Bond University

# Bond Law Review

Volume 38 Issue 2

2026

## Reconsidering the Epistemological Foundations of Economic Analysis of Law in the Age of Artificial Intelligence: From Posnerian Efficiency to Algorithmic Rationality and the Transformation of Legal Teleology

Mostafa Bikaranbehesht

Private and Economic Law, Allameh Tabataba'i University

# RECONSIDERING THE EPISTEMOLOGICAL FOUNDATIONS OF ECONOMIC ANALYSIS OF LAW IN THE AGE OF ARTIFICIAL INTELLIGENCE: FROM POSNERIAN EFFICIENCY TO ALGORITHMIC RATIONALITY AND THE TRANSFORMATION OF LEGAL TELEOLOGY

MOSTAFA BIKARANBEHESHT*

Abstract

*This article reconsiders the epistemological foundations of the economic analysis of law in the context of artificial intelligence. It argues that AI reconfigures the core assumptions of law and economics under which these assumptions are operationalised. In classical frameworks, rationality functions as a simplifying device for modelling how individuals respond to legal rules. The increasing use of data-driven systems reshapes the informational and institutional environments within which such behaviour is analysed.*

*The article examines how this transformation affects four central concepts: rationality, efficiency, legitimacy, and justice. It shows that efficiency is increasingly mediated by technical metrics such as predictive accuracy and system performance. Similarly, legal legitimacy continues to depend on normative justification, even as computational processes influence how decisions are evaluated. In the domain of justice, individualised reasoning coexists with system-level approaches based on statistical patterns, generating new tensions within legal practice.*

*Through a series of case studies, the article demonstrates how these dynamics emerge in concrete legal contexts. It concludes that the impact of artificial intelligence on law is best understood not as a process of substitution, but as a reconfiguration of the epistemic and institutional conditions of legal reasoning.*

* Private and Economic Law, Allameh Tabataba'i University

## I Introduction

Economic analysis of law, since its consolidation in Posner's works, has rested on the assumption that law can be explained as a system designed to optimise human behaviour—a system whose primary function is to reduce social costs through rules that guide actors toward efficient conduct. Within this framework, two theoretical foundations play a central role: first, human rationality as the basis of legal decision-making; second, efficiency as the ultimate criterion for evaluating legal rules, institutions, and policies. These foundations, which made economic analysis of law one of the most influential legal approaches of the twentieth century, rely on the assumptions of neoclassical economics, including the premise that human beings are calculating, self-interested, and capable of making efficient choices. Nevertheless, a vast body of literature over the past four decades has demonstrated that these foundations face a range of cognitive, behavioural, and informational limitations that prevent the full realisation of the theoretical claims of economic analysis of law.[1]

With the advent of artificial intelligence in decision-making structures, data processing, and the interpretation of social and legal phenomena, the prevailing picture has been transformed in unprecedented ways. For the first time in the history of law, a non-human entity has emerged that not only compensates for the cognitive limitations of human agents but is also capable of producing decisions that are predictable, uniform, and data-driven—decisions that, from a computational perspective, are more optimal and, from a statistical standpoint, more stable than human judgment.[2]

Importantly, this article does not assume a simple replacement of human decision-making by artificial intelligence. Contemporary legal systems increasingly operate through hybrid human–machine assemblages, in which algorithms support, guide, and sometimes constrain human judgment. Understanding this hybridity is essential for analysing how economic reasoning operates in practice, as rationality remains formally attributed to human actors even when it is materially shaped by computational systems.

The rise of artificial intelligence has made it possible for neoclassical assumptions—which until now remained largely theoretical and abstract—to move closer to empirical reality. Today, deep learning algorithms, generative neural networks, graph-based

---

1 Cass R Sunstein, *How Change Happens* (MIT Press, 2019) 44–58.

2 Cary Coglianese and David Lehr, 'Regulating by Robot: Administrative Decision-Making in the Machine Learning Era' (2017) 105 *Georgetown Law Journal* 1147, 1153–62.

models, and multilayered predictive systems possess the capacity to analyse millions of legal, economic, and behavioural data points and to forecast legal outcomes through advanced computational methods.[3]

The more pressing question is not whether artificial intelligence replaces human rationality in the economic analysis of law, but how AI reconfigures the operational environment within which the assumption of rational behaviour functions. In the classical framework, rationality operates as a simplifying assumption about how individuals respond to legal rules by maximising their subjective utility. The emergence of AI does not displace this assumption. Rather, it introduces new forms of informational infrastructure and decision-support systems that reshape how such utility-maximising behaviour is expressed, predicted, and governed.

Accordingly, the central claim of this article is not that artificial intelligence replaces human rationality, but that it alters the epistemic conditions under which rationality is operationalised in legal analysis. In this transformed environment, economic analysis of law increasingly relies on computational systems that optimise decisions at a systemic level, thereby shifting attention from individual choice to the architecture of decision-making systems.

Within this framework, the issues of legitimacy and justice also enter a new stage. Previously, critiques of economic analysis of law largely focused on the claim that efficiency is not a sufficient normative justification for legal rules and cannot adequately address the justice-oriented demands of law.[4] With the emergence of artificial intelligence, however, these challenges have not been resolved but rather intensified. Algorithms are structurally exposed to data biases, informational asymmetries, computational determinism, and weakened explainability—features that strike directly at the core of the legitimacy of legal decisions. As a result, the reliance of law on algorithmic models within the framework of economic analysis of law transforms both the foundation of efficiency and the foundation of legitimacy at an epistemological level.

The existing literature shows that scholars have conducted extensive research on issues such as algorithmic discrimination, data-driven justice, explainability, accountability, and algorithmic legitimacy.[5]

---

3 Jonah B Gelbach and David Marcus, 'Rethinking Judicial Review of High-Volume Agency Adjudication' (2017-2018) 96 *Texas Law Review* 1097.

4 Louis Kaplow and Steven Shavell, *Fairness versus Welfare* (Harvard University Press, 2006) 10–18.

5 Sandra Wachter, Brent Mittelstadt and Chris Russell, 'Bias Preservation in Machine Learning: The Legality of Fairness Metrics Under EU Non-Discrimination Law' (2021) 4 *West Virginia law review.* 735; Karen Yeung, 'Algorithmic Regulation' (2017) 12 *Regulation & Governance* 1, 8–16; Andrew D Selbst and Solon Barocas, 'The Intuitive Appeal of Explainable AI' (2019) 87 *Fordham Law Review* 1085, 1093–1104; Ben Greene et al, 'Algorithmic Realism: Expanding the Boundaries of Algorithmic Thought' (2020) Conference on Fairness,

While a growing body of literature has begun to explore the intersection between artificial intelligence and legal decision-making—including work on algorithmic governance, platform regulation, and computational law—these contributions have not yet been systematically integrated into the theoretical framework of economic analysis of law. This article builds on these emerging strands while seeking to provide a more explicit epistemological account of their implications.

In contrast to much of the contemporary literature on law and artificial intelligence, which has concentrated on improving efficiency, accuracy, or speed in legal decision-making, this article advances a different claim. Its central argument is that the expansion of algorithmic and data-driven rationalities does not merely provide law with new tools but fundamentally disrupts the epistemological and normative foundations of economic analysis of law. Within this framework, artificial intelligence is not a gradual technological development but a revelatory force that exposes the limitations of efficiency-based reasoning in legal analysis. The significance of artificial intelligence in this article lies in its demonstration of how the replacement of human judgment with predictive rationality transforms the very assumptions that once enabled economic analysis of law to assert normative authority.

This study adopts a conceptual and philosophical approach grounded in theoretical analysis. Its line of argument unfolds in three interconnected steps. First, it undertakes a conceptual reconstruction by revisiting the classical assumptions of the economic analysis of law, particularly those concerning human rationality, efficiency, and legitimacy. Second, it engages in a philosophical reconsideration of these assumptions, critiquing and redefining them through insights drawn from epistemology and the philosophy of technology. Finally, it explores the normative implications of this redefinition, articulating key principles such as epistemological substitution and the transition toward algorithmic legitimacy.

The aim of the article is to provide a foundational theoretical framework that can serve as the basis for future empirical and doctrinal studies, rather than to examine practical instances.

This article introduces several conceptual tools that are central to its analysis. "Algorithmic rationality" refers not to a replacement of human rationality, but to the role of computational systems in structuring decision-making processes. The "epistemic substitution principle" is

---

Accountability, and Transparency, January 27–30, 2020, Barcelona, Spain. ACM, New York, NY, USA, 13 pages. https://doi.org/10.1145/3351095.3372840, 3–9; Jenna Burrell, 'How the Machine "Thinks"' (2016) 3 *Big Data & Society* 1, 5–8; Michael Veale and Lilian Edwards, 'Clarity, Surprises, and Further Questions in Algorithmic Decision-Making' (2018) 31 *Computer Law & Security Review* 1–6;

used heuristically to describe the increasing reliance on computational processes in the production of legal knowledge, without implying a complete displacement of human judgment. Similarly, the "algorithmic efficiency paradigm" captures the growing importance of system-level optimisation metrics alongside traditional welfare-based approaches. Finally, the notion of an "algorithmic legitimacy shift" refers to the increasing reliance on statistical outputs in assessing the acceptability of legal decisions.

The article proceeds in five parts. The first section examines how AI reshapes the role of rationality in law and economics. The second revisits the concept of efficiency in light of computational optimisation. The third analyses the implications for legal legitimacy. The fourth explores the transformation of justice under automated decision-making. The final section considers practical implications through selected case studies.

## II From Human Rationality to Algorithmic Rationality

The transformation brought about by the advent of artificial intelligence in economic analysis of law first emerges at the level of the concept of "rationality"—a notion that, from the outset in Posner's theory, has served as both the epistemic and practical cornerstone of economic analysis of law. Posner conceived rationality as a characteristic of the human actor: an agent who makes decisions by calculating costs and benefits, maintains stable preferences, and adopts optimal behaviour when confronted with legal rules.[6] Although this rationality was abstract and idealised, it consistently presupposed that legal decision-making was grounded in human mental calculation—a calculation carried out under relatively simple conditions and with limited but manageable information.

In economic analysis of law, rationality is typically framed through the model of "rational choice." This means that the human actor, given stable preferences and available information, selects the option that maximises expected utility. Symbolically, this assumption can be expressed as:

$$a^* = \arg\max_{a \in A} \; U(a \mid I, P)$$

$\boldsymbol{a}^*$= the chosen action

$\boldsymbol{A}$= the set of available actions

$\boldsymbol{U(a \mid I, P)}$= expected utility of action $a$, given information $I$ and preferences $P$

$\boldsymbol{I}$= information accessible to the actor

$\boldsymbol{P}$= stable preference structure

6 Richard Posner, *Economic Analysis of Law* (7th ed, Aspen 2007) 15-22.

With the advent of machine-learning technologies and deep neural networks, this conception of rationality has undergone a fundamental transformation. Artificial intelligence has revealed that rationality is not merely a property of the human mind, but rather the function of a "computational system" capable of generating efficient, consistent, and stable decisions from data patterns that remain incomprehensible to humans.[7] In behavioural economics, extensive critiques had already been directed at Posner's assumptions concerning the stability of preferences, the coherence of mental calculation, and the cognitive capacities of human actors. Yet these critiques never shifted the process of rationality from the human subject to the computational system.[8]

The transformation associated with artificial intelligence in the economic analysis of law first becomes visible at the level of how "rationality" is operationalised, rather than at the level of its formal definition. In the Posnerian framework, rationality functions as a simplifying assumption that enables the modelling of how individuals respond to legal rules by maximising their utility. This assumption does not depend on the cognitive accuracy of decision-making, but on its predictive usefulness. The introduction of artificial intelligence does not displace this assumption; rather, it alters the informational and institutional context within which it operates.

In the Posnerian version of law and economics, the theory rests on the assumption that human beings are rational actors seeking to maximise their utility.[9] In his early works, Posner conceived rationality as the essence of individual behaviour within the legal system—behaviour that renders individuals' responses to rules, sanctions, and incentives predictable.[10] From this perspective, law is a set of rules that, by altering costs and benefits, shapes the trajectory of rational human conduct. This conception of rationality relied on three fundamental human capacities: the ability to process information; the ability to anticipate consequences; and the ability to make optimal choices under constraints.

In the algorithmic era, these three capacities are no longer exclusive to human beings; rather, machine-learning systems perform them in many domains with far greater accuracy, speed, and consistency.[11] Generative neural networks, reinforcement-learning algorithms, and graph-based models are capable of processing millions of unstructured data points, extracting hidden relationships among economic and legal

---

7 Coglianese and Lehr (n 2).

8 Daniel Kahneman, *Thinking, Fast and Slow* (Farrar, Straus and Giroux, 2011) 19-97.

9 Richard A Posner, 'Rational Choice, Behavioral Economics, and the Law' (1998) 50 *Stanford Law Review* 1551.

10 Richard Posner, *The Problems of Jurisprudence* (Harvard University Press, 1990) 19-40.

11 Stuart Russell and Peter Norvig, *Artificial Intelligence: A Modern Approach* (4th ed, Pearson 2020) 695–850.

variables, and generating predictions that even human experts cannot achieve.[12] Rationality is thus transferred from the individual level to the level of "systemic architecture"—an architecture in which decisions are produced not through mental calculation but through mathematical and statistical processes. In other words, whereas in Posner's model rationality resided "within the human," in the age of artificial intelligence rationality resides "within the algorithm." In this context, rationality remains attributed to human actors within legal-economic analysis, but its empirical modelling increasingly depends on computational systems that process large-scale data and generate predictive insights.

This transformation carries a significant implication: economic analysis of law, which previously treated human behaviour as the foundation of explanation, must now incorporate the behaviour of computational systems into the domain of legal understanding. In this new structure, law is no longer merely a set of incentives shaping behaviour; it becomes part of the data upon which the algorithm constructs its optimal system.

To understand the necessity of this transition, one must attend to the limitations of human beings in executing optimisation under constraints. Economic theory assumes that individuals are capable of making optimal choices given temporal, informational, and cost constraints. Yet extensive research in cognitive science and behavioural economics has demonstrated that, in many legal, economic, and judicial contexts, humans lack the capacity for systematic information processing. [13]

Human rationality is constrained by several limitations. Firstly, the human mind cannot simultaneously analyse millions of behavioural data points, judicial precedents, economic variables, and social factors. [14] Additionally, from confirmation bias to availability bias, human decision-making is permeated by psychological distortions that obstruct full rationality. [15] Moreover, contrary to the neoclassical assumption, human preferences are dynamic, context-dependent, and often inconsistent. [16] Individuals also tend to focus on short-term consequences and lack the ability to compute complex future

---

12 Yann LeCun, Yoshua Bengio and Geoffrey Hinton, 'Deep Learning' (2015) 521 *Nature* 436.

13 Cass Sunstein, 'Boundedly Rational Borrowing' (2006) 73 *University of Chicago Law Review* 249.

14 Herbert Simon, 'A Behavioral Model of Rational Choice' (1955) 69(1) *TheQuarterly Journal of Economics* 99.

15 Amos Tversky and Daniel Kahneman, 'Judgment under Uncertainty: Heuristics and Biases' (1974) 185(4157) *Science* 1124.

16 Shane Frederick, George Loewenstein and Ted O'Donoghue, 'Time Discounting and Time Preference: A Critical Review' (2002) 40(2) *Journal of Economic Literature* 351.

cost-benefit streams with precision.[17] Finally, unlike algorithms, humans cannot systematically reconfigure their decision pathways in light of past errors.[18]

In contrast, artificial intelligence systems directly overcome precisely those points at which human rationality is constrained. Algorithms are capable of processing millions of structured and unstructured data points in fractions of a second, extracting hidden relationships among legal variables through predictive models, and executing optimisation under constraints with mathematical precision.[19]

Optimisation available to algorithms does not merely mean faster calculation; it signifies the ability to reconstruct and re-engineer the entire "decision space"—a task fundamentally beyond human capacity. For example, in high-volume judicial systems, humans cannot compare thousands of similar cases or model the behavioural patterns of judges and defendants. Yet AI systems can perform these tasks with high accuracy and even simulate the systemic consequences of implementing a criminal policy.[20]

This capability enables, for the first time, the realisation of rationality at a level beyond the individual—at the level of the legal system itself—what may be termed systemic rationality.

At this stage, a profound epistemic transformation occurs: artificial intelligence is no longer merely a decision-support tool or auxiliary system but becomes a producer of rationality. This shift can be articulated through an innovative principle—the Epistemic Substitution Principle. According to this principle, rationality is the product of algorithmic computation, while human perception is treated merely as one variable within the model, rather than its starting point.[21]

In this sense, human judgment is not eliminated but increasingly mediated by algorithmic systems that structure how legal decisions are analysed and predicted.

This transformation has several implications for economic analysis of law. First, legal knowledge is increasingly informed by predictive models that analyse large-scale data, complementing traditional behavioural assumptions. Second, the evaluation of legal rules becomes more closely tied to system-level performance metrics, rather than

---

17 Edward L Glaeser, 'Paternalism and Psychology' (Winter, 2006) 73(1) *University of Chicago Law Review* 133.

18 Sendhil Mullainathan and Eldar Shafir, *Scarcity: Why Having Too Little Means So Much* (Times Books, 2013) 31–45.

19 Jon Kleinberg et al, 'Human Decisions and Machine Predictions' (2018) 133(1) *Quarterly Journal of Economics* 237.

20 Jonah Gelbach and David Marcus (n 3).

21 Mireille Hildebrandt, 'Law as Computation in the Era of Artificial Legal Intelligence' (2018) 68 *University of Toronto Law Journal* 12.

solely to individual behavioural responses. Third, questions of legitimacy and justification are increasingly shaped by the interaction between normative reasoning and computational outputs.[22]

The transformation can be clearly observed in the visa cancellation mechanism under section 501 of the *Migration Act 1958* (Cth). This transformation became particularly pronounced after the 2014 amendments, which expanded the scope of mandatory cancellation and introduced rigid criminal thresholds, thereby shifting decision-making from the individualised assessment of an administrative officer to the automatic application of data-driven criteria.[23]

Within this framework, visa cancellation is no longer the outcome of human evaluation of an individual's circumstances, but rather the result of processing criminal records, judicial histories, and risk indicators within an automated decision-making system. Although this mechanism does not necessarily rely on "strong AI," it is nonetheless fully grounded in algorithmic logic, administrative automation, and statistical prediction. Consequently, law ceases to function merely as an instrument for guiding human behaviour and instead becomes part of the input data for computational systems. Concepts such as "risk to the community" are reduced to quantifiable indicators, and the legitimacy of legal decisions depends less on human assessment than on the coherence and predictive reliability of algorithmic outputs.[24]

Academic critiques have demonstrated that this mechanism, by weakening natural justice and individualised evaluation, constitutes a clear instance of *epistemic substitution*: legal rationality no longer arises from human understanding but is produced through automated decision-making and data processing.[25] This example illustrates how the entry of algorithmic logic into law structurally challenges the human-centered foundations of economic analysis of law.

This example should not be understood as evidence of the replacement of human rationality, but rather as an instance of how decision-making processes become structured by rule-based and data-driven criteria. The legal framework continues to assume rational compliance and response from individuals, yet the operationalisation of decision-making increasingly reflects systemic and administrative optimisation concerns.

---

22 Sandra Wachter, Brent Mittelstadt and Chris Russell (n 5).

23 *Migration Amendment (Character and General Visa Cancellation) Act 2014 (Cth)*; *Migration Act 1958 (Cth)* s 501.

24 Mary Crock and Laurie Berg, *Immigration, Refugees and Forced Migration: Law, Policy and Practice in Australia* (Federation Press, 2011) ch 12.

25 Australian Human Rights Commission, *Background Paper: Human Rights Issues Raised by Visa Refusal or Cancellation under Section 501 of the Migration Act* (June 2013).

## III Redefining Efficiency within the Computational Order

The transformation discussed in the previous section does not displace the concept of efficiency in economic analysis of law, but invites a reconsideration of how efficiency is defined, measured, and operationalised in increasingly data-driven environments. Within the framework of law and economics, efficiency is conventionally associated either with the minimisation of social costs or with the maximisation of aggregate welfare. In condensed form, efficiency may be expressed as follows:

$$\max W = \sum_{i=1}^{n} U_i$$

$\boldsymbol{W}^*$= the maximised level of social welfare
$\boldsymbol{U_i}$= utility of individual $i$in society
$\boldsymbol{n}$= the total number of individuals considered
$\sum_{\boldsymbol{i=1}}^{\boldsymbol{n}} \boldsymbol{U_i}$ = aggregate welfare function, representing the sum of individual utilities

This formulation captures efficiency as an aggregate measure of individual welfare. Importantly, it remains a normative construct, even when used in positive analysis.

In classical law and economics, rationality and efficiency were conceived as coextensive notions: human rationality was regarded as the instrument through which welfare efficiency could be realised. In the context of artificial intelligence, this relationship becomes more complex. While rational behaviour continues to be modelled at the individual level, efficiency is increasingly assessed through system-level performance metrics. The introduction of artificial intelligence does not eliminate the welfare-based understanding of efficiency. However, it introduces additional layers through which efficiency is evaluated. In many contemporary systems, efficiency is increasingly operationalised through metrics such as predictive accuracy, error reduction, and systemic stability. These metrics do not replace welfare considerations, but they can reshape how welfare is approximated, measured, and sometimes prioritised within decision-making processes.

Within the classical framework of law and economics, efficiency was invariably derived from human welfare. Even across its diverse formulations—whether in the Kaldor–Hicks criterion, Pareto optimality, or wealth maximisation—the foundational assumption was that the legal system ought to be designed to enhance the welfare conditions of individuals.[26] Efficiency, in this sense, was not an

---

26 Guido Calabresi, *The Costs of Accidents: A Legal and Economic Analysis* (Yale University Press, 1970) 25.

autonomous technical construct but rather an index for measuring social utility.

Within this tradition, legal rules were deemed "efficient" insofar as they either increased the aggregate of individual utilities or, at the very least, enabled winners to compensate losers.[27] Even when certain theorists sought to disentangle efficiency from justice, efficiency nonetheless remained a criterion for the organisation of human interests.[28] Law, under this conception, functioned as an instrument for regulating human behaviour, and efficiency served as the measure of success in that regulation.

Crucially, this understanding of efficiency implicitly rested upon three foundational assumptions: that human welfare is measurable, or at least amenable to economic articulation; that legal outcomes can be predicted through relatively simple models; and that the legal decision-maker—whether judge, legislator, or regulatory authority—can reasonably perform such calculations.[29]

Yet these three assumptions are precisely the points that collapse with the advent of artificial intelligence. Once legal decision-making is entrusted to deep learning systems, human welfare no longer necessarily constitutes the target variable; rather, it becomes merely one possible parameter within a far more complex optimisation space.

Algorithmic systems are typically designed to optimise specific objective functions—such as accuracy, efficiency, or stability—which may serve as proxies for broader welfare considerations. However, the relationship between these technical objectives and human welfare is neither direct nor guaranteed and requires careful interpretation within legal contexts. Within algorithmic architectures, what matters is the reduction of error, the enhancement of predictive accuracy, the stability of outputs, and the maximisation of overall system performance.[30]

In this framework, efficiency no longer signifies that "human beings are better off," but rather that "the system functions more effectively." For instance, within a criminal risk-prediction regime, efficiency may be defined in terms of reducing error rates, rather than necessarily improving the lived experience of defendants or advancing distributive justice.[31] In an algorithmic financial regulatory system, efficiency may be equated with systemic stability, even if certain human actors incur

---

27 Robert Cooter and Thomas Ulen, *Law and Economics* (6th ed, Pearson 2016) 45.

28 Jules L Coleman, 'Efficiency, Utility, and Wealth Maximization' (1980) 8 *Hofstra Law Review* 509.

29 Duncan Kennedy, 'Cost-Benefit Analysis of Entitlement Problems' (1981) 33 *Stanford Law Review* 387.

30 Pedro Domingos, *The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World* (Basic Books, 2015) 87.

31 Bernard Harcourt, *Against Prediction: Profiling, Policing, and Punishing in an Actuarial Age* (University of Chicago Press, 2007) 45.

losses.[32] In some contexts, the increasing reliance on system-level optimisation may shift attention away from direct assessments of human welfare toward the performance of decision-making systems. Thus, law is no longer designed to "improve human conditions," but is reconfigured to "optimise the performance of decision-making systems." Within such an order, human welfare may be sacrificed to the imperatives of stability, predictability, or statistical efficiency.[33]

One tangible indicator of this transformation is the shift in the instruments used to calculate efficiency. In neoclassical economics, efficiency was primarily assessed through analytical equations, utility functions, and equilibrium models.[34] These models were necessarily simplified and, as a result, ignored many real-world variables.

By contrast, artificial intelligence calculates efficiency through deep learning models—models that are not only nonlinear but also frequently lack interpretive transparency.[35] In such models, efficiency emerges from the interaction of millions of parameters that are dynamically updated. In such models, efficiency is often derived from complex statistical relationships that may be difficult to interpret, rather than from transparent analytical formulations. This can make the justification of efficiency-based decisions more challenging, even if their predictive performance improves.

This new formalisation carries significant implications: first, the possibility of normative critique of efficiency is diminished, since it is no longer clear why a given outcome is deemed "efficient";[36] second, the scope for human intervention in setting efficiency criteria is restricted, as the system itself optimises those criteria;[37] and third, efficiency, rather than serving as a tool for evaluating rules, becomes the very metric by which rules are generated.[38]

Thus, once efficiency becomes the product of systemic optimisation, its connection to human values—justice, dignity, equality—weakens.[39]

In such a situation, efficiency is presented not as a "reason" but as a "statistical fact." Legal decisions are no longer justified because they are fair or equitable, but because the model has demonstrated that they

---

32 Dan Awrey, 'Complexity, Innovation and the Regulation of Modern Financial Markets' (2012) 2 *Harvard Business Law Review* 235, 240.

33 Frank Pasquale, *The Black Box Society: The Secret Algorithms That Control Money and Information* (Harvard University Press, 2015) 101.

34 Hal Varian, *Microeconomic Analysis* (3rd ed, W W Norton 1992) 213.

35 Cathy O'Neil, *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy* (Crown, 2017) 21.

36 Jenna Burrell (n 5).

37 Nick Seaver, 'Algorithms as Culture: Some tactics for the ethnography of algorithmic systems' (2017) *Big Data & Society.* Doi: 10.1177/2053951717738104

38 Stephen E Margolis, 'Two Definitions of Efficiency in Law and Economics' (1987) 16(2) *Journal of Legal Studies* 471, 475.

39 Thomas Nagel, *Equality and Partiality* (Oxford University Press, 1991) 62.

are less error-prone or more stable.[40] This shift profoundly undermines the normative foundations of economic analysis of law.

The RoboDebt program (2016–2020)[41] illustrates how efficiency metrics based on automation and large-scale data processing can shape administrative decision-making. Relying on automated decision-making, the system calculated and claimed individuals' debts without human intervention by cross-matching tax data with Centrelink welfare information.[42] From an administrative perspective, the system was able to process an enormous volume of cases at minimal cost and maximum speed—a metric that, within algorithmic logic, is regarded as success.

The significance of RoboDebt lies not merely in its failure, but in how it reveals the tension between different conceptions of efficiency. From a system-level perspective, the program achieved scale and cost reduction. From a welfare perspective, however, it produced significant harm due to flawed assumptions and lack of human oversight. This tension demonstrates that computational efficiency and welfare-based efficiency may diverge, rather than one replacing the other.

Yet systemic optimisation based on flawed computational assumptions led to the mass production of erroneous debts and erased human responsibility from the decision-making chain. The Federal Court found this system to lack a valid legal foundation and declared it in violation of fundamental principles of administrative fairness.[43] In this way, RoboDebt demonstrated that algorithmic efficiency, once severed from human welfare, not only undermines legal legitimacy but also challenges the welfare foundations of economic analysis of law.

From this perspective, one may say that economic analysis of law in the age of artificial intelligence faces a paradox: as efficiency becomes increasingly mediated by computational systems, its normative interpretation may become less transparent and more contested.

The critique advanced in this article is not directed against economic reasoning as an analytical tool, nor does it deny its descriptive or instrumental usefulness. What is contested is the normative authority of law and economics—particularly in its efficiency-centred version—as a framework for justifying legal decisions. The central problem is that, with the entry of algorithmic rationality into legal decision-making processes, foundational assumptions such as the rational actor, welfare maximisation, and equilibrium can no longer adequately explain or justify the actual mechanisms of decision.

---

40 Danielle Allen, *Justice by Means of Democracy* (University of Chicago Press, 2023) 87.

41 *Royal Commission into the Robodebt Scheme* (Final Report (July 2023) vol 1, 45.

42 Terry Carney, 'Robo-debt Illegality: The Seven Veils of Failed Guarantees of the rule of law' (2019) 44(1) *Alternative Law Journal* 106, 108.

43 *Amato v Commonwealth of Australia* [2019] FCA 780, [45]; *Prygodicz v Commonwealth of Australia [2021]* FCA 634, [52].

This development does not eliminate the normative role of efficiency but complicates its relationship with legal justification. As efficiency becomes intertwined with technical optimisation, questions arise regarding how such metrics can be translated into legitimate legal reasoning. This tension sets the stage for the analysis of legal legitimacy in the following section.

## IV The Evolution of Legal Legitimacy under Non-Human Rationality

The central claim of this section is not that artificial intelligence complicates the relationship between legal justification and decision-making processes increasingly shaped by computational systems. Rather than replacing normative reasoning, algorithmic decision-making introduces new conditions under which legitimacy must be articulated and defended.

Within this framework, legitimacy is no longer assessed primarily through the question of whether a rule is just or fair; but also, by the question of whether the output of a computational system is "efficient," "low-error," or "stable." This displacement is what the present article designates as the Algorithmic Legitimacy Shift. It does not imply that legitimacy is replaced by computational outputs. Rather, it refers to the increasing reliance on statistical and predictive systems in shaping how legal decisions are evaluated, thereby creating tension between traditional forms of normative justification and emerging forms of technical validation

Prior to the introduction of artificial intelligence, the economic analysis of law had always been confronted with an unresolved tension concerning legitimacy: while the approach sought to justify legal rules on the basis of efficiency, it was unable to offer a persuasive account of why efficiency ought to function as the ultimate criterion of normative justification. This crisis was not historical but structural, arising from within the internal logic of the economic analysis of law itself.

The first axis of these critiques focused on the reduction of justice to aggregate social welfare. Critics demonstrated that even if a legal rule is efficient from the standpoint of total social welfare, it may nevertheless be profoundly unjust or in violation of individuals' fundamental rights. From this perspective, efficiency lacks independent justificatory force and cannot substitute for concepts such as rights, equality, and dignity. The core argument was that law is not merely a mechanism for the allocation of resources, but a normative institution endowed with autonomous moral commitments.[44]

---

44 Amartya Sen, *The Idea of Justice* (Harvard University Press, 2009) 30-35.

The second axis of critique focused on the indifference of the economic analysis of law to distributive concerns. Even ostensibly more moderate criteria, such as the Kaldor–Hicks standard, presuppose that the *possibility* of compensation is sufficient for legitimacy, without requiring that such compensation actually occur. As a result, legal legitimacy becomes contingent upon an abstract and empirically unreal assumption.[45]

The third axis constituted an epistemological critique: the economic analysis of law justified legal decisions on the basis of models that themselves rested on overly simplistic assumptions about human behaviour. Within this framework, legitimacy was grounded in a form of "reasoning" whose empirical and cognitive foundations were themselves open to serious doubt.[46]

These critiques remain relevant in the age of artificial intelligence, but they are now reframed within a context in which decision-making processes are increasingly mediated by computational systems.

The crucial point is that, across all these critiques, legitimacy continued to remain a concept dependent on human reasoning. Even the most radical critics of the economic analysis of law assumed that a legal decision must be answerable to human rationality—that is, it must be capable of explaining why a decision is justified despite its costs.[47]

The introduction of artificial intelligence fundamentally alters this equation. Whereas the earlier problem lay in the unpersuasive character of efficiency-based reasoning, the present issue is the difficulty of translating computational processes into forms of reasoning that are recognisable within legal doctrine. Algorithmic decisions, even when they generate accurate outcomes, are not necessarily grounded in a reconstructible chain of reasons.[48]

In some contexts, there is a tendency to rely on computational outputs as indicators of reliability or accuracy, which may, in practice, influence perceptions of legitimacy even when the underlying reasoning is not fully transparent. The system's output is accepted because "the model has said so," rather than because it can be articulated in terms of traditional legal concepts—such as proportionality, fairness, or responsibility.[49]

This dynamic may shift attention from the reasoning process toward the reliability and consistency of outcomes, without fully eliminating the importance of procedural justification. In the algorithmic order,

---

45 Louis Kaplow and Steven Shavell (n 4).

46 Ronald Dworkin, *Law's Empire* (Harvard University Press, 1986) 151-175.

47 Jeremy Waldron, 'The Concept and the Rule of Law' (2008) 43 *Georgia Law Review* 1, 6-12.

48 Frank Pasquale (n 34) 6.

49 Mireille Hildebrandt, 'Algorithmic Regulation and the Rule of Law' (2018) 376(2128) *Philosophical Transactions* 1, 3-8.

however, the decision-making process occurs within a computational black box, and legitimacy becomes tied solely to the final output.[50]

At this point, the concept of the Algorithmic Legitimacy Shift can be articulated: the transition of legal legitimacy from normative justification grounded in human reasoning to technical acceptance based on statistical computation.

This transition is not merely an instrumental change, but a transformation in the logic of legal authority. Law becomes binding because the system has deemed it optimal.[51]

In response to this crisis, the concept of explainability has been proposed as a techno-ethical solution. The core idea is that if algorithms can be made explainable, the problem of legitimacy will be resolved. Explainability can provide insight into how a model operates, but it does not necessarily fulfill the requirements of legal reasoning, which involves normative justification, addressability, and accountability. The distinction between explanation and justification therefore remains crucial. Explaining which variables carried the most weight is not equivalent to answering the question of why an individual should normatively be subject to the decision.[52]

Legal reasoning, in contrast to technical explanation, possesses three distinct characteristics:

- Normativity: it refers to values, not merely statistical correlations.
- Addressability: it is capable of persuading a human audience.
- Accountability: it allows the decision to be attributed to a responsible agent.[53]

Algorithmic explainability guarantees none of these three characteristics. Even if a model is fully transparent, it still cannot explain why a statistical pattern should serve as the basis for legal obligation.[54] Consequently, explainability cannot substitute for legal reasoning. The two operate on entirely different levels: one at the level of computation, the other at the level of justification.[55]

The peak of the legitimacy crisis lies in the fact that data carry no inherent values. Data merely reflect past patterns—patterns that may

---

50 Karen Yeung, (n 5) 505.

51 Julie E Cohen, 'Law for the Platform Economy' (2017) 51 *UC Davis Law Review* 133, 140-146.

52 Andrew L Selbst and Solon Barocas (n 5) 1085.

53 Neil MacCormick, *Rhetoric and the Rule of Law* (Oxford University Press, 2005) 54-60.

54 Sandra Wachter, Brent Mittelstadt and Chris Russell, 'Counterfactual Explanations without opening the black box: Automated decisions and the GDPR' (2018) 31 *Harvard Journal of Law & Technology* 841, 852–857.

55 Joseph Raz, *The Authority of Law: Essays on Law and Morality* (Oxford University Press, 1979) 26–31.

themselves be the product of existing power structures, inequalities, and prior decisions.[56] When the legal system relies on data as the basis for decision-making, it effectively reproduces these embedded historical values without any normative reflection. This is particularly perilous in areas such as criminal law, financial regulation, and social policy.[57]

The use of predictive policing systems provides an example of how computational tools can influence decision-making within legal institutions. These programs rely on analyses of criminal, social, and spatial data to identify "high-risk" areas and guide the allocation of police resources based on statistical outputs. Within this framework, the legitimacy of police intervention derives not from individual-centred legal justification, but from predictive accuracy and algorithmic coherence.

The significance of these systems lies in how they reshape the relationship between justification and prediction. While traditional legal legitimacy depends on reasons addressed to individuals, predictive systems operate by identifying statistical patterns at the population level. This creates a tension: decisions may be supported by strong predictive evidence while remaining difficult to justify in individual normative terms.

This shift carries significant institutional implications. Traditional criteria of legitimacy, such as transparency and accountability, gradually become overshadowed by technical metrics, like statistical accuracy. At the same time, algorithms, by reproducing existing patterns in the data, increase the risk of entrenching social inequalities under the guise of seemingly neutral and scientific decisions.[58] In such a context, the algorithmic output itself becomes a source of legitimacy, even when the underlying data are questionable.

Law, unlike algorithms, deals not only with "what is" but also with "what ought to be." Machine learning models, however, fundamentally lack the capacity to distinguish between the two. They can detect patterns, but they cannot evaluate their normative desirability.[59] In this sense, reliance on non-human rationality in law carries the risk of stripping legal norms of their value-laden content. Legal legitimacy, if reduced to data-driven outputs, will no longer rest on justice, but on statistical repetition.[60]

---

56 Cathy O'Neil (n 34) 84.

57 Bernard Harcourt (n 30) 43.

58 Sarah Brayne, 'Big Data Surveillance: The Case of Policing' (2017) 82(5) *American Sociological Review* 977, 980–986; Andrew Goldsmith, 'Policing's New Visibility' (2010) 50(5) *British Journal of Criminology* 914, 920–926.

59 Thomas Nagel (n 38) 10.

60 Danielle Allen (n 40) 29.

## V Justice and the Challenge of Automated Rationality

Justice, more than any other concept in law, has resisted reduction. Even as law has become increasingly technocratic, economic, or managerial, justice has remained a point against which the law has been compelled to justify itself. Automated rationality introduces new pressures on this historical balance, particularly by emphasising consistency, scalability, and pattern-based decision-making within legal systems.

The central claim of this section is that automated rationality transforms justice not only at the level of implementation but also at the conceptual level. contemporary legal systems increasingly incorporate statistical and system-level considerations into how justice is operationalised. This creates a tension between individualised and systemic conceptions of justice.

In the modern legal tradition, justice has always been linked to the identifiable individual. Even theories that focus on structure or distribution ultimately return to the question of whether a particular person can regard themselves as the addressee of a just decision.[61] In this sense, justice required that the individual be able to understand the decision, hear its reasons, and attribute it to a responsible agent.[62]

Automated rationality gradually displaces this framework. Algorithmic systems are not designed to respond to individual claims; they are built to regulate collective patterns. What matters is not the condition of the individual, but the performance of the system at the population level. [63] justice may, in certain contexts, be partially reframed as a property of decision-making systems, without eliminating its relational and individual dimensions.[64]

This transition from individual justice to systemic justice may at first appear as progress. Systemic justice promises to reduce randomness, personal discretion, and even unconscious human bias.[65]

In systemic justice, a decision may be manifestly unjust for an individual yet still be defended as "just" because the system as a whole performs satisfactorily. [66] This logic transforms justice from a

---

61 John Rawls, *Justice as Fairness: A Restatement* (Harvard University Press, 2001) 17–19.

62 Neil MacCormick, *Legal Reasoning and Legal Theory* (Oxford University Press, 1978) 45–46.

63 Antoinette Rouvroy and Thomas Berns, 'Algorithmic Governmentality and prospects of emancipation: Disparateness as a precondition for individuation through relationships?' (2013) 177(1) *Réseaux* 163, 165–168.

64 Axel Honneth, *The Struggle for Recognition: The Moral Grammar of Social Conflicts* (Polity Press, 1995) 172.

65 Frederick Schauer, *Playing by the Rules: A Philosophical Examination of Rule-Based Decision-Making in Law and in Life* (Oxford University Press, 1991) 26.

66 Jeremy Waldron, 'The Rule of Law and the Importance of Procedure' (2011) 50 *Nomos* 3.

commitment to the individual into a criterion for evaluating the system—a shift that empties the human core of legal justice.[67]

One of the most innovative outcomes of analysing justice in the age of automated rationality is the formulation of the paradox of algorithmic justice. Contrary to the prevailing literature, the central problem of algorithmic justice is not "bias" understood as a correctable error, but rather the excessive success of algorithms in representing social reality.[68]

Algorithms are trained on data that are themselves the product of history—a history saturated with structural inequalities, discriminatory decisions, and unjust distributions of opportunities.[69] The more accurately a model identifies patterns in these data, the more faithfully it reproduces those very patterns.[70] Consequently, algorithmic justice neither creates nor remedies discrimination; rather, it normalises and entrenches it. This is the paradox: the greater the efficiency of prediction, the greater the durability of inequality.

This inequality is no longer perceived as injustice, for it appears in the guise of statistical necessity.[71] What once required moral justification is now accepted as the "natural" outcome of data.[72] In this condition, justice shifts from a standard for critiquing reality to a tool for optimally representing it—and this is precisely the point at which legal justice loses its critical function.[73]

Economic analysis of law, even prior to the advent of artificial intelligence, has always been confronted with the tension between efficiency and distributive justice.[74] Yet in its classical version, this tension remained manageable, since the human decision-maker could still incorporate normative considerations into the decision. Automated rationality structurally restricts this possibility. In the algorithmic order, efficiency is not merely one criterion among others, but the objective

---

67 Gerald J Postema, '*Law's Rule: The Nature*' (2022) 1 Oxford University Press, 45.

68 Marion Fourcade and Kieran Healy, 'Seeing Like a Market' (2017) 15 *Socio-Economic Review* 9.

69 Ruha Benjamin, *Race After Technology: Abolitionist Tools for the New Jim Code* (Polity Press, 2019) 100.

70 Virginia Eubanks, *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor* (St Martin's Press, 2018) 20-21.

71 Bernard Williams, *Moral Luck in Moral Luck: Philosophical Papers 1973–1980* (Cambridge University Press, 1981) 20–39.

72 Sheila Jasanoff, *The Ethics of Invention: Technology and the Human Future* (W W Norton, 2016) 7–9.

73 Roberto Mangabeira Unger, *Law in Modern Society: Toward a Criticism of Social Theory* (Free Press, 1976) 58.

74 Jules L Coleman (n 28).

function of the system.[75] Whatever cannot be formalised and optimised is gradually eliminated from the decision-making process.[76]

Distributive justice, if it cannot be translated into the language of data and models, is reduced to a minor technical constraint.[77] Consequently, in its algorithmic version, law and economics transforms from a theory for justifying rules into a technology of governance based on optimisation.[78]

The deepest consequence of automated rationality for justice is the displacement of the moral subject of decision. In classical law, a decision—even an unjust one—was attributed to a human agent. This attribution made possible blame, remorse, and correction.[79] In the algorithmic order, decision is attributed to the "system." Yet the system possesses neither conscience, nor moral responsibility, nor the capacity for normative doubt.[80] This condition transforms justice from a human virtue into a technical output. The decline of legal conscience is not merely an ethical concern; it undermines the very foundation of law's legitimacy. Law without human responsibility becomes a non-accountable mechanism.[81]

The introduction of automated decision-making systems into judicial and administrative processes in Australia signals a redefinition of justice from an individualised judgment to a structural and statistical feature. Recent institutional reports indicate that AI tools, particularly in small claims and case management, are employed as solutions to reduce costs, increase speed, and standardise decisions.[82] Within this framework, justice is defined not by the hearing of individuals and contextual evaluation of cases, but by the consistency of outputs and the minimisation of decision-making deviations.

From the perspective of law and economics, this transformation signifies a shift in the criterion of justice from individual fairness to systemic efficiency. A just decision is conceived as one that is produced at the lowest cost and with the greatest uniformity. The so-called "AI judge," in handling simple cases, reduces justice to indicators such as

---

75 Pedro Domingos (n 31) 280.

76 Julie E Cohen, *Between Truth and Power: The Legal Constructions of Informational Capitalism* (Oxford University Press, 2019) 145.

77 Elizabeth Anderson, *Private Government: How Employers Rule Our Lives (and Why We Don't Talk about It)* (Princeton University Press, 2017) 37–40.

78 Frank Pasquale (n 33) 101.

79 Hannah Arendt, *Responsibility and Judgment* (Schocken Books, 2003) 17-48.

80 Onora O'Neill, *Justice Across Boundaries: Whose Obligations?* (Cambridge University Press, 2016) 29-31.

81 Joseph Raz, above n 58, 3-36.

82 Australasian Institute of Judicial Administration, *AI Decision-Making and the Courts: A Guide for Judges, Tribunal Members and Court Administrators* (rev ed, 2023) 12-15.

speed, scalability, and predictability, while relegating the role of human judgment to the supervision of algorithmic outputs.[83]

The normative implication of this transition is that algorithmic justice can be both efficient and unjust at the same time. Efficient, insofar as it reduces the burden on courts and accelerates decision-making; and unjust, insofar as it restricts attention to individual differences, the hearing of voices, and the exercise of contextual fairness. Thus, justice, instead of being a human and responsible relationship, becomes the statistical output of a decision-making system—a transformation that carries the risk of eroding human responsibility and undermining normative legitimacy.

This section has shown that justice, when confronted with automated rationality, faces a crisis deeper than error or bias. The central issue is the transformation of the very nature of justice: from a normative, individual-centred commitment into a systemic and computational property. If this transformation is accepted without theoretical reconsideration, law risks being reduced from a moral institution to a regulatory technology. This conclusion directs the trajectory of the article toward its final claim: the theoretical reconstruction of law in the age of non-human rationality.[84] There is a risk that an uncritical reliance on automated systems may narrow the scope of law's normative commitments, shifting emphasis toward regulatory efficiency at the expense of broader moral considerations.

## VI Practical Implications in the Age of AI

The case of *O'Donnell v Commissioner of Police* (NSW)[85] provides a prominent example of how Australian courts have confronted data-driven decision-making in the criminal justice context. In this case, the system for determining pre-trial liberty or detention was structurally based on risk assessments and predictions of future offender behaviour—assessments that relied not on the judge's case-specific judgment but on statistical patterns and historical data.

In O'Donnell, the use of risk assessment mechanisms in determining pre-trial detention introduces a form of reasoning that is structurally distinct from classical judicial deliberation. While formally the judge remains the decision-maker, the architecture of decision-making is reorganised around predictive outputs derived from statistical models. From the perspective of classical economic analysis of law, such an arrangement can initially be justified: predictive accuracy enhances the

---

83 Tania Sourdin, *Judge v Robot? Artificial Intelligence and Judicial Decision-Making* (Judicial Commission of NSW, 2022) 1-5.

84 Commonwealth Ombudsman, *Automated Decision-Making – Better Practice Guide* (March 2025) 15–17.

85 *O'Donnell v Commissioner of Police (NSW)* 2015 NSWSC 1063.

allocation of judicial resources, reduces the social costs of crime, and aligns with an efficiency-oriented framework that seeks to minimise error in legal outcomes. On this reading, the integration of risk assessment appears as a refinement of rational adjudication rather than its transformation.

However, once examined through the lens developed in the earlier sections, this interpretation proves insufficient. The issue is not simply that additional information is introduced into the decision-making process, but that the structure of reasoning itself is displaced. The predictive model does not merely inform judicial reasoning; it reorganises it. Judicial deliberation becomes oriented toward the validation, interpretation, or acceptance of probabilistic outputs whose internal logic is not fully accessible within the traditional framework of legal reasoning. At this point, the Epistemic Substitution Principle becomes empirically visible: rationality is no longer generated through the interpretive activity of the judge but emerges from a computational process that the judge can neither fully reconstruct nor normatively justify within established doctrinal categories.

This transformation exposes a tension that the classical economic analysis of law is not equipped to resolve. While that tradition assumes that improved information leads to more efficient and therefore more desirable outcomes, it presupposes that the decision-maker retains epistemic authority over the evaluative process. In O'Donnell, by contrast, epistemic authority is partially relocated to the predictive system itself. The judge's role shifts from that of a rational optimiser of legal outcomes to that of an intermediary between legal norms and algorithmic outputs. This shift complicates the assumption that efficiency remains a human-centered evaluative criterion, as the determination of what counts as "efficient" is increasingly embedded in the design and operation of the predictive model.

The second dimension of analysis concerns the concept of efficiency as it operates within the case. At first glance, the use of predictive tools aligns closely with the economic logic of minimising both Type I and Type II errors in pre-trial detention decisions. Yet the case reveals that efficiency is no longer measured in relation to human welfare in any direct sense. Instead, it is operationalised through metrics such as predictive accuracy, risk calibration, and systemic consistency. These metrics reflect what has been described in Section 2 as the shift from welfare-based efficiency to systemic optimisation. The critical point is that this shift is not externally imposed but arises from within the logic of the decision-making system itself. The consequence is that efficiency ceases to function as a transparent normative criterion and instead becomes a technical property of the system. The judge does not evaluate efficiency in terms of competing human interests but relies on the model's capacity to produce stable and low-error predictions. In this

configuration, the economic analysis of law risks reinforcing a conception of efficiency that it can no longer normatively ground. The case thus reveals an internal tension within the law and economics framework: the more it embraces predictive accuracy as a proxy for efficiency, the more it distances itself from its own welfare-based foundations.

The third analytical dimension concerns legal legitimacy. In classical doctrine, legitimacy is tied to the capacity of the legal system to provide reasons that are intelligible, contestable, and attributable to human decision-makers. In O'Donnell, however, the increasing reliance on predictive assessments shifts the basis of legitimacy toward the perceived reliability of the system's outputs. The justification for detention is no longer grounded solely in the interpretation of legal norms or the evaluation of individual conduct, but in the statistical likelihood of future behaviour as generated by the model.

This development exemplifies the Algorithmic Legitimacy Shift identified earlier. Legitimacy becomes partially detached from normative reasoning and reattached to computational performance. Importantly, the case does not suggest that normative reasoning disappears entirely; rather, it becomes secondary to the validation of predictive outputs. The judge's reasoning is structured around the model's assessment, and the space for contestation shifts accordingly. What is challenged is no longer only the application of legal standards, but the reliability, design, and assumptions of the predictive system itself. This reconfiguration places significant strain on existing legal doctrines, which are not designed to accommodate challenges to epistemic infrastructures.

At this point, the methodological value of the case study becomes clear. Rather than merely illustrating the article's theoretical claims, O'Donnell reveals the points at which those claims encounter resistance and complexity. For instance, while the Epistemic Substitution Principle captures the relocation of rationality, the case shows that this substitution is neither complete nor uncontested. Judicial authority persists, but in a transformed and constrained form. Similarly, while the concept of algorithmic efficiency explains the operational logic of the system, the case demonstrates that this logic coexists uneasily with residual commitments to individualised justice and procedural fairness.

Engaging more directly with contemporary debates in AI governance further clarifies the significance of this tension. European regulatory frameworks, most notably the emerging literature surrounding the EU AI Act, emphasise principles such as transparency, accountability, and human oversight as mechanisms for preserving legal

legitimacy in algorithmic systems.[86] From that perspective, the issues raised in O'Donnell are not anomalies but symptomatic of a broader structural transformation in legal decision-making. Yet the case also suggests the limits of these regulatory responses. Even where oversight and explainability are introduced, they do not fully restore the conditions of normative justification that underpin classical legal reasoning. Instead, they operate within the same computational paradigm, seeking to render it more acceptable rather than fundamentally reconfiguring it. This observation reinforces the article's broader claim regarding the gap between AI & Law and law and economics. While AI governance literature has developed sophisticated tools for addressing issues of bias, transparency, and accountability, it often treats efficiency and optimisation as unproblematic objectives. Conversely, law and economics continues to rely on a conception of efficiency that presupposes human-centered rationality. The analysis of O'Donnell demonstrates that these two strands of literature intersect but do not fully integrate. The case thus serves as an empirical site in which this theoretical gap becomes visible.

Ultimately, the significance of O'Donnell lies in its capacity to transform the article's argument from an abstract claim into an analytically grounded account of legal change. The case does not simply confirm that algorithmic rationality is reshaping law; it shows how that reshaping occurs in practice, through incremental shifts in reasoning, evaluation, and justification. It also reveals that these shifts generate tensions that cannot be resolved within existing frameworks, whether doctrinal, economic, or regulatory.

In this sense, the case study fulfills a dual function. It substantiates the central thesis of the article by demonstrating the practical operation of epistemic substitution and algorithmic efficiency, while simultaneously exposing the limits of those concepts when confronted with the complexities of legal practice. The result is not a closure of the argument, but its deepening: the recognition that the transformation of economic analysis of law in the age of artificial intelligence is not merely theoretical but is already embedded in the institutional fabric of legal decision-making, where it continues to generate unresolved and, at times, irreducible tensions.

## VII Conclusion

This article has examined the implications of artificial intelligence for the economic analysis of law, with a particular focus on its epistemological foundations.

---

86 European Commission, 'Proposal for a Regulation of the European Parliament and of the Council Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act) and Amending Certain Union Legislative Acts' COM (2021) 206 final.

At the level of rationality, the article has clarified that economic analysis continues to rely on the assumption that individuals respond to incentives in utility-maximising ways. Artificial intelligence does not displace this assumption. However, it significantly alters the informational and institutional conditions under which such behaviour is modelled and predicted. Rationality remains formally attributed to human actors, but its operationalisation is increasingly shaped by data-driven systems that structure decision-making environments.

A similar pattern emerges in relation to efficiency. Traditionally grounded in welfare-based conceptions, efficiency remains a central evaluative criterion in economic analysis of law. Yet, in computational contexts, it is increasingly operationalised through technical metrics such as predictive accuracy, scalability, and system performance. These metrics do not replace welfare considerations, but they complicate the relationship between technical optimisation and normative evaluation, making the assessment of efficiency more contingent and context-dependent.

The analysis of legal legitimacy has further demonstrated that artificial intelligence does not eliminate the need for normative justification. Legal decisions must still be justified through reasons that are accessible, contestable, and grounded in institutional authority. However, the growing role of computational systems introduces new challenges, particularly where decision-making processes become difficult to translate into conventional forms of legal reasoning. This creates a tension between explanation and justification, requiring legal systems to reconsider how legitimacy is articulated in technologically complex environments.

In the domain of justice, the article has argued that the rise of automated decision-making introduces competing conceptions of justice within legal systems. Individualised justice, grounded in relational and contextual reasoning, continues to coexist with systemic approaches that emphasise consistency, scalability, and pattern recognition. The resulting tension highlights the difficulty of reconciling individual fairness with system-level optimisation in contemporary legal practice.

The case studies examined in this article have illustrated how these theoretical tensions manifest in practice. Across different legal domains, they reveal recurring patterns: the structuring of discretion through data-informed systems, the divergence between system-level efficiency and individual welfare, the challenge of maintaining legitimacy in the absence of transparent reasoning, and the coexistence of individual and systemic conceptions of justice. These examples confirm that the impact of artificial intelligence is neither uniform nor unidirectional but characterised by complexity and contestation.

Taken together, these findings suggest that the significance of artificial intelligence for law lies not in the replacement of human rationality with algorithmic processes, but in the transformation of the conditions under which legal reasoning operates. Legal concepts retain their formal structure, yet their practical meaning is increasingly shaped by the interaction between human judgment, institutional frameworks, and computational systems.

This reconfiguration has important implications for legal theory and practice. It calls for a more precise understanding of the relationship between technical systems and normative reasoning, as well as a critical engagement with the metrics and models that underpin contemporary decision-making. Rather than treating artificial intelligence as either a neutral tool or a disruptive force, legal analysis must account for its role in structuring the epistemic environment of law.

Future research may build on this analysis by further examining how hybrid human–machine systems can be designed to preserve core legal values, including accountability, transparency, and fairness. In doing so, the challenge is not to resist technological change, but to ensure that its integration into legal systems remains aligned with the normative commitments that define law as a distinct institutional practice.

## VIII Acknowledgements

An initial version of this article was prepared in Persian. For the purposes of international publication, the author employed artificial intelligence tools to assist in the translation process. The author accepts full responsibility for any remaining errors or omissions. Sincere gratitude is extended to all those who supported the development of this work, both during its initial drafting and subsequent translation.