ORIGINAL ARTICLE

INTERNATIONAL JOURNAL OF ARTIFICIAL INTELLIGENCE

ISSN: 2692-5206, Impact Factor: 12,23
American Academic publishers, volume 05, issue 10,2025

Journal: https://www.academicpublishers.org/journals/index.php/ijai

THE TRAJECTORY OF AI-DRIVEN CREDIT SCORING AND THE REFINEMENT OF LEGAL MECHANISMS FOR A DIGITAL FUTURE

Amirjon Mardonov,
Lecturer, Cyber Law Department, Tashkent State University of Law
Email: amirmardonov39@gmail.com
Orcid: 0009-0007-1997-6870

Abstract: This paper examines how modern machine-learning techniques are reshaping consumer credit scoring and what legal-regulatory guardrails are required to keep accuracy accountable. It argues that, in tabular lending contexts, ensemble models (e.g., gradient-boosted trees) deliver meaningful lifts over traditional logistics while remaining compatible with faithful, applicant-level explanations (e.g., SHAP). The core risks are opacity, drift, and disparate impacts—problems that cannot be solved by technical fixes alone but must be governed through clear model-risk controls, documented fairness choices, auditable data lineage, and reason-code specificity. Surveying leading regimes, the paper highlights the EU AI Act's "high-risk" obligations for creditworthiness assessment, U.S. expectations under ECOA/Reg B and SR 11-7/OCC model-risk guidance, MAS FEAT/Veritas operationalization, and BIS/BCBS prudential concerns about correlated model failures. Against this backdrop, Uzbekistan's existing statutes on personal data, consumer credit, and credit histories provide an enabling legal scaffold. The paper proposes targeted secondary rules—explainability standards tied to actual computations, ex-ante fairness metrics and monitoring, mandatory drift detection with challenger models, and full vendor oversight—to align innovation with due process. Done well, AI scoring can expand inclusion, strengthen dignity through actionable explanations, and enhance system resilience without trading off consumer protection.

Keywords: Adverse-action reasons; Basel/BCBS prudential risk; Consumer protection; Credit histories; Data governance; Drift monitoring; ECOA/Reg B; EU AI Act (high-risk); Explainability (SHAP); Fairness metrics (equal opportunity); Gradient-boosted trees; Model calibration; Model risk management (SR 11-7/OCC); Monotonic constraints; Personal data law (Uzbekistan); Tabular ML; Transparency & accountability; Vendor model oversight; Veritas/FEAT (MAS); Uzbekistan—consumer credit law.

Artificial intelligence in credit scoring feels like a story we've been trying to tell for years but only recently learned the language for. Banks once leaned on tidy, linear models and a handful of inputs to decide who deserved a loan; now they can draw on sprawling behavioral histories, cash-flow traces, alternative data, and non-linear interactions that classical regressions simply glossed over. Gradient-boosted trees, random forests, and occasionally neural networks find faint patterns in tabular data that humans would miss, pushing default prediction a few percentage points better—small on paper, but decisive at scale. Yet the better these models predict, the more they tempt us into opacity: variable interactions become too tangled to explain, drift creeps in as macro conditions shift, and historical inequities can resurface as "statistical facts." The regulatory response, worldwide, is converging on a common intuition: when an algorithm says yes or no to money, the answer must be auditable, explainable, fair, and safe—not just clever. In this article, I make the case for a practical, rights-respecting path that embraces the accuracy gains of modern ML while imposing governance, explainability, data

INTERNATIONAL JOURNAL OF ARTIFICIAL INTELLIGENCE

ISSN: 2692-5206, Impact Factor: 12,23
American Academic publishers, volume 05, issue 10,2025

Journal: https://www.academicpublishers.org/journals/index.php/ijai

discipline, and enforceable consumer protections. I anchor the analysis in widely cited supervisory texts (EU, U.S., Singapore, BCBS) and situate Uzbekistan's maturing legal base (personal data, consumer credit, credit histories) as fertile ground for a high-trust AI credit ecosystem.[1]

If we start with the technology rather than the law, we see that most winning credit-risk pipelines today are not deep end-to-end neural systems but ensembles tuned for tabular data, often with meticulous feature engineering and regularization. Tree-boosting methods (like XGBoost/LightGBM) typically outperform baseline logistics, especially when interactions and thresholds matter; they calibrate reasonably well with post-processing, and they play nicely with local explanation tools such as SHAP that translate complex splits into human-interpretable reason codes for individual applicants. This is the understated revolution: not only better AUCs, but the ability to look an applicant in the eye—figuratively—and say, "These were the most influential factors in your decision," in a way that isn't fictional. Modelers still wrestle with leakage, target drift, and stability under shifting applicant pools, but we've learned to monitor data drift, run challenger models, and set retrain triggers before performance or fairness decays into customer harm. The open debate is whether to prefer truly interpretable models (monotonic GBMs, scorecards, rule lists) over black boxes with post-hoc explanations; in many tabular lending settings, the accuracy gap is slimmer than people think, which strengthens the argument for simpler, audit-friendly approaches where stakes are high and reasons matter.[2]

Fairness is where technical neatness ends and policy choices begin. Because base rates differ across groups, some fairness definitions simply conflict: equalizing false-negative rates can worsen positive predictive value gaps, and vice versa. The research community has proven "impossibility" trilemmas that mean lenders must choose and justify which fairness notion they will target, and regulators must be explicit about the normative goal to prevent well-meaning teams from optimizing into legal ambiguity. A responsible workflow treats fairness like any other measurable requirement: define metrics (e.g., equal opportunity on approvals above a given threshold), test on training and out-of-time windows, document trade-offs, and monitor in production alongside calibration. The point is not to pretend bias can be engineered away instantly, but to make every step visible and contestable, so that policy, compliance, and consumer protection can see the same dashboards product sees. When these choices stay tacit, "accuracy" becomes the alibi for historical patterns we'd never defend in plain words.[3]

The European Union took a decisive step by classifying AI systems used to evaluate creditworthiness as "high-risk," which pulls them into a robust set of obligations: risk management, data quality and governance, technical documentation, transparency, human oversight, robustness testing, and post-market monitoring. Even outside the EU, that template now shapes boardroom conversations and vendor contracts because global banks and fintechs

[1] EU AI Act—high-risk creditworthiness; EBA Big Data & Advanced Analytics; SR 11-7 & OCC Model Risk; CFPB adverse-action specificity; MAS FEAT; BCBS on prudential AI risks; Uzbekistan: Law on Personal Data (Lex.uz), Law on Consumer Credit (Lex.uz), Law on Exchange of Credit Information (Lex.uz).
[2] Applied reviews showing ML gains in credit risk; Dastile et al., 2020; survey literature; Rudin, 2019 on interpretable models; SHAP methodology for local explanations.
[3] Foundational fairness work: Hardt et al., 2016; Kleinberg et al., 2016; industry fairness testing guidance.

ISSN: 2692-5206, Impact Factor: 12,23
American Academic publishers, volume 05, issue 10,2025
Journal: https://www.academicpublishers.org/journals/index.php/ijai

can't afford divergent ethics stories across jurisdictions. The EBA's loan origination guidance and analytics reports complement the AI Act by describing what sound governance looks like in day-to-day modeling—clear model inventories, validation independence, and data lineage discipline. If you imagine a world where a regulator or a court asks, "What exactly did you train on, when, with what checks, and how do you know it didn't drift into discriminatory behavior?" these European texts provide the binder dividers and the checklists that make a coherent answer possible.[4]

The United States, without a single omnibus AI law, arrives at similar expectations through the Equal Credit Opportunity Act and Regulation B's requirement to give specific, accurate adverse-action reasons—even when the decision came from a complex model. The Consumer Financial Protection Bureau has been unusually direct: generic boilerplate won't cut it if an ML system's actual pivots were, say, high utilization on recent revolving accounts and unusual cash-flow volatility. At the prudential end, SR 11-7 and OCC guidance on model risk management demand institution-wide frameworks: inventories, development standards, independent validation, effective challenge, change control, and ongoing monitoring. What's crucial is that vendor models and "AI-as-a-service" are not loopholes; they're squarely in scope, with the supervised institution owning the risk. Culturally, this nudges banks away from "mystique-as-strategy" and toward "documentation-as-asset," which is exactly where consumer trust is built.[5]

Singapore's MAS framed the FEAT principles—Fairness, Ethics, Accountability, Transparency—and then did something pragmatic: it released Veritas tooling so institutions could operationalize those values rather than just nod to them. This combination of principles and implementation is worth emulating because it bridges the common gap between board-level slogans and engineer-level tests. Meanwhile, global prudential voices (BIS/BCBS) are warning about correlated model risks and macroprudential opacity—how a thousand black boxes, each locally "good," can still synchronize badly under stress. The message here is sober: governance is not just about individual rights but also about system stability. If many lenders learn the same patterns from the same data, then drift, regime changes, or data quality shocks could propagate in ways we don't intend.[6]

Turning to Uzbekistan, the legal scaffolding already in place is more enabling than many realize. The Law "On Personal Data" establishes lawful grounds for processing, security safeguards, and important subject rights—cornerstones for model training and automated decision-making conducted with dignity and legality. The Law "On Exchange of Credit Information (Credit Histories)" defines how credit data can move and be used—a vital substrate for any scoring approach that aspires to be both accurate and fair. The Law "On Consumer Credit," updated in recent years, frames the obligations around lending to individuals and aligns with a broader consumer-protection posture that automated decision-making must respect. With these anchors, the next step is not to start from scratch, but to issue secondary regulations and supervisory guidance tailored to AI: what counts as adequate explainability, how fairness

[4] EU AI Act (Annex/Articles on high-risk), EBA Guidelines on Loan Origination & Monitoring, EBA Big Data/Advanced Analytics report.
[5] CFPB circulars/press guidance on adverse action and specificity; FRB SR 11-7; OCC Model Risk Handbook.
[6] MAS FEAT; MAS Veritas; BIS/BCBS speeches and notes on AI and systemic risk.

ISSN: 2692-5206, Impact Factor: 12,23
American Academic publishers, volume 05, issue 10,2025
Journal: https://www.academicpublishers.org/journals/index.php/ijai

should be measured and reported, what vendor access is required for audits, how frequently drift must be checked, and how consumers can meaningfully contest a decision.[7]

What, then, should a practical, risk-based framework look like in day-to-day banking operations? First, scope the problem honestly: credit scoring and pricing that materially affect access to finance should be deemed "high-impact" automated decision-making, which triggers stronger duties. That designation should not be punitive; it is a permission structure—"Yes, you can use modern ML, but here are the rails." Second, require institution-wide model-risk management akin to SR 11-7: a current model inventory; clear ownership; standardized development documentation (data sources, sampling, feature engineering, hyperparameters, training windows); truly independent validation with authority to block deployment; and change control that records why a model was updated and what testing preceded the change. Third, insist on explanation that is faithful to the actual computation, not post-hoc storytelling that feels plausible but isn't. If lenders use SHAP or monotonic constraints to make reason codes stable and plain-language, they should back-test the communication itself with user research: did consumers understand the main drivers, and could they act on them? Fourth, define fairness metrics ex ante—e.g., equal opportunity at a decision threshold—and publish how they're monitored and what remediation steps follow when gaps exceed tolerances. Finally, treat vendor models as first-class citizens in oversight: contractual rights to documentation, to sandbox replicas for validation, to fairness and drift metrics, and to termination terms if the model cannot meet supervisory expectations.[8]

On the technical side, institutions can keep their house in order with a predictable blueprint. Begin with a living model inventory that maps every scoring and pricing model to a clear business purpose, input data sets, and a retraining cadence. Build lineage tracking so that every feature has a provenance chain back to raw data, with data quality checks that fail loudly rather than degrade silently. Establish a "challenger" lane where a second model (simpler or differently regularized) runs in shadow to spot drift early. Implement calibration and stability checks in monthly or quarterly cycles, and couple them with fairness dashboards that track disparities in approval, pricing, and error rates across legally protected and contextually sensitive groups. Most importantly, make the dashboards intelligible across roles: compliance, model risk, business owners, and—through carefully designed communications—consumers who deserve to know why the model reacted the way it did. When teams can sit around a single page of metrics and agree on what "good" looks like, governance stops being a cost center and becomes part of the product.[9]

There are predictable failure modes to guard against. One is the "template reason" fallacy: the bank uses standardized adverse-action codes that don't match the model's true drivers, which misleads consumers and violates legal specificity requirements. Another is "silent drift," where macroeconomic shifts change applicant behavior and a model trained on a sunny regime becomes brittle; you see it in calibration plots first, then in complaint rates,

[7] Uzbekistan: Personal Data Law (Lex.uz, ZRU-547); Consumer Credit Law (Lex.uz, as amended); Law on Exchange of Credit Information (Lex.uz, ZRU-301).
[8] EU AI Act—high-risk logic as a template; SR 11-7 playbook; CFPB specific-reason expectations; EBA governance; MAS FEAT operationalization.
[9] Industry model-risk and MLOps patterns; EBA/EU documentation expectations; OCC handbook practices.

ISSN: 2692-5206, Impact Factor: 12,23
American Academic publishers, volume 05, issue 10,2025

Journal: https://www.academicpublishers.org/journals/index.php/ijai

then—too late—in higher defaults. A third is the "proxy creep": variables that look harmless in isolation (geography, device type, browsing behavior) combine to reconstruct sensitive attributes in practice. Avoiding these failure modes is not about forbidding variables wholesale but about making proxy detection and sensitivity analysis a routine part of validation, documenting what was tested and why a variable remains justified. Each of these guardrails is worth articulating in secondary regulation so that the industry hears one voice: better models are welcome, but ghost stories are not.[10]

What does all of this buy the public? First, inclusion: by reading behavior more finely and treating "thin-file" applicants as knowable rather than unknowable, AI can expand access responsibly. Second, dignity: faithful reason codes let declined applicants take corrective action rather than feel shut out by a machine. Third, resilience: governance that anticipates drift and demands independent challenge reduces the chance that many lenders make the same mistake at once. And fourth, legitimacy: when regulators can see into models through documentation and testing rather than headlines, the sector earns room to experiment without lurching from hype to backlash. The future of credit scoring is not a choice between innovation and rights; it is a discipline that merges them so that accuracy is accountable, and accountability is informed by evidence rather than slogans.[11]

Finally, a word about voice and responsibility. AI in finance attracts grand claims and equally grand fears. The sober path is to treat the model like any other critical system: we specify what it should do (predict risk), what it must not do (produce unjustified disparities or inscrutable decisions), how we will test that over time (validation, fairness, calibration, stress), and what rights the affected person has (understand, correct, contest). Uzbekistan's existing laws already point in this direction; by adding targeted AI-specific guidance—definitions, explainability standards, fairness reporting, vendor duties—the country can both accelerate digital lending and harden its commitment to the rule of law. If we do this well, the next time a loan decision arrives by algorithm, it will not feel like the verdict of a black box but like the output of a system that took the applicant seriously. That is what trustworthy AI looks like in real life: not magic, just responsibility at scale.[12]

REFERENCES:

1. EU AI Act—creditworthiness as "high-risk" (European Commission/Parliament texts); EBA Big Data & Advanced Analytics; U.S. SR 11-7 (Federal Reserve) and OCC Model Risk Management; CFPB guidance on adverse-action specificity; MAS FEAT principles; BCBS/BIS notes on AI prudential risk; Uzbekistan: Law on Personal Data; Law on Consumer Credit; Law on Exchange of Credit Information (all via Lex.uz). ↩
2. Systematic reviews of ML in credit risk (e.g., Dastile et al., Applied Soft Computing, 2020); survey papers on ML for default prediction; Rudin, "Stop explaining black box models…" (Nature Machine Intelligence, 2019); SHAP explainability for tabular credit models. ↩

[10] CFPB on specificity; SR 11-7 validation standards; research on proxy variables and fairness auditing.
[11] BCBS prudential concerns; MAS/EBA transparency aims; CFPB consumer-rights lens; EU AI Act accountability.
[12] Uzbek statutes via Lex.uz; international supervisory texts cited above; academic evidence on explainability and fairness.

INTERNATIONAL JOURNAL OF ARTIFICIAL INTELLIGENCE
ISSN: 2692-5206, Impact Factor: 12,23
American Academic publishers, volume 05, issue 10,2025
Journal: https://www.academicpublishers.org/journals/index.php/ijai

3. Hardt, Price, & Srebro (2016) on equal opportunity; Kleinberg, Mullainathan, & Raghavan (2016) on inherent trade-offs; additional industry guidance on fairness testing in lending. ↩
4. EU AI Act (obligations for high-risk systems); EBA Guidelines on Loan Origination & Monitoring; EBA report on Big Data/Advanced Analytics—governance, data quality, validation expectations. ↩
5. CFPB circulars/press materials clarifying adverse-action specificity for AI-driven denials under ECOA/Reg B; FRB SR 11-7 on model risk; OCC Comptroller's Handbook sections on validation and independent challenge. ↩
6. MAS FEAT principles and Veritas initiative (tooling for responsible AI in finance); BIS/BCBS speeches and notes warning about systemic risks from opaque, correlated ML models. ↩
7. Republic of Uzbekistan: Law "On Personal Data" (ZRU-547); Law "On Consumer Credit" (as amended); Law "On Exchange of Credit Information (Credit Histories)" (ZRU-301)—all published on Lex.uz. ↩
8. High-impact classification logic inspired by EU AI Act; institution-wide controls per SR 11-7/OCC; CFPB's expectation of faithful, specific adverse-action reasons; EBA/MAS guidance on documentation and transparency. ↩
9. Practical MLOps/validation practices aligned with EBA documentation norms and OCC/FRB expectations: model inventories, lineage, challenger frameworks, calibration/fairness dashboards. ↩
10. CFPB specificity as a guardrail against "template reasons"; SR 11-7 on independent validation; literature on proxy variables and sensitivity testing for fairness in credit. ↩
11. BCBS prudential framing (AI and systemic resilience); MAS/EBA transparency for consumer trust; CFPB consumer-rights benefits; accountability regime in EU AI Act. ↩
12. Synthesis of Uzbek legal anchors (Lex.uz) and international supervisory materials; academic sources on explainability and fairness cited above. ↩