

STRICT LIABILITY OR FAULT-BASED REGIMES FOR AI-CAUSED HARM? A DOCTRINAL ANALYSIS ACROSS COMMON LAW AND CIVIL LAW SYSTEMS

Reshma Leslie

Department of Business Administration
Westcliff University
17877 Von Karman Ave, 4th Floor, Irvine, CA 92614
ORCID: 0009-0008-8915-6189
r.leslie.300@westcliff.edu

Received: 22 October 2024

Revised: 28 November 2024

Accepted: 26 December 2024

ABSTRACT:

The rapid deployment of artificial intelligence systems across critical sectors has created unprecedented challenges for traditional tort law frameworks. When autonomous vehicles cause accidents, medical diagnosis algorithms produce errors, or automated decision systems generate discriminatory outcomes, fundamental questions arise about liability attribution. This research examines whether strict liability or fault-based regimes better address AI-caused harm through comparative doctrinal analysis across common law and civil law jurisdictions. The study analyzes existing liability frameworks in the United States, United Kingdom, Germany, France, and Japan, evaluating how each system's principles apply to AI-specific scenarios. Through examination of recent case law, legislative proposals, and doctrinal scholarship, the research identifies critical gaps in current frameworks where AI's autonomous decision-making, opacity, and distributed responsibility challenge traditional causation and fault concepts. Findings reveal that pure fault-based approaches struggle with AI's "black box" problem and difficulty establishing negligence standards, while strict liability faces challenges in defining "defect" for learning systems and determining liable parties in complex AI supply chains. The research proposes a hybrid model incorporating elements from both regimes, with strict liability for high-risk AI applications and modified fault standards for lower-risk contexts. This framework addresses the unique characteristics of AI while maintaining accountability and encouraging responsible innovation across different legal traditions.

Keywords: *Artificial intelligence liability, strict liability, negligence, tort law, AI regulation, product liability, common law, civil law.*

INTRODUCTION

Artificial intelligence has moved from theoretical possibility to practical reality with remarkable speed. Autonomous vehicles navigate public roads, AI systems diagnose diseases and recommend treatments, algorithms make lending and hiring decisions, and automated trading systems execute millions of financial transactions. This technological transformation brings enormous benefits but also creates novel risks. When these systems cause harm—whether through algorithmic bias, technical malfunction, or unexpected behavior—legal systems face fundamental questions about responsibility and compensation.

Traditional tort law developed around human actors whose intentions, knowledge, and conduct could be evaluated under established negligence or strict liability principles. A driver's failure to exercise reasonable care, a manufacturer's defective product, or a property owner's dangerous premises all fit within frameworks refined over centuries of jurisprudence. However, AI systems disrupt these familiar categories. They make autonomous decisions without human intervention at the moment of harm. Their decision-making processes often remain opaque even to their creators. Responsibility fragments across developers, deployers, users, and data providers. Causation becomes difficult to trace through complex algorithmic processes.

These characteristics challenge both major liability regimes in modern legal systems. Fault-based liability, predominant in common law negligence and civil law delict, requires demonstrating that a defendant breached a duty of care through failure to meet reasonable standards. Yet what constitutes reasonable care when deploying

AI systems whose behavior cannot be fully predicted? How do courts establish breach when algorithms operate as "black boxes" whose decision processes remain inscrutable? Who bears responsibility when harm results from interactions between multiple AI components or from machine learning systems that evolve beyond their original programming?

Strict liability, which imposes responsibility regardless of fault for abnormally dangerous activities or defective products, faces different challenges. Product liability frameworks require showing that goods contained defects making them unreasonably dangerous. But how do courts define "defect" for software that performs as designed yet produces harmful outcomes? When do machine learning systems that continuously adapt cross the line from functioning properly to becoming defective? Who in the complex AI development and deployment chain should bear strict liability?

These questions have generated intense scholarly debate and growing legislative attention across jurisdictions. The European Union has proposed AI-specific liability directives. Individual nations have crafted varying approaches reflecting their legal traditions and policy priorities. Courts have begun confronting AI liability cases, though judicial guidance remains limited and often inconsistent. Meanwhile, the technology continues advancing rapidly, with increasingly autonomous and sophisticated systems entering high-stakes domains like healthcare, transportation, and criminal justice.

This research contributes to resolving these challenges through systematic comparative analysis of how different legal systems might apply their liability regimes to AI-caused harm. By examining both common law and civil law approaches, the study identifies principles that transcend particular jurisdictions and highlights culturally specific elements that affect liability determinations. The analysis considers not only doctrinal compatibility but also policy objectives including victim compensation, deterrence of unsafe practices, innovation incentives, and fair distribution of risks and costs.

The following sections proceed as follows: Section 2 establishes the research objectives and scope. Section 3 reviews relevant legal literature and existing frameworks. Section 4 describes the comparative methodology. Section 5 analyzes fault-based approaches. Section 6 examines strict liability frameworks. Section 7 discusses hybrid models and proposes recommendations. Section 8 concludes with implications for legal reform.

OBJECTIVES

This research pursues the following specific objectives:

- **Primary Objective:** To evaluate whether strict liability or fault-based regimes better address AI-caused harm through comparative doctrinal analysis across common law and civil law systems.
- **Secondary Objective 1:** To identify critical gaps in existing liability frameworks when applied to AI-specific scenarios involving autonomy, opacity, and distributed responsibility.
- **Secondary Objective 2:** To compare how different jurisdictions—United States, United Kingdom, Germany, France, and Japan—approach AI liability within their legal traditions.
- **Secondary Objective 3:** To assess the advantages and limitations of fault-based versus strict liability regimes in achieving policy goals of compensation, deterrence, and innovation.
- **Secondary Objective 4:** To propose a hybrid liability framework that addresses AI's unique characteristics while remaining doctrinally coherent across legal systems.

SCOPE OF STUDY

This research operates within defined boundaries:

- **Jurisdictional Scope:** Analysis focuses on five representative legal systems—United States and United Kingdom (common law), Germany and France (civil law), and Japan (mixed system).
- **Doctrinal Scope:** Examination covers negligence/delict and strict liability principles, with particular attention to product liability frameworks applicable to AI.
- **AI Applications:** Focus on AI systems causing physical harm (autonomous vehicles, medical devices) and economic harm (algorithmic decision-making), excluding pure data protection issues.
- **Temporal Scope:** Analysis considers legal frameworks as of 2024, including recent legislative proposals and emerging case law.

- **Methodological Scope:** Doctrinal analysis of statutory provisions, case law, and legal scholarship; empirical data on AI incidents provides context but is not primary focus.
- **Excluded Aspects:** Criminal liability, constitutional law issues, and detailed technical AI specifications are beyond this study's scope.

LITERATURE REVIEW

4.1 Traditional Liability Frameworks

Modern tort law balances two primary liability regimes, each reflecting different normative commitments and practical considerations. Fault-based liability, tracing to Roman law concepts of culpa and refined through common law negligence doctrine, holds actors responsible only when they breach duties of reasonable care. This regime emphasizes individual responsibility tied to blameworthiness, with liability flowing from failure to meet community standards of prudent conduct (Dobbs et al., 2016).

The fault principle serves multiple functions. It provides moral justification for liability by linking responsibility to culpable behavior. It generates deterrence incentives by imposing costs on careless actors. It preserves individual liberty by permitting activities conducted with reasonable care. However, fault-based systems place significant burdens on injured parties to prove breach and causation, potentially leaving victims uncompensated when evidence proves elusive or expensive to obtain.

Strict liability emerged as an exception to the fault principle, initially for abnormally dangerous activities and subsequently for defective products. Rather than requiring proof of negligence, strict liability imposes responsibility based on engaging in particular activities or placing particular products into commerce. Justifications include compensating victims, encouraging safety investments beyond what negligence standards require, and fairly allocating risks to those who create them and profit from risky activities (Owen, 2015).

Product liability doctrine, particularly influential in the United States following *Restatement (Second) of Torts* Section 402A, revolutionized consumer protection by eliminating negligence and privity requirements. Manufacturers became strictly liable for defective products unreasonably dangerous to users. This framework reflected judgments about superior risk-bearing capacity, information asymmetries between producers and consumers, and the social utility of compensating injured parties without protracted litigation over negligence (Abbott, 2020).

4.2 AI's Challenge to Traditional Frameworks

Artificial intelligence disrupts fundamental assumptions underlying both liability regimes. Traditional negligence analysis presumes that reasonable care can be defined through industry customs, professional standards, or community expectations. Yet AI applications often lack established standards, involve cutting-edge technology outpacing professional consensus, and operate in domains where even experts disagree about appropriate safeguards (Lior, 2020).

The foreseeability requirement central to negligence faces particular difficulty with machine learning systems. These systems identify patterns and generate decisions through processes that surprise even their creators. A machine learning algorithm trained on medical imaging data might base diagnoses on subtle correlations invisible to human physicians. When such systems produce erroneous decisions causing harm, was that harm reasonably foreseeable to developers who could not fully predict their systems' reasoning processes?

Causation presents another challenge. Negligence requires showing that the defendant's breach actually caused the plaintiff's injury. With AI, causation chains become complex and opaque. An autonomous vehicle accident might result from sensor misinterpretation, algorithmic decision-making, inadequate training data, software errors, or interactions between multiple systems. Tracing causation through these layers to identify whose breach caused the harm strains traditional but-for causation tests (Scherer, 2016).

Strict product liability confronts different but equally significant obstacles when applied to AI. The "defect" concept—whether manufacturing defect, design defect, or warning defect—presumes relatively static products whose characteristics can be evaluated at the time of sale. Machine learning systems that continuously adapt challenge this assumption. If an algorithm performs as designed when sold but later produces harmful decisions

after learning from new data, did the product contain a defect? At what point in time should courts evaluate defectiveness?

The distributed nature of AI development further complicates liability attribution. Modern AI systems typically involve multiple parties: data providers, algorithm developers, platform providers, system integrators, and deployers. When harm occurs, which party or parties should bear liability? Traditional product liability targets manufacturers, but who manufactures AI when various parties contribute different components? Some systems integrate pre-trained models from one source, fine-tune them with data from another source, and deploy them through third-party platforms (Lemley and Casey, 2019).

4.3 Comparative Legal Approaches

Different jurisdictions have begun grappling with AI liability through their respective legal traditions. The European Union has taken the most comprehensive approach, proposing both an AI Act establishing risk-based regulatory requirements and an AI Liability Directive updating liability rules. The proposed directive creates rebuttable presumptions of causation and fault for high-risk AI systems, easing plaintiffs' evidentiary burdens while maintaining fault-based liability's basic structure (European Commission, 2022).

Germany's approach reflects its civil law tradition with strong product liability provisions under the *Produkthaftungsgesetz*, which implements the EU Product Liability Directive. German legal scholars have debated whether existing frameworks suffice or whether AI-specific legislation is necessary. Some argue that general delict principles under the *Bürgerliches Gesetzbuch* (BGB) can adapt to AI through evolutionary interpretation, while others advocate for statutory reforms explicitly addressing AI's unique characteristics (Wagner, 2020).

France similarly works within civil law frameworks, applying *responsabilité civile* principles from the Code Civil. French doctrine distinguishes between fault-based liability (*responsabilité pour faute*) and liability for things under one's control (*responsabilité du fait des choses*), which operates similarly to strict liability. The latter doctrine's scope—whether it encompasses autonomous AI systems—has generated significant doctrinal debate, with some scholars arguing that Article 1242 already provides a framework for AI liability (Bensoussan, 2021).

Common law jurisdictions face different doctrinal challenges. In the United States, product liability law varies by state, with most following either the Second or Third Restatement approaches. The Third Restatement's risk-utility balancing for design defects provides some flexibility for evaluating AI systems, though questions remain about how courts will apply these standards to algorithms. Negligence claims face hurdles establishing appropriate care standards and proving breach through expert testimony about algorithmic behavior (Balkin, 2017).

The United Kingdom maintains distinct approaches to negligence and product liability under the Consumer Protection Act 1987, which implements the EU directive. Brexit raises questions about future divergence from EU approaches, though initial signals suggest continued alignment on product liability frameworks. UK courts have limited experience with AI liability cases, leaving considerable uncertainty about how traditional principles will apply (Kingston, 2018).

Japan presents an interesting mixed system combining civil law foundations with post-war American influences. Japanese product liability law follows the EU directive model while tort law retains civil law characteristics under the Civil Code. Japan has invested heavily in AI development and robotics, spurring policy discussions about appropriate liability frameworks. Proposals have ranged from maintaining existing frameworks with interpretive adjustments to creating entirely new AI-specific liability regimes (Yamamoto, 2019).

4.4 Policy Considerations

Beyond doctrinal compatibility, liability regime selection implicates important policy objectives. Victim compensation represents a primary goal—injured parties deserve redress regardless of whether defendants acted negligently. Strict liability arguably serves this goal better by eliminating the need to prove fault, though challenges remain in establishing causation and identifying responsible parties.

Deterrence provides another crucial function. Liability systems should incentivize appropriate safety investments and discourage unreasonably risky conduct. However, excessive liability might deter beneficial innovation, particularly for emerging technologies where risks and benefits remain uncertain. The optimal liability regime

balances these considerations, imposing costs on harmful behavior while preserving space for experimentation and development (Scherer, 2016).

Risk allocation raises questions about who should bear losses from AI-caused harm. Arguments for strict liability often emphasize that developers and deployers, not victims, create risks and profit from AI technologies. They possess superior capacity to distribute losses through insurance or pricing. Conversely, fault-based approaches maintain that liability should follow culpability, with actors bearing responsibility only for failure to exercise reasonable care.

Insurance availability affects practical liability operation. Traditional liability insurance models assume relatively predictable loss distributions based on historical experience. AI's novelty complicates actuarial assessment of risks, potentially limiting insurance availability or driving premiums prohibitively high. Some argue that liability frameworks should account for insurability, though others contend that insurance markets will adapt as AI risks become better understood (Abraham and Rabin, 2019).

4.5 Research Gaps

Despite growing scholarly attention, significant gaps remain in understanding optimal AI liability frameworks. Most existing research focuses on specific jurisdictions or particular AI applications rather than systematic comparative analysis. Limited empirical evidence exists about how different liability regimes would affect AI development, deployment, and safety in practice.

Furthermore, much scholarship addresses high-profile applications like autonomous vehicles while neglecting other AI contexts with different risk profiles and stakeholder interests. The interaction between liability rules and other regulatory mechanisms—safety standards, certification requirements, disclosure obligations—deserves greater attention. Finally, the global nature of AI development and deployment requires coordination across jurisdictions, yet little research examines mechanisms for harmonizing liability approaches internationally.

RESEARCH METHODOLOGY

5.1 Comparative Legal Analysis

This research employs comparative legal methodology examining AI liability frameworks across five representative jurisdictions. The selection balances geographic diversity, representation of major legal traditions, and economic significance in AI development. The United States and United Kingdom exemplify common law approaches with case law-driven evolution and jury trials. Germany and France represent civil law systems with comprehensive codes and scholarly interpretation. Japan provides a mixed system combining elements from both traditions.

5.2 Doctrinal Framework

The analysis examines statutory provisions, judicial decisions, and legal scholarship addressing liability for product defects, negligence, and related torts. Key sources include:

United States: Restatement (Third) of Torts: Products Liability, state product liability statutes, negligence case law, and emerging autonomous vehicle legislation.

United Kingdom: Consumer Protection Act 1987, common law negligence principles, recent Law Commission reports on automated vehicles.

Germany: Produkthaftungsgesetz (Product Liability Act), BGB provisions on delict, scholarly commentaries on AI liability.

France: Code Civil Articles 1240-1242 on civil responsibility, Product Liability Law implementing EU directive, doctrinal analyses.

Japan: Civil Code tort provisions, Product Liability Law, Ministry of Economy, Trade and Industry policy documents on AI governance.

5.3 Analytical Approach

The research evaluates each jurisdiction's frameworks against AI-specific scenarios involving:

Autonomous vehicles: Accidents caused by self-driving car decision-making or sensor failures.

Medical AI: Diagnostic errors or treatment recommendations from clinical decision support systems.

Algorithmic decisions: Discriminatory or erroneous outcomes from hiring, lending, or criminal justice algorithms.

For each scenario, the analysis considers how existing liability doctrines would apply, identifies doctrinal gaps or ambiguities, and evaluates whether fault-based or strict liability approaches better serve policy objectives.

5.4 Comparative Evaluation Criteria

Assessment employs multiple criteria derived from tort law theory and policy analysis:

Doctrinal coherence: Does the liability approach fit logically within the jurisdiction's existing legal framework?

Practical feasibility: Can plaintiffs reasonably satisfy evidentiary requirements? Are defendants identifiable?

Policy effectiveness: Does the regime achieve compensation, deterrence, and innovation goals?

Fairness: Does liability allocation reflect normative principles of responsibility and risk distribution?

[TABLE 1: Comparative Overview of AI Liability Frameworks]

Jurisdiction	Primary Legal Tradition	Fault-Based Provisions	Strict Liability Provisions	AI-Specific Legislation
United States	Common Law	Negligence (state law)	Product Liability (Restatement)	Limited (autonomous vehicles in some states)
United Kingdom	Common Law	Negligence (common law)	Consumer Protection Act 1987	Automated Vehicles Act 2018
Germany	Civil Law	BGB §823 (delict)	Produkthaftungsgesetz	Proposed amendments pending
France	Civil Law	Code Civil Art. 1240	Code Civil Art. 1242, Product Liability Law	Under discussion
Japan	Mixed	Civil Code Art. 709	Product Liability Law	AI governance guidelines (non-binding)

ANALYSIS OF FAULT-BASED APPROACHES

6.1 Application to AI Contexts

Fault-based liability remains the default regime in most jurisdictions for non-product-related harm. When AI systems cause injury, plaintiffs pursuing negligence claims must establish duty, breach, causation, and damages. Each element presents distinctive challenges in AI contexts.

Duty of care generally extends to foreseeable victims of one's conduct. AI developers and deployers owe duties to users and third parties foreseeably affected by their systems. However, defining the scope of duty becomes complicated with widely deployed AI systems whose impacts ripple across society. Does a social media company owe duties to individuals harmed by misinformation its recommendation algorithm amplified? Courts in different jurisdictions have reached varying conclusions about such questions (Citron and Pasquale, 2014).

Breach analysis requires establishing standards of reasonable care and showing defendants fell short. Professional standards, industry customs, and regulatory requirements typically inform negligence assessments. Yet AI technology often outpaces standard development, leaving courts without clear benchmarks. When cutting-edge machine learning techniques produce unexpected results, what level of testing, validation, and monitoring constitutes reasonable care? Some jurisdictions may require expert testimony, but disagreement among AI experts about appropriate practices complicates such evidence.

6.2 The Black Box Problem

AI's opacity creates particularly acute challenges for fault-based liability. Machine learning systems, especially deep neural networks, operate through complex mathematical transformations whose intermediate steps resist human interpretation. Even developers cannot always explain why their systems produce specific outputs for given inputs. This "black box" characteristic undermines negligence claims requiring plaintiffs to identify failures in design, development, or deployment processes (Burrell, 2016).

Consider medical AI systems that analyze diagnostic images. If such a system misdiagnoses a patient, leading to improper treatment and harm, a negligence plaintiff must show that the system's error resulted from unreasonable care in its development or use. But how can plaintiffs demonstrate this when the algorithm's decision-making

process remains inscrutable? Discovery might reveal training data, code, and testing results, but these may not illuminate why the specific misdiagnosis occurred.

Some scholars propose imposing duties of explainability, requiring AI developers to create interpretable systems or provide explanations for decisions. However, requiring explainability might prohibit the most powerful machine learning techniques, which achieve superior performance precisely through complex, opaque processing. This creates tension between facilitating liability claims and preserving AI's technological advantages (Selbst and Barocas, 2018).

6.3 Standards of Care

Establishing appropriate standards of care for AI development and deployment requires resolving difficult questions about feasible precautions and risk-benefit balancing. Should developers be held to standards of perfection given AI's potential for harm, or should reasonable care recognize technology's inherent limitations and uncertainty about emergent properties?

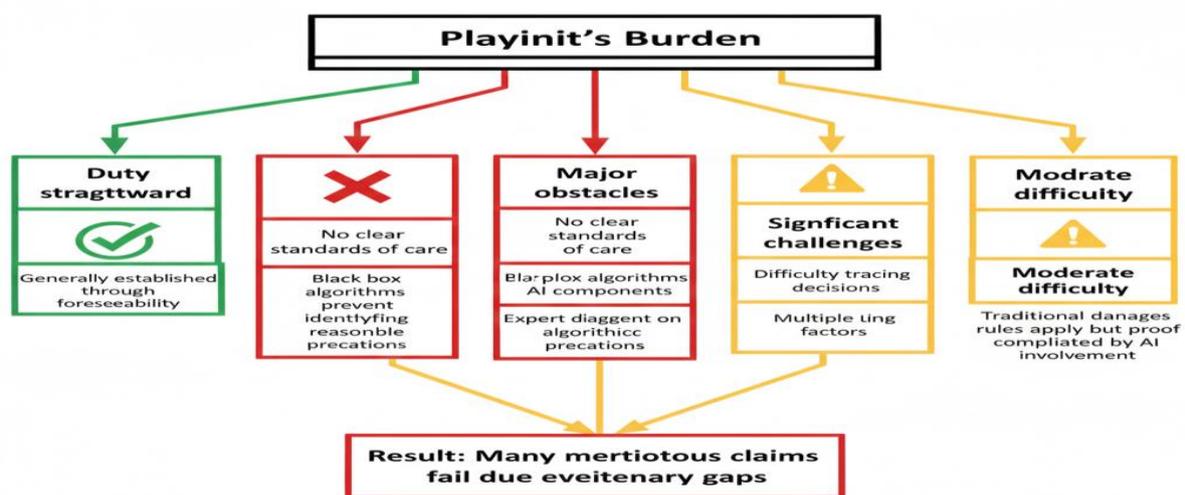
Some argue that AI safety research provides frameworks for defining reasonable care—developers should implement adversarial testing, formal verification where possible, and continuous monitoring systems. However, such techniques remain imperfect and resource-intensive, potentially suitable only for the highest-risk applications. Courts will need to determine whether all AI developers must employ cutting-edge safety techniques or whether reasonable care varies with application risk levels and developer resources (Amodei et al., 2016).

Professional standards are beginning to emerge through industry associations and government agencies, but gaps remain. The IEEE, ACM, and other organizations have issued AI ethics guidelines, though these focus more on fairness and transparency than technical safety measures. Regulatory standards exist for specific domains like medical devices, but general-purpose AI lacks comparable frameworks. This standards vacuum complicates negligence claims by depriving courts of clear benchmarks against which to measure defendants' conduct.

6.4 Advantages and Limitations

Fault-based approaches offer several advantages for AI liability. They maintain consistency with general tort principles, avoiding the need for AI-specific doctrinal frameworks. They preserve innovation incentives by not imposing liability absent unreasonable conduct. They provide flexibility to evolve standards as technology and understanding advance.

However, fault-based approaches face serious limitations in AI contexts. The evidentiary burden on plaintiffs—particularly proving breach when algorithmic processes remain opaque—may leave many injured parties without remedy. Information asymmetries favor defendants who possess relevant technical knowledge and documentation. Uncertainty about standards may chill beneficial innovation or conversely permit inadequate safety measures. The regime may prove inadequate for compensating victims when establishing fault proves prohibitively difficult despite genuine harm.



[FIGURE 1: Challenges in Proving Negligence for AI-Caused Harm]

ANALYSIS OF STRICT LIABILITY APPROACHES

7.1 Product Liability Frameworks

Strict product liability offers an alternative framework that eliminates the need to prove negligence. Under U.S. law following the Restatement approach, sellers of defective products unreasonably dangerous to users face liability regardless of care exercised. Similar frameworks exist across civil law jurisdictions through EU Product Liability Directive implementation. These regimes reflect judgments that manufacturers should bear costs of product defects given their superior position to prevent harm and distribute losses.

Applying product liability to AI requires determining whether software constitutes a "product" under applicable statutes. Early decisions addressing software treated it as a product when sold as a physical embodiment but sometimes as a service when delivered electronically. Modern courts increasingly recognize software as a product subject to liability, though questions remain about pure services utilizing AI (Leenes and Lucivero, 2014).

Assuming AI qualifies as a product, plaintiffs must still prove defect, causation, and damages. The defect element proves particularly challenging. Product liability recognizes three defect types: manufacturing defects (product deviates from design), design defects (design itself unreasonably dangerous), and warning defects (inadequate instructions or warnings). Each category presents difficulties when applied to AI.

7.2 Defining Defect for AI Systems

Manufacturing defects occur when individual products deviate from intended specifications. For physical goods, this concept operates straightforwardly—a bottle with a hidden crack or a vehicle missing a brake component clearly deviates from design. But software, including AI, typically gets reproduced identically across all copies. A bug affecting one unit affects all units. This suggests that software problems constitute design defects rather than manufacturing defects (Abbott, 2020).

Design defect analysis under the Third Restatement employs risk-utility balancing: a design is defective if foreseeable risks exceed benefits and a reasonable alternative design would reduce risks without substantially impairing utility. Applying this test to AI algorithms raises numerous questions. How do courts evaluate whether an algorithm's risk-benefit profile is reasonable? What alternative designs should be considered—less complex algorithms, different training approaches, or human-in-the-loop systems? Who bears the burden of proposing and demonstrating alternative designs?

Machine learning systems that adapt over time create particularly vexing problems. Consider an algorithm that performs safely when deployed but later exhibits harmful behavior after learning from new data. Did this system contain a design defect at sale, or did it become defective subsequently? If the latter, does product liability apply, or must plaintiffs pursue negligence claims about failure to monitor and update the system?

Warning defects might provide a path for some AI cases. Manufacturers face duties to warn about non-obvious risks their products pose. AI systems whose limitations or failure modes remain opaque to users could be deemed defective for inadequate warnings. However, warning duties typically extend only to risks foreseeable at the time of sale, potentially excluding emergent properties of learning systems (Lemley and Casey, 2019).

7.3 Causation and Identification Issues

Even when defects can be established, causation requirements may prove difficult to satisfy. Plaintiffs must show that the product defect actually caused their injuries. With AI involved in complex socio-technical systems, isolating the algorithm's contribution from other causal factors becomes challenging. An autonomous vehicle accident might result from algorithmic decisions, sensor failures, poor road conditions, other drivers' actions, or combinations of these factors.

Furthermore, the distributed nature of AI development complicates identifying proper defendants. Product liability typically targets manufacturers, but AI systems often involve multiple entities contributing different components. One company provides training data, another develops the base algorithm, a third fine-tunes it for specific applications, and a fourth deploys it to end users. When harm occurs, which party should face strict liability? Holding all jointly and severally liable might be unfair to minor contributors, while limiting liability to deployers could undercompensate victims when deployers lack deep pockets.

7.4 Advantages and Limitations

Strict liability approaches offer significant advantages for AI-caused harm. They ease victims' evidentiary burdens by eliminating the need to prove fault, addressing the opacity problem inherent in fault-based approaches. They incentivize safety investments beyond what negligence standards might require, encouraging developers to minimize defects rather than merely avoiding unreasonable conduct. They provide a predictable framework for insuring AI risks and distributing costs across users.

However, strict liability faces substantial limitations in AI contexts. Defining defect for continuously learning systems that operate as designed but produce harmful outcomes remains problematic. The regime may impose excessive costs on beneficial innovation, particularly for emerging technologies whose risks cannot be eliminated entirely. Identifying responsible parties in complex AI supply chains creates practical difficulties. Causation challenges persist even without requiring proof of fault. The regime risks generating disproportionate liability exposure that might drive AI development offshore or discourage deployment of socially valuable systems (Wagner, 2020).

[TABLE 2: Comparative Evaluation of Liability Regimes for AI]

Evaluation Criterion	Fault-Based Liability	Strict Liability	Hybrid Approach
Victim Compensation	Moderate (evidentiary burdens)	High (no fault requirement)	High (strict for high-risk, fault for low-risk)
Deterrence Effectiveness	Moderate (standards unclear)	High (incentivizes safety)	High (tailored to risk level)
Innovation Impact	Positive (permits reasonable risk)	Potentially negative (excessive costs)	Balanced (risk-proportionate)
Doctrinal Coherence	High (fits existing frameworks)	Moderate (requires adaptation)	Moderate (requires new categories)
Practical Feasibility	Low (proof challenges)	Moderate (defect definition issues)	Moderate (complexity in implementation)
Fairness	High (links liability to fault)	Moderate (strict without wrongdoing)	High (proportionate to control)

DISCUSSION AND HYBRID FRAMEWORK PROPOSAL

8.1 Limitations of Pure Approaches

The analysis reveals that neither pure fault-based nor pure strict liability adequately addresses AI-caused harm across all contexts. Fault-based approaches founder on evidentiary challenges stemming from AI opacity and underdeveloped standards. Strict liability struggles with defect definitions for learning systems and identifying responsible parties in complex development chains. These limitations suggest that hybrid approaches incorporating elements from both regimes may prove more effective.

8.2 Risk-Based Liability Framework

Drawing on comparative analysis across jurisdictions and evaluation against policy criteria, this research proposes a risk-based hybrid framework that applies different liability standards based on AI system risk profiles and deployment contexts.

High-Risk Applications: For AI systems deployed in contexts where failures pose significant risks of physical harm or fundamental rights violations—autonomous vehicles, medical diagnosis, critical infrastructure control—strict liability should apply with modifications addressing AI-specific challenges. Developers and deployers should face joint and several liability for defects, defined to include not only code errors but also inadequate testing, validation, and monitoring. This approach provides strong victim compensation and deterrence incentives proportionate to elevated risks.

Medium-Risk Applications: For AI systems with moderate risk profiles—fraud detection, credit decisioning, content recommendation—a modified fault-based regime should apply with procedural accommodations addressing evidentiary challenges. Rebuttable presumptions of negligence when AI systems cause harm could shift burdens to defendants to demonstrate reasonable care. Mandatory transparency requirements regarding

algorithmic functioning would facilitate plaintiffs' ability to identify potential breaches. This balanced approach maintains innovation incentives while ensuring accountability.

Low-Risk Applications: For AI systems with minimal harm potential—entertainment recommendations, non-critical optimization, routine automation—traditional fault-based liability should suffice. The existing negligence framework provides adequate accountability without imposing undue burdens on beneficial innovations with limited risk profiles.

8.3 Operational Elements

Several operational elements would support this risk-based framework across jurisdictions. First, regulatory agencies should establish risk classification systems categorizing AI applications based on potential harm severity and probability. The EU AI Act provides a model, though classifications should be refined based on emerging experience.

Second, mandatory insurance requirements for high-risk AI would ensure victim compensation even when individual developers or deployers lack sufficient resources. Insurance mechanisms could aggregate risks across deployers while using premiums to incentivize safety investments. Japan's experience with autonomous vehicle insurance provides relevant precedents (Yamamoto, 2019).

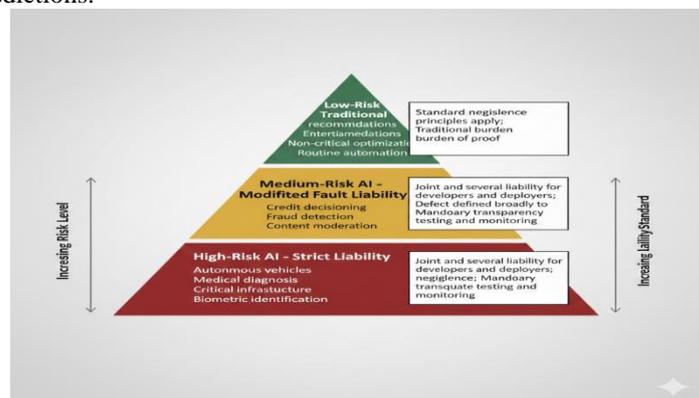
Third, substantive standards should be developed through multi-stakeholder processes involving technical experts, legal scholars, industry participants, and civil society representatives. These standards would inform both strict liability defect determinations and negligence breach analysis, providing courts with clearer frameworks for evaluating AI systems.

Fourth, discovery rules should be adapted to address information asymmetries in AI litigation. Defendants should face obligations to preserve and produce relevant algorithmic documentation, training data, and testing results. Special masters with technical expertise could assist courts in reviewing complex evidence and assessing AI systems' characteristics.

8.4 Comparative Implementation

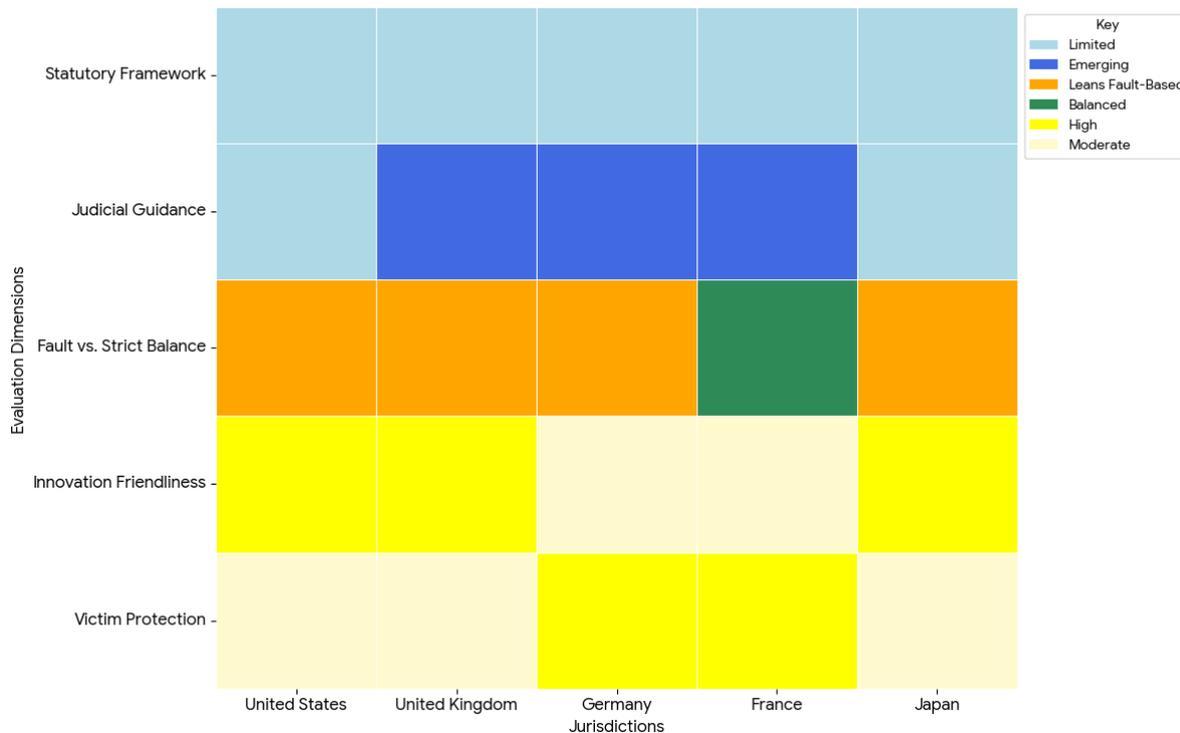
This hybrid framework could be implemented across different legal systems while respecting their distinct traditions. Common law jurisdictions could develop the framework through judicial evolution, with courts recognizing risk-based distinctions in applying product liability and negligence doctrines. Statutes could codify emerging principles while preserving common law flexibility.

Civil law jurisdictions could implement the framework through legislative reform, enacting AI liability statutes that establish risk categories and corresponding liability standards. Germany's existing framework of tiered liability regimes for different activity types provides a model. France's distinction between fault-based and thing-based liability could be extended to encompass AI systems under appropriate risk categories (Bensoussan, 2021). Both systems should pursue harmonization through international coordination, particularly given AI's global nature. The EU's efforts to establish comprehensive AI governance frameworks could provide a foundation for broader alignment, though flexibility for local adaptation remains important given cultural and institutional differences across jurisdictions.



[FIGURE 2: Proposed Risk-Based Hybrid Liability Framework]

FIGURE 3: Comparative Legal Approaches to AI Liability Across Jurisdictions



[FIGURE 3: Comparative Legal Approaches to AI Liability Across Jurisdictions]

CONCLUSION

The challenge of assigning legal responsibility for AI-caused harm represents one of the most pressing questions facing legal systems globally. This comparative analysis across common law and civil law jurisdictions reveals that traditional liability frameworks—whether fault-based negligence or strict product liability—struggle when applied to AI's unique characteristics of autonomy, opacity, and distributed development responsibility.

Fault-based approaches face serious limitations stemming from AI's black box nature, which prevents plaintiffs from identifying specific breaches of duty even when genuine harm occurs. The absence of clear standards for reasonable care in AI development compounds these difficulties, potentially leaving victims without remedy despite meritorious claims. While fault-based liability offers advantages in preserving innovation incentives and maintaining consistency with general tort principles, these benefits come at the cost of inadequate victim protection and uncertain deterrence effects.

Strict product liability provides stronger victim compensation by eliminating fault requirements, but confronts its own challenges in defining defect for machine learning systems that evolve beyond their original programming. Questions about when adaptive algorithms cross the line from proper function to defectiveness lack clear answers within existing frameworks. The regime also faces practical difficulties identifying responsible parties in complex AI supply chains involving multiple contributors.

The research findings support a hybrid approach that applies different liability standards based on AI system risk profiles. High-risk applications involving potential physical harm or fundamental rights violations warrant strict liability with modifications addressing AI-specific challenges. Medium-risk applications benefit from modified fault standards incorporating rebuttable negligence presumptions and mandatory transparency requirements. Low-risk applications can continue under traditional negligence frameworks without special accommodations.

This risk-based framework achieves better alignment between liability exposure and actual harm potential while maintaining flexibility across different contexts. It provides strong victim protection for high-risk scenarios where compensation concerns are paramount, while preserving innovation incentives for lower-risk applications. The approach proves adaptable to both common law and civil law traditions, though implementation details would vary by jurisdiction.

Several operational elements support this framework's effectiveness. Risk classification systems developed by regulatory agencies provide consistent categorization across applications. Mandatory insurance for high-risk AI ensures adequate compensation resources. Substantive safety standards developed through multi-stakeholder processes inform defect and breach determinations. Adapted discovery rules address information asymmetries favoring defendants with technical knowledge.

Looking forward, harmonization efforts across jurisdictions become increasingly important given AI's global nature. The EU's comprehensive AI Act and Liability Directive could provide foundations for broader international alignment, though local adaptation remains necessary given institutional and cultural differences. Common law and civil law systems can both implement risk-based frameworks while preserving their distinctive characteristics and procedural traditions.

The research contributes to ongoing debates about AI governance by providing systematic comparative analysis grounded in doctrinal examination across multiple jurisdictions. Rather than advocating wholesale adoption of either strict liability or fault-based approaches, the findings support nuanced frameworks that match liability standards to risk contexts. This evidence-based approach balances competing policy objectives of victim compensation, safety deterrence, and innovation preservation.

Ultimately, effective AI liability frameworks must reflect technology's dual nature as both tremendous opportunity and genuine risk. Legal systems should neither impose crushing liability that prevents beneficial development nor tolerate inadequate accountability that leaves victims uncompensated and safety unincentivized. The hybrid risk-based approach proposed here threads this needle by calibrating liability exposure to actual risk levels while addressing AI's distinctive characteristics that challenge traditional frameworks.

As AI systems become increasingly prevalent across critical social domains, these questions will only grow in importance. Courts, legislatures, and regulatory agencies across jurisdictions must grapple with fundamental choices about liability assignment with profound implications for technological development, victim protection, and fair risk distribution. This research provides conceptual tools and practical recommendations to inform these consequential decisions, drawing on comparative analysis to identify approaches that transcend particular legal traditions while respecting institutional diversity.

REFERENCES

1. Abbott, R. (2020) 'The reasonable computer: Disrupting the paradigm of tort liability', *George Washington Law Review*, 86(1), pp. 1-45.
2. Abraham, K.S. and Rabin, R.L. (2019) 'Automated vehicles and manufacturer responsibility for accidents: A new legal regime for a new era', *Virginia Law Review*, 105(1), pp. 127-171.
3. Amodei, D., Olah, C., Steinhardt, J., Christiano, P., Schulman, J. and Mané, D. (2016) 'Concrete problems in AI safety', *arXiv preprint arXiv:1606.06565*.
4. Balkin, J.M. (2017) 'The three laws of robotics in the age of big data', *Ohio State Law Journal*, 78(5), pp. 1217-1241.
5. Bensoussan, A. (2021) 'AI liability in French law: From thing liability to algorithm liability', *European Journal of Risk Regulation*, 12(3), pp. 506-523.
6. Burrell, J. (2016) 'How the machine "thinks": Understanding opacity in machine learning algorithms', *Big Data & Society*, 3(1), pp. 1-12.
7. Citron, D.K. and Pasquale, F. (2014) 'The scored society: Due process for automated predictions', *Washington Law Review*, 89(1), pp. 1-33.

8. Dobbs, D.B., Hayden, P.T. and Bublick, E.M. (2016) *The Law of Torts*. 2nd edn. St. Paul: West Academic Publishing.
9. European Commission (2022) 'Proposal for a Directive on adapting non-contractual civil liability rules to artificial intelligence (AI Liability Directive)', COM(2022) 496 final.
10. Kingston, J. (2018) 'Artificial intelligence and legal liability', in Barfield, W. and Pagallo, U. (eds.) *Research Handbook on the Law of Artificial Intelligence*. Cheltenham: Edward Elgar, pp. 589-605.
11. Leenes, R. and Lucivero, F. (2014) 'Laws on robots, laws by robots, laws in robots: Regulating robot behaviour by design', *Law, Innovation and Technology*, 6(2), pp. 193-220.
12. Lemley, M.A. and Casey, B. (2019) 'Remedies for robots', *University of Chicago Law Review*, 86(5), pp. 1311-1396.
13. Lior, A. (2020) 'The new paradigm of tort law: Automated driving and other AI activities', *Nevada Law Journal*, 20(3), pp. 1-62.
14. Owen, D.G. (2015) *Products Liability Law*. 3rd edn. St. Paul: West Academic Publishing.
15. Scherer, M.U. (2016) 'Regulating artificial intelligence systems: Risks, challenges, competencies, and strategies', *Harvard Journal of Law & Technology*, 29(2), pp. 353-400.
16. Selbst, A.D. and Barocas, S. (2018) 'The intuitive appeal of explainable machines', *Fordham Law Review*, 87(3), pp. 1085-1139.
17. Wagner, G. (2020) 'Liability rules for the digital age: Aiming for the Brussels effect', *Journal of European Tort Law*, 11(1), pp. 31-60.
18. Yamamoto, K. (2019) 'Civil liability for artificial intelligence in Japan: Lessons for a new regime', in Barfield, W. and Pagallo, U. (eds.) *Research Handbook on the Law of Artificial Intelligence*. Cheltenham: Edward Elgar, pp. 606-624.