Safety Science 200 (2026) 107228

Contents lists available at ScienceDirect

Safety Science

journal homepage: www.elsevier.com/locate/safety

Holding the egg: ethical tensions of care in safety-critical technologies

Oscar Oviedo-Trespalacios *

Department of Values, Technology and Innovation, Faculty of Technology, Policy and Management, Delft University of Technology, Jaffalaan 5, 2628 BX Delft, The Netherlands

ARTICLE INFO

Keywords:
Ethical care after deployment
Safety-critical technologies
Technological persistence
Responsibility
Legacy systems
Safety Science

ABSTRACT

Safety-critical technologies frequently persist for decades, outliving their original design assumptions, organisational arrangements, and regulatory contexts. Yet ethical and safety scholarship remains predominantly oriented toward early lifecycle phases, emphasising design assurance, deployment decisions, and prevention, while offering limited conceptual resources for understanding responsibility once systems are already in operation, ageing, and increasingly difficult to modify or withdraw. This paper develops a conceptual account of ethical care after deployment, referring to forms of responsibility that persist once safety–critical technologies become operationally embedded and remain in use while practical steerability, authority, epistemic certainty, and exit options gradually erode. The animated film *Angel's Egg* (Oshii, 1985) is used as a heuristic device—not as empirical evidence or film analysis—to surface ethical tensions associated with sustaining fragile systems over time. On this basis, the paper proposes a preliminary conceptual framework for examining ethical conditions that remain weakly captured by prevailing safety concepts. The resulting CARE-TECH framework identifies six interrelated tensions: care without repair, responsibility without authority, care under epistemic uncertainty, fragility without exit, inherited burden, and irreversibility of ethical failure. The framework further clarifies how these tensions cluster as limits of agency and temporal–epistemic burdens, sustained by structural conditions of persistence and shaped by an ethical horizon of irreversibility. The contribution is intentionally conceptual, offering a vocabulary to recognise and analyse the accumulation and distribution of ethical burden across the operational life of safety–critical technologies, and outlining directions for empirical engagement in domains where long-lived systems must be sustained under constraint.

1. Introduction

Safety-critical technologies often stay in use for a long time, sometimes for decades. Across transport, healthcare, and industry, many safety-critical systems outlive the conditions they were designed for. In the public sector, this is also evident in large-scale information systems: a U.S. Government Accountability Office report (U.S. Government Accountability Office, 2016) shows that federal agencies rely on decades-old legacy IT systems—some more than 50 years old—that absorb the majority of operational budgets yet remain difficult to modernise, replace, or retire. Once embedded, such systems become resistant to change while continuing to shape patterns of risk, responsibility, and potential harm in everyday practice. Consequently, safety cannot be understood solely as a matter of sound design or initial deployment, but must also be considered in terms of sustained care, vigilance, and ethical responsibility throughout the long operational life of safety-critical technologies.

Research in Safety Science has long examined the challenges associated with ageing infrastructures, asset management, and long-term risk governance. Studies of infrastructure systems highlight how safety and reliability must be sustained over extended operational lifetimes through maintenance practices, organisational arrangements, and evolving regulatory environments (Komljenovic et al., 2016; Little, 2012). Empirical analyses of major hazard installations show that ageing assets and material degradation can contribute directly to accidents (Hansler et al., 2022). These systems are often treated as complex adaptive systems in which risks emerge and evolve over time from interactions among technical components, organisational processes, and external conditions. While this body of work provides valuable insights into the management of ageing technological systems, it has only occasionally framed these challenges in terms of the ethical character of responsibility that persists once technologies remain in operation despite diminishing steerability and limited exit options.

The relative absence of explicit discussion of ethical care after

* Corresponding author.
E-mail address: O.OviedoTrespalacios@tudelft.nl.

https://doi.org/10.1016/j.ssci.2026.107228
Received 4 January 2026; Received in revised form 2 April 2026; Accepted 2 April 2026
Available online 20 April 2026

deployment is notable given the prevalence of long-lived systems across many safety-critical domains. Ethical reflection on such systems is not entirely absent from the literature—Bowen (2000), for example, examines responsibility, competence, and professional caution in the development and use of safety-critical software. However, these contributions are largely oriented toward design-time decisions and professional conduct during development. Consequently, the ethical challenges that arise once technologies enter extended operation—when systems age, design choices become difficult to revise, and organisations grow dependent on legacy infrastructures—remain comparatively under-theorised within Safety Science. This limits our ability to explain how ethical responsibility is experienced, distributed, and negotiated over the long operational life of safety–critical technologies.

Against this backdrop, the present paper draws conceptual inspiration from an unexpected source in popular culture. *Angel's Egg* (1985) is an experimental animated film in which a girl carries and protects a fragile egg in a desolate world, briefly bonds with a boy who challenges the meaning of what she protects, and witnesses the irreversible rupture of that fragile object. The film is introduced here neither as evidence about real-world systems nor as an object of film criticism, but as a heuristic prompt: a narrative device that helps make visible and discussable ethical phenomena—endurance, obligation under ambiguity, and responsibility sustained beyond initial novelty—that are widely encountered in safety-critical domains yet remain weakly articulated in prevailing frameworks. The film is considered in its entirety, but the discussion refers only to a small number of illustrative moments that most clearly foreground these ethical conditions. These scenes are used descriptively as narrative prompts rather than analysed as empirical material, helping to articulate tensions that are already recognisable in long-lived safety-critical systems.

Using popular culture in this way has a precedent in scholarship. Across the social sciences and humanities, films, novels, and other cultural artefacts are routinely used as *sensitising devices* to surface tacit assumptions, sharpen conceptual distinctions, and render abstract ethical dynamics more intelligible, without treating the artefact itself as empirical evidence. Within Safety Science, Walker et al. (2015) provide a methodological precedent by using the fictional Death Star from *Star Wars* as an analytic case to compare component-based and systems-oriented safety methods, demonstrating how different analytical approaches reveal—or obscure—critical vulnerabilities in complex sociotechnical systems. In parallel, some empirical research has examined films themselves as objects of safety inquiry, systematically analysing how movies depict safety-related behaviours and risk practices—for example, studies of children's films have quantified portrayals of seatbelt use, helmet wearing, speeding, and pedestrian behaviour, highlighting how cinematic representations model both safe and unsafe practices (e.g., Tongren et al., 2010; Boppana et al., 2016). Taken together, this literature establishes popular culture as both an analytic resource and a legitimate site of safety-relevant insight. In this paper, *Angel's Egg* is positioned within that tradition: not as an object of cinematic critique, but as an analytic trigger that helps articulate post-deployment ethical tensions—burdens and frictions of care—that shape how safety is sustained when systems persist, age, and become difficult to alter or withdraw.

Building on this gap, the present paper develops a conceptual account of ethical care after deployment, using *Angel's Egg* as a heuristic lens to reflect on responsibility under conditions of technological persistence. The film is not treated as evidence about real-world systems, but as a narrative prompt that foregrounds ethical experiences that are difficult to articulate within dominant technical and organisational vocabularies of Safety Science. By linking moments in the film to recurring ethical tensions encountered in long-lived safety-critical systems, the paper connects these reflections to technologies that persist and become difficult to modify or withdraw. The contribution is intentionally conceptual: rather than offering an empirical case study, the paper proposes a preliminary framework for thinking about post-deployment care, intended to support Safety Science in addressing forms of responsibility that arise beyond novelty, design, and initial control in an increasingly technology-dependent world. In doing so, the paper contributes to Safety Science by introducing a conceptual vocabulary for analysing responsibility under technological persistence, by articulating recurring ethical tensions that arise when systems remain operational despite declining steerability, and by reframing everyday safety practices—such as monitoring, maintenance, and vigilance—as forms of ethical care sustained under constraint.

To develop this argument, the paper proceeds in three steps. Section 2 conceptualises ethical care after deployment as a problem for Safety Science, clarifying how responsibility can persist when safety–critical technologies remain in operation despite diminishing steerability, authority, and realistic exit options. Building on this foundation, Section 3 introduces the CARE-TECH framework, which articulates six interrelated ethical tensions that arise under conditions of technological persistence: care without repair, responsibility without authority, care under epistemic uncertainty, fragility without exit, inherited burden, and the irreversibility of ethical failure. These tensions do not represent discrete failure modes or normative prescriptions, but recurring conditions that shape how responsibility is experienced and sustained in long-lived sociotechnical systems. The final sections discuss the implications of this framework for Safety Science and outline directions for future empirical research.

2. Ethical care after Deployment: A conceptual problem for Safety Science

This section introduces the concept of ethical care after deployment and situates it within existing Safety Science and technology ethics scholarship, clarifying the conceptual problem that motivates the framework developed later in the paper. Ethical care after deployment refers to a form of responsibility that arises when safety–critical technologies remain in operation long after their introduction and become deeply embedded in infrastructures, organisational routines, and regulatory arrangements. Under these conditions of technological persistence, responsibility for system outcomes continues even as practical capacities for redesign, withdrawal, or meaningful steering gradually erode. In this paper, deployment refers primarily to the introduction and early embedding of a technology into operational practice. Design and assurance refer to ex ante activities through which acceptable risk is specified and justified prior to this introduction, while control occupies an intermediate zone in which actors retain the practical capacity to steer, suspend, or redirect system behaviour during operation. Post-deployment care becomes salient once technologies are sufficiently embedded that responsibility persists despite declining opportunities for redesign, intervention, or exit. The distinction is therefore analytical rather than temporal: post-deployment care refers not to a specific moment in a technological lifecycle, but to conditions in which systems remain operationally embedded while the capacity to meaningfully modify or withdraw them becomes constrained.

This condition can be observed in contemporary AI-enabled safety-critical systems such as machine-learning–based clinical decision-support tools used in hospitals. These systems analyse patient data—such as medical records, laboratory results, and imaging—to generate predictions or recommendations that assist clinicians in tasks such as diagnosis, treatment planning, or risk assessment. Once these systems enter routine use, however, opportunities for system-level modification are often limited. Control is therefore exercised primarily through human oversight—for example, clinicians reviewing or overriding recommendations—rather than through redesign or withdrawal of the system. As a result, operators may remain responsible for outcomes even when they lack the practical ability to modify core system behaviour or decisively halt its operation. Empirical studies of AI-based safety-critical systems have documented this pattern: responsibility for outcomes persists even as the practical ability to influence system behaviour

declines. In some cases, this produces "liability sink" dynamics, where operational actors carry ethical and legal responsibility despite having limited control over the system (Ryan et al., 2023).

Responsibility in technology and engineering has long been recognised as temporally complex, encompassing both forward-looking and backward-looking dimensions that unfold across collective and uncertain action contexts (Doorn & van de Poel, 2012). In the context of infrastructures and safety–critical systems, responsibility is also distributed across multiple actors and institutions and often extends beyond moments of design or failure, encompassing ongoing obligations to maintain, restore, and safeguard systems upon which societies depend (Van der Bruggen, 2008). However, existing safety concepts tend to frame responsibility primarily in technical, procedural, or organisational terms rather than as a distinct ethical condition that persists when control over systems diminishes (Dekker & Breakey, 2016). Ethical care after deployment builds on this insight by focusing on situations in which responsibility endures even as practical control, redesign capacity, and clear authority become constrained within complex sociotechnical infrastructures characterised by distributed risks and responsibilities (Grady et al., 2021). In this context, responsibility refers to the enduring obligation associated with the safety consequences of technologies that remain in operation over time, while care refers to the sustained attentiveness, stewardship, and practical engagement through which this responsibility is enacted when redesign, withdrawal, or full control are no longer feasible. Within this framework, ethical burden refers to the experiential weight that arises when actors must sustain such care under conditions of constrained agency, uncertainty, and limited exit options. This operational use of care resonates with, but is not identical to, the notion of care developed in parts of science and technology studies and technology ethics.

Work in technology ethics has explored the relevance of care for understanding sociotechnical practices and responsibilities. For example, Puig de la Bellacasa (2011) introduces the notion of "matters of care" to emphasise that technoscientific assemblages are sustained not only through technical design and governance but also through ongoing practices of attentiveness, maintenance, and responsibility toward fragile sociotechnical arrangements. Building on related insights, Baas, Metselaar, and Klaassen (2022) develop a care-ethics perspective on Safe-by-Design, arguing that safety should not be treated solely as a technical property of technologies but as a relational and ongoing practice in which multiple actors share responsibility for caring for safety throughout a technology's life cycle. Their concept of "circles of care" highlights how safety emerges through sustained attentiveness, responsibility, and responsiveness among stakeholders involved in the design, governance, and use of technologies. However, the present paper employs the notion of care primarily as a conceptual and descriptive lens rather than as a direct application of care ethics as a normative philosophical framework. In this context, care refers operationally to the sustained attentiveness, stewardship, and practical engagement through which responsibility for safety is enacted when systems persist in operation even as opportunities for redesign, withdrawal, or meaningful control become constrained. While existing care-ethics discussions focus primarily on design and innovation processes, the present paper extends the discussion to the post-deployment phase of safety–critical technologies, where attentiveness and stewardship continue to be required long after systems have entered routine operation.

Several established safety concepts capture important aspects of sustaining safety under conditions of constrained control, but they rarely address its ethical implications directly. Maintenance focuses on the technical and organisational work required to preserve system functionality over time, yet its normative orientation remains largely instrumental: systems are kept running relative to predefined performance expectations. Compliance establishes accountability through adherence to standards, regulations, and certification regimes, but presumes that responsibility can be discharged through conformity to formal requirements. Resilience foregrounds adaptive capacity and recovery under disturbance, emphasising how systems cope, adjust, and continue to function. Safety management integrates these elements within organisational structures, framing responsibility as allocable, auditable, and subject to optimisation. Taken together, these approaches are indispensable for sustaining safety. Yet they often treat responsibility as something that can ultimately be resolved through intervention, adaptation, or procedural accountability (Dekker & Breakey, 2016; McCall & Pruchnicki, 2017), rather than examining situations in which responsibility persists even when the practical capacity to steer, redesign, or withdraw a system becomes limited.

Similar tensions have begun to surface in safety and risk governance scholarship addressing long-lived sociotechnical infrastructures and contested transitions, where safety–critical systems remain operational despite growing ethical or societal pressure for transformation. Recent work has highlighted how ethical disengagement from harmful systems, including fossil fuel infrastructures, may be normatively compelling while responsibility nevertheless persists during prolonged phases in which infrastructures, risks, and institutional dependencies remain active (Oviedo-Trespalacios et al., 2025). Under such conditions, responsibility cannot be discharged through withdrawal alone, because safety–critical systems continue to operate and generate potential harm during non-ideal periods of transformation. These analyses point to a structural condition in which responsibility endures under institutional lock-in and constrained exit, yet they stop short of articulating the ethical character of this persistence. Ethical care after deployment builds on this emerging recognition by providing a conceptual vocabulary for examining responsibility as a sustained burden when redesign, withdrawal, or resolution are not practically available.

Ethical care after deployment refers to a condition in which safety–critical technologies persist in operation beyond the horizons of design assumptions, while responsibility for their consequences continues even as practical capacities for redesign, withdrawal, or meaningful steering become increasingly constrained. It concerns situations in which systems continue to matter ethically—because they shape risk, vulnerability, and potential harm—yet cannot be meaningfully redesigned, decisively improved, or cleanly withdrawn. In such contexts, responsibility is not exhausted by procedural conformity or organisational optimisation, but remains as a moral burden borne by actors who must continue to attend to something fragile and consequential without clear authority, resolution, or prospects for closure.

One influential attempt to address responsibility in complex technological systems focuses on preserving meaningful human control. This framework grounds responsibility in the preservation of tracking—system behaviour remaining responsive to relevant human reasons and values—and tracing—identifiable human agents retaining knowledge, role awareness, and answerability for system behaviour (Santoni de Sio & van den Hoven, 2018; Santoni de Sio & Mecacci, 2021). While this approach powerfully clarifies the conditions under which responsibility can be sustained, it remains oriented toward maintaining or restoring control over technological systems. Ethical care after deployment becomes most visible in a different and under-theorised situation: when systems persist despite the gradual erosion of these control conditions. In such cases, responsibility endures even as the practical capacities to steer, revise, or withdraw the system decay over time.

This erosion of practical steerability is not only an operational or governance problem; it also generates a distinct ethical burden that remains only partially articulated within Safety Science. When long-lived systems become infrastructural—organisationally depended upon, costly to suspend, and difficult to redesign—responsibility does not disappear; it accumulates, is inherited across roles, and is often unevenly distributed to those closest to operations and maintenance. This burden is intensified when responsibility relations themselves are structurally degraded. Related work in the philosophy of technology has similarly argued that responsibility failures in advanced sociotechnical systems cannot be understood solely in terms of diminished control or oversight. Vallor and Vierkant (2024), for example, describe a

"vulnerability gap" in which institutional and technical arrangements weaken reciprocal relations of answerability between system actors and those affected by system action. While developed in the context of AI governance, this diagnosis reinforces the need to examine responsibility and care beyond design and deployment, particularly where systems persist and responsibility is sustained under constraint.

Taken together, these strands of scholarship reveal an emerging but still fragmented recognition that responsibility in sociotechnical systems can persist under conditions of declining control. Research in Safety Science has examined the organisational and technical work required to sustain ageing systems; governance scholarship has highlighted the persistence of responsibility during contested infrastructural transitions; and work in technology ethics has sought to preserve responsibility through frameworks such as meaningful human control or through analyses of degraded answerability relations. Yet these contributions address different aspects of the problem and rarely converge on the ethical condition that arises when safety–critical systems continue to operate despite diminishing capacities to redesign, suspend, or meaningfully steer them. What remains insufficiently articulated is how responsibility is experienced, distributed, and sustained in such circumstances, when actors must continue to care for systems that remain consequential but increasingly resistant to intervention. Ethical care after deployment names this condition. It foregrounds the ethical burdens that arise when responsibility extends beyond the practical horizons of design, control, and withdrawal, thereby opening a conceptual space for examining how responsibility is sustained in long-lived sociotechnical systems. The following section builds on this conceptualisation by introducing the CARE-TECH framework, which articulates the specific ethical tensions that emerge under conditions of technological persistence.

3. The CARE-TECH Framework: Ethical care after deployment in Safety-Critical technologies

3.1. Framing ethical care under technological persistence

This section develops a conceptual framework for ethical care after deployment (hereafter, the CARE-TECH framework) by articulating a set of ethical tensions that arise under conditions of technological persistence, where safety–critical systems remain operationally embedded while opportunities for redesign, withdrawal, or effective control become increasingly constrained. In this paper, these tensions are understood not as discrete moral dilemmas or binary trade-offs, but as persistent conditions in which responsibility remains active while the practical, epistemic, and institutional conditions for sustaining it are progressively constrained. The CARE-TECH framework is not intended as a taxonomy of failure modes, a set of normative prescriptions, or a guide for intervention. Rather, it provides a structured way of naming and examining ethical conditions that are already present in practice but remain weakly captured by existing safety concepts, particularly in contexts of technological persistence and institutional lock-in.

The CARE-TECH framework is developed heuristically through selected scenes from Angel's Egg, used to foreground ethical experiences that are difficult to access through technical or organisational analysis alone. Importantly, the tensions themselves are not derived from the film; rather, they emerge from reflection on recurring conditions observed in long-lived safety–critical systems, with the film serving only as a heuristic lens that renders these ethical dynamics more conceptually visible. These scenes are treated descriptively, focusing on what is enacted rather than what is symbolised, and function as narrative prompts that help render visible tensions associated with sustained care under conditions of fragility, inheritance, and constraint. Each tension is articulated through a brief scene-based description, followed by an account of the ethical condition it surfaces and its resonance with real-world situations recognisable across safety–critical domains. Taken together, the tensions do not resolve into a unified normative stance but describe the frictions, burdens, and asymmetries that shape ethical care when obligations remain despite eroding practical control, limited exit options, and constrained prospects for repair.

3.2. Ethical tensions

The tensions articulated below describe distinct dimensions of ethical care that arise in long-lived safety–critical systems operating under conditions of technological persistence and constrained steerability. They are not alternative labels for the same phenomenon, but analytically different situations through which post-deployment responsibility becomes constrained. The six tensions presented here are not intended as an exhaustive taxonomy, but as recurring patterns that emerge when responsibility persists under conditions of constrained agency, technological persistence, and limited exit options. Some tensions concern limits of agency, where actors remain responsible for safety outcomes despite restricted capacity to repair systems or exercise authority over their evolution. Others reflect temporal and epistemic burdens, where responsibility must be sustained across generations of practitioners or under conditions of irreducible uncertainty. Additional tensions arise from structural persistence, where fragile systems cannot be safely abandoned even when their limitations are recognised because viable exit options are foreclosed. Finally, the possibility that certain failures would produce irreversible consequences operates as an ethical horizon that intensifies the weight of care across all other tensions. The subsections that follow describe how these tensions surface in practice and shape the experience of responsibility in long-lived safety–critical systems.

3.2.1. Care without repair

> *"The girl moves slowly through a ruined city, holding the egg close to her body. The environment is static and desolate: buildings are collapsed, machinery is inert, and no signs of renewal are visible. The girl repeatedly checks the egg for cracks, cleans it, and adjusts how she carries it. Nothing in her actions alters the surrounding environment or changes the condition of the egg itself."*

This scene foregrounds a form of care oriented toward preservation rather than improvement. The girl's actions are attentive and continuous, yet they do not aim at repair, optimisation, or progress. Care is enacted as the effort to hold something together in its current, fragile state, despite the absence of realistic prospects for restoration or transformation. Importantly, the scene does not depict negligence or passivity; it depicts responsibility exercised under conditions where repair is no longer available as a meaningful option. Care here is not inertia, but sustained moral effort under constraint.

In safety-critical systems, this tension arises when technologies remain in operation despite being fundamentally constrained by legacy design, obsolescence, or infrastructural dependence. In the rail industry, for example, signalling systems continue to operate as long-lived, safety–critical infrastructures composed of multiple generations of mechanical, electrical, and computer-based interlockings, many of which remain safety-certified decades after their introduction (Office of Rail and Road, 2021). Replacement or major renewal is costly and often unaffordable at scale, and local change is constrained by the need to interface with surrounding legacy systems, making withdrawal or full replacement disruptive at the network level (Office of Rail and Road, 2021). Research on contemporary rail signalling further shows that safety is typically sustained through stability and incremental adaptation layered onto existing architectures, rather than through wholesale redesign, reinforcing the persistence of legacy systems in everyday operation (Aoun et al., 2024). As a result, organisations remain responsible for maintaining safety in systems that they cannot easily modernise or decisively retire, illustrating a form of post-deployment care shaped by technological persistence rather than active design

control.

In such cases, practitioners are not primarily tasked with system-level repair or transformation, but with preventing degradation and sustaining acceptable performance within known constraints. Ethical responsibility persists even though the technology cannot be repaired in any substantive sense. While Safety Science has traditionally framed these situations in terms of maintenance, reliability, and risk containment, the tension of care without repair draws attention to the ethical burden of sustaining safety when meaningful improvement or withdrawal is no longer achievable.

3.2.2. Responsibility without authority

> *"The girl continues to protect the egg even as the boy repeatedly questions its value. He handles the egg without her consent and openly states his desire to break it. Despite her responsibility for the egg, the girl has no authority to prevent his interference beyond physical proximity."*

Here, responsibility is clearly located with the girl, yet authority over the fate of the object she cares for is not. The scene renders visible a misalignment between who bears responsibility and who holds decision-making power. The girl's obligation to protect the egg remains intact even as her capacity to control its future is undermined by another actor. This tension maps onto safety-critical contexts in which responsibility for outcomes is assigned to those closest to operation, even when they lack meaningful authority to govern, modify, or withdraw the technologies involved.

This misalignment can be observed in safety-critical AI deployments in healthcare. Habli et al. (2020) describe how clinicians using AI-based clinical decision-support systems remain morally and professionally accountable for patient outcomes, despite having limited control over the system's internal decision-making processes and only partial understanding of how recommendations are generated. Although clinicians retain formal authority over the final decision, key elements of interpretation and judgement are delegated to opaque systems that they cannot meaningfully alter once operationally embedded. Ethical care after deployment captures this persistence of responsibility under asymmetric power relations, where responsibility (and, in practice, accountability) remains local while authority is distributed elsewhere. This condition exceeds standard notions of human oversight or accountability allocation, pointing instead to a structural ethical imbalance that safety management and regulatory frameworks tend to stabilise rather than critically interrogate.

3.2.3. Care under epistemic uncertainty

> *"The contents of the egg are never shown. The girl states that she believes something with flapping wings will hatch, while the boy doubts that anything exists inside at all. No evidence is provided to resolve the disagreement. The girl's care continues without confirmation or validation."*

This tension concerns care exercised under conditions where knowledge remains structurally incomplete and cannot be resolved through additional data, modelling, or verification. Ethical responsibility persists despite the absence of epistemic closure, requiring sustained vigilance, restraint, and stewardship rather than informed control. A contemporary analogue can be found in the governance of deep-sea mining in areas beyond national jurisdiction, such as the Clarion–Clipperton Zone. Regulatory authorities are required to make decisions about environmental protection, thresholds of harm, and acceptable risk in systems where baseline ecological conditions remain only partially known. Authoritative scientific assessments indicate that deep-sea mining may generate wide-ranging, long-lasting, and potentially irreversible impacts, including biodiversity loss, sediment plume dispersal, and disruption of interconnected deep-sea ecosystems, while simultaneously emphasising the profound uncertainty surrounding these effects (Scientific Advisory Board on Deep-Sea Mining, 2024). From a Safety Science perspective, this uncertainty is not a temporary deficit but a defining feature of the system. Many species remain undocumented, recovery times may span centuries, and long-term ecosystem dynamics cannot be reliably predicted. Yet international law requires effective protection of the marine environment, placing responsibility on regulators and states to act before definitive evidence of harm is available. This has led to increasing calls for precautionary pauses or moratoria on deep-sea mining, justified not by proven damage but by the inability to ensure adequate protection under current scientific conditions (Singh et al., 2025).

As in Angel's Egg, care is sustained without confirmation. Responsibility persists even when the object of care cannot be fully known, empirically validated, or repaired. Ethical care after deployment therefore shifts attention away from achieving complete knowledge toward sustaining precaution and vigilance, recognising that safety–critical practice must sometimes proceed by guarding against irreversible harm even when uncertainty cannot be resolved.

3.2.4. Fragility without exit

> *"The girl never sets the egg down permanently, nor does she abandon it. Even when exhausted, she rests beside it. There is no moment in which leaving the egg behind is presented as a viable option."*

This scene makes visible a condition in which care is required precisely because exit is not available. Unlike care without repair, which refers to situations where meaningful improvement or restoration of a system is no longer feasible, fragility without exit concerns cases in which vulnerabilities are recognised but withdrawal or replacement of the system is not a viable option. The egg is fragile, and abandoning it would constitute a moral failure, yet continuing to carry it imposes an ongoing burden. The absence of an acceptable exit option renders care obligatory rather than chosen. Care persists not because it is optimal or sustainable, but because withdrawal would produce immediate and unacceptable harm.

In safety–critical systems, exit is often similarly constrained by infrastructural dependence and risk displacement. Certain technologies cannot be safely withdrawn without creating new hazards that exceed those associated with continued operation. A clear example can be found in wide-area smoke detectors deployed in critical infrastructure facilities such as airports, transport hubs, and large public buildings. These systems are known to exhibit fragilities, including susceptibility to false alarms and vulnerabilities that can suppress genuine fire detection under specific conditions. Empirical security analyses have shown that optical beam smoke detectors can be remotely manipulated using relatively simple external interference, with potentially severe consequences for evacuation safety and emergency response (Shin et al., 2020).

Despite these documented vulnerabilities, such systems cannot simply be decommissioned. Removing or disabling wide-area smoke detection would create an immediate safety void in environments where early fire detection is critical and where alternative detection infrastructures are not readily available. Operators and facility managers therefore remain ethically responsible for continuing to rely on systems whose fragility is well understood, precisely because no safe exit exists. Continued use is not an endorsement of adequacy, but a necessity imposed by the risks associated with absence. Ethical care after deployment foregrounds this asymmetry: responsibility persists not because the system is robust, but because withdrawal would be more dangerous than continued, carefully managed reliance.

3.2.5. Inherited burden

> *"The girl does not create the egg. She inherits it already intact, already fragile, already meaningful. No origin story is provided, and no prior choice is depicted."*

This scene highlights responsibility that is received rather than chosen. The girl's obligation does not stem from authorship, consent, or initial decision-making, nor is it accompanied by instruction or explanation about the nature of what she has inherited. Care is assumed as a condition of the present, independent of past intentions, design rationales, or opportunities for preparation. In long-lived safety–critical systems, responsibility is similarly inherited across generations of engineers, operators, regulators, and organisations. Contemporary actors are accountable for systems designed under historical assumptions, standards, and constraints they did not set, and for which authoritative knowledge may no longer be fully available.

An example of this can be found in cross-generational space exploration missions such as NASA's Voyager programme. The Voyager spacecraft were designed and launched in the late 1970s and have remained in operation far beyond their originally anticipated mission lifespan. Many of the engineers who designed and built the systems are no longer available, and the spacecraft were developed at a time when documentation practices, archiving standards, and assumptions about system longevity differed substantially from contemporary engineering norms. Responsibility for the continued safe operation and scientific integrity of Voyager has therefore been inherited by successive generations of engineers who did not design the systems and cannot fully reconstruct their original rationale. This condition became visible during recent anomalies in Voyager 1′s flight data subsystem, where engineers relied on fragmented, decades-old documentation and cautious remote interventions to address corrupted onboard memory in hardware designed under historical constraints and without the possibility of physical repair (National Aeronautics and Space Administration [NASA], 2024). The Voyager mission thus exemplifies inherited ethical burden: care exercised under extreme temporal distance, eroded instruction, and irreversible constraints, where responsibility persists even as authorship, authority, and complete knowledge have long since dissipated. Ethical care after deployment names this inherited burden of responsibility, which persists beyond obsolete design rationales and eroded instruction, reframing responsibility as temporally distributed and morally cumulative rather than located solely with designers or deployers.

3.2.6. Irreversibility of ethical failure

> *"When the egg is broken, it cannot be restored. The environment does not reset. The girl's fall into the water marks a point of no return; the consequences unfold without repair or reversal."*

Once the egg is broken, no corrective action can restore the prior ethical condition. This scene foregrounds the irreversibility of certain failures and the temporal fragility of care: when care fails, its consequences persist into futures that cannot be undone. The tension at stake is therefore not about sudden catastrophe, acute negligence, or singular moments of collapse, but about responsibility sustained over time in systems where failure would permanently alter environmental, social, or institutional conditions. It is not concerned with assigning blame after disaster, but with the ethical weight of maintaining care in contexts where margins for error progressively shrink. Unlike fragility without exit, which concerns situations in which vulnerable systems cannot be safely withdrawn, irreversibility of ethical failure concerns the temporal consequences of failure itself, where certain harms would permanently alter future conditions of safety and cannot be undone once they occur.

An example of this tension can be found in the Netherlands' long-term dependence on its primary flood-defence system. For a country in which large areas lie below sea level, failure of dikes or storm-surge barriers would not constitute a repairable incident but an irreversible state change involving large-scale inundation, saltwater intrusion, soil subsidence, and long-term ecological and societal disruption. This illustrates the tension of irreversibility because such failure would permanently alter environmental and societal conditions in ways that cannot be restored through subsequent intervention. Dutch water authorities manage an inherited infrastructure developed under historical climatic assumptions, while sea-level rise and land subsidence increasingly constrain safety margins. National assessments explicitly recognise that existing protection strategies have finite horizons and must nevertheless be sustained with great care until alternative approaches can be developed and implemented (Friocourt et al., 2025). Ethical care after deployment, in this context, emphasises that responsibility is not only about preventing failure, but about recognising that the costs of failure would be borne in futures that cannot be repaired, only endured.

3.3. Interactions among ethical tensions

The ethical tensions articulated in the preceding subsections are analytically distinct, yet they do not arise in isolation. In practice, they are closely related and tend to co-occur in long-lived safety–critical technologies, where responsibility persists even as design control, organisational authority, and realistic withdrawal options erode. Ethical care after deployment is therefore best understood not as a series of discrete dilemmas, but as a configurational condition in which multiple burdens of responsibility interact and compound over time. Fig. 1 schematically represents this configuration by showing how different tensions cluster and interact under structural and boundary conditions, rather than as a set of parallel categories.

A first set of interactions concerns **limits of agency**, most clearly expressed through the conjunction of *care without repair* and *responsibility without authority*. As technologies age and become embedded within broader infrastructures, meaningful redesign or improvement may no longer be feasible. At the same time, responsibility for safety outcomes is frequently assigned to actors who lack authority over system modification, procurement, or decommissioning. These conditions reinforce one another: the inability to repair heightens the ethical strain of being held responsible without decision-making power, while constrained authority narrows the scope of care to preservation, monitoring, and containment rather than transformation. Ethical responsibility thus persists even as the practical capacity to act on that responsibility progressively diminishes.

A second cluster of interactions foregrounds **temporal and epistemic burdens**, linking *inherited burden* with *care under epistemic uncertainty*. Responsibility for safety–critical systems is often inherited across generations of practitioners, organisations, and regulators who neither authored the original design nor fully control its ongoing evolution. This inherited responsibility is frequently accompanied by epistemic uncertainty that cannot be resolved through additional analysis or data, but remains a structural feature of complex sociotechnical systems. When responsibility is inherited rather than chosen, and uncertainty cannot be closed, care becomes oriented toward stewardship under incomplete knowledge rather than informed optimisation. In this way, temporal distance and epistemic limits jointly intensify the ethical weight of post-deployment care.

These interacting tensions are sustained by **structural conditions** that render ethical care unavoidable. *Fragility without exit* captures situations in which continued operation of a safety–critical system is ethically required precisely because withdrawal would introduce greater harm, risk displacement, or social disruption. Here, care persists not because systems are robust or satisfactory, but because abandonment would generate unacceptable consequences. This condition binds together the limits of agency and temporal-epistemic burdens: when repair is infeasible, authority is constrained, and responsibility is inherited under uncertainty, exit may also be foreclosed, transforming care from a strategic choice into an ethical obligation. Finally, all of

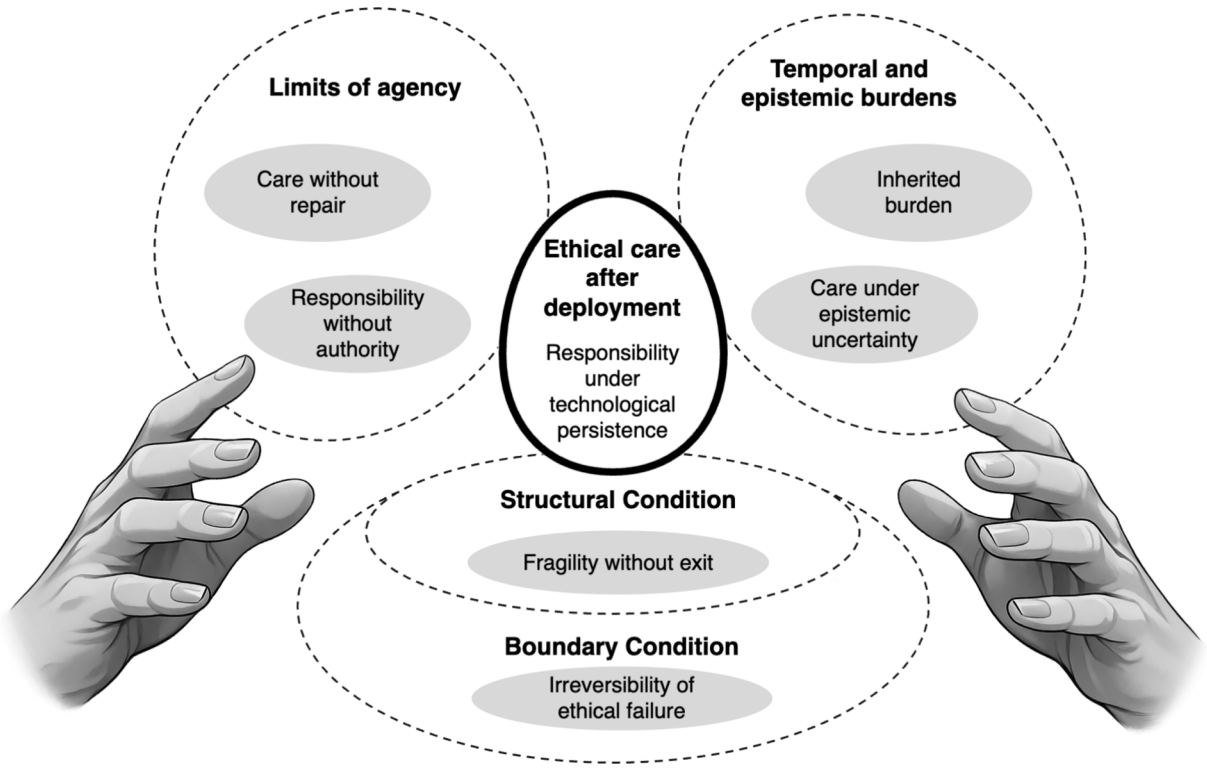

Fig. 1. The CARE-TECH framework: Ethical care after deployment in safety–critical technologies.

these interactions unfold under a **boundary condition**: the *irreversibility of ethical failure*. The recognition that certain failures would permanently alter environmental, social, or institutional futures intensifies the moral significance of care under constraint. Responsibility is sustained not because safe outcomes can be guaranteed, but because the costs of failure would be borne by future users, communities, or ecosystems in ways that cannot be repaired or reversed. Irreversibility thus does not constitute an additional tension, but an ethical horizon that amplifies and stabilises all others.

Taken together, these interactions suggest that ethical care after deployment emerges from the convergence of two clusters of tensions—limits of agency and temporal–epistemic burdens—sustained by structural conditions of persistence and unfolding under the boundary condition of the irreversibility of ethical failure. Ethical responsibility in such contexts cannot be reduced to individual virtues, organisational practices, or governance mechanisms alone. Rather, it is experienced as a fragile and ongoing practice of sustaining safety when control, repair, certainty, and exit are progressively limited. The CARE-TECH framework synthesises these relations, offering a structured way of recognising how ethical tensions cluster, reinforce one another, and shape the lived experience of responsibility in long-lived safety–critical systems.

4. Limits of the heuristic and conceptual scope

While *Angel's Egg* provides a powerful heuristic for articulating ethical care under persistence, it also introduces important limitations. Most notably, the film foregrounds a single fragile object sustained by a single caretaker, encouraging an analytical focus on bounded responsibility and localised care. Safety-critical technologies, by contrast, are rarely singular or isolated. They are typically embedded in distributed, networked, and highly coupled systems, where risks emerge from interactions across components, organisations, and temporal scales. As a result, the tensions articulated in this paper translate most directly to situations of component-level stewardship or frontline operational care, and less cleanly to systemic dynamics such as cascading failures, cross-sectoral dependencies, or collective coordination breakdowns. Accordingly, the CARE-TECH framework is intended primarily to illuminate ethical conditions associated with long-lived safety–critical sociotechnical systems whose persistence constrains opportunities for redesign, repair, or withdrawal. The logic of ethical care holds most clearly where responsibility can be experienced as proximate and embodied; its extension to fully networked sociotechnical systems therefore requires caution rather than direct transposition.

A second limitation concerns the asymmetry foregrounded throughout the framework. By emphasising ethical burdens borne by operators, maintainers, and local actors who sustain care without authority, exit, or repair, the analysis risks underplaying upstream forms of responsibility. Designers, regulators, procurers, and organisational leaders remain ethically implicated in the persistence of safety–critical systems, even when responsibility is experienced most acutely at the operational level. The film's narrative structure—centred on a single figure carrying an inherited burden—offers little scope for representing distributed accountability or institutional decision-making. Without explicit clarification, this asymmetry could be misread as relocating responsibility downward rather than analysing how ethical burden is unevenly experienced within broader responsibility relations. The framework should therefore be read as complementary to, rather than substitutive of, established accounts of upstream accountability in safety and responsible innovation.

Finally, some of the tensions articulated—particularly the notion that protection itself can accumulate risk—require careful handling to avoid conceptual overreach. In uncertain worlds, almost any act of preservation can be reframed as a source of future harm, raising the risk of tautology. The intention here is not to claim that all care is ethically suspect, but to highlight a specific post-deployment condition in which continued preservation forecloses alternative futures and amplifies long-term vulnerability. Moreover, the framework remains intentionally conceptual and internally coherent rather than empirically validated. While the tensions resonate with well-documented historical experiences in safety–critical domains, they are not derived from case studies

and do not claim empirical generalisation. Their value lies in offering a language for recognising and discussing ethical responsibility under persistence, rather than in adjudicating particular technological decisions. As such, the framework invites empirical engagement and refinement, rather than presenting itself as a complete or exhaustive account.

5. Discussion

5.1. Implications for safety Science

The CARE-TECH framework has implications for how Safety Science conceptualises responsibility, care, and ethical burden in long-lived safety–critical systems. It shifts ethical attention away from moments of design, deployment, or acute failure and toward the prolonged operational periods in which technologies persist, age, and become increasingly resistant to modification or withdrawal. In doing so, the framework complements established safety concepts—such as maintenance, resilience, and safety management—by foregrounding ethical responsibility as an ongoing and fragile practice rather than a condition that can be fully discharged through design assurance, procedural compliance, or organisational optimisation. Activities commonly treated as technical or managerial—monitoring, maintenance, compensatory action, and vigilance—are thus reinterpreted as forms of ethical care enacted under conditions of constrained agency, limited authority, and incomplete knowledge.

The framework also reframes how responsibility is distributed and experienced across sociotechnical systems. Rather than treating responsibility as something that can be cleanly allocated, resolved, or exhausted, CARE-TECH emphasises its cumulative and inherited character in long-lived technological systems, particularly where exit options are limited and failures would irreversibly shape future patterns of safety and harm. This perspective does not relocate responsibility downward or substitute ethical concern for governance, regulation, or design accountability. Instead, it provides a conceptual vocabulary for examining how ethical burden is sustained and negotiated in practice, especially by actors closest to operation and maintenance. As such, the framework offers a lens for empirical inquiry into post-deployment safety contexts—directing attention not to whether systems meet optimal or idealised standards, but to how ethical care is maintained when control, repair, certainty, and withdrawal are progressively constrained.

5.2. Future research

While the CARE-TECH framework is intentionally conceptual, it generates a set of empirically tractable questions about how ethical care is enacted, experienced, and distributed in long-lived safety–critical systems. Rather than seeking to validate the framework as a predictive model or to operationalise its components as measurable variables, the framework is intended to function as an analytic lens for examining how responsibility persists when technologies remain in operation despite declining steerability, constrained authority, and limited exit options. The research questions presented in Table 1 follow directly from the ethical tensions articulated in the CARE-TECH framework. For each tension, the questions focus on three interrelated empirical dimensions: the reasoning and interpretations through which practitioners understand their responsibilities, the organisational practices through which ethical care is enacted in everyday safety work, and the institutional conditions that shape how ethical burden is distributed across actors and over time. In this way, the framework does not prescribe specific empirical methods, but provides a structured vocabulary for investigating how post-deployment responsibility is justified, negotiated, and sustained in practice. Empirical research guided by this perspective would be particularly valuable for illuminating how ethical burden becomes recognised, normalised, or obscured within routine safety practices, and how different organisational and governance arrangements shape the capacity of actors to sustain care under technological persistence and irreversible risk.

6. Conclusion

Safety-critical technologies increasingly persist beyond the horizons of effective design control, organisational authority, and realistic withdrawal, yet responsibility for their consequences does not dissipate with time. This paper has advanced a conceptual account of ethical care after deployment to address this under-theorised condition in Safety Science, using Angel's Egg as a heuristic device to render visible ethical experiences that are often difficult to articulate through technical and managerial vocabularies alone. The resulting CARE-TECH framework names a set of interrelated tensions—care without repair, responsibility without authority, care under epistemic uncertainty, fragility without exit,

Table 1
Empirical issues and research questions generated by the CARE-TECH framework.

CARE-TECH component	Empirical issue to be explored	Illustrative empirical research questions*
Care without repair	Sustaining safety in systems where meaningful redesign or improvement is infeasible	What organisational and professional rationales lead practitioners to prioritise preservation rather than system improvement? How does the recognition that repair is not feasible influence everyday safety practices, vigilance, and risk management strategies?
Responsibility without authority	Misalignment between accountability and decision-making power	How do frontline actors interpret and justify their responsibility when authority over system modification or withdrawal lies elsewhere? What effects does this responsibility–authority misalignment have on decision-making, professional judgement, and perceived ethical burden?
Care under epistemic uncertainty	Acting responsibly under conditions of persistent and irreducible uncertainty	How do practitioners reason about responsibility when key aspects of system behaviour or risk remain uncertain? How does persistent uncertainty shape precautionary practices, decision thresholds, and organisational attitudes toward risk?
Inherited burden	Cross-generational responsibility for ageing systems	How do organisations explain and legitimise inherited responsibility for technologies designed under earlier assumptions? What effects does this inherited burden have on knowledge transfer, operational conservatism, and the preservation of institutional memory?
Fragility without exit	Continued reliance on systems that cannot be safely withdrawn	How do practitioners justify continued reliance on systems recognised as fragile or limited? What organisational and ethical considerations sustain the decision to continue operating systems for which viable exit options are unavailable?
Irreversibility of ethical failure	Anticipation of permanent consequences shaping safety practices	How does awareness of potentially irreversible harm influence everyday safety decisions and organisational risk tolerance? How do practitioners adjust operational practices when failures could permanently alter environmental, societal, or infrastructural conditions?
CARE-TECH (integrative)	Accumulation and interaction of ethical burdens over time	How do multiple CARE-TECH tensions co-occur within specific safety–critical domains, and how do organisations manage the resulting accumulation of ethical burden? What organisational mechanisms amplify, redistribute, or mitigate the accumulation of ethical burden over the operational life of technological systems?

Notes: (*) The research questions are intended as exploratory, theory-informed prompts to guide empirical inquiry, rather than hypotheses for formal testing.

inherited burden, and the irreversibility of ethical failure—and shows how these tensions cluster into limits of agency and temporal–epistemic burdens, sustained by structural conditions of persistence and intensified under the boundary condition of the irreversibility of ethical failure. By foregrounding care as a sustained practice of attending to fragile and consequential systems under constraint, the framework complements established safety concepts without substituting for governance, regulation, or upstream accountability, and provides a vocabulary for examining how ethical burden is experienced and negotiated throughout the life of safety–critical systems once operationally embedded. Future empirical research can use CARE-TECH as an analytic lens to investigate how these tensions co-occur, how responsibility is distributed and normalised in practice, and how institutional arrangements shape the capacity to sustain ethical care when control, repair, certainty, and exit are progressively constrained.

CRediT authorship contribution statement

Oscar Oviedo-Trespalacios: Writing – review & editing, Writing – original draft, Visualization, Validation, Supervision, Software, Resources, Project administration, Investigation, Funding acquisition, Formal analysis, Data curation, Conceptualization.

Declaration of competing interest

The author declare that he has no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

No data was used for the research described in the article.

References

Aoun, J., Goverde, R.M.P., Nardone, R., Quaglietta, E., Vittorini, V., 2024. Analysis of safe and effective next-generation rail signalling systems. Transp. Res. Part C Emerging Technol. 162, 104573. https://doi.org/10.1016/j.trc.2024.104573.

Baas, L., Metselaar, S., Klaassen, P., 2022. Circles of care for safety: a care ethics approach to safe-by-design. NanoEthics 16 (2), 167–179. https://doi.org/10.1007/s11569-022-00419-w.

Boppana, S., Shen, J., Schwebel, D.C., 2016. Transportation-related safety behaviors in top-grossing children's movies from 2008 to 2013. J. Dev. Behav. Pediatr. 37 (4), 306–313. https://doi.org/10.1097/DBP.0000000000000273.

Bowen, J., 2000. The ethics of safety-critical systems. Commun. ACM 43 (4), 91–97. https://doi.org/10.1145/332051.332078.

Doorn, N., Van de Poel, I., 2012. Editors' overview: Moral responsibility in technology and engineering. Sci. Eng. Ethics 18 (1), 1–11. https://link.springer.com/article/10.1007/s11948-011-9285-z.

Dekker, S.W., Breakey, H., 2016. 'Just culture:'Improving safety by achieving substantive, procedural and restorative justice. Saf. Sci. 85, 187–193. https://doi.org/10.1016/j.ssci.2016.01.018.

Friocourt, Y., Blaas, M., Bonte, M., Vos, R., Slomp, R., Wilmink, R., Lodder, Q., Brakenhoff, L., van Gool, S., 2025. The impact of extreme sea level rise on the national strategies for flood protection and freshwater in the Netherlands. Water 17 (7), 919. https://doi.org/10.3390/w17070919.

Grady, C., Rajtmajer, S., Dennis, L., 2021. When smart systems fail: the ethics of cyber–physical critical infrastructure risk. IEEE Trans. Technol. Soc. 2 (1), 6–14. https://doi.org/10.1109/TTS.2021.3058605.

Habli, I., Lawton, T., Porter, Z., 2020. Artificial intelligence in health care: Accountability and safety. Bull. World Health Organ. 98 (4), 251–256. https://doi.org/10.2471/BLT.19.237487.

Hansler, R.J., Bellamy, L.J., Akkermans, H.A., 2022. Ageing assets at major hazard chemical sites–the dutch experience. Saf. Sci. 153, 105788. https://doi.org/10.1016/j.ssci.2022.105788.

Komljenovic, D., Gaha, M., Abdul-Nour, G., Langheit, C., Bourgeois, M., 2016. Risks of extreme and rare events in Asset Management. Saf. Sci. 88, 129–145. https://doi.org/10.1016/j.ssci.2016.05.004.

Little, R.G., 2012. Managing the risk of aging infrastructure. In *Public sector governance of emerging risks: Infrastructure case*. Geneva: International Risk Governance Council (IRGC). Prepared for the IRGC Project on Public Sector Governance of Emerging Risks.

McCall, J.R., Pruchnicki, S., 2017. Just culture: a case study of accountability relationship boundaries influence on safety in HIGH-consequence industries. Saf. Sci. 94, 143–151. https://doi.org/10.1016/j.ssci.2017.01.008.

National Aeronautics and Space Administration. (2024, April 4). Engineers pinpoint cause of Voyager 1 issue, are working on solution. NASA Science. https://science.nasa.gov/blogs/voyager/2023/12/12/engineers-working-to-resolve-issue-with-voyager-1-computer-2/.

Office of Rail and Road. (2021). *Signalling market study: Final report.* Office of Rail and Road. https://www.orr.gov.uk/sites/default/files/2021-12/signalling-market-study-final-report.pdf.

Oviedo-Trespalacios, O., Bouchaut, B., van Nunen, K., Afghari, A.P., Yang, M., Yarveisy, R., Reniers, G., 2025. Safety, Security, and the Unintended Implications of cutting the Academic Ties with the Fossil fuel Industry. Environmental Development 101264. https://doi.org/10.1016/j.envdev.2025.101264.

Puig de la Bellacasa, M., 2011. Matters of care in technoscience: Assembling neglected things. Soc. Stud. Sci. 41 (1), 85–106. https://doi.org/10.1177/0306312710380301.

Ryan, P., Porter, Z., Al-Qaddoumi, J., McDermid, J., Habli, I., 2023. What's my role? Modelling Responsibility for AI-Based Safety-Critical Systems. https://doi.org/10.48550/arXiv.2401.09459.

Santoni de Sio, F., Van den Hoven, J., 2018. Meaningful human control over autonomous systems: a philosophical account. Front. Rob. AI 5, 323836. https://doi.org/10.3389/frobt.2018.00015.

Santoni de Sio, F., Mecacci, G., 2021. Four responsibility gaps with artificial intelligence: why they matter and how to address them. Philosophy & Technology 34 (4), 1057–1084. https://doi.org/10.1007/s13347-021-00450-x.

Scientific Advisory Board on Deep-Sea Mining. (2024). *Deep-sea mining: What it is, why it matters* (Brief). United Nations system. https://www.un.org/scientific-advisory-board/sites/default/files/2025-05/250403-DSM-Brief-Rev-7.pdf.

Shin, H., Noh, J., Kim, D., Kim, Y., 2020. The system that cried wolf: Sensor security analysis of wide-area smoke detectors for critical infrastructure. ACM Trans. Privacy Secur. 23 (3).

Singh, P.A., Jaeckel, A., Ardron, J.A., 2025. A pause or moratorium for deep seabed mining in the Area? the legal basis, potential pathways, and possible policy implications. Ocean Dev. Int. Law 56 (1), 18–44. https://doi.org/10.1080/00908320.2024.2439877.

Tongren, J.E., Sites, A., Zwicker, K., Pelletier, A., 2010. Injury-prevention practices as depicted in G- and PG-rated movies, 2003–2007. Pediatrics 125 (2), 290–294. https://doi.org/10.1542/peds.2009-1282.

U.S. Government Accountability Office. (2016). *Information technology: Federal agencies need to address aging legacy systems* (GAO-16-468). https://www.gao.gov/products/gao-16-468.

Vallor, S., Vierkant, T., 2024. Find the gap: AI, responsible agency and vulnerability. Minds and Machines 34 (3), 20. https://doi.org/10.1007/s11023-024-09674-0.

van der Bruggen, K., 2008. Critical infrastructures and responsibility: a conceptual exploration. Saf. Sci. 46 (7), 1137–1148. https://doi.org/10.1016/j.ssci.2007.06.003.

Walker, G., Salmon, P. M., & Stanton, N. A. (2015). What the Death Star can tell us about system safety. In D. Harris (Ed.), *Engineering psychology and cognitive ergonomics* (EPCE 2015) (Lecture Notes in Computer Science, Vol. 9174, pp. 297–306). Springer. Doi:10.1007/978-3-319-20373-7_28.