

Spalević Žaklina*

<https://orcid.org/0000-0001-8502-2038>

Ivković Ratko**

<https://orcid.org/0000-0002-6557-4553>

Ječmenić Milan***

<https://orcid.org/0009-0003-6109-3490>

UDC: 347.412:34|004.8(4-672EU)

Review article

DOI: 10.5937/ptp2603134S

Received on: June 3, 2026

Approved for publication on:

July 29, 2026

Pages: 134–154

BEYOND DETERMINISM – LEGAL RESPONSIBILITY IN THE AGE OF PROBABILISTIC AI SYSTEMS

ABSTRACT: This paper examines the fundamental distinction between classical deterministic programming and contemporary artificial intelligence systems, with particular emphasis on the legal consequences arising from the non-deterministic behavior of AI models. By analyzing the basic postulates of programming, the *Softmax* function as a key mechanism for output generation, and stochastic factors in training and generative processes, the paper demonstrates that artificial intelligence deviates from expected norms of predictability—an aspect that has far-reaching implications for the law of obligations, consumer protection, and tort liability regimes. In addition to a theoretical overview, the paper provides a legal analysis of the European Union’s AI Act in the context of non-determinism, highlighting regulatory efforts to address this characteristic of AI models through a risk-classification system as well as technical and organizational requirements. The paper aims to contribute to an interdisciplinary understanding of AI-related challenges by positioning legal scholarship as a key factor in developing a new model of liability in the digital age.

*LLD, Associate Professor, Singidunum University, Faculty of Tourism and Hospitality Management, Belgrade, Serbia, e-mail: zspalevic@singidunum.ac.rs

**PhD, Associate Professor, State University of Novi Pazar, Department of Technical Sciences, Novi Pazar, Serbia, e-mail: ivkovic@np.ac.rs

***LLD, Associate Professor, MB University Belgrade, Business and Law Faculty, Belgrade, Serbia, e-mail: mjecmenic@mbuniverzitet.edu.rs



© 2026 by the authors. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

Keywords: *artificial intelligence, non-determinism, law of obligations, EU AI Act, consumer protection.*

1. Introduction

Programming as a discipline has always rested on firm foundations of logic, mathematics, and predictability. At the core of this discipline lies the postulate of determinism—the idea that, for any valid program code that receives the same set of input data, the execution result must be identical in every run. This property enables the formal verification of software systems, guarantees their reliability, and makes the development of complex information systems possible in domains where accuracy and predictability are critical—from aviation and medicine to finance and the management of critical infrastructure. However, with the emergence of advanced artificial intelligence (AI) models—especially those based on deep learning and statistical inference—this foundational postulate is being seriously called into question. AI systems, unlike classical deterministic programs, do not execute instructions according to a strictly defined algorithm; rather, they model probabilities, drawing inferences from vast amounts of data and complex functions such as the softmax function. This means that the same input may produce different outputs across different instances of execution, depending on the model’s state, the temperature of the softmax function, or even the semantic context of the task.

For example, in traditional programming, a function $f(x)$ would always return the same value y if x is the same. In the AI context, a function $f_{AI}(x)$ uses mechanisms such as softmax normalization to produce a probability distribution over possible outputs, where the model then “decides” which output is most appropriate, without thereby guaranteeing identical results on each execution. This concept is fundamentally different from the logical rigor of traditional programming and introduces an element of uncertainty into the behaviour of the software system. The difference between the traditional and AI approach can also be explained philosophically: classical programming relies on an extensional understanding (what the program computes), whereas AI systems increasingly incline toward intensional properties (how the system arrives at a result), as discussed in philosophical analyses of programming such as Quinon’s work (Quinon, 2025), which highlights the distinction between expressing and specifying procedures and the formal correctness of program behaviour. In addition, Bertrand Meyer and Reto Weber, in their work on reconstructing the foundations of programming

through the pure mathematical theory of programming PRISM (Meyer & Weber, 2025), emphasize that all classical postulates of programming—such as composition, restriction, and correctness—are grounded in zero axioms, i.e., they are formally provable from set theory. This paradigm is difficult to transfer to AI systems, which are not formally provably correct, but rather empirically “good enough”. It is precisely the softmax function that amplifies this problem. As part of the output layer of many neural networks, softmax converts unordered logits into probabilities that together form a distribution whose maximum corresponds to the model’s most likely answer. However, softmax also accentuates differences in scores, such that small shifts in input logits can lead to radically different responses. Moreover, models also use so-called temperature-scaled softmax, which further influences how “decisive” the model will be in its response. Higher temperatures lead to a more uniform distribution and lower confidence, whereas lower temperatures lead to more decisive (but potentially premature) conclusions. For this reason, generative AI systems cannot be considered classical programs—they are probabilistic oracles, not exact executors. In such systems, even with the same input, the result may differ due to:

- the different state of internal memory (the so-called context window),
- stochastic elements of the sampling algorithm (e.g., top-k sampling, nucleus sampling),
- or changes in model weights during fine-tuning.

This inherent non-determinism introduces profound challenges into legal frameworks that have hitherto treated software as a deterministic entity. For example, within the legal framework of liability for software defects—under classical doctrine—it is possible to establish causation between software behaviour and the harm incurred. But if software (e.g., an LLM) is not deterministic and instead yields the “most probable” solution, how can liability be proven—and, more importantly, who is liable? This issue becomes even more complex when we consider that most AI systems are not designed, implemented, and controlled by a single party. We have:

- the model developer (e.g., OpenAI),
- the system implementer (e.g., Microsoft, Google),
- the end user (e.g., a lawyer, physician, student) who uses an AI tool for decision-making.

A need thus arises to redefine the concepts of liability, foreseeability, and users’ rights under conditions in which the traditional binary logic of

software code is replaced by the probabilistic decisions of a neural network. At the legal level, the question becomes: does the user have the right to expect deterministic behaviour from a system that technically cannot behave deterministically?

This is why the present manuscript matters—it does not aim to answer technical questions about the performance of AI models, but rather to establish a legal-philosophical foundation for understanding a new form of program behaviour. The need to examine AI systems legally from the perspectives of the law of obligations, consumer protection, and tort liability is no longer theoretical. There is a growing number of examples in which AI provides harmful, incorrect, or non-compliant information, while no clear legal doctrine exists as to how liability should be addressed. Through this paper, we will show how the classical premises of deterministic programming, grounded in theories such as PRISM (Moine, Westrick & Tassarotti, 2025), can no longer serve as the basis for understanding and regulating modern AI systems. Furthermore, by analyzing software mechanisms such as the softmax function, we will demonstrate how the mathematical logic of a model’s output is itself fundamentally different from the traditional, strictly deterministic output of program functions.

In legal terms, this difference opens new space for considering concepts such as:

- acceptable deviation instead of accuracy,
- probabilistic liability instead of binary fault,
- and conditional reliability as a basis for contractual obligation.

Through this interdisciplinary analysis, the aim of this manuscript is to show that what is at stake is not merely a “technical difference” in how an AI system functions, but a fundamental paradigm shift that demands a new understanding of law, responsibility, and trust in twenty-first-century information systems.

2. AI Models and the Undermining of Determinism

The basic postulates of classical programming clearly define determinism as a central attribute: for any specific input to a program, it must always generate the same, predictable, and reproducible result. This principle is grounded in the idea that software functions as a mathematical function or an algorithmic transformation, in which inputs and outputs are deterministically connected through strictly defined operations and formal

semantics. Such predictability enables programmers, engineers, and legal professionals to prove correctness, test behavior, and establish responsibility for the outcome of software execution. However, with the development and accelerated adoption of contemporary artificial intelligence models—especially those based on statistics and neural networks—this fundamental postulate of programming is being called into question. New AI systems—including large language models (Large Language Models – LLMs) and other generative architectures—do not function as classical deterministic functions. Instead, they employ probabilistic approaches in which the output value is not precisely determined in advance, but rather results from a complex process of estimating probabilities among possible outcomes, thereby undermining the foundational assumption of predictability that has governed traditional software engineering.

2.1 The Deterministic Approach in Classical Programming

In traditional software, code is most often executed by directly interpreting explicitly written commands, without internal elements of randomness or statistical evaluation. If the same program is executed multiple times with identical input data, the result will always be the same, provided that external conditions (hardware, environment) remain the same. This behavior enables formal methods for proving correctness and validating software systems—for example, through the use of static analysis tools, proof-theoretic methods, or model checking. Deterministic systems are therefore reproducible, testable, and formally predictable. This determinism is also integrated into legal domains: when software is treated as a reliable product or service, the law of obligations and consumer protection often proceed from the assumption that functionality must be predictable and error-free. For example, if software forms part of a service contract, it is expected to satisfy the agreed specifications—which implies that it will always produce the same result for the same input data.

2.2 The Probabilistic Approach in Modern AI

By contrast, modern AI systems often use probabilistic models grounded in learning from data and statistical inference, rather than in the explicit execution of rules. In such models, the result is not deterministically fixed, but is generated on the basis of an assessment of the probabilities of different possible answers. This approach is particularly pronounced in deep neural networks and generative models that use techniques such as Bayesian

inference, Gaussian process models, or other probabilistic methods to estimate uncertainty and multiple possible outputs (Drabent, 2025). Within these frameworks, the model does not return a single strictly predetermined value for a given input, but rather a probability distribution—which is especially evident in generative tasks such as producing text, images, or other forms of content. Such a probabilistic approach enables models to learn patterns from enormous datasets and to adapt to new, unknown situations. However, it simultaneously introduces non-determinism into model behavior, so that the same input may yield different outputs across different runs of the model, even when all other conditions are nominally the same (Hagen et al., 2012). By way of illustration, LLMs such as GPT are trained to generate, for each next token, the most probable choice based on context and statistical patterns in the training data, using functions such as softmax for normalization and for estimating the probability of each possible output token. Consequently—even with identical input data—the model may behave differently on each invocation, because the generation process is based on probabilistic sampling and contextual learning, rather than on fixed rules (Franke & Degen, 2023).

2.3 The Role of Softmax and Probabilistic Decision-Making

The softmax function in neural networks converts unnormalized values (logits) into probabilities that represent the relative likelihood of different possible responses. This process is inherently probabilistic, because there is no single “certain” result—rather, the model relies on a probability distribution to select the next element of the output. The accuracy and repeatability of responses thus become functions of statistical estimation, model design, and the training set, rather than of deterministic execution of a predefined algorithm. This probabilistic approach is further complicated when sampling techniques and other parameters are used that directly affect the probability distribution. For example, in text generation, a change in the softmax distribution can change token-selection probabilities, enabling the model to provide different answers to the same input in different situations (Krause & Hübötter, 2025).

2.4 Factors That Cause Non-Determinism in AI Models

There are multiple reasons why modern AI systems exhibit non-deterministic characteristics:

- Stochasticity of the training process: Models are typically trained using optimization algorithms such as stochastic gradient descent (SGD),

which introduce randomness into weight initialization and the ordering of data across training epochs. This randomness can lead to different local minima and different model parameters after training (Atil, Aykent, Chittams, Fu & Passonneau, 2025).

- Sampling during output generation: In generative tasks, output may be sampled from a probability distribution rather than selecting the maximally probable result each time, which directly affects reproducibility of results.
- Hardware and system differences: Even when models are configured “deterministically” (e.g., temperature set to zero), variability in hardware or execution modes (e.g., GPU parallelization) can lead to different results due to operation ordering and the accumulation of numerical errors (Fu, 2026).

Because of these factors, AI behaviors such as “black-box” decision-making and unpredictability become an essential feature rather than merely an artifact of implementation or an error.

2.5 Implications of Undermined Determinism within Legal, Ethical, and Regulatory Frameworks

This probabilistic nature of modern AI systems has significant implications—not only for domestic and international deployment of the technology, but also for legal, ethical, and regulatory frameworks. In the context of law, particularly within the law of obligations and consumer protection, determinism is often implicitly assumed. When software produces different outcomes for the same input, this raises questions about reliability, verification, and liability. In the traditional paradigm, if software fails to meet the contractually agreed performance standard, the user may claim damages or terminate the contract. However, for AI systems there is no clear line between normal variation in model behaviour and a software defect, because probabilistic outputs are not technically “wrong”—they are an inherent product of the model. This phenomenon is referred to as the Model Variability Problem (MVP) and relates to the fact that LLMs and similar systems can produce different classifications or outcomes for the same input information due to inherent variability in their design and data-interpretation mechanisms (Herrera-Poyatos et al., 2025). This leads to legal uncertainty, because liability for harm or consequences cannot always be clearly attributed to a “program error” in the traditional sense.

3. Legal Consequences of the Non-Determinism of AI Systems

The non-determinism of contemporary AI systems is not merely a technical or philosophical issue; it generates profound and multi-layered legal consequences, particularly in the fields of the law of obligations, consumer protection, and tort liability. Traditional legal orders—developed over decades in symbiosis with deterministic software and mechanical systems—proceed from the assumption that the behaviour of a technological product is predictable, repeatable, and verifiable. Artificial intelligence, as a probabilistic system, directly undermines these assumptions and thereby calls into question existing legal categories of liability, causation, and fault.

3.1 Determinism as an Implicit Legal Assumption

In classical obligations law, especially in contracts for the provision of services and the sale of software, there is an implicit assumption that a product or service will meet a certain standard of performance. That standard presupposes that, under the same set of circumstances, the same system behavior is expected. In other words, the law proceeds from an assumption of determinism in the technical instrument. This assumption is visible across several areas:

- in rules on material defects of goods;
- in the doctrine of non-performance or improper performance of obligations;
- in standards of professional diligence in the provision of services.

With deterministic software, a deviation in behaviour can readily be qualified as an error or malfunction. With AI systems, however, variability of outputs for the same input does not necessarily signify an error—it is often an anticipated and systemic characteristic of the model (Quinon, 2025). This creates a fundamental legal dilemma: is non-deterministic behaviour of an AI system legally acceptable, or does it constitute a breach of obligation?

3.2 Contractual Liability and the Problem of Performance

In the context of contractual relationships (Law of Obligations), the key questions are:

- whether an AI system performs the contractual obligation;
- how to define the object of the obligation when outcomes are probabilistic;
- who bears the risk of result variability.

With classical software solutions, the contractual object is clearly defined: the software must perform a specific function in a precisely specified manner. With AI systems—particularly those offered as “intelligent assistants”—the object of the contract is often not a concrete result, but rather the best estimate or recommendation. Accordingly, the contractual obligation shifts from outcome to process. This transformation has serious consequences. If an AI system produces an incorrect or harmful outcome, the provider may argue that the system operated in accordance with its probabilistic nature and that no contractual breach occurred. Conversely, the user may argue that they legitimately expected a consistent and reliable response, especially if the system was marketed as “accurate,” “reliable,” or “safe” (Devetzis & Samaras, 2024). This leads to legal uncertainty and indicates the need for more precise regulation of contracts involving AI systems, including:

- an explicit statement of the level of reliability;
- the definition of acceptable deviation;
- the allocation of risks arising from probabilistic behaviour.

A framework for assessing the conformity of digital services with the contract—including services based on artificial intelligence—is provided by Directive (EU) 2019/770 on certain aspects concerning contracts for the supply of digital content and digital services, which governs the conformity of digital content and services and the remedies available to consumers in the event of a lack of conformity (Directive (EU) 2019/770).

3.3 Non-Contractual Liability and Causation

Even more complex issues arise in the domain of non-contractual (tort) liability. The classical model of liability for damage requires the existence of:

- damage;
- unlawful conduct;
- a causal link;
- fault (or strict liability).

For AI systems, the element of causation is particularly problematic. If the same input can generate different outputs, the question arises whether a specific harmful outcome can be regarded as a necessary consequence of the system’s behaviour, or merely as one among statistically possible outcomes (Wachter, Mittelstadt & Russell, 2021).

For example, if an AI system recommends a particular medical or financial course of action that leads to harm, can it be claimed that the harm was “caused” by the system, or is it a legitimate but unfavourable outcome of probabilistic inference? Traditional causation tests (*conditio sine qua non*) are difficult to apply in such situations. This problem gives rise to the idea of facilitated proof or presumed causation for AI systems, which has also been recognized in contemporary EU regulatory initiatives (European Commission, 2024).

3.4 Consumer Protection and Legitimate Expectations

In consumer law, a key concept is the legitimate expectation of the average consumer. Consumers are entitled to expect that a product or service will function in a reasonably predictable manner and in accordance with its intended purpose. However, the average consumer is often unaware of the probabilistic nature of AI systems. When using an AI tool, the consumer may assume that the system will provide the same answer to the same question, as is the case with classical software tools. If this does not occur, the consumer may consider themselves deceived or misled. This raises issues of unfair commercial practices and duties of transparency. If the manufacturer or provider of an AI service fails to clearly explain that the system does not guarantee consistent results, consumer rights may be infringed (Devetzis & Samaras, 2024).

3.5 Product Liability and the Concept of a “Defective Product”

Another important aspect concerns product liability. Traditionally, a product is considered defective if it does not provide the safety that can reasonably be expected. With AI systems, the question is whether non-determinism can be regarded as a defect. If non-determinism is an essential characteristic of the system, the manufacturer may argue that the product is not defective, but rather functions as designed.

However, if that characteristic leads to harm or serious legal uncertainty, the legislature may conclude that the product fails to meet standards of safety and reliability (European Parliamentary Research Service, 2025). This issue is particularly relevant in the context of the new EU framework—Directive (EU) 2024/2853 on liability for defective products, which repeals the earlier Directive 85/374/EEC. The new Directive retains strict (no-fault) liability as its core principle, but it expressly broadens the notion of a “product” to

encompass software, including artificial intelligence systems, and introduces presumptions that ease the burden of proof for the injured person in cases of technical or scientific complexity (Directive (EU) 2024/2853).

3.6 The Need to Redefine Legal Standards

All of the above indicates that the legal system must evolve to address the challenges posed by the non-determinism of AI systems. Classical concepts such as complete predictability, binary correctness, and clear causation are no longer sufficient. Instead, legal theory increasingly refers to:

- standards of reasonable probabilistic behaviour;
- risk-based liability rather than fault-based liability;
- functional responsibility of actors who exercise control over an AI system.

This transformation has a particularly strong impact on the Law of Obligations, because it requires a redefinition of the very notion of an obligation: the obligation is no longer a guaranteed outcome, but rather due diligence in managing probabilities.

4. EU AI ACT – Legal Analysis in the Context of the Non-Determinism of AI Systems

By adopting the EU Artificial Intelligence Act (AI Act), the European Union has established the first comprehensive normative framework aimed at regulating the development, placing on the market, and use of artificial intelligence systems in order to manage the risks these systems generate. The AI Act, published in the Official Journal of the EU on 12 July 2024 after years of negotiations (Regulation (EU) 2024/1689), represents an attempt to harmonize rules across the Union and has the potential to become a global standard for AI regulation due to the so-called “Brussels effect”—the phenomenon whereby EU regulations often influence non-European jurisdictions through economic and legal channels (EU Artificial Intelligence Act, 2024). In the legal context of AI system non-determinism, the AI Act pursues several objectives: reducing risks through a harmonized classification of systems by risk level; imposing obligations on operators and providers; improving transparency; and enabling oversight and accountability mechanisms. However, its application and interpretation with respect to non-deterministic systems—such as generative large language models (LLMs)—raises numerous legal questions concerning

how traditional legal institutes of the law of obligations and liability adapt to probabilistic outputs that are not deterministic.

The AI Act is based on a risk-based approach, under which systems are classified into several categories: unacceptable risk, high risk, limited risk, and low or minimal risk. Systems that present an unacceptable risk are prohibited, while those considered high-risk are subject to stringent requirements before they may be placed on the EU market. This approach aims to ensure a high level of protection for citizens' rights and safety, with particular emphasis on transparency, documentation, and responsible governance of models that may significantly affect fundamental rights. In this way, the AI Act sets legal standards oriented toward risk minimization, while simultaneously having to address inherent properties of AI systems—such as non-determinism—that conflict with traditional legal assumptions regarding the predictability of software behaviour.

One of the key questions in analyzing the AI Act in the context of non-determinism is how the law treats systems whose results are not always the same for the same input, as a product of statistical and probabilistic learning methods. The AI Act imposes obligations on providers of high-risk systems to establish and maintain systems for collecting, managing, and retaining technical documentation and activity logs so that system behaviour and its impact on users and third parties can be traced (Ivković, 2025). These obligations are intended to enable evaluation of model performance and behaviour, but they do not guarantee automatic legal liability for every possible deviating outcome. The Act requires providers to implement procedures for risk assessment and risk mitigation, including testing prior to deployment and continuous monitoring during operation. At the organizational level, these obligations can be operationalized through auditable AI governance and management system structures (ISO/IEC 42001:2023). This aligns with established AI risk governance approaches that operationalize risk identification, measurement, and mitigation across the AI lifecycle (NIST, 2023). This framework indicates a shift from the traditional understanding of software as a deterministic product toward understanding AI systems as entities whose performance must be assessed and supervised through empirical indicators, statistical tests, and compliance assessment against risk standards (Gasparri, 2025). In addition, AI-specific risk management guidance provides practical structure for identifying, assessing, and treating risks throughout the AI lifecycle (ISO/IEC 23894:2023).

Transparency and explainability are two further components of the AI Act that are directly connected to the problem of non-determinism. For high-risk

systems, the Act requires that users be enabled to understand the rationale of a decision, i.e., to receive an explanation of outcomes in an intelligible format, particularly where decisions affect people's rights (e.g., employment, credit scoring, health decisions). These requirements are significant because they provide a legal framework for reducing the negative consequences of non-determinism—although they cannot eliminate the probabilistic nature of the model itself, they can ensure that embedded uncertainty, expected outcome variability, and the limits of model reliability are clearly communicated to the user. Transparency is thus legally positioned as a core component of responsible AI development and use, thereby partially addressing non-determinism by requiring providers to supply clear information on performance and on the risks associated with managing AI systems (Ivković, 2025).

Another important aspect of the AI Act is its extraterritorial application—clearly articulated in provisions that allow the Act to apply not only to providers and users within the EU, but also to entities outside the EU whose products have market impact or affect the rights of EU citizens. This is particularly important in light of the global nature of AI technology, where large models and services often originate beyond the EU's legal boundaries. The AI Act's extraterritorial reach reflects the EU's legal ambition to set global norms, but it also raises the question of how standards of liability for non-deterministic model behaviour are aligned across jurisdictions when other legal systems do not recognize or regulate AI in a comparable manner. This feature creates an imperative to align domestic legal frameworks with EU standards and may have consequences for international trade and legal practice outside the EU.

Nevertheless, despite its comprehensive framework, the EU AI Act—as many legal scholars commenting on it have noted—faces criticism regarding its capacity to effectively address legally relevant challenges posed by probabilistic AI systems. For instance, analyses of the AI Act indicate that, although it introduces strict obligations for high-risk systems and a risk-based approach, it does not fully resolve questions relating to legal liability in cases of harm or negative outcomes resulting from inherent model variability. These critics argue that the AI Act should include more detailed mechanisms that would provide legal predictability and security for users with respect to probabilistic outcomes, especially where a model may generate objectively inaccurate results that harm the rights or interests of data subjects.

Another paper in *Pravo – teorija i praksa* addressing the AI Act highlights the legal implications of the risk-based approach to regulating artificial intelligence. The paper points out that the AI Act recognizes specific characteristics of machine learning and autonomous behaviour, and that the risk-based approach

enables differentiated treatment of systems according to the potential danger they may cause. The author also notes that a broad definition of AI may lead to over-regulation, because certain characteristics of artificial intelligence do not appear in all types of algorithms. At the same time, the AI Act imposes technical documentation and record-keeping obligations that are important for oversight, yet a legal gap remains regarding what happens when a system operates within its technical limits but nevertheless produces a harmful outcome due to the inherent logic of a probabilistic model (Mladenov, 2023).

A particular legal challenge concerns liability for damage in situations where AI non-determinism leads to a harmful outcome—for example, when an AI model provides an inaccurate diagnosis in a medical context or inadequate legal advice. Traditional rules of liability, both contractual and non-contractual, proceed from the assumption that the precise cause and responsibility can be clearly determined. However, probabilistic models with inherent non-determinism make it difficult to determine whether harm is the result of a “defect” in the system or of legitimately foreseeable variation in performance. The AI Act seeks to mitigate this problem through obligations of risk assessment and mitigation, as well as requirements for system oversight, but it still lacks specific legal mechanisms that directly address the problem of associating a harmful outcome with a model’s probabilistic behaviour.

On the other hand, the AI Act also contains provisions relating to fundamental rights, including protection against discrimination, data privacy, and the transparency of decisions made by automated systems. These provisions are particularly important where AI systems, such as machine-learning models, make decisions affecting individuals’ living conditions or rights. The EU legislature has attempted to capture these aspects normatively as well, yet the challenge remains of how to implement and supervise protection standards when inherent model non-determinism complicates evaluation and the attribution of responsibility for negative outcomes to legal subjects.

5. Liability and Sanctions under EU Law in Cases of AI System Errors

Within the legal system of the European Union, liability for damage caused by artificial intelligence—including damage arising from the non-deterministic behaviour of AI systems—is regulated through a combination of the AI Act and the separate proposal for an Artificial Intelligence Liability Directive (AI Liability Directive). The key legal question is: who is responsible when an AI system produces an unpredictable or inaccurate result that leads to harm? Under

the EU approach, primary responsibility rests with the “provider” and the “user of a high-risk AI system,” where the provider is defined as the entity that develops or places the system on the market, and the user as the party that employs it for professional purposes. Where damage occurs, and it is established that the AI system failed to comply with statutory obligations—such as transparency, risk management, oversight, and safety—the provider may be directly liable, even in the absence of intent or negligence. The AI Act also contains a tiered penalty regime in which the level of the administrative fine depends on the nature of the infringement (Article 99). For the most serious infringements, namely the placing on the market or use of prohibited AI practices under Article 5, administrative fines of up to EUR 35 million or, where the offender is an undertaking, up to 7% of its total worldwide annual turnover for the preceding financial year may be imposed, whichever amount is higher (Article 99(3)). Non-compliance with obligations applicable to high-risk systems and to other operators—other than the obligations laid down in Article 5 - is subject to lower fines of up to EUR 15 million or up to 3% of total worldwide annual turnover (Article 99(4)), while the supply of incorrect, incomplete, or misleading information to the competent authorities is subject to fines of up to EUR 7.5 million or up to 1% of turnover (Article 99(5)). Non-determinism as a technical characteristic does not exempt actors from liability; rather, it entails stricter obligations for providers regarding control and documentation of system behavior, precisely in order to ensure legal certainty and protect users and third parties from the harmful consequences of AI-driven decisions. However, the proposed AI Liability Directive was subsequently withdrawn by the Commission—first announced in the Commission Work Programme for 2025 (11 February 2025) and then formally, through a notice published in the Official Journal of the EU (OJ C/2025/5423, 6 October 2025)—so that specific rules on non-contractual liability for damage caused by artificial intelligence were not adopted. Liability questions are thereby left to be addressed through other instruments, primarily the strict-liability regime established by Directive (EU) 2024/2853 on liability for defective products, as well as through national law mechanisms (European Parliament Legislative Train, 2025).

6. The Legal Framework of AI Non-Determinism in Serbia

Although the Republic of Serbia still does not have a specific statute that explicitly regulates artificial intelligence, the legal consequences of the use of AI systems—especially those exhibiting non-deterministic behaviour—can be assessed within the existing normative framework, primarily through the Law on

Obligations (1974) and the Law on Consumer Protection (2026). These statutes, although enacted in an era of deterministic technological systems, continue to constitute the basis for legal assessment of liability, damage, and the protection of the weaker contracting party in the context of contemporary AI technologies.

The Law on Obligations (1974) proceeds from the assumption that the object of an obligation is determinable, predictable, and verifiable, and that the debtor is liable for non-performance or improper performance where the result deviates from what was agreed. However, in AI systems whose behaviour depends on probabilistic models and statistical estimation, the outcome is not always identical for the same input data. This calls into question the classical understanding of performance of an obligation, because deviations in results do not necessarily constitute an error, but may instead be an inherent feature of the system. This opens the need for a new interpretation of the concept of due care and the allocation of risk in obligations relationships that involve AI systems.

On the other hand, the Law on Consumer Protection (2026) provides a basis for considering the legitimate expectations of users of AI products and services. If the consumer is not clearly informed that an AI system may produce different results for the same inputs, this may lead to a violation of the right to information and the prohibition of misleading commercial practices. A particularly significant question is whether an AI system may be regarded as a defective product if its non-determinism leads to harmful or unreliable outcomes.

In that sense, the topic of this manuscript directly points to the need to align domestic legislation with contemporary technological realities, as well as to the necessity of future normative intervention which would—following the model of the EU AI Act—more precisely regulate the legal consequences of deploying probabilistic AI systems in the Republic of Serbia.

7. Conclusion

This paper has shown that the development of probabilistic artificial intelligence systems seriously challenges the basic postulates upon which classical programming rests, particularly the principle of determinism. This fact has a direct impact on legal systems that traditionally rely on the predictability and verifiability of the behaviour of technical systems. Through an analysis of the functional architecture of AI models—especially the role of the softmax function, sampling strategies, and response temperature—it has been established that these systems produce variable results even under

identical inputs, which raises legal problems concerning liability, causation, and reliability. The legal analysis of the EU AI Act confirms that the European legislature has recognized this specificity and responded with a normative framework based on risk classification, technical documentation, and transparency obligations. However, the question remains open as to how non-determinism should be treated legally in cases where it proves harmful to users or third parties. In the context of the Republic of Serbia, an analysis of existing legislation—particularly the Law on Obligations and the Law on Consumer Protection—indicates that these statutes, although formally applicable, were not designed for technologies operating on principles of stochastic inference. The absence of specific rules governing liability for probabilistic outcomes points to the need for reform of the domestic legal framework, following European approaches. The paper opens space for further research aimed at defining new legal constructs—such as “probabilistic liability,” “technical legitimate expectations,” and “functional accuracy”—which could serve as a foundation for legal adaptation in the age of artificial intelligence. In this regard, legal scholarship has a key role in developing a model that enables a balance between technological progress and legal certainty.

Conflict of Interest

The authors declare no conflict of interest.

Author Contributions

Conceptualization, Ž.S. and R.I.; formal analysis, Ž.S. and R.I.; investigation, Ž.S. and M.J; methodology, R.I., Ž.S. and M.J; validation, Ž.S. and M.J; writing – original draft, R.I., Ž.S. and M.J; writing – review & editing, Ž.S., R.I. and M.J. All authors have read and agreed to the published version of the manuscript.

Funding

No external financial support was provided for this research.

Data Availability Statement

The original contributions presented in the paper are included in the manuscript itself. Additional inquiries may be addressed to the author.

Informed Consent / Approval of the Competent Committee

Not applicable.

Spalević Žaklina

Univerzitet Singidunum, Fakultet za turistički i hotelijerski menadžment, Beograd, Srbija

Ivković Ratko

Državni Univerzitet u Novom Pazaru, Departman za tehničko-tehnološke nauke, Novi Pazar, Srbija

Ječmenić Milan

MB Univerzitet, Poslovni i pravni fakultet, Beograd, Srbija

IZVAN DETERMINIZMA – PREISPITIVANJE PRAVNE ODGOVORNOSTI U DOBU PROBABILISTIČKIH VI SISTEMA

APSTRAKT: Ovaj rad razmatra suštinsku razliku između klasičnog determinističkog programiranja i savremenih sistema veštačke inteligencije, sa posebnim naglaskom na pravne posledice koje proizlaze iz nedeterminističkog ponašanja VI sistema. Analizom osnovnih postulata programiranja, *Softmax* funkcije kao ključnog mehanizma generisanja izlaza, kao i stohastičkih faktora prisutnih u procesima treniranja i generativne produkcije, pokazuje se da veštačka inteligencija odstupa od uobičajenih očekivanja predvidivosti, što ima dalekosežne implikacije na obligaciono pravo, zaštitu potrošača i režime vanugovorne (deliktne) odgovornosti. Pored teorijskog pregleda, rad pruža i pravnu analizu Akta Evropske unije o veštačkoj inteligenciji (EU AI Act) u kontekstu nedeterminisanosti, ukazujući na regulatorne napore da se ova osobina VI modela normativno obuhvati kroz sistem klasifikacije rizika, kao i kroz tehničke i organizacione zahteve. Rukopis ima za cilj da doprinese interdisciplinarnom razumevanju izazova povezanih sa veštačkom inteligencijom, pozicionirajući pravnu nauku kao ključnog aktera u izgradnji novog modela odgovornosti u digitalnom dobu.

Ključne reči: *veštačka inteligencija, nedeterminisanost, obligaciono pravo, EU AI Act, zaštita potrošača.*

References

1. Atil, B., Aykent, S., Chittams, A., Fu, L., & Passonneau, R. J. (2025). *Non-Determinism of “Deterministic” LLM Settings*. arXiv preprint arXiv:2408.04667v5. DOI: 10.48550/arXiv.2408.04667
2. Devetzis, D., & Samaras, S. (2024). Consumer Protection Safeguards after the AI Act. *Perspectives of Law and Public Administration*, 13(2), pp. 298–309. DOI: 10.62768/PLPA/2024/13/2/12
3. Directive (EU) 2019/770 of the European Parliament and of the Council of 20 May 2019 on certain aspects concerning contracts for the supply of digital content and digital services. *Official Journal of the European Union (OJ)*. Downloaded 2026, January 20, from <https://eur-lex.europa.eu/eli/dir/2019/770/oj/eng>
4. Directive (EU) 2024/2853 of the European Parliament and of the Council of 23 October 2024 on liability for defective products and repealing Council Directive 85/374/EEC (Product Liability Directive). *Official Journal of the European Union (OJ)*. Downloaded 2026, January 26, from <https://eur-lex.europa.eu/eli/dir/2024/2853/oj/eng>
5. Drabent, W. (2025). *On systematic construction of correct logic programs*. arXiv preprint arXiv:2508.16782. DOI: 10.48550/arXiv.2508.16782
6. EU Artificial Intelligence Act. (2024). *The EU Artificial Intelligence Act: Up-to-date developments and analyses of the EU AI Act*. Downloaded 2025, December 30, from <https://artificialintelligenceact.eu/>
7. European Commission. (2024). *Proposal for a Directive on Adapting Non-Contractual Civil Liability Rules to Artificial Intelligence* (COM (2024)). Downloaded 2026, January 20, from [https://www.europarl.europa.eu/RegData/etudes/STUD/2024/762861/EPRS_STU\(2024\)762861_EN.pdf](https://www.europarl.europa.eu/RegData/etudes/STUD/2024/762861/EPRS_STU(2024)762861_EN.pdf)
8. European Parliament. (2025). AI liability directive – Legislative Train Schedule (includes reference to withdrawal notice in OJ C/2025/5423, 6 October 2025). Downloaded 2026, January 26, from <https://www.europarl.europa.eu/legislative-train/theme-a-europe-fit-for-the-digital-age/file-ai-liability-directive>
9. European Parliamentary Research Service. (2025). *Artificial Intelligence and Civil Liability – A European Perspective*. Downloaded 2026, January 21, from [https://www.europarl.europa.eu/RegData/etudes/STUD/2025/776426/IUST_STU\(2025\)776426_EN.pdf](https://www.europarl.europa.eu/RegData/etudes/STUD/2025/776426/IUST_STU(2025)776426_EN.pdf)
10. Franke, M., & Degen, J. (2023). The softmax function: Properties, motivation, and interpretation. PsyArXiv. DOI: 10.31234/osf.io/vsw47

11. Fu, T. (2026). Beyond Reproducibility: Token Probabilities Expose Large Language Model Nondeterminism. arXiv preprint arXiv:2601.06118v1. DOI: 10.48550/arXiv.2601.06118
12. Gasparri, W. (2025). Artificial intelligence and AI Act: from the individual to the algorithm? *Zbornik radova Pravnog fakulteta u Novom Sadu*, 59(1), pp. 291–314. DOI: 10.5937/zrpfns59-57519
13. Hagen, E. H., Chater, N., Gallistel, C. R., Houston, A., Kacelnik, A., Kalenscher, T., Nettle, D., Oppenheimer, D., & Stephens, D. W. (2012). Decision making: What can evolution do for us? In: Hammerstein P., & Stevens, J., R. (eds.), *Evolution and the mechanisms of decision making* (pp. 97–126). MIT Press
14. Herrera-Poyatos, D., Peláez-González, C., Zuheros, C., Herrera-Poyatos, A., Tejedor, V., Herrera, F., & Montes, R. (2025). An overview of model uncertainty and variability in LLM-based sentiment analysis: challenges, mitigation strategies, and the role of explainability. *Frontiers in Artificial Intelligence*, 8, 1609097. DOI: 10.3389/frai.2025.1609097
15. ISO/IEC. (2023). ISO/IEC 23894:2023 – *Artificial intelligence – Guidance on risk management*. International Organization for Standardization (ISO). Downloaded 2026, January 27, from <https://www.iso.org/standard/77304.html>
16. ISO/IEC. (2023). ISO/IEC 42001:2023 – *Artificial intelligence – Management system*. International Organization for Standardization (ISO). Downloaded 2026, January 26, from <https://www.iso.org/standard/42001>
17. Ivković, R. (2025). Transformation and economic aspects of software engineering through the implementation of the EU AI ACT. *Pravo - teorija i praksa*, 42(1), pp. 186–200. DOI: 10.5937/ptp2501186I
18. Krause, A., & Hübötter, J. (2025). *Probabilistic Artificial Intelligence*. arXiv preprint arXiv:2502.05244v1. DOI: 10.48550/arXiv.2502.05244
19. Meyer, B., & Weber, R. (2025). *Programming Really Is Simple Mathematics*. arXiv preprint arXiv:2502.17149. DOI: 10.48550/arXiv.2502.17149
20. Mladenov, M. (2023). Human vs. artificial intelligence – EU’s legal response. *Pravo - teorija i praksa*, 40(suppl.), pp. 32–43. DOI: 10.5937/ptp2300032M
21. Moine, A., Westrick, S., & Tassarotti, J. (2025). *All for One and One for All: Program Logics for Exploiting Internal Determinism in Parallel Programs*. arXiv preprint arXiv:2511.23283. DOI: 10.48550/arXiv.2511.23283

22. National Institute of Standards and Technology (NIST). (2023). AI Risk Management Framework (AIRMF 1.0). NIST. Downloaded 2026, January 22, from <https://www.nist.gov/itl/ai-risk-management-framework>
23. Quinon, P. (2025). Intensional Differences Between Programming Languages: A Conceptual and Practical Analysis. *Philosophies*, 10(6), p. 129. DOI: 10.3390/philosophies10060129
24. Regulates on artificial intelligence (Artificial Intelligence Act - 2024/1689). *Official Journal of the European Union (OJ)*. Downloaded 2026, January 25, from <https://eur-lex.europa.eu/eli/reg/2024/1689/oj>
25. Wachter, S., Mittelstadt, B., & Russell, C. (2021). Why fairness cannot be automated: Bridging the gap between EU non-discrimination law and AI. *Computer Law & Security Review*, 41, 105567. DOI: 10.1016/j.clsr.2021.105567
26. Zakon o obligacionim odnosima. [Law on Obligations]. *Službeni list SFRJ*, br. 29/78, 39/85, 45/89 – odluka USJ i 57/89, *Službeni list SRJ*, br. 31/93, *Službeni list SCG*, br. 1/03 – *Ustavna povelja i Službeni glasnik RS*, br. 18/20
27. Zakon o zaštiti potrošača [Law on Consumer Protection]. *Službeni glasnik RS*, br. 35/26