



Architecting Organisational Intent: A Nudge-Based Behavioural Engineering Framework for Resilient Safety Management

Abubakar Alhaji Janga^{1,*}, Alibe Mustapha Ali¹, Umar Shamsuddeen Abubakar¹

¹ Department of Mechanical Engineering, School of Engineering Technology, The Federal Polytechnic Damaturu, Damaturu, Nigeria.

* Corresponding author: jangatagwai@gmail.com

Abstract

Purpose: This study addresses the Cognitive Friction Paradox, where the accumulation of administrative controls triggers non-compliance by exceeding operator cognitive resources. The proposed Behavioural Engineering Framework integrates dual-process psychology and resilience engineering to make the safest action the path of least resistance. **Theory and Design:** High-tempo environments favour automatic System 1 processing over deliberate reasoning. By integrating choice architecture and Safety-II, the framework treats System 1 as a stable structural feature to be designed around rather than a defect to be eliminated. **The Framework:** Four operational pillars are proposed: Defaults and Constraints, Friction Engineering, Salience and Priming, and Social Proof. **Methodology:** Intervention design follows a four-stage Diagnose–Architect–Pilot–Calibrate protocol mapped to the resilience potentials of monitoring, anticipating, responding, and learning. **Practical Implications:** Positioned as an administrative-layer enhancement within the Hierarchy of Controls, the framework shifts safety measurement from lagging metrics to leading behavioral indicators while establishing ethical safeguards to prevent nudge-washing. **Conclusion:** High-reliability performance is best achieved by engineering systems that work in harmony with natural human cognition, providing a scalable method to close the gap between work-as-imagined and work-as-done.

Keywords: Choice Architecture, Nudge Theory, Safety Management Systems, Dual-Process Theory, Work-As-Done, Resilience Engineering.

Article Information:

Received: 25 August 2026
Revised: 26 September 2026
Accepted: 3 October 2026
Published: 2026

Vol. 16, No. 1, 2026

© MRN Publishing

Introduction

Industrial safety management has evolved through successive, additive layers of control. Early twentieth-century practice relied almost exclusively on physical and engineering barriers—guarding, containment, and interlocks designed to remove hazards from the point of human contact. As operations grew more complex, safety management expanded into administrative territory: written procedures, permits, checklists, and compliance audits intended to standardize behavior across large, distributed workforces. This administrative layer is now embedded in formal management-system standards such as ISO 45001, which frame occupational health and safety as a matter of documented process and legal compliance (Liu *et al.*, 2023; Dolgunyürek & Ada, 2025).

However, this evolution has been cumulative rather than substitutive. Each new administrative control is typically layered atop existing ones rather than replacing them, producing a phenomenon of *administrative stacking* where an operator must simultaneously satisfy physical safeguards, written procedures, and behavioral expectations in the same moment of work. This pattern is often driven by bureaucratic entrepreneurship, wherein safety departments and regulatory bodies seek to demonstrate organizational value and professional competence through the proliferation of documentation. In many industrial contexts, administrative controls have become a defensive mechanism; because they are significantly cheaper and faster to implement than engineering redesigns, they are deployed as the default response to every identified near-miss.

This creates a cycle where procedures are increasingly written to satisfy the requirements of an auditor or a legal framework rather than to support the functional needs of the operator at the point of work.

This layering of controls creates what this paper terms the *Cognitive Friction Paradox*: while each individual rule or checklist is introduced to reduce risk, their cumulative cognitive burden increases the likelihood that operators will bypass, simplify, or reinterpret those same controls (Figure 1). This results in *safety clutter*—the accumulation of excessive paperwork, duplicative licensing, and administrative tasks that hinder frontline action rather than supporting it (Breda *et al.*, 2026; Raheemy *et al.*, 2026). When formal systems add friction without adding perceived operational value, operators naturally construct informal workarounds as a means of maintaining productivity (Ashraf *et al.*, 2026; Wong *et al.*, 2022).

The urgency of addressing this paradox is amplified by the transition to Industry 4.0. While digitalization, ubiquitous sensing, and the Internet of Things (IoT) were heralded as tools to eliminate human error, they have frequently introduced a secondary layer of cognitive demand. Real-time data streams and predictive alerts often translate into more *digital paperwork* and monitoring tasks, adding to the operator's mental load rather than simplifying the decision environment. Instead of subtracting human variance, high-tech environments often exacerbate cognitive friction by demanding that workers manage both traditional administrative protocols and complex digital interfaces simultaneously. The result is a persistent and largely institutionalized gap between *Work-as-Imagined* (WAI)






Figure 1. The Cognitive Friction Paradox: Divergent Pathways from Prescribed Work to Work-as-Done

—the procedure as written—and *Work-as-Done* (WAD)—the procedure as actually performed under real operational constraints. Rather than treating this gap as a compliance failure to be punished, this paper treats it as a *design signal*: a measurable indication of where the safety system's demands exceed the operator's available cognitive resources.

The gap persists because conventional Safety Management Systems (SMS) implicitly model the operator as a *System 2* reasoner who deliberately weighs risks, consults manuals, and selects the compliant action. However, dual-process accounts of cognition suggest this model breaks down precisely in the high-tempo conditions where safety matters most. Under time pressure, fatigue, and multitasking, human decision-making shifts toward *System 1*—fast, automatic, low-effort processing dominated by heuristics rather than explicit reasoning (Hao *et al.*, 2025; Zhylin *et al.*, 2026). Empirical work confirms that high workloads narrow attention and degrade situational awareness in complex control tasks (Causse *et al.*, 2024; Liu & Ham, 2025). SMS frameworks that fail to account for this automatic mode are, in effect, designing for an operator who does not exist under real operational conditions.

The central thesis of this paper is that safety performance can be improved not by adding further deliberative burden to the operator, but by re-engineering the choice environment itself so that the safest available action is also the path of least cognitive and physical resistance. This reframes choice architecture—the deliberate design of the context in which decisions are made—as a legitimate engineering discipline within SMS, positioned alongside, rather than as a replacement for, existing physical and procedural controls. To advance this thesis, the paper pursues three objectives: (1) to synthesize dual-process psychology, nudge theory, and resilience engineering into a coherent theoretical foundation for industrial behavioral design; (2) to formalize a four-pillar Behavioral Engineering Framework with associated testable propositions; and (3) to translate that framework into a practical diagnostic, design, and measurement protocol that safety engineers can apply and researchers can empirically validate.

The remainder of the paper is organized as follows. Section 2 develops the theoretical foundations, covering dual-process psychology, procedural drift, and the alignment between choice architecture and Safety-II. Section 3 presents the core

Behavioral Engineering Framework, its system topology, and its mapping onto the Hierarchy of Controls. Section 4 details a diagnostic and validation protocol, including a behavioral safety audit method and leading behavioral indicators. Section 5 establishes boundary conditions and ethical safeguards to prevent the framework from being used to substitute for primary engineering controls. Section 6 discusses theoretical and practical implications, and Section 7 concludes.

Theoretical and Literature Foundations

Dual-Process Psychology and Cognitive Load in High-Hazard Workplaces

Traditional safety management systems (SMS) are largely designed under the assumption of deliberate, analytical reasoning. However, high-hazard operations rely heavily on fast, automatic processing to conserve cognitive resources for novel or high-stakes moments (Table 1). The central failure of conventional SMS design is its demand for constant analytical attention on tasks that, through repetition and habituation, inevitably shift toward heuristic-based shortcuts (Reale *et al.*, 2023; Zhylin *et al.*, 2026).

This cognitive shift is better understood through the lens of Cognitive Load Theory (CLT). CLT posits that the human working memory has a limited capacity, and when this capacity is exceeded, learning and error-free performance collapse. In industrial settings, we must distinguish between intrinsic load—the inherent difficulty of the mechanical or operational task itself—and extraneous load, which is the mental effort required by the way information is presented or the environment is structured (Holden *et al.*, 2015).

A significant portion of safety clutter translates directly into extraneous cognitive load. When an operator is required to navigate 20-page permit or cross-reference multiple redundant checklists while performing a high-risk lockout, the mental energy consumed by the paperwork subtracts from the mental energy available for hazard recognition. If the extraneous load of the safety system is too high, cognitive overflow occurs. In this state, the brain instinctively sheds the most taxing load—analytical System 2 compliance—and defaults to System 1 heuristics to survive the shift. Thus, excessive safety documenta-

Table 1. Cognitive Characteristics and Behavioral Drivers of System 1 vs. System 2 Processing in Industrial Work Environments

Characteristic	System 1 Processing	System 2 Processing
Speed and effort	Fast, automatic, low perceived effort	Slow, deliberate, high perceived effort
Primary trigger conditions	Time pressure, fatigue, multitasking, cognitive saturation	Novel, unfamiliar, or explicitly flagged high-stakes decisions
Governing mechanism	Heuristics, habit, associative memory	Rule-based reasoning, explicit procedure consultation
Typical failure mode	Inattentional blindness, sign-blindness, habituation	Decision fatigue, analysis delay under time pressure
Industrial example	Reaching for a familiar tool without re-checking a changed permit status	Conducting a formal pre-task hazard review or JSA
Design implication	Requires default-safe states and salient, low-effort cues (Pillars A & C)	Benefits from structured information and explicit feedback (Pillar D)

tion does not merely add time; it actively induces the very automatic processing that leads to missed hazard cues (Causse *et al.*, 2024; Su *et al.*, 2025).

This cognitive shift is driven by structural environmental triggers rather than individual negligence. Time pressure, multitasking, and fatigue reliably degrade situational awareness and slow response times (Gondale *et al.*, 2025; Liu & Ham, 2025). Under such cognitive saturation, inattentional blindness causes operators to miss even salient hazards (Matias *et al.*, 2022; Yu & Jiang, 2026). Effective safety engineering must therefore accommodate these natural cognitive constraints through the choice environment rather than relying on training to enforce analytical compliance.

The Mechanisms of Procedural Drift and the Law of Least Effort

Procedural drift is the directional, cumulative divergence between documented procedures and everyday practice (Hollnagel, 2014). This divergence is often described in organizational literature as a *slow erosion of margins* or *normalized deviance*. As Sidney Dekker (2011) argues, drift is not a sudden collapse but a steady, incremental movement toward the edge of safety, driven by the need to balance safety with production pressures. Each unpunished deviation—such as skipping a minor verification step to save five minutes—becomes the new standard for the next task.

However, viewing drift merely as an organizational failure overlooks a fundamental biological imperative: the *Law of Least Effort*. From an evolutionary perspective, the human brain is an energy-expensive organ that has evolved to achieve goals using the minimum possible expenditure of cognitive and physical resources. In the workplace, if a compliant path requires high effort (e.g., walking 200 meters to retrieve a specific tool) and an adapted path requires low effort (e.g., using a nearby but non-standard tool) with no immediate negative consequence, the brain will naturally select the adapted path.

This is not a character flaw or a lack of safety culture; it is a biological optimization. When operators create workarounds, they are often solving a friction problem that the designers of the safety system failed to see. Because drift is a structural adaptation to operational friction rather than random error, it requires behavioral engineering to nudge the path of least resistance back toward safety-critical intent (Tedone & Lanz, 2024). Traditional, snapshot-based audits are structurally poor at detecting this drift because they sample rare moments of ceremonial compliance where workers temporarily re-engage System 2 processing for the benefit of the auditor (Dekker, 2011; Kim *et al.*, 2024).

Synthesis: Behavioral Economics and High-Reliability Organizing (HRO)

The integration of behavioral economics and High-Reliability Organizing (HRO) provides the scholarly bridge necessary for a modern, resilient SMS. HRO theory focuses on the organizational traits required to maintain safety in complex,

high-risk environments, such as a *preoccupation with failure* and a *reluctance to simplify* (Bertoni *et al.*, 2021). While HRO provides the high-level philosophical goals, it often lacks the granular mechanics for frontline implementation.

Behavioral economics, specifically Nudge Theory, provides these mechanics. By understanding the *predictable irrationality* of human operators, safety engineers can design a choice architecture that aligns System 1 habits with HRO goals. For example, the HRO principle of *sensitivity to operations* can be operationalized through Contextual Salience (Pillar C), ensuring that the most critical hazard information is brought to the operator's attention at exactly the moment they are most likely to overlook it.

This synthesis moves safety from a discipline-based model to a design-based model. In a design-based model, resilience is not just an organizational culture but a tangible feature of the physical and digital workspace. By architecting the choice environment, we ensure that even when an operator is fatigued, under pressure, or cognitively saturated, the system's defaults guide them toward high-reliability outcomes. This prevents the downward drift into failure by making safe behavior the most efficient and natural choice for the human brain (Hollnagel, 2017; Trancoso *et al.*, 2024).

The Behavioral Engineering Framework

Framework Architecture and System Topology

The proposed framework (Figure 2) models safety behavior as a systemic pipeline running from environmental inputs, through cognitive processing modes, to observable actions. At the point of work, environmental triggers (ranging from extreme time pressure and task complexity to administrative load and fatigue) act as a *cognitive switch*. Under high-load conditions, an operator's decision-making is naturally pushed into System 1 processing.

Rather than attempting to retrain the operator out of this automatic mode, the framework treats the processing mode as a design parameter. Choice architecture is engineered to ensure that when System 1 is active, the safe action is the path of least cognitive and physical resistance. The topology is governed by a continuous feedback loop: Behavioral Trace data (e.g., override frequency, checklist quality) is measured via leading indicators and fed back into the diagnostic protocol to refine the friction profile of the task.

The Four Pillars of Industrial Nudging

The framework utilizes four core pillars to align the choice environment with human intent. By architecting these pillars, safety engineers can ensure that the WAD matches WAI even under operational strain.

Pillar A: Defaults and Constraints

- **The Cognitive Mechanism:** At the neuroergonomic level, default options leverage status quo bias and decision fatigue




Figure 2. Framework Architecture and System Topology

to preserve limited neurocognitive resources (Bechara, 2008). High mental workload and sustained cognitive exertion deplete executive resources in the dorsolateral prefrontal cortex (dlPFC) and anterior cingulate cortex (ACC), which increases the subjective cost of active deliberation (Shenhav *et al.*, 2017). When cognitive fatigue sets in, the neurobiological decision network shifts away from resource-intensive analytical processing toward metabolic conservation, defaulting to the path of least cognitive effort (Bechara, 2008; Shenhav *et al.*, 2017). By configuring the safe state as the structural default or applying physical interlocks as hard constraints, the system bypasses the requirement for active prefrontal override. Compliance is thus achieved naturally even when operators experience high extraneous load, severe time pressure, or attentional saturation.

Industrial Scenario (Lock-out Tag-out):

- ✓ **Before:** An operator must manually cross-reference a 10-page paper isolation list with physical energy points. Under production pressure and cognitive fatigue, they rely on memory and fail to isolate a secondary hydraulic line.
- ✓ **After:** A digital LOTO interface is implemented where the Safe-to-Work permit is pre-populated with required energy points based on the specific asset ID. The system defaults to a locked status, and the machine-startup sequence is digitally constrained until sensors confirm all isolations are engaged. The operator achieves compliance by following the pre-selected path provided by the interface.

Evidence Review: Neuroergonomic and human factors research confirms that optimizing decision cost-benefit tradeoffs reduces input errors and procedural steps under high mental workload (Bechara, 2008; Shenhav *et al.*, 2017). Pre-populating safety workflow parameters drastically minimizes cognitive load at the point of action. When safe choices are embedded as pre-

selected defaults or system constraints, initial compliance significantly outperforms interventions that rely on post-hoc warning reminders or continuous analytical vigilance.

Pillar B: Friction Engineering

- **The Cognitive Mechanism:** Friction engineering is grounded in the Cost of Action theory. By asymmetrically increasing the physical or cognitive effort required for an unsafe path, we force a momentary System 2 interruption. neurologically, adding a physical delay or a minor obstacle creates a choice point that breaks the automaticity of habit-driven shortcuts. While Pillar A removes effort from the safe path, Pillar B adds thoughtful friction to risky ones. This shifts the operator's internal cost-benefit calculation; if the shortcut becomes more effortful than the safe procedure, the brain's optimization drive naturally reverts to the safer route (Ashraf *et al.*, 2026; Thaler & Sunstein, 2008).
- **Industrial Scenario (Bypassing Guards):**
 - ✓ **Before:** A machine guard is missing a bolt, allowing it to be swung open easily to clear a jam without stopping the machine. The shortcut saves two minutes, making it the low-effort choice.
 - ✓ **After:** The guard is redesigned such that clearing a jam requires a specialized tool located at a workstation ten meters away. While the unsafe action is still technically possible, the physical effort required to bypass the safety protocol now exceeds the effort of following the standard Stop-Clear-Start procedure.

Evidence Review: Costa *et al.* (2025) demonstrated that in suspended-load operations, adding physical barriers (friction) on unsafe walking paths produced more consistent safe-distance behavior than supervisory verbal instructions. Similarly, Ashraf *et al.* (2026) found that adding digital skip-order friction—requiring a justification for out-of-sequence steps—reduced procedural drift in process-industry workflows.

Pillar C: Salience and Priming

- **The Cognitive Mechanism:** Static warning signs suffer from Sign Blindness caused by sensory habituation. Neurologically, the brain's reticular activating system (RAS) filters out repetitive environmental stimuli that it deems non-threatening. Contextual salience addresses this by delivering novel, action-triggered stimuli (visual, haptic, or spatial) at the exact moment of decision. By priming the brain with a salient cue just before a risky action, we increase the probability that the safe option will be registered by System 1. This orienting response forces the brain to re-evaluate the immediate environment, breaking the autopilot mode that often precedes struck-by incidents (Hussain *et al.*, 2024; Kim *et al.*, 2021).
- **Industrial Scenario (Confined Space Entry):**
 - ✓ **Before:** A static Danger: Confined Space sign is bolted to a tank hatch. After months of exposure, operators walk past it daily without cognitive registration.
 - ✓ **After:** A smart beacon is installed. When an operator's wearable sensor approaches the hatch without an active permit, the hatch handle vibrates (haptic priming) and a floor projection flashes a safe-zone boundary. The novelty and timing of the cue force the brain to re-register the hazard.

Evidence Review: Research in construction safety has shown that action-triggered cues—such as floor projections marking safe zones—raise hazard-recognition skill by 29% compared to static signage (Albert *et al.*, 2017). Hussain *et al.* (2024) used eye-tracking in virtual reality to confirm that the *timing* of a salient prompt is more critical than its content; cues delivered even seconds after a decision point are often ignored.

Pillar D: Social Proof and Commitment

- **The Cognitive Mechanism:** Humans are an obligately social species with highly evolved neurological pathways for social mirroring and reputation management. Pillar D harnesses the Bandwagon Effect and the desire for social consistency. When operators perceive that their peers are consistently following safe protocols, System 1 mimics that behavior to maintain social belonging. Furthermore, making a visible commitment to a safety goal activates the Consistency Principle, where the brain seeks to avoid the cognitive dissonance of behaving in a way that contradicts one's stated social identity (Cialdini, 2021; Vo & Kim Quy, 2026).
- **Industrial Scenario (PPE Compliance):**
 - ✓ **Before:** PPE compliance is monitored through periodic, punitive audits. Workers wear the gear when a supervisor is present but remove it when alone, viewing it as a private burden.
 - ✓ **After:** A non-punitive, team-based dashboard is installed in the breakroom. It displays the team's collective Safety Streak and highlights peer-to-peer recognition comments. Safety is transformed into a visible social norm, and the Social Proof of others' compliance primes individuals to follow suit without direct oversight.

Evidence Review: Sparer *et al.* (2015) documented that recognition-based feedback programs in construction increased voluntary hazard reporting. Recent telematics research also indicates that non-punitive peer-comparison dashboards reduce risky driving behaviors more effectively than individual coaching, provided the culture supports psychological safety (Ghamari *et al.*, 2022; Veerasimman, 2026).

Formal Research Propositions (P_1 – P_4)

The framework advances four testable propositions for empirical validation, shifting the focus from lagging injuries to proximal behavioral outcomes.

1. P_1 (Defaults): *Safe defaults reduce unsafe first-attempt actions more effectively than informational prompts, especially under high workloads.*
 - **Theoretical Rationale:** Based on the work of Thaler and Sunstein (2008), this proposition assumes that cognitive inertia is the dominant driver of behavior during fatigue. Because System 2 reasoning is metabolically expensive, operators will follow the Status Quo unless a highly salient reason to override it is presented.
2. P_2 (Friction): *Introducing a friction differential between safe and unsafe paths reduces procedural drift and near-miss frequency.*
 - **Theoretical Rationale:** Grounded in Kahneman's (2011) Law of Least Effort, this hypothesis posits that operators skip steps not due to negligence, but because the brain is an optimizer. By making the unsafe path costly in terms of time or effort, we align the brain's optimization drive with safety requirements.
3. P_3 (Salience): *Action-triggered cues produce significantly smaller habituation decay than static signage over repeated exposures.*
 - **Theoretical Rationale:** This proposition is grounded in the psychological principle of Attentional Capture (Theeuwes, 1992). Habituation occurs when a stimulus is static. By making the cue contextual and action-triggered, we ensure the brain's novelty-detection pathways remain active over time.
4. P_4 (Social Proof): *Non-punitive peer-comparison data increases near-miss reporting velocity, provided the workgroup possesses high psychological safety.*
 - **Theoretical Rationale:** Based on Cialdini's (2021) Social Proof mechanism and Bandura's (1977) Social Learning Theory, this proposition assumes that behavior is shaped by observed peer norms. If reporting is seen as a group-wide value, reporting velocity will increase as a function of social reinforcement.

Mapping to the Hierarchy of Controls and ISO 45001

Choice architecture is positioned strictly within the administrative-controls layer of the Hierarchy of Controls (Figure 3). It serves as a reinforcement for, rather than a substitute for, primary engineering controls (Clarke *et al.*, 2026; Dolgunyurek & Ada, 2025). Within ISO 45001, these interventions support *Operational Planning and Control* by providing practical mechanisms to translate documented procedures into consistent field behavior. This positioning resists downward drift, where organizations might otherwise adopt low-cost behavioral nudges to accommodate unresolved physical hazards (Liberati *et al.*, 2018; Stackhouse & Stewart, 2017). Framing choice architecture as an engineering discipline ensures management remains accountable for the work environment, resisting the tendency to blame the worker for deviations structurally encouraged by a poorly designed interface.

Diagnostic Protocol, Measurement, and Validation

The Behavioral Safety Audit and Friction Differential

Before designing a nudge, the framework requires a structured Behavioral Safety Audit to map actual WAD and locate systemic friction points (de Freitas *et al.*, 2026; Woltjer *et al.*, 2015). A central function of the audit is differential diagnosis: distinguishing whether a deviation stems from skill deficits, equipment failure, or genuine cognitive friction. Only the latter—evidenced by consistent workarounds performed by competent, adequately equipped operators—is the specific target of behavioral engineering (Ashraf *et al.*, 2026; Johnson *et al.*, 2025).






Figure 3. Positioning Choice Architecture within the Hierarchy of Controls

Operational Intervention Cycle (inner ring) mapped to Resilience Potentials (outer ring)




Figure 4. Iterative Intervention Cycle and Resilience-Potential Mapping

To operationalize this, managers utilize a Friction Scorecard to calculate the Friction Differential (F_{Δ}) between the compliant path (P_c) and the shortcut/workaround (P_s). The scorecard evaluates three variables: Time (T), Physical Effort (P), and Cognitive Steps (C). The differential is expressed as:

$$F_{\Delta} = (T_c + P_c + C_c) - (T_s + P_s + C_s)$$

A high positive F_{Δ} value indicates that the safe path is significantly more expensive to the operator than the workaround, creating a structural invitation for procedural drift. The objective of Pillars A and B is to reduce F_{Δ} to zero or a negative value ($F_{\Delta} \leq 0$), ensuring that the intended safe adaptation is also the path of least resistance.

Nudge Design and Calibration Protocol

Intervention design follows a four-stage iterative loop (Figure 4) that maps directly to the four resilience potentials (Bertoni *et al.*, 2021; Trancoso *et al.*, 2024). While Diagnose and Architect focus on planning, the Respond and Learn phases are critical for addressing the common organizational Failure to Learn.

In the Respond phase (Pilot), the organization tests the system’s reaction to a bounded change. This is not merely an observation of worker compliance, but a test of sociotechnical compatibility. If a nudge (e.g., a new digital default) is bypassed within the first week, it indicates that the Architect stage failed to account for a specific operational constraint.

The Learn phase (Calibrate) addresses the tendency for safety interventions to become ceremonial or static. Most organizations suffer from intervention decay because they treat safety fixes as one-time solutions. Iterative calibration uses leading indicators to adjust nudge volume—such as the salience of a visual cue or the magnitude of a friction point—before the intervention becomes background noise. This ensures that the SMS functions as an evolving feedback loop rather than a static log (Ankamah-Lomotey, 2025).

Leading Indicators and Statistical Challenges

To support the transition from lagging metrics to proactive design signals, the framework utilizes four leading behavioral

indicators (Table 2). These metrics are intended to provide real-time feedback on the health of the choice architecture and the level of alignment between WAI and WAD.

However, moving from simple Incident Counts to Near-miss Reporting Velocity presents distinct statistical challenges. Traditional metrics like Total Recordable Incident Rate (TRIR) are zero-inflated and lagging; a low score can signify either a safe site or a suppressed reporting culture. In contrast, Near-miss Reporting Velocity (V_{nm}) measures the rate and speed at which hazards are surfaced relative to operational volume. The primary statistical challenge here is the latency-bias: the time delta between an event occurring and its entry into the SMS. A resilient system aims for high V_{nm} with low latency. High-reliability safety requires analyzing these patterns through a system-dynamics lens, recognizing that a sudden drop in reporting velocity is often a more accurate predictor of an impending failure than a stable incident rate (Bakshi & Peura, 2022; Thoroman *et al.*, 2018).

Recommended Validation Methods by Pillar

To empirically validate the propositions in Section 3.3, researchers must select methodologies capable of distinguishing sustained change from transient novelty effects (Table 3). Because the four pillars target distinct mechanisms [ranging from structural defaults to social learning (Bandura, 1977)], validation must be contextually tailored to the intervention type.

Boundary Conditions, Ethics, and Risk Mitigation

The Primary Safety Layer Rule and the Nudge-Washing Debate

The framework’s non-negotiable boundary condition is the Primary Safety Layer Rule: choice architecture must never substitute for elimination, substitution, or engineering controls, only reinforce them. This prevents *nudge-washing*—a moral hazard where low-cost behavioral interventions are deployed as a façade for under-investment in fundamental system redesign (Arya & Hindolia, 2025).

Table 2. Leading Behavioral Indicators for Nudge-Based Safety Management

Leading Indicator	Metric: What it Measures	Predictive Rationale	Key Supporting Literature
Voluntary Safeguard Adoption	PPE usage rates, permit-to-work adherence, and hazard self-correction without supervisory oversight.	High adoption indicates that the safe path is successfully engineered as the path of least resistance (Pillars A & B).	Veerasimman (2026); Ammad <i>et al.</i> (2025)
Near-miss Reporting Velocity (V_{nm})	The volume of surfaced hazards and the time delta between an event and its entry into the safety system.	Faster, more frequent reporting indicates high hazard-awareness and a non-punitive learning culture (Pillar D).	Bakshi & Peura (2022); Thoroman <i>et al.</i> (2018)
Safety-voice Behavior	Frequency of preventive speaking up and peer-to-peer stop-work interventions.	Reflects organizational psychological safety; identifies procedural blind spots before they lead to drift.	Mathisen & Tjora (2023); Tedone <i>et al.</i> (2026)
Friction-Reduction Score	The measured decrease in the Friction Differential (F_{Δ}) after a nudge implementation.	Directly validates that the architectural intervention has made the safe action cognitively cheaper for the operator.	Ashraf <i>et al.</i> (2026); Johnson <i>et al.</i> (2025)

Table 3. Methods and Rationale for target Pillar Validation

Methodology	Target Pillars	Rationale	Key Reference
Quasi-experiments	A, B	Compares nudged vs. non-nudged sites to test organizational-level uptake.	Clapier <i>et al.</i> (2026)
Interrupted Time-Series	C, D	Tracks compliance persistence or decay across extended quarters.	Bilgin <i>et al.</i> (2023)
Immersive VR Simulation	C	Isolates habituation mechanisms and measures attentiveness in a safe environment.	Kim <i>et al.</i> (2024)
Bayesian Belief Networks	General	Estimates intervention effects retrospectively using existing operational data.	Cooper <i>et al.</i> (2023)

The economic incentive for nudge-washing is substantial. Behavioral nudges primarily incur operational expenditures (OPEX)—such as software configuration, dynamic signage, or procedural redesign—whereas primary physical engineering controls (e.g., automated interlocks, machine enclosures, high-capacity ventilation) require significant capital expenditure (CAPEX). When safety budgets face fiscal constraints, risk management risks degrading into cost-avoidance, substituting CAPEX isolation barriers with cheap OPEX choice architecture.

To prevent this capital substitution, organizations must implement a Financial and Safety Control Audit (FSCA) alongside the standard Hierarchy of Controls (Figure 3). The FSCA requires safety steering committees and internal auditors to evaluate proposed nudge interventions against three specific financial and operational indicators:

- **The CAPEX-Avoidance Ratio (R_{ca}):** Auditors must calculate the ratio of projected OPEX for a behavioral nudge to the estimated CAPEX required for physical elimination or engineering control. An R_{ca} significantly below baseline thresholds triggers a mandatory engineering review to confirm whether physical isolation is technically infeasible rather than merely financially deferred.
- **Control Layer Expenditure Matching:** Any administrative choice-architecture project budgeted under OPEX must include a formal engineering justification sign-off. The sign-off certifies that higher-order controls (Elimination, Substitution, Engineering Interlocks) were evaluated using a formal Risk Reduction Cost-Benefit Analysis (CBA) and deemed physically unviable before behavioral support was approved.
- **Decoupling Audit Trails:** Financial audit cycles must cross-reference CAPEX deferral logs with field-level friction scorecards. If an engineering project to eliminate a physical hazard is delayed or rejected on budget grounds, management is explicitly prohibited from deploying a nudge as a permanent substitute. Any interim nudge deployed under such conditions must be flagged as a temporary risk accommodation with a hard expiration date tied to the next CAPEX budget cycle.

Management must resist treating human cognition as a low-cost filter for technical system deficiencies. If a nudge is used to mitigate a risk that could have been eliminated through higher-order controls, it constitutes a managerial workaround—a top-down version of the very procedural drift the framework seeks to solve.

Auditors must independently verify these financial and operational entries against the boundary matrix (Table 4) to prevent façade adoption where a polished behavioral program masks a systemic failure to address structural hazards (Arya & Hindolia, 2025; von Thiele Schwarz *et al.*, 2021).

Ethical Choice Architecture: Liability and the Architecture of Failure

Because nudges can operate below the threshold of conscious deliberation, their legitimacy depends on transparency and systemic accountability. Organizations must explicitly communicate the rationale for engineered defaults or friction points; opacity risks the intervention being perceived as manipulation, undermining trust in the SMS (Grelle *et al.*, 2024; Rantala *et al.*, 2023).

A critical frontier for management research involves the legal implications of nudge failure. Because nudges are explicitly non-compulsory, a liability loophole may emerge: organizations might attempt to shift blame to the operator by claiming the safe option was available and the individual simply chose the unsafe path. However, in high-hazard operations, the distinction between voluntary choice and environmental coercion is thin. If

a nudge fails due to poor architectural design [such as inadequate contextual salience or miscalibrated friction] legal and ethical liability should remain with the organization. Management acts as the system architect; thus, designing the failure through a poorly calibrated choice environment constitutes a breach of the duty of care. Responsibility for the digital and physical environment remains with the firm, and interventions must be tailored to local cultural symbolism to ensure ethical fit (Alshamy, 2025; Wu & Paluck, 2021).

Table 4. Boundary Matrix: Distinguishing Primary Controls from Auxiliary Choice Architecture

Control Layer	Representative Examples	Role of Choice Architecture
Elimination /Substitution	Removing hazardous materials entirely	Out of scope; resolve before considering nudges.
Engineering Controls	Guarding, containment, interlocks	None as replacement; nudges may only prompt correct use.
Adminis-trative Controls	Procedures, ISO 45001 processes	Primary operating zone; pillars enhance reliability.
PPE/ Behavioral Layer	Protective equipment, compliance	Supported by Pillars A–D; never used to justify omitting higher controls.

Managing Nudge Decay and Long-Term Habituation

Nudge decay is an expected lifecycle event where cues eventually become background noise due to habituation (Kim *et al.*, 2024). The framework addresses this via the Calibrate stage, which involves periodic refreshing of visual cues and scheduled reviews of default settings. The four pillars are not equally susceptible to decay: Pillar A (defaults) is highly resistant as it requires no sustained attention, whereas Pillars C (salience) and D (social proof) are decay-prone and require frequent rotation to remain effective (Causse *et al.*, 2024; Lee & Kim, 2022).

Discussion and Future Research Agenda

Theoretical Implications

The Behavioral Engineering Framework relocates the unit of design from the individual worker’s judgment to the decision environment itself. By treating System 1 processing as a stable structural feature of high-tempo work rather than a defect to be trained away, the framework asks the engineering system, rather than the worker, to adapt. This reconciles the historical tension between Safety-I and Safety-II: it uses the measurable, auditable rigor of choice architecture to operationalize and support the human safety factor—the adaptive capacity of workers to produce safe outcomes under pressure (Ham, 2021; Hollnagel, 2017). Furthermore, by linking these interventions directly to higher worker engagement, the framework suggests that designing for human intent carries motivational benefits beyond simple compliance.

Managerial Implications: The ROI of Behavioral Engineering

For facility managers, the transition from Lagging Indicator Management to behavioral engineering represents a fundamental shift in resource allocation. Traditional safety management often incurs massive hidden costs through Lagging Indicator Management—the reactive cycle of post-incident investigations, legal fees, stop-work orders, and insurance premiums. While administrative stacking (adding more rules) is often perceived as a low-cost fix, it typically yields a negative Return on Investment (ROI) because it induces ceremonial compliance, where docu-

mentation exists but behaviors remain unchanged, leaving the organization vulnerable to catastrophic failure.

In contrast, behavioral engineering requires an upfront investment in diagnostic audits and digital/physical architectural changes (CAPEX), but it yields a measurable Behavioral Trace that functions as an early warning system. By reducing the Friction Differential, managers can lower the cognitive tax on their workforce, which correlates with higher productivity and lower burnout. The ROI of this framework is found in the prevention of the drift into failure; by making the safe path the path of least resistance, organizations reduce the frequency of unpunished workarounds that eventually aggregate into major accidents.

Limitations: The Boundaries of Nudging

Despite its utility, choice architecture is not a panacea. A critical limitation of this framework is its reliance on System 1 processing. Nudges are designed for tasks that are frequent, repetitive, and stable—environments where autopilot behavior is the primary risk. In contrast, during novel, high-complexity emergencies, operators must transition back to System 2 (deliberate, analytical reasoning). In such scenarios, a nudge designed for routine operations may actually become an attentional trap or induce confirmation bias, leading an operator to default to a standard response when a creative, System 2 intervention is required.

Furthermore, choice architecture cannot fix a broken technical system. If a hazard is so severe that it requires elimination or engineering interlocks, attempting to nudge around it is a violation of the Primary Safety Layer Rule. Behavioral engineering is an auxiliary support layer; it assumes that the underlying equipment is fit for purpose. It is ineffective in organizations where the safety culture is overtly punitive, as the Social Proof and Safety Voice pillars (Pillar D) will be suppressed by fear.

Global and Cultural Context: Re-Architecting for Diversity

The effectiveness of the four pillars is not culturally universal. Management research suggests that behavioral interventions are culturally contingent (Alshamy, 2025). For example, in High Power Distance (HPD) cultures, Pillar D (Social Proof) may function differently; workers may be more responsive to Vertical Social Proof (feedback from authority figures) than to the Horizontal Social Proof (peer telemetry) favored in Western contexts.

Similarly, the visual and haptic cues used in Pillar C (Salience) must be re-architected to fit local symbolic meanings. As noted by Wu & Paluck (2021) in their study of Chinese manufacturing, using culturally specific symbols (like Golden Coin decals) significantly outperformed standard Western hazard icons. In environments with diverse, multi-national workforces, safety engineers must ensure that salience cues and defaults do not inadvertently create cultural friction, where a nudge that feels supportive to one group feels coercive or confusing to another.

Practical Roadmap for Safety Engineers and Facility Managers

Integrating the framework involves layering the iterative loop (Figure 5) onto existing workflows rather than replacing current systems. Digital HSE platforms offer a natural entry point for Social Proof (Pillar D) through non-interruptive dashboards (Grouse *et al.*, 2023). For daily safety walks, the roadmap recommends replacing generic checklists with prompts that ask: “Where did a worker choose the harder path today, and why?” (Stauffer & Chastain-Knight, 2021).

Future Empirical Research Agenda

Five domains warrant priority attention to mature the framework:

1. *Human–AI Collaboration*: Determining how choice architecture functions when operators play supervisory roles in automated plants (Gondale *et al.*, 2025).
2. *Smart Wearables*: Testing real-time haptic or visual nudges (Pillar C) responsive to an individual’s location and task state.
3. *Long-horizon Durability*: Conducting multi-year research to address the nudge-decay concerns raised by quarter-scale studies (Huang & Boudjema, 2021).
4. *Cultural Contingency*: Systematic study of how power distance and cultural symbolism affect pillar uptake (Alshamy, 2025; Wu & Paluck, 2021).
5. *Measurement Standardization*: Validating a shared, cross-industry indicator battery to replace inconsistently defined metrics (Sadeghi *et al.*, 2025; Mathisen & Tjora, 2023).

Conclusion

The persistent gap between Work-as-Imagined and Work-as-Done is a design problem, not a discipline problem. The Behavioural Engineering Framework developed here responds by making the safe action the path of least resistance through its four pillars: Defaults and Constraints, Friction Engineering, Salience and Priming, and Social Proof. Positioned beneath primary engineering controls and bounded by the Primary Safety Layer Rule, the framework provides a missing administrative layer that interfaces directly with human cognition. High-reliability performance is achieved not by forcing operators to fit rigid systems, but by engineering systems that work in harmony with the way humans actually think and act.

Looking ahead, the most transformative application of this framework lies in its integration with Artificial Intelligence (AI) and wearable technology. As industrial environments transition toward full digitalization, the Diagnose–Architect–Pilot–Calibrate loop can be automated to provide real-time, personalized behavioural support. AI-driven systems can monitor physiological markers (such as eye-tracking for distraction or heart-rate variability for fatigue) to predict when an operator has shifted into a vulnerable System 1 state. At the micro-moment of risk, wearable devices can then deliver Dynamic Priming via haptic vibrations or augmented reality overlays, automating Pillar C (Contextual Salience) with a precision previously impossible. This vision reframes the safety system as a cognitive co-pilot that proactively shields the operator from habituation and sign-blindness. The ultimate goal is to move beyond static manuals toward a responsive sociotechnical landscape where safety is an emergent, engineered property of the work itself. By embracing these technologies, organizations can finally ensure that safe intent is always translated into safe action.

References

- Albert, A., Hallowell, M. R., Skaggs, M., & Kleiner, B. (2017). Empirical measurement and improvement of hazard recognition skill. *Safety Science*, 93, 1–8. <https://doi.org/10.1016/j.ssci.2016.11.007>
- Alshamy, Y. (2025). Mind the cultural gap: Cultural contingency of behavioral interventions. *Mind & Society*, 24(1), 339. <https://doi.org/10.1007/s11299-025-00339-8>
- Ammad, S., Mostafa, S., & Stewart, R. A. (2025). Development of a safety oriented framework for fall prevention in construction projects using smart PLS-SEM analysis. *Journal of Safety and Sustainability*, 1(1), 1004. <https://doi.org/10.1016/j.jsasus.2025.10.004>






Figure 5. Practical Implementation Roadmap for Safety Engineers and Facility Managers

- Ankamah-Lomotey, S. A. (2025). Developing a safety culture through effective management systems: Considering outcomes from various management system models. *International Journal of Research and Innovation in Social Science*, 9(10), 8794-8803. <https://doi.org/10.47772/IJRISS.2025.910000723>
- Arya, J., & Hindolia, A. (2025). Nudging towards excellence: Harnessing behavioral insights for organizational learning. *Development and Learning in Organizations: An International Journal*, 39(1), 25-27. <https://doi.org/10.1108/DLO-01-2024-0019>
- Ashraf, A. M., Peres, S. C., & Sasangohar, F. (2026). Beyond compliance: Investigating work-as-done in procedural work with the systematic Skip-Order-Action (SOA) framework. *Safety Science*, 200, 107209. <https://doi.org/10.1016/j.ssci.2026.107209>
- Bakshi, N., & Peura, H. (2022). Prevent or report? Managing near misses for safer operations. *Manufacturing & Service Operations Management*, 24(4), 2118-2136. <https://doi.org/10.1287/msom.2022.1100>
- Bandura, A. (1977). *Social learning theory*. Prentice-Hall.
- Bechara, A. (2008). The Neurology of emotions and feelings, and their role in behavioral decisions. In: Parasuraman, R., & Rizzo, M. (Eds.), *Neuroergonomics: The Brain at Work* (pp. 178-192). Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780195177619.003.0012>
- Bertoni, V. B., Saurin, T. A., Fogliatto, F. S., & Patriarca, R. (2021). Monitor, anticipate, respond, and learn: Developing and interpreting a multilayer social network of resilience abilities. *Safety Science*, 136, 105148. <https://doi.org/10.1016/j.ssci.2020.105148>
- Bilgin, H., Sili, U., Pazar, N., Karaaslan, A., Soylemez, K., Akkoç, G., Kadayifci, E. K., & Korten, V. (2023). Effect of video camera monitoring feedback on hand hygiene compliance. *American Journal of Infection Control*, 51(8), 983-989. <https://doi.org/10.1016/j.ajic.2022.12.011>
- Breda, P., Di Pietrantonio, V., & Schneider, S. (2026). Spaceport selection, licensing complexity and launch operator constraints in NewSpace. *Journal of Space Safety Engineering*, 13(1), 45-56. <https://doi.org/10.1016/j.jsse.2026.02.012>
- Causse, M., Mouratille, D., Rouillard, Y., & Hidalgo-Muñoz, A. (2024). How a pilot's brain copes with stress and mental load? Insights from the executive control network. *Behavioural Brain Research*, 459, 114698. <https://doi.org/10.1016/j.bbr.2023.114698>
- Cialdini, R. B. (2021). *Influence: The Psychology of Persuasion (New and Expanded)*. Harper Business.
- Clapier, A., Herrbach, O., Lérat-Pytlak, J., & Lombardot, E. (2026). Nudges and risk prevention in the construction sector. *European Management Journal*, 44(2), 101345. <https://doi.org/10.1016/j.emj.2026.05.005>
- Clarke, R. J., Bickerton, R., Butcher, T., & Wyborn, B. (2026). Laser safety design considerations for the EPAC 10Hz petawatt laser system. *Journal of Laser Applications*, 38(1), Article 012001. <https://doi.org/10.2351/7.0001775>
- Cooper, A., Mazzeo, F., Waterson, P., & Louis, D. (2023). The use of Bayesian Belief Networks (BBNs) to probe deeper into railway safety management systems – Two studies from Great Britain and Italy. *Applied Ergonomics*, 111, 103968. <https://doi.org/10.1016/j.apergo.2023.103968>
- Costa, S., Duyck, W., & Dirix, N. (2025). Preventing hazardous exposure to suspended loads with nudge interventions. *Journal of Loss Prevention in the Process Industries*, 94, 105665. <https://doi.org/10.1016/j.jlp.2025.105665>
- de Freitas, J. H., Wachs, P., Henriqson, E., & Saurin, T. A. (2026). Moving from description to actionable improvement recommendations: Supporting resilient performance in oil and gas safety-critical operations. *Safety Science*, 182, 107192. <https://doi.org/10.1016/j.ssci.2026.107192>
- Dekker, S. (2011). *Drift into failure: From hunting scapegoats to understanding complex systems*. Ashgate Publishing.
- Dolgunyurek, O., & Ada, S. (2025). Evaluation of the effectiveness of implementation of ISO 45001 occupational health and safety management system in educational institutions. *Milli Egitim*, 54(245), 185-210. <https://doi.org/10.37669/milliegitim.1500544>
- Ghamari, A., Rezaei, N., Malekpour, M.-R., & Farzadfar, F. (2022). The effect of non-punitive peer comparison and performance feedback on drivers' behavior using telematics. *Journal of Safety Research*, 82, 312-322. <https://doi.org/10.1016/j.jsr.2022.07.010>
- Gondale, A., Sachdeva, L., & Gupta, S. D. (2025). Neurocognitive indicators of automation readiness in Industry 4.0. *TPM - Testing*,

- Psychometrics, *Methodology in Applied Psychology*, 32(1), e105014. <https://tpmap.org/submit/index.php/tpm/article/view/672>
- Grelle, S., Kuhn, S., Fuhrmann-Riebel, H., & Hofmann, W. (2024). The role of framing and effort in green nudging acceptance. *Behavioural Public Policy*, 8(3), 345–367. <https://doi.org/10.1017/bpp.2024.8>
- Grouse, C. K., Waung, M. W., Holmgren, A. J., & Khanna, R. R. (2023). Behavioral “nudges” in the electronic health record to reduce waste and misuse: 3 interventions. *Journal of the American Medical Informatics Association*, 30(4), 624–630. <https://doi.org/10.1093/jamia/ocac238>
- Ham, D.-H. (2021). Safety-II and resilience engineering in a nutshell: An introductory guide. *Safety and Health at Work*, 12(1), 3–10. <https://doi.org/10.1016/j.shaw.2020.11.004>
- Hao, S., Lifei, X., Ruiying, X., Jian, W., Jiaxin, J., Siping, F., Xiaoqin, W., Xin, Q., & Yuhang, Z. (2025). The impact of mind-wandering on driller’s situation awareness in drilling operations. *Human Factors and Ergonomics in Manufacturing & Service Industries*, 35(1), e70010. <https://doi.org/10.1002/hfm.70010>
- Holden, R. J., Brown, R. L., Scanlon, M. C., & Karsh, B. T. (2015). Micro-and macroergonomic changes in mental workload and medication safety. *International Journal of Industrial Ergonomics*, 46, 29–43. <https://doi.org/10.1016/j.ergon.2014.04.003>
- Hollnagel, E. (2014). *Safety-I and Safety-II: The past and future of safety management*. Ashgate Publishing.
- Hollnagel, E. (2017). *Safety-II in practice: Developing the resilience potentials*. Routledge. <https://doi.org/10.4324/9781315201023>
- Huang, F., & Boudjema, S. (2021). Three-year hand hygiene monitoring and impact of real-time reminders on compliance. *Journal of Hospital Infection*, 118, 129–135. <https://doi.org/10.1016/j.jhin.2021.08.015>
- Hussain, R., Zaidi, S. F. A., Pedro, A., & Park, C. (2024). Exploring construction workers’ attention and awareness in diverse virtual hazard scenarios. *Safety Science*, 171, 106526. <https://doi.org/10.1016/j.ssci.2024.106526>
- Johnson, K., Morais, C., & Patelli, E. (2025). Enhancing procedure quality: Advanced language tools for identifying ambiguity and high-potential violation triggers. *Reliability Engineering & System Safety*, 253, 111308. <https://doi.org/10.1016/j.ress.2025.111308>
- Kahneman, D. (2011). *Thinking, fast and slow*. Farrar, Straus and Giroux.
- Kim, N., Anderson, B. A., & Ahn, C. R. (2021). Reducing risk habituation to struck-by hazards in a road construction environment using virtual reality behavioral intervention. *Journal of Construction Engineering and Management*, 147(11), 04021147. [https://doi.org/10.1061/\(ASCE\)CO.1943-7862.0002187](https://doi.org/10.1061/(ASCE)CO.1943-7862.0002187)
- Kim, N., Grégoire, L., Razavi, M., & Anderson, B. A. (2024). Protocol to evaluate the effectiveness of a virtual-reality-based behavioral intervention in enhancing sensory responses to real-world warning. *STAR Protocols*, 5(2), 103262. <https://doi.org/10.1016/j.xpro.2024.103262>
- Lee, B., & Kim, H. (2022). Measuring effects of safety-reminding interventions against risk habituation. *Safety Science*, 154, 105857. <https://doi.org/10.1016/j.ssci.2022.105857>
- Liberati, E. G., Peerally, M. F., & Dixon-Woods, M. (2018). Learning from high risk industries may not be straightforward: A qualitative study of the hierarchy of risk controls approach in healthcare. *International Journal for Quality in Health Care*, 30(1), 39–44. <https://doi.org/10.1093/intqhc/mzx163>
- Liu, D., & Ham, Y. (2025). Multistage cognitive dynamics for team-based teleoperation in challenging environments. *Automation in Construction*, 171, 106537. <https://doi.org/10.1016/j.autcon.2025.106537>
- Liu, X., Liu, Y., Li, H., & Wen, D. (2023). Identification and analysis of barriers to the effectiveness of ISO 45001 certification. *Journal of Cleaner Production*, 384, 135447. <https://doi.org/10.1016/j.jclepro.2022.135447>
- Mathisen, G. E., & Tjora, T. (2023). Safety voice climate: A psychometric evaluation and validation. *Journal of Safety Research*, 86, 285–294. <https://doi.org/10.1016/j.jsr.2023.05.008>
- Matias, J., Pelletier, C., Izaute, M., & Silvert, L. (2022). The role of perceptual and cognitive load on inattention blindness. *Quarterly Journal of Experimental Psychology*, 75(11), 2002–2022. <https://doi.org/10.1177/17470218211064903>
- Raheemy, Y., Hollowell, M. R., Lingard, H., & Sherratt, F. (2026). Safety clutter: An attribute-based conceptualization and taxonomy for diagnosis. *Journal of Safety Research*, 96, 107–118. <https://doi.org/10.1016/j.jsr.2026.06.005>
- Rantala, E., Vanhatalo, S., Perez-Cueto, F. J. A., & Absetz, P. (2023). Acceptability of workplace choice architecture modification for healthy behaviours. *BMC Public Health*, 23, 2146. <https://doi.org/10.1186/s12889-023-17331-x>
- Reale, C., Salwei, M. E., Militello, L. G., Weinger, M. B., Burden, A., Sushereba, C., Torsher, L. C., Andreae, M. H., Gaba, D. M., McIvor, W. R., Banerjee, A., Slagle, J., & Anders, S. (2023). Decision-making during high-risk events: A systematic literature review. *Journal of Cognitive Engineering and Decision Making*, 17(2), 188–212. <https://doi.org/10.1177/15553434221147415>
- Sadeghi, H., Mohandes, S. R., Yunusa-Kaltungo, A., Cheung, C., & Manu, P. (2025). A state-of-the-art review of safety leading indicators across diverse industries. *Journal of Safety Science and Resilience*, 100272. <https://doi.org/10.1016/j.jnlssr.2025.100272>
- Shenhav, A., Musslick, S., Lieder, F., Kool, W., Griffiths, T. L., Cohen, J. D., & Botvinick, M. M. (2017). Toward a rational decision-making framework for mental effort investment. *Neuron*, 94(4), 690–702. <https://doi.org/10.1016/j.neuron.2017.04.030>
- Sparer, E. H., Herrick, R. F., & Dennerlein, J. T. (2015). Development of a safety communication and recognition program for construction. *New Solutions*, 25(1), 42–58. <https://doi.org/10.1177/104829115569025>
- Stackhouse, M. R. D., & Stewart, R. (2017). Failing to fix what is found: Risk accommodation in the oil and gas industry. *Risk Analysis*, 37(2), 316–328. <https://doi.org/10.1111/risa.12583>
- Stauffer, T., & Chastain-Knight, D. (2021). Do not let your safe operating limits leave you SO-L (out of luck). *Process Safety Progress*, 40(2), 125–131. <https://doi.org/10.1002/prs.12163>
- Tedone, A. M., & Lanz, J. J. (2024). Staying silent during a crisis: How workplace factors influence safety decisions in U.S. nurses. *Social Science & Medicine*, 342, 116548. <https://doi.org/10.1016/j.socscimed.2023.116548>
- Tedone, A. M., Mesmer-Magnus, J., Viswesvaran, C., & Newkirk, E. (2026). See something, say something: A meta-analytic review of safety voice. *International Journal of Management Reviews*, 28(1), 45–67. <https://doi.org/10.1111/ijmr.12401>
- Thaler, R. H., & Sunstein, C. R. (2008). *Nudge: Improving Decisions About Health, Wealth, and Happiness*. Yale University Press.
- Theeuwes, J. (1992). Perceptual selectivity for color and form. *Perception & Psychophysics*, 51(6), 599–606. <https://doi.org/10.3758/BF03211656>
- Thoroman, B., Goode, N., & Salmon, P. (2018). System thinking applied to near misses: A review of industry-wide near miss reporting systems. *Theoretical Issues in Ergonomics Science*, 20(4), 469–490. <https://doi.org/10.1080/1463922X.2018.1484527>
- Trancoso, R., Patriarca, R., & Henriqson, É. (2024). Assessing resilience in everyday work through observations. *Safety Science*, 176, 106461. <https://doi.org/10.1016/j.ssci.2024.106461>
- Veerasimman, A. (2026). A department level performance metric architecture for operationalizing behavior-based safety in high-risk manufacturing. *Safety Science*, 182, 107330. <https://doi.org/10.1016/j.ssci.2026.107330>
- Vo, B., & Kim Quy, H. T. (2026). Safe skies begin with safe voices: Leadership and reporting in aviation. *Journal of Air Transport Management*, 114, 103072. <https://doi.org/10.1016/j.jairtraman.2026.103072>
- von Thiele Schwarz, U., Nielsen, K., Edwards, K., Hasson, H., Ipsen, C., Savage, C., Abildgaard, J. S., Richter, A., Caroline Lornudd, Pamela Mazzocato, C. L.P., & Reed, J. E. (2021). How to design, implement and evaluate organizational interventions for maximum impact: the Sigtuna Principles. *European Journal of Work and Organizational Psychology*, 30(3), 415–427. <https://doi.org/10.1080/1359432X.2020.1803960>
- Woltjer, R., Pinska-Chauvin, E., Laursen, T., & Josefsson, B. (2015). Towards understanding work-as-done in air traffic management safety assessment and design. *Reliability Engineering & System Safety*, 141, 64–74. <https://doi.org/10.1016/j.ress.2015.03.010>
- Wong, L. H. M., Hurbean, L., Davison, R. M., & Muntean, M. (2022). Working around inadequate information systems in the workplace. *International Journal of Information Management*, 64, 102471. <https://doi.org/10.1016/j.ijinfomgt.2022.102471>
- Wu, S. J., & Paluck, E. L. (2021). Designing nudges for the context: Golden coin decals nudge workplace behavior in China. *Organizational Behavior and Human Decision Processes*, 163, 1–13. <https://doi.org/10.1016/j.obhdp.2018.10.002>
- Yu, P., & Jiang, Y. (2026). The mediating role of inattention blindness between risk propensity and risk perception. *Sustainability*, 18(11), 5498. <https://doi.org/10.3390/sui18115498>
- Zhylin, M., Morozova, T., Malyshev, V., Shcherba, V., & Medianova, O. (2026). The use of heuristics in quick decision-making as a manifestation of bounded rationality under stress. *Revista Cubana de Medicina Militar*, 55(1), e105036. <https://revmedmilitar.sld.cu/index.php/mil/article/view/77139>

