Fintech 2040

AI Ethics and Accountability in Hyperindividualized Lending: Pathways to Optimized Consumer Credits

A Position Paper by Anita Klingel, Jan Richard & Florian Haubold

FIVERTY

Fintech 2040
Imprint

Publisher:
Riverty Group GmbH
Rheinstraße 99
76532 Baden-Baden
Germany

www.riverty.com

Registered Office: Baden-Baden
Commercial Register: Amtsgericht
Mannheim, HRB 202547
VAT ID No.: DE 814491108

Management: Andreas Barth,
Volker Bornhöft, Oliver Kuhaupt,
Dr. Teresa Schlichting

Authors:
Anita Klingel,
Jan Richard,
Florian Haubold

Media and Press Contact:
Kristina Hunter Nilsson
Corporate Affairs Lead
E-mail: press@riverty.com

Editorial Design & Layout:
Dominik Ziller

Project Website:
www.riverty.com/Fintech-2040

Responsible according
to § 55 (2) RStV:
Riverty Group GmbH
Rheinstraße 99
76532 Baden-Baden
Germany

Design and Creation Notes:
AI was used to lecture and proofread
the final text version (GPT 5.2).
Images in this publication were created
by humans using artificial intelligence.
Typeface used: GT Ultra (Grilli Type)

Regulatory Notice
This publication presents hypothetical and analytical future scenarios relating to AI-based credit systems. It does not describe existing, deployed, or planned operational systems of Riverty Group GmbH or its affiliated entities. References to regulatory classifications or technological architectures are made solely for discussion purposes and do not imply current implementation.

RIVERTY

With Fintech 2040, we aim to contribute to a structured and fact-based discussion about the future of financial services. Following our first paper, which examined the broader evolution of the Fintech ecosystem, this publication explores a specific and highly relevant topic: AI-driven, hyperindividualized lending.

Artificial intelligence is increasingly shaping discussions around credit processes across the financial sector. It offers the potential for more granular risk assessment, faster decision-making, and the consideration of broader categories of data. These developments may improve efficiency and could expand access to credit, particularly for consumers with limited traditional credit histories.

At the same time, they raise important questions regarding transparency, accountability, and fairness. As credit assessment frameworks become more data-intensive and automated, governance structures, process design, and regulatory standards must continue to evolve.

This paper examines the perceived trade-off between profitability and ethical responsibility in AI-based credit markets. It outlines conceptual approaches to addressing this tension, including data sovereignty, interoperability, and structured process governance, and highlights the importance of a consistent regulatory framework in supporting responsible innovation.

Our objective with Fintech 2040 is not to promote a specific technological model or operational approach, but to encourage informed dialogue across industry, academia, and policymaking. Sustainable financial innovation depends on both technological progress and credible safeguards.

We hope this contribution supports a constructive debate on how AI-related technologies can be developed and governed in credit markets in ways that support market stability and consumer confidence.

Andreas Barth
CEO, Riverty

Management Summary

This position paper explores how new technologies and processes can reduce risks for consumers and businesses in hyperindividualized, AI-based credit markets. AI is reshaping the credit value chain by accelerating decision-making, expanding the data used in underwriting, and enabling increasingly personalized pricing and product design. These developments can theoretically improve efficiency and broaden access to credit, including thin-file consumers. But they also heighten risks around discriminatory outcomes through proxies, error propagation and the concentration of informational power. Borrowers often experience AI-driven credit as difficult to understand and hard to contest, especially when decisions depend on data that does not feel legitimate or transparent. Meanwhile, lenders may view ethical and legal safeguards as costly and as sources of operational and liability risk, particularly when models and data pipelines rely on third parties.

As part of Riverty's Fintech 2040 thought-leadership initiative, the position paper serves as a guideline for assessing how AI-powered financial innovation could benefit borrowers and lenders alike. The project Fintech 2040 is designed to look beyond the near technological borders in credit. It explores a future scenario of finance shaped not only by new capabilities, but by deliberate choices around fairness, inclusion, and human-centered design (Riverty, 2025). This paper on hyperindividualized lending captures that tension especially well: it can help enable more accurate and potentially fairer credit access, while simultaneously lowering the danger of hidden bias, hard-to-explain decisions, and widening asymmetries between institutions and individuals.

All scenarios described in this paper are hypothetical and intended to contribute to academic and regulatory dialogue. They do not represent current governance frameworks, deployed systems, or strategic roadmaps of Riverty Group GmbH.

Keywords: Hyperindividualized lending, AI Ethics, Credit scoring, Explainable AI, XAI, AI Bots, Accountability, Future of Credit, Data Sovereignty, Open Banking, Open Finance, Interoperability

Content

Introduction

AI is personalizing credit, but responsibility remains opaque.

Why hyperindividualized lending demands new answers on fairness, transparency, and consumer protection.

A migrant taxi driver in her mid-30s – Josephine – has been living in Germany for two years. She works long night shifts, sends money home regularly, and is proud that she has never missed a rent payment. What she does not have, however, is a long local credit history: no old German credit cards, no long-standing bank relationship, no decades of trackable borrowing.

One morning, after a particularly busy week, Josephine buys a pair of Converse sneakers[2] online with her credit card. It is a small purchase – practical shoes for a job that keeps her on her feet. A few days later, she applies for a modest consumer loan to repair her car, which she needs to keep working. But the interest rate offered is unexpectedly high. When she asks why, she receives a generic answer: "Our assessment of your credit score indicates elevated risk."

Only later – through a friend who works in Fintech – Josephine hears a worrying rumor: in some modern credit assessment systems that rely on alternative data, certain purchasing patterns can influence the individual credit score. In Josephine's case the Converse Sneakers triggered her worse credit interest rates. It all happened inside the algorithms. The brand, the type of shoe, the timing of the purchase – some combination of features in a model – have contributed to the outcome. The decision feels arbitrary, almost surreal. Yet the core dynamic is real: lenders and scoring providers increasingly explore data beyond classic bureau records and income statements, precisely to assess people who look thin-file in traditional systems (Federal Reserve, 2025).

The fictive example of Josephine demonstrates the problem of accountability in hyperindividualized lending. The usage of large datasets and AI have changed the credit process between borrowers and lenders substantially. What used to be a comparatively slow, document-heavy process that relied on a limited set of standardized indicators is increasingly becoming a real-time, data-rich decision pipeline. Artificial intelligence and machine learning enable lenders to process large volumes of heterogeneous information, identify complex patterns in repayment behavior, and automate decisions at a scale and speed that would be infeasible with traditional methods.

The underlying logic from the perspective of the lenders is economically intuitive: Better prediction of default risk and more granular segmentation can reduce losses, improve capital efficiency, and enable products that match borrowers' needs more closely. At the same time, hyper-individualization can reduce traditional forms of mutualization and cross-subsidization, amplifying distributional effects: some consumers may gain access or receive better terms, while others – like Josephine – may face higher prices, exclusion, or volatile credit conditions.

The societal stakes of this debate are high because credit is not merely another consumer product. It is a gateway to economic participation: housing, education, mobility, entrepreneurship, and resilience against shocks. AI therefore creates a double-edged dynamic. On one edge, it can widen access and improve efficiency; on the other, it can harden existing inequalities through biased data, opaque models, and feedback loops that make disadvantage self-reinforcing.

2 Illustrative example only. The described purchasing pattern does not reflect actual Riverty data practices or model inputs.

Within Riverty's Fintech 2040 thought-leadership platform, the position paper on hyperindividualized lending serves as a critical lens through which to understand the long-term societal, economic, and regulatory implications of AI-driven financial innovation (Riverty, 2025). Fintech 2040 aims to anticipate the next era of finance – one defined not only by technological progress but by the ethical, inclusive, and human-centric design choices that shape it. Hyperindividualized lending exemplifies this future: a powerful development with the potential to broaden access to fair credit, but also one that carries inherent risks of algorithmic discrimination, opacity, and power imbalances, developed in a theoretical scenario analysis.

This paper examines these opportunities and risks. It outlines how AI-based credit assessment models function, how bias can emerge within them, and which ethical dilemmas arise for lenders, regulators, and society. Anchored in the Fintech 2040 perspective, the paper ultimately looks into the current market conditions to illustrate the best – and worst – possible outcomes of the choices being made today.

The paper is written for a Fintech and technology-oriented audience and provides an argumentative guideline in the ethical debate around AI-based lending. The central question of the paper is:

How can new technologies and processes help mitigate risks for consumers and businesses in a hyper-individualized, AI-based credit market?

Rather than treating ethics as an external constraint that merely limits innovation, the paper explores how ethical and economic objectives can be aligned through concrete technical and organizational measures, and how regulation can create the conditions in which this alignment is feasible and competitively neutral.

Systems & Risks

From single scores to interacting decision systems.

How hyperindividualized credit systems work and where risk and bias emerge.

Highlighting the role of AI in the hyperindividualized credit business

In practice, AI in credit rarely means a single model replacing a single score. Hyper-individualization emerges from the interaction of multiple models and decision rules across the lifecycle: from marketing and onboarding, to creditworthiness assessment, pricing and monitoring. Many EU banks report using a mix of methods ranging from more established statistical approaches to decision trees, natural language processing and neural networks. They also commonly combine in-house development with third-party tools, including open-source components and cloud-based APIs (European Banking Authority, 2024). This modular reality matters, because the ethical and economic properties of hyper-individualized credit are often not in any model but in the system-level design: which data sources are permitted, which objectives are optimized, where in the process is human discretion allowed, and what recourse exists when automated decisions harm consumers?

The emerging frontier: future scenarios for hyper-individualized credit

The near-term trajectory of AI in credit is about re-architecting the credit process into a real-time, data-interoperable pipeline. One prominent scenario is cash-flow and transaction-data underwriting, enabled by open banking access to bank transaction data and other verified streams. The World Bank's landscape study on alternative data emphasizes that integrating structured and unstructured alternative data can expand access for underserved populations, particularly thin- or no-file consumers, while also improving approval rates and risk outcomes. (World Bank, 2024).

A second scenario is continuous underwriting, where eligibility, pricing, and credit limits are periodically updated as new information arrives. This can support earlier detection of financial stress and more timely forbearance offers, but it can also create borrower uncertainty if product terms feel unstable or if consumers cannot anticipate how everyday behavior affects future access. System design choices – such as considered variables, frequency of updates – become core determinants of market trust in such a development (Financial Stability Board, 2024).

A third scenario is the deeper integration of unstructured information. Generative AI and advanced NLP can extract signals from documents and communications, summarize borrower-provided materials, and support credit analysts with faster review – particularly relevant in credit card lending and complex financing where documentation is extensive (Krivorotov, 2023). The operational upside is real: faster cycle times and lower per-loan processing costs can make smaller-ticket credit more viable, potentially supporting inclusion.

Taken together, the central question is not whether AI will be used in credit markets in the future, but how the data architecture and process design surrounding AI can make outcomes auditable, contestable, and aligned with both welfare and profitability.

The role of data and possible bias

The fictional example of Josephine in the first chapter demonstrates a central problem of the ethical discussion: a model can be statistically right and still be conceptually wrong. The sneakers might correlate with higher default rates in the historical training data – perhaps because that purchase pattern is more common in certain age groups, neighborhoods, occupations, or income situations. But correlation is not causality. The shoes do not cause anyone to miss repayments. At best, they are a proxy – a shortcut that stands in for other factors the model cannot directly observe. In the worst case, they become an opaque indicator of sensitive characteristics, producing discriminating outcomes while remaining difficult to detect and contest (Barocas and Selbst, 2016).

Discrimination in credit refers to unequal treatment of individuals or groups based on characteristics such as race, gender, age, or sexual orientation. In credit assessment, discrimination can arise intentionally, but more often results from biases – systematic distortions in data, feature design, model training, or human oversight.

These biases may be technical, such as under-representation of certain groups in training data or reliance on flawed proxy variables. But also psychological or organizational, such as insufficient human oversight or unchallenged assumptions embedded in product design.

Where bias enters: the main mechanisms along the credit lifecycle

One major channel of discrimination is historical or label bias. Credit datasets typically learn from past decisions and outcomes. But past decisions may already reflect unequal access, different terms, or differential treatment. If the target variable is defined as repayment or default, it is influenced not only by borrower ability and willingness to pay, but also by product terms, servicing practices, and macroeconomic conditions. This creates the risk that models encode patterns of disadvantage as if they were natural risk signals, especially when credit histories are missing precisely because groups were historically excluded (World Bank, 2025).

A second channel for discrimination is sampling and coverage bias. Hyper-individualized lending often relies on digital channels, which can under-represent consumers with limited connectivity, unstable housing, or low digital literacy. When the training population differs from the target population, models can systematically underperform on groups that were less present in the data. This risk is explicitly recognized in governance frameworks that demand attention to dataset representativeness and statistical properties relative to the groups a system is intended to affect (European Commission, 2025).

A third channel of discrimination is measurement and data quality bias. Even when the same variable exists across groups – income, employment stability, expenses – its measurement can differ due to informal work, cash-based livelihoods, or inconsistent documentation. Alternative data can sometimes reduce these gaps by observing cash-flow directly, but it can also introduce new measurement artifacts that reflect platform access, device type, or local infrastructure rather than creditworthiness.

A fourth channel is feedback loops. When models influence who receives credit and on what terms, they also influence the future data that models learn from. If a system systematically offers worse terms or tighter limits to a group, it may increase delinquency rates in that group, which then confirms the model's original pattern and entrenches it. (FinRegLab, 2023).

A fifth channel is objective misalignment. Hyperindividualization is often economically motivated by profitability: risk-based pricing, limit optimization, and targeted marketing can increase returns. But optimizing for profit can select strategies that are privately rational while socially harmful, for example by exploiting behavioral vulnerabilities, steering consumers toward expensive products, or amplifying exclusion. The fact that AI can increase predictive accuracy does not guarantee welfare-improving outcomes (Gambacorta et al., 2019).

Josephine's fictional case illustrates the practical dilemma of hyperindividualized lending: the promise is more precision and broader inclusion, but the danger is that individuals become legible to models in ways they cannot understand, challenge, or meaningfully correct. When the consumer's credit score decreases because of an everyday behavior – like buying shoes – without a clear, human-interpretable justification, trust may erode quickly. This paper tackles exactly that tension: how can consumer credit be optimized with advanced analytics while ensuring ethical use of data, robust accountability, and explanations that make sense not only to data scientists, but also to customers, regulators, and credit professionals?

The Dilemma

From efficiency gains to legitimacy concerns.

Why hyperindividualized lending generates tension between profit and trust.

Identification of the perceived trade-off between lenders and borrowers

Public debate about AI-driven credit mainly revolves around a single intuitive suspicion: that tools like AI and ML that improve profitability and efficiency for lenders can simultaneously erode fairness, autonomy, and trust for borrowers.

From the borrower side, the perceived situation is sharpened by the black-box feeling of AI decisioning. People tend to be more skeptical of algorithmic tools in high-stakes contexts when they anticipate privacy invasion, unfair treatment, or an inability to challenge the outcome.

Experimental research that includes bank loan approvals similarly examines public preferences between algorithmic and human decision-makers in high-stakes settings. The results of these studies reflect the concern that better performance alone does not resolve legitimacy questions when decisions deeply affect individuals' lives (Bansak and Paulson, 2024). This correlation may even lead to a harmful perception: It almost seems that humans prefer worse credit decisions from other humans than better decisions from AI-based credit-scoring algorithms.

In societal perception, therefore, the AI-based decision making in the credit value chain becomes a moralized topic. Profit appears to demand maximum data extraction, optimization, and personalization, while ethical behavior appears to demand restraint, transparency, and due process. The perceived trade-off is therefore not simply banks versus customers, it is a clash between two plausible narratives of what AI is doing to the credit relationship: one narrative of efficiency and inclusion, and another of surveillance, manipulation, and automated discrimination.

The economic roots of the societal perception: information asymmetry and the logic of segmentation

The shift toward hyperindividualization introduces a series of dilemmas for financial services, many of which require balancing commercial logic with broader societal values.

One dilemma is inclusion versus risk minimisation. Consider two applicants with similar incomes, but different financial footprints. A long-established customer has years of bureau history, a stable bank relationship, and predictable repayment patterns. A newcomer – like Josephine – may be economically reliable but thin-file. Hyperindividualization promises to include the thin-file borrower by using cash-flow patterns or other non-traditional signals. BIS research using Fintech data illustrates how non-traditional data and machine learning can improve predictive performance relative to traditional approaches, especially when richer behavioural data is available (Gambacorta et al., 2019).

A secon dilemma is efficiency versus concentration. Hyperindividualization rewards scale: the more customers you have, the more data you collect. The more data you collect, the better your models can become. The better your models, the cheaper and faster you can underwrite. Over time, this can push markets toward a "winner-takes-most" dynamics, where a few institutions or platforms dominate because they can out-learn and out-price smaller competitors. OECD analysis on data-driven markets discusses how data use shapes market structure and can concentrate advantages in ways that affect competition and consumer outcomes (OECD, 2022). For bankers, the risk is not only competitive pressure but dependency pressure: if critical underwriting signals flow through a few large data intermediaries, credit markets become more

vulnerable to vendor lock-in, outages, and correlated model behaviour.

These banking-specific dilemmas are embedded in broader cross-sector tensions that shape how society interprets AI-based credit. Data privacy versus predictive accuracy is the most visible. Many signals that improve prediction – location patterns, device metadata, granular transaction details –also feel like surveillance when used to decide who gets affordable credit. The more the model knows, the more it can predict. But the more it knows, the less autonomy and dignity borrowers may feel they have in the relationship (NIST, 2023).

However, the deepest cross-sector tension is efficiency incentives versus fairness requirements. Markets reward what can be measured and monetized quickly: lower cost per loan, higher conversion, and tighter risk selection. Explainability, and contestability are harder to price into competition because their benefits are diffuse and long-term: fewer disputes, fewer scandals, more durable trust, and better legitimacy. In practice, this can create a structural race-to-the-bottom risk. A lender that invests heavily in fairness testing, conservative feature policies, and robust recourse may carry higher short-term costs than a competitor that pushes automation and expansive data use more aggressively.

When these dilemmas accumulate, society often compresses them into one simple storyline: a social dilemma between ethical behaviour and profit (Figure 1). Borrowers see hyperindividualization as a system that can optimize against them – using data they did not understand, making decisions they cannot contest, and charging prices that feel like personalized penalties. Lenders, in turn, may see ethical safeguards as costs, liability risks, and sources of operational friction – especially when models and data pipelines involve third parties and complex tooling. The result is a coordination problem: even well-intentioned actors fear being competitively punished for restraint (Pies, 2009).

Possible worst-case scenarios for both sides

From the perspective of the societal debate, lenders and borrowers are seen as counterparts in a zero sum game: what one side wins is the direct loss for the other side – and vice versa. In this chapter we look into the worst possible outcome for both market sides in order to make the perceived dilemmas even more tangible.

For borrowers, the worst case is exclusion by design. Hyperindividualization promises to bank the unbanked, but without guardrails it can flip into a requirement to be legible to machines. People who cannot or do not want to generate the right data – newcomers without local history, workers paid in cash, older consumers with low digital usage, or anyone intentionally minimizing their digital footprint - become structurally disadvantaged. They may be treated as unknown risk, which in automated underwriting often means rejection, extremely conservative limits, or punitive pricing. This is not a hypothetical edge case: even in mature credit markets, a meaningful share of adults have limited or no credit file (CFPB, 2015).

FIGURE 1
Perceived Trade-off between Ethical Behaviour and Profit in Credit

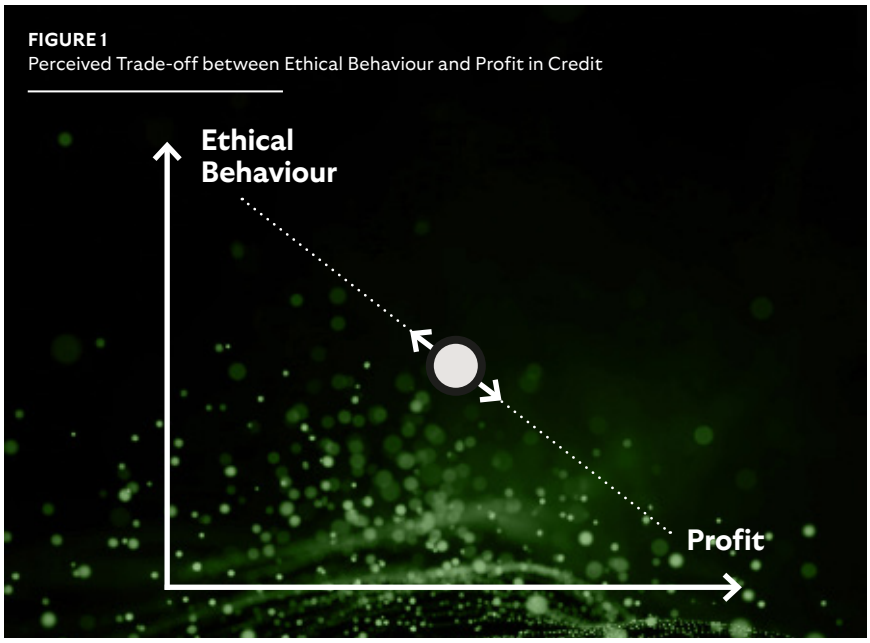

A simple example makes the mechanism tangible. Imagine Elena, who recently moved for work and shares a flat with roommates. She has stable income, but she avoids credit cards, pays many expenses in cash, and uses prepaid mobile plans. When she applies for a small loan to furnish her apartment, the automated system cannot confidently place her into any profitable segment. In a market optimized for margin, Elena gets filtered out. The exclusion is not framed as discrimination, but as insufficient signals. The result feels similar: the people who most need mainstream credit terms are pushed away from them.

Added to this worst case for the borrower is the loss of agency when something goes wrong. In a high-speed ecosystem, borrowers may not know which data source harmed them, which model used it, or how to correct it before real-life opportunities expire. Credit reporting already shows how damaging errors can be and how difficult disputes can be to navigate. Hyperindividualization multiplies the points of failure: more data furnishers, more derived attributes, more model updates, more automated limit changes. The borrower-side worst case is not only a wrong decision, but a practical impossibility of diagnosis and timely recourse. In such a scenario data richness that fuels prediction expands the blast radius of breaches and misuse. Because compromised identities and polluted files can haunt borrowers for years.

Put together, the borrower-side worst case is a market that feels exclusionary and coercive in daily life. Credit becomes hyperpersonalized in the same way advertising is personalized: optimized for conversion and profitability, not for mutual benefit (see Figure 2). People who are outside the system do not merely have less access, they become structurally less financeable, because the market architecture treats lack of data as lack of worthiness. The social consequence is predictable: inequality is reinforced, because the groups already disadvantaged in labor, housing, and digital participation are also the groups most likely to be data-poor, proxy-penalized, or vulnerability-targeted. In this scenario, the trade-off between profit and ethical behavior is not an abstract debate, it becomes a distributional outcome baked into the credit market's operating system.

Worst-case for lenders: paralysis through overregulation, or being punished for acting alone

The lender-side worst case is less often described in public debates, but it matters for the credibility of any governance proposal: if rules make underwriting impossible or commercially irrational, credit supply will contract, risk will concentrate, and the market can become unstable.

The first lender worst case arises in a world where data usage is so tightly restricted that lenders cannot reliably distinguish between low-risk and high-risk applicants. In credit markets, this is not a minor inconvenience. It strikes at the basic function of lending, which is to price and allocate capital under uncertainty. The classic market for lemons logic explains why: when one side of the market has better information and the other side cannot verify quality, trade can unravel because prices converge toward an average that drives the good participants away. Akerlof originally illustrated this with used cars, but he explicitly connects the principle to markets such as credit, where quality uncertainty and asymmetric information can prevent mutually beneficial transactions (Akerlof, 1970).

FIGURE 2
Worst Case Scenario from the perspective of the Borrowers

A simple example makes the collapse mechanism tangible. Suppose two applicants look identical under the limited dataset: both report similar monthly income and have short credit histories. One is a careful saver who would almost certainly repay, the other is already overextended but hides it in ways the restricted dataset cannot detect. If the bank must price both applicants similarly, it either sets a conservative price to cover the hidden risk – making the good borrower walk away – or it prices aggressively to stay competitive – absorbing losses from the bad borrower. Over time, the good borrowers disproportionately exit (they self-finance, borrow from employers, use secured products, or go to lenders with better information), leaving a pool that is increasingly dominated by the harder cases. That is the lemons spiral in a credit setting: the average quality of the applicant pool deteriorates because the market cannot credibly differentiate quality.

In this lender worst case, the market does not necessarily collapse overnight; it can shrink and harden. Credit becomes scarcer, more expensive, and more collateralized. Innovation slows because experimentation requires data feedback loops. Risk migrates to less regulated corners of the market that can still price with richer information. For policymakers, this is the warning: regulation that focuses only on restricting data, without creating workable alternatives for risk assessment and accountability, can unintentionally increase exclusion and concentrate systemic risk rather than reduce it (Stiglitz and Weiss, 1981).

The second lender worst case is competitive, not legal: what happens if one bank behaves ethically but others do not? Here, ethical behaviour means more than marketing claims. It means the bank voluntarily limits invasive data sources, rejects opaque proxy features, invests in explanations and recourse, avoids borderline fee strategies, and accepts tighter internal constraints on profit optimization. The problem is that these choices can create a cost disadvantage and, even more importantly, an information disadvantage (Song and Thakor, 2022).

A concrete illustration is the instant credit race. Bank A builds a slower but safer credit process: it adds friction where uncertainty is high, requires manual review for edge cases, and invests in consumer-friendly explanations and dispute pathways. Bank B advertises instant approvals with minimal documentation, using richer behavioral data to automate risk selection and pricing. In the short run, Bank B wins on conversion and customer acquisition, especially among low-risk customers who enjoy favorable terms. Bank A looks uncompetitive, loses volume, and sees its unit economics deteriorate because fixed compliance and governance costs are spread over fewer loans. In extreme form, Bank A either exits the segment or is forced to adopt Bank B's practices to survive (See Figure 3).

These worst case scenarios explain why the ethical behaviour versus profit dilemma feels so plausible to society. Borrowers imagine a market that is optimized against them, and lenders imagine a market in which innovation creates new liabilities and coordination failures. Therefore, the design challenge is not to deny the risks but to re-engineer the data and process environment in credit so that the probability and severity of these failures drop sharply.

FIGURE 3
Trade-off between Ethical Behaviour and Profit from the lender's perspective

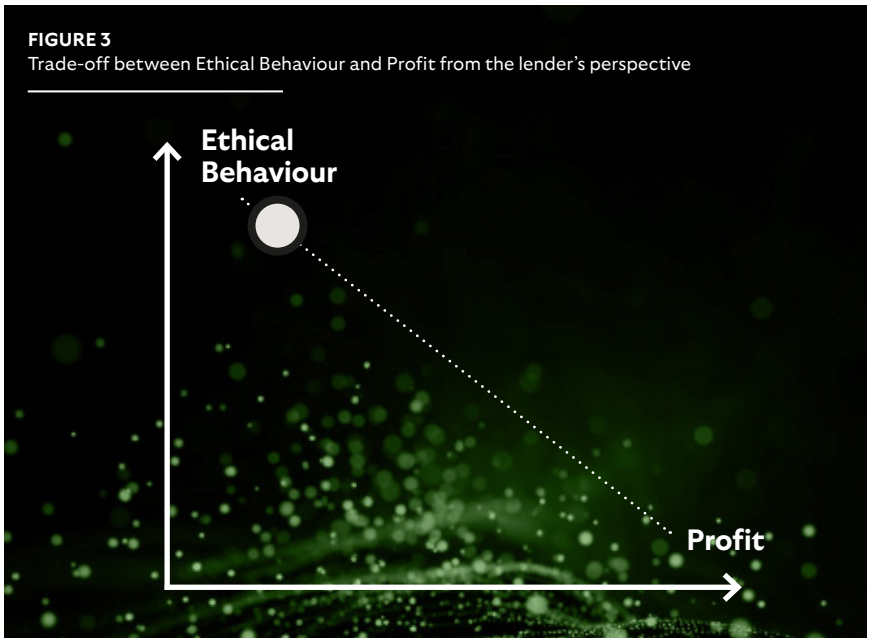

Alignment

From dilemma to governance.

How market design and regulation enable responsible hyperindividualization.

The role of data sovereignty, interoperability and process design in overcoming the perceived dilemma

The previous chapters offer readers little reason for hope. From the perspective of the public debate credit market seems to be caught in an insurmountable dilemma between profit and ethical behavior. But is that really the case?

The answer is: No! If all actors involved take their responsibility seriously, hyperindividualized lending promises win-win scenarios for lenders and borrowers alike. This chapter aims to create a vision of which tangible results we would see that allow us to fully use the potential of AI in hyperindividualized lending: the ability to process more diverse and even incomplete data sets to define the credit worthiness of any given individual, while respecting both ethical and legal boundaries.

Artificial intelligence does not exclusively strengthen the position of lenders. It also dramatically lowers the cost of financial sophistication on the borrower side. AI-powered credit advisors, comparison agents, and affordability simulators can help individuals understand trade-offs between price, privacy, speed, and flexibility before entering a credit contract. Such tools can alert borrowers to unstable repayment structures, identify exploitative pricing patterns, and simulate the long-term consequences of short-term liquidity decisions. When widely accessible, borrower-side AI becomes a market-based counterweight to institutional optimization, reducing asymmetries without requiring paternalistic intervention.

While public debate often frames borrowers primarily as passive recipients of AI-driven decisions, this perspective risks underestimating the adaptive capacity of market participants. Borrowers are not only subjects of algorithmic assessment, but also economic actors who learn, compare, and adjust their behaviour when meaningful information and choice architectures are available. In competitive credit markets, transparency and explainability can enable borrowers to discipline providers through switching, product selection, and reputational feedback. From this perspective, market mechanisms are not opposed to ethical safeguards; rather, they can amplify them when borrowers are empowered to act on understandable information.

Therefore, the argumentation of this chapter follows a straight line: The described dilemma cannot be solved by better models alone. Additionally, the defining elements of the credit value chain have to be included into the equation. The dilemma can be solved by the merging of three pillars:

- Data Sovereignty
- Interoperability
- Process Design

Data sovereignity

Data sovereignty means that individuals and businesses retain meaningful control over how their data is accessed, shared, and used for credit decisions. It is not identical to privacy, and it is not simply a legal checkbox. In a hyper-individualized credit market, sovereignty is the mechanism that turns more data from a threat into a negotiatable resource: borrowers can share data to improve access and terms, while lenders gain higher-quality signals without relying on opaque surveillance-like inputs.

In this win–win architecture, the starting point changes: underwriting becomes a collaboration. Instead of the lender silently assembling a profile from scattered brokers and inferred proxies, the borrower can actively choose a verified data package for this specific loan decision – such as consented transaction history, verified income flows, and on-time utility payments. The borrower understands that more relevant, verifiable data typically improves the lender's ability to price risk precisely, but the outcome is not guaranteed to be lower rate. Precision can also mean a higher, risk-adjusted rate if the data reveals fragility. The difference is that the borrower sees the logic as a transparent trade: privacy and disclosure choices on one side, uncertainty and price on the other.

Technically, AI assistants in the form of Haggling Bots could operationalize data sovereignty by turning data sharing into a controlled, goal-oriented negotiation between borrower and lender: instead of a borrower handing over a broad, static data package, the bot acts as the borrower's agent, disclosing only the minimum information needed (and only with explicit permission) to unlock better loan conditions, while the lender can transparently request additional signals when they materially reduce uncertainty. Conceptually, these bots are specialized AI agents built to negotiate deals on behalf of their users by comparing many possible contract options and iteratively improving offers and counteroffers (Frank, 2025).

This is where data sovereignty becomes more than a slogan. Sovereignty means the borrower can decide what is shared, for which purpose, for how long, and with whom—and can withdraw or correct information through usable mechanisms rather than theoretical rights.

A real-world illustration of sovereignty-as-infrastructure is India's Account Aggregator framework, which operationalizes consent-based financial data sharing through regulated entities and explicit customer instruction. Official descriptions stress that data moves only with explicit consent and that participation is voluntary (Government of India, 2025). That does not automatically mean that Europe should copy India, but that the design lesson should be transferred: when consent and purpose limitation are built into the rails, lenders can access higher-quality, borrower-authorized signals without relying on shadow data practices that undermine trust.

In a European context, a comparable sovereignty-oriented approach can be implemented through consent dashboards, granular permissioning, purpose limitation embedded into data-sharing contracts, and technical enforcement measures such as auditable consent receipts. The point is not that sovereignty eliminates risk, but that it changes the bargaining position. This is how sovereignty begins to dissolve the perceived trade-off: it allows inclusion and personalization to be earned through consent and demonstrable benefit.

Interoperability

Sovereignty on its own, however, quickly hits a practical wall: friction. If every lender needs bespoke integrations and every borrower needs to repeatedly upload documents, re-verify information, and restart the process from scratch, then the market drifts back toward lock-in and data monopolies. That is why interoperability is the economic enabler of responsible personalization. PSD2 is an instructive precedent: it created a regulated access model for payment account data and payment initiation via licensed third parties, tied to customer permission and security requirements (European Union, 2015). Industry standards such as the Berlin Group's NextGenPSD2 framework show how such interoperability can be operationalized through shared API patterns and consent models (The Berlin Group, 2025). In credit terms, interoperable access to permissioned account and transaction data can reduce the temptation to compensate with more intrusive, unstructured alternative data – because the meaningful signals become cheaper and more portable.

The credit-relevant implication is direct. Interoperable access to transaction and account information enables cash-flow-based affordability and risk assessment, which can expand access for consumers and SMEs who are poorly served by bureau-centric scoring. At the same time, interoperability can reduce the incentive to use unstructured or intrusive alternative data because higher-quality financial data becomes portable, permissioned, and standardised. This is one reason the European Commission has pushed beyond open banking toward open finance. The Commission's framework for financial data access (FiDA) is intended to establish rights and obligations for data sharing beyond payment accounts, across a wider range of financial services (European Commission, 2023). The associated legislative proposal explicitly aims to set rules on access, sharing, and use of certain categories of customer data in financial services. In a hyper-individualized credit market, such a framework can reduce coordination failures by setting baseline interoperability expectations and reducing the friction of responsible data sharing.

Interoperability also has a second, often overlooked function: it improves auditability. When data flows are standardized – through consistent schemas, field definitions, and provenance metadata –institutions can more readily explain which information drove a decision and can more readily detect drift, bias, and errors.

In practical terms, interoperability is also how ethical constraints become affordable. If lenders must implement fairness checks, explainability tooling, and compliant disclosures, they benefit from shared rails that reduce bespoke integration costs. Without interoperability, only the largest players can amortize compliance investments. With interoperability, responsible practices can scale across the market and become competitively neutral.

Process design as the trust and safety layer across the credit value chain

Even with sovereignty and interoperability, hyper-individualized credit will not be trusted if decisions remain opaque and unchallengeable. Process design is therefore the third lever: it embeds accountability into the end-to-end lifecycle, from data ingestion and model development to decisioning, servicing and monitoring. The core principle of the new process design is that the system must be governable not only when it works, but also when it fails.

A first practical example of process design that reduces the perceived dilemma is explainable adverse action and meaningful recourse. Translating that principle into credit operations means building a decision pipeline where models produce not only a score or recommendation, but also stable, policy-aligned reason codes that can be communicated in plain language, audited internally, and contested by the consumer. The value for lenders is not only compliance. It is operational resilience and reputational protection: when customers can understand and challenge outcomes, errors surface earlier, escalation costs fall, and trust is less fragile.

From a technical perspective an Explainable AI Assistant could sit at the last mile of the credit decision pipeline and turn an otherwise intransparent model decision into a stable, plain-language explanation that is both consumer-facing and audit-ready: instead of returning only "credit declined," the underwriting model is instrumented to produce policy-aligned reason codes (for example, derived from well-established explanation methods like SHAP or LIME) that are then translated by the assistant into a short, specific adverse-action narrative the applicant can actually act on (Zytek et al., 2024).

Concretely, imagine someone applies for an instalment loan to replace a broken car; the model declines, and the assistant immediately generates: "We couldn't approve this request because (1) your current revolving credit utilization is high relative to your limits, and (2) your stated monthly expenses leave too little remaining income after existing obligations; if either of these is inaccurate, you can submit updated income documents or corrected account information, and you may request a manual review." Designing the process this way directly supports adverse-action requirements that expect creditors to give applicants accurate, specific reasons tied to factors actually considered, even when complex models are used.

A second example of robust process management is model risk management and change control. The PRA's SS1/23 sets out supervisory expectations for banks' management of model risk, aiming for a strategic approach to model risk management as a risk discipline across the model lifecycle (Bank of England, 2023). In hyper-individualized credit, where many models interact and update frequently, disciplined governance becomes a profitability enabler. It reduces silent drift, prevents uncontrolled feature creep into proxy discrimination, and forces organizations to document what a model is for, what it is not for, and under what conditions it should be overridden or retired. This kind of process discipline directly counters the lender-side case of compliance shock or model failure in stress, while also countering the borrower-side fear that decisions are arbitrary and unaccountable.

A third example of new ideas in process design is purpose-bound data handling embedded into workflows. Sovereignty becomes real only when processes enforce it: data is accessed for a declared purpose, retained for a defined period, and segregated from unrelated marketing or cross-selling unless renewed consent is granted. PSD2's access model, which ties third-party access to express customer permission, illustrates how consent conditions can be operationalized at scale through regulated access patterns. The same logic can be extended through open finance-style architectures where permissions, authentication, and access logs become standard features of the market rather than bespoke implementations.

A final piece of process design is industry-level testing and benchmarking. Standardized testing protocols and shared synthetic benchmark datasets that mimic edge cases and vulnerable demographics would allow institutions to stress-test the disparate impact without sharing proprietary data or trade secrets. In practice, this is how an ethical competition becomes possible: firms can demonstrate performance and fairness on comparable tests, regulators can evaluate systems more consistently, and the market is less reliant on alternative data sources.

Why the three levers work together to dissolve the perceived trade-off

Each of these three levers addresses a different root of the perceived dilemma. Data sovereignty reduces power asymmetry by making data sharing conditional, revocable, and legible to the individual. Interoperability reduces coordination failure by making responsible data access cheap, standardised, and portable, preventing lock-in and enabling competition. Process design reduces legitimacy failure by making decisions explainable, auditable, and contestable, and by ensuring that failures trigger disciplined responses rather than silent propagation.

Crucially, the levers reinforce each other. Sovereignty without interoperability becomes friction. Interoperability without sovereignty becomes surveillance. Both without process design become unaccountable automation. Together, they enable a credit market in which inclusion can be pursued without sacrificing agency, and in which profitability can be pursued without relying on opaque discrimination or exploitative personalization. This is the concrete pathway through which a perceived trade-off can become a win-win equilibrium: lenders gain better verified signals, lower operational uncertainty, and reduced compliance and reputational risk, while borrowers gain control, transparency, and meaningful recourse.

Where interoperability and explainability are standardized, ethical conduct does not remain a pure cost factor. Instead, trust, transparency, and fair process design can become dimensions of competition. Lenders may differentiate themselves by offering clearer explanations, more predictable product behaviour, and more robust recourse mechanisms, reducing customer acquisition costs and increasing long-term loyalty. In such an environment, market incentives begin to align with ethical objectives, reinforcing rather than undermining responsible innovation.

The next chapter builds on this engineering logic and turns explicitly to the legislature. A level playing field is not only about banning bad outcomes. It is about ensuring that sovereignty, interoperability, and robust processes are not optional competitive disadvantages, but market-wide requirements and shared infrastructure.

The role of the legislature as the creator of a level playing field: Political, economic, legal, organizational framework conditions

If we take the arguments of the previous chapter seriously, the legislature's role is not limited to consumer protection in the narrow sense. Instead, the legislator becomes the architect of a level playing field that aligns incentives, prevents a race to the bottom in data exploitation and risk shifting, and makes the win-win potential of technological progress politically and economically feasible. This matters because the core dilemma described earlier – ethical behaviour versus profit – often persists not because single firms prefer unethical outcomes, but because competitive pressure and asymmetric information make ethical restraint individually costly and collectively fragile.

From an economic perspective, the case for public rule-setting is strongest where individual market actors cannot credibly commit to standards that would benefit all parties, where harms are externalized, or where informational power is systematically concentrated. In AI credit, these conditions appear simultaneously: lenders and platforms can observe and infer much more than borrowers; outcomes can create spillovers (over-indebtedness, exclusion, social and macroeconomic instability); and the technical opacity of models undermines contestability and trust. The result is a coordination problem that private ordering alone struggles to solve. The state's level playing field function therefore consists of three mutually reinforcing tasks: defining enforceable rights and duties around data sovereignty, mandating or enabling interoperability to reduce lock-in and information rents, and setting process and governance requirements that make compliance auditable, comparable, and scalable across the market.

A useful way to think about legislation here is as a market design layer that sets the boundary conditions for innovation. Instead of choosing between innovation and protection, the state can legislate technical and organizational guardrails that make responsible innovation cheaper than irresponsible innovation. This is precisely the point where data sovereignty, interoperability, and robust process design stop being nice-to-have ethics initiatives and become the infrastructure of a functioning AI credit market.

Hard-law baselines: aligning AI, credit, data protection, and operational resilience

In the EU context, the emerging regulatory stack already sketches the outline of such a level playing field. Most importantly, the EU AI Act classifies AI systems used to evaluate the creditworthiness of natural persons or establish

their credit score as high-risk systems, reflecting the fundamental-rights impact and the high stakes of access to financial resources (European Union, 2024).[2] This classification is not symbolic. It is designed to trigger concrete obligations (risk management, data governance, documentation, transparency, human oversight, robustness), and it is paired with mechanisms such as registration and conformity assessment to make compliance observable rather than merely declaratory.

Consumer-credit regulation increasingly converges with this logic. The revised EU Consumer Credit Directive (Directive (EU) 2023/2225) strengthens expectations around creditworthiness assessment and explicitly addresses automated processing in credit decisions. It requires that the assessment be based on relevant and accurate information, prohibits the use of special categories of data for this purpose, and – importantly for a hyper-individualized market – states that social networks shall not be considered an external source for the creditworthiness assessment under the Directive. Where automated processing of personal data is used, the Directive provides the consumer with rights to human intervention, an explanation of the assessment (including the logic and risks involved), the ability to express a viewpoint, and the right to request a review of the decision. It also requires that rejected applicants be informed without delay and, where relevant, be referred to accessible debt advisory services—an institutional bridge between credit-market decisions and protective social infrastructure.

Data protection law provides an additional baseline that shapes process design. The GDPR establishes the right to data portability, which – when operationalized – supports switching and interoperability incentives in data-driven markets. It also sets out a right not to be subject to decisions based solely on automated processing, including profiling, where such decisions produce legal effects or similarly significant effects, subject to defined exceptions and safeguards.

Finally, even ethically perfect credit models can become socially harmful if the surrounding infrastructure is fragile or compromised. Here, the Digital Operational Resilience Act (DORA) sets a market-wide baseline for ICT risk management, incident reporting, and third-party risk in the financial sector – requirements that become especially relevant when credit decisions rely on outsourced cloud services, external data pipelines, and AI tooling.

Taken together, these instruments demonstrate a key legislative pattern: the level playing field is created not by specifying a single ethical model, but by mandating verifiable governance, controllable data practices, and robust operational processes that all actors must implement.

Data sovereignty as a governed choice architecture, not a one-time consent screen

In hyper-individualized credit, data sovereignty cannot be reduced to consent banners. The core issue is that increasingly powerful inferences can be drawn from data that borrowers never experienced as credit data and, in many cases, never knowingly provided for that purpose. A borrower's explicit right to accept or reject the use of inferred data and synthetic data in credit assessment, combined with transparency about the pricing and decision consequences of those choices could address this problem.

Legally and technically, this kind of sovereignty becomes more feasible when it is tied to interoperable identity and attribute infrastructures. The amended eIDAS framework establishing the European Digital Identity Wallet is directly relevant because it is intended to enable individuals and businesses to authenticate and share identity-linked attributes across the EU in a harmonized way, with user control and strong security and certification expectations. For credit markets, this creates a credible pathway toward attribute-based underwriting where the borrower can selectively present verified proofs like income range, employment status, residency, education credential and repayment history attestations (Austenaa et al., 2024).

Interoperability as anti-lock-in infrastructure: open finance, data governance, and standardized reporting

Interoperability is the economic counterpart to sovereignty. Rights to control data have limited practical value if data cannot move, cannot be understood by recipients, or cannot be verified. The EU's broader data strategy offers building blocks here. The Data Act is designed to increase access to and use of data across sectors and to reduce frictions such as switching costs and contractual imbalances, which can otherwise entrench dominant platforms and limit consumer and business choice. Complementing this, the Data Governance Act seeks to increase trust in data sharing and to strengthen mechanisms (including data intermediation) that can make data availability and reuse more reliable and less transactionally costly.

In financial services specifically, the European Commission's proposal for a Framework for Financial Data Access (FiDA) is a pivotal reference point because it aims to extend data-sharing capabilities beyond payments toward a broader open finance environment, enabling customers to share their financial data with third-party providers under standardized conditions. If implemented effectively,

2 Reference to EU AI Act high-risk classification included for analytical context only.

such an approach can reduce information asymmetries and stimulate competition on service quality rather than on proprietary data hoarding.

Interoperability also has a supervisory dimension. If the legislature wants a genuine level playing field, it needs comparability of compliance and risk across firms. This is where standardized reporting models become relevant. The European Banking Authority's work on data reporting standardization (including common data models and machine-readable reporting concepts) illustrates the direction of travel toward harmonized supervisory data infrastructures. Building on such foundations real-time regulatory reporting APIs on lending platforms would allow continuous monitoring of fairness metrics by supervisory bodies (Crisanto et al., 2020).

Process design obligations: from "trust us" to "show us" via audits, certification, and testing environments

Even with strong rights and interoperability, a hyper-individualized credit market remains fragile unless process obligations make the system inspectable. A framework for mandatory, regular independent third-party model audits – explicitly analogized to a German TÜV certification – could help to detect technical bias and to establish minimum safety expectations in the market. In EU law, the AI Act's conformity assessment logic provides a clear institutional anchor for this idea, because high-risk systems are expected to meet defined requirements and to be assessed in ways that enable market surveillance and accountability (European Union, 2024).

Another element of this strategy could be the provision of open datasets for training and testing to enable private or third-party verification of discrimination risks and to increase comparability between systems. This is a governance choice with trade-offs: open datasets can improve benchmarking and accountability, but they must be designed to avoid privacy leakage, re-identification risks, or the creation of new standardized biases. A legislative approach could therefore focus on curated, privacy-preserving benchmark suites and secure evaluation access rather than unrestricted publication of sensitive datasets, aligning the goal of comparability with data protection constraints.

Market integrity and antitrust: preventing "the machine did it" defenses

AI-driven personalization also creates new competition risks. If multiple platforms use algorithmic pricing and monitoring, markets can drift toward tacit coordination without explicit agreements, undermining price competition while making enforcement difficult. The proactive expanding of antitrust regulation could help to cover algorithmic pricing and monitoring and prevent machine-created price agreements or tacit collusion. This concern is reflected in policy research: the OECD has explicitly analyzed how algorithms can facilitate collusion and complicate detection and enforcement in competition policy (OECD, 2017). In Germany, the Monopolkommission has also discussed how algorithms can affect competition and collusion dynamics, reinforcing the point that modern market power can be exercised through data and optimization rather than through traditional price-fixing meetings (Monopolkommission, 2018).

A level playing field approach here would combine competition enforcement with interoperability policies that reduce data-driven lock-in. If switching is feasible, data access is standardized, and decision processes are auditable, dominant players have fewer structural advantages – and smaller lenders can compete on service, ethics, and product design.

Financial Education and advice and representation services

A central promise of AI credit markets is that better data and models can widen access – especially for thin-file consumers and underserved groups. But inclusion without safeguards can become extraction, particularly in short-term credit and BNPL-like environments. The EU Consumer Credit Directive explicitly addresses the need for financial education and support for consumers in difficulty, requiring Member States to promote measures for responsible borrowing and debt management and to ensure dissemination of information about guidance from consumer organizations and authorities.

Borrower empowerment, therefore, also depends on financial education, not as a moral obligation but as an economic productivity factor. Better-informed borrowers tend to select products that fit their repayment capacity more accurately, leading to lower default rates, fewer disputes, and more stable portfolio performance. In this sense, financial literacy complements data sovereignty: transparency and choice only translate into better outcomes when borrowers can interpret and act on them. Investments in accessible financial education therefore support both consumer protection and market efficiency.

Imagine a single parent, Paul, whose washing machine breaks on the 28th of the month. Two kids, no help available and the laundromat would cost time and money he does not have. His account is at zero because rent and childcare just went out, and his salary arrives in three days. Paul is not asking for a long-term loan. He needs a small, short-term bridge so daily life does not fall apart.

In a market without safeguards, this is exactly the kind of moment hyperpersonalized credit systems can exploit. Paul searches online for "washing machine delivery tomorrow" and "pay later." Within minutes, he gets offers for "instant approval" microloans or BNPL plans. The lender's algorithm can infer urgency and vulnerability from signals like timing (end of month), prior overdrafts, location, device data, shopping behavior, and the fact that he's looking at essentials rather than discretionary goods. The risk is not just that he pays some interest, but that the price becomes tailored to his weakest moment: a high fee for immediate cash, short repayment deadlines, penalties if he's late by even one day, and automatic rollovers that turn a washing machine emergency into a repeating, high-cost cycle.

In such a situation an advice and representation service can help Paul understand options, challenge unfair terms, and intervene early if no credits are available or repayment becomes difficult. In this view, the state's level playing field role is not to eliminate fast credit, but to ensure that speed and flexibility do not become a tool for extracting maximum profit from people precisely when they have the least bargaining power.

Organizational governance: multi-stakeholder policy processes, sandboxes, and enforceability under uncertainty

Because AI credit markets evolve quickly, legislative design must balance stability with adaptability. The eIDAS framework explicitly encourages innovation-friendly implementation instruments such as sandboxes to test controlled solutions, while maintaining strong requirements for trust, security, and certification. This points to a broader institutional lesson: level playing fields are built not only through prohibitions, but also through shared infrastructure and trusted mechanisms for testing, certification, and iteration.

In practical terms, an effective legislative strategy for AI credit markets therefore combines enforceable baseline rights and obligations (data protection, consumer credit safeguards, AI governance, operational resilience) with targeted enabling infrastructure (digital identity, interoperable data access frameworks, standardized reporting) and credible accountability mechanisms (audits, certification, supervisory analytics, and accessible redress). When these elements are aligned, the perceived trade-off between ethical behaviour and profit can shift: responsible practices become a competitive advantage and a compliance necessity rather than a discretionary cost.

Table 1 gives an overview over the proposed mechanisms to ensure Accountability in Hyperindividualized Lending and show pathways to optimized consumer credits:

TABLE 1
Proposed Measurments to overcome the ethical behaviour vs. profit dilemma in credit

Proposed Measurements
Aligning AI, credit, data protection, and operational resilience
Data sovereignty as a governed choice architecture, not a one-time consent screen
Interoperability as anti-lock-in infrastructure: open finance, data governance, and standardized reporting
Process design obligations: from "trust us" to "show us" via audits, certification, and testing environments
Market integrity and antitrust: preventing "the machine did it" defenses
Financial Education and advice and representation services
Organizational governance: multi-stakeholder policy processes, sandboxes, and enforceability under uncertainty

Outlook

From perceived dilemma to a sustainable equilibrium.

How shared standards and infrastructure can stabilize AI-based credit markets.

An important promise associated with hyperindividualized credit is "banking the unbanked" and, more broadly, expanding financial inclusion. This is a powerful hope: it could mean that people are pulled away from unregulated credit markets towards regulated ones, provinding a transparent process with a risk-adequate pricing. The World Bank estimates that 1.4 billion adults globally could profit from such a development as they remain unbanked in 2021 (Demirgüç-Kunt et al., 2021).

In principle, digital delivery channels and alternative data can help bridge gaps for the unbanked with thin traditional credit histories. Digital footprints – such as transaction histories, cash-flow patterns, or other verifiable behavioral signals – can reduce information asymmetries and enable lenders to offer appropriately sized, responsibly priced credit to individuals and micro-businesses previously excluded from formal finance. Yet inclusion is not automatically beneficial. Digital finance can also expose vulnerable groups to new forms of harm, including aggressive targeting, over-borrowing, opaque pricing, and digital financial exclusion when people lack the skills, connectivity, or agency to navigate increasingly complex products. The fictive example of Josephine in the first chapter is a worrying example of this development.

This position paper set out to clarify how new technologies and processes can mitigate risks for consumers and businesses in a hyper-individualized, AI-based credit market. The preceding chapters have shown why the perception of a structural trade-off between profit maximisation and ethical lending remains so persistent: when credit decisions are delegated to complex data pipelines and opaque models, lenders can gain speed and margin, while borrowers may experience reduced transparency, heightened vulnerability to error, and diminished practical contestability.

The described dilemma cannot be solved by better AI models alone. It can be solved by following a clear pathway to overcome the perceived dilemma between ethical behaviour and profit in credit, the paper propses a clear pathway: The win-win potential becomes tangible once Data Sovereignty, Data Interoperability and the development of transparent credit processes are treated as operational capabilities.

The central conclusion of this paper is that the "profit versus ethical behaviour" tension is best understood as a governance deficit rather than an immutable economic law. It is not AI per se that creates the perceived dilemma. Rather, the dilemma emerges when data is captured without meaningful user agency, when interoperability is weak and locks consumers into information asymmetries, and when process design fragments accountability across the credit value chain.

Concrete examples show that these ideas can be implemented at scale when governance and architecture coevolve. Consent-based data-sharing infrastructures such as India's Account Aggregator ecosystem are frequently cited as a form of rails that combine interoperable APIs with user-mediated consent artefacts, enabling secure, purpose-limited data flows between regulated entities. The European context is different, but the underlying design lesson is transferable: Digital sovereignity, interoperability and process standards are not just technical standards. They are governance instrument that can make consumers' data agency legible, enforceable, and economically valuable without collapsing into uncontrolled data extraction.

The outlook of the paper, therefore, is optimistic: the hyper-individualized, AI-based credit market can evolve toward a regime where profitability and ethical behaviour

reinforce each other and the publicly perceived dilemma between lenders and borrowers in credit can be overcome (Figure 4).

The possible future equilibrium of hyperindividualized credit is therefore not one in which regulation replaces markets, but one in which markets operate under conditions that empower both sides technologically. When borrowers are supported by education and decision tools, and lenders compete under shared governance standards, market mechanisms can reinforce accountability rather than erode it. In this configuration, ethical behaviour and profitability are no longer opposing forces, but mutually reinforcing drivers of sustainable credit markets.

> The future of AI in finance will not be defined by how much data we can extract, but by how responsibly we design the systems that use it.

FIGURE 4
Overcoming the Perceived Dilemma in the Credit Market

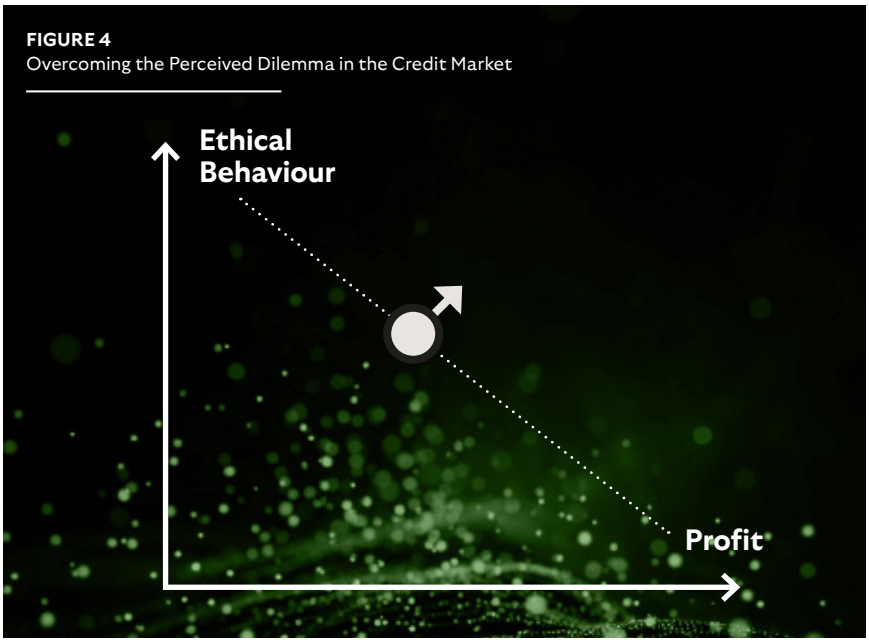

Akerlof, G. A. (1970). The Market for "Lemons": Quality Uncertainty and the Market Mechanism, he Quarterly Journal of Economics, Volume 84, Issue 3, URL: [https://www.sfu.ca/~wainwrig/Econ400/akerlof.pdf].

Austenaa, M., L. Bailly, S. Kauhaus, R. Prasad and J. van Vonno (2024): What does it take to use the European Digital Identity Wallet for payment, White paper from the European Digital Identity Wallet Consortium, published in: www.eudiwalletconsortium.org, URL: [https://eudiwalletconsortium.org/wp-content/uploads/2024/11/241101_Eudiw_white-paper_web.pdf].

Bank of England (2023): SS1/23 – Model risk management principles for banks, published in: www.bankofengland.co.uk, URL: https://www.bankofengland.co.uk/prudential-regulation/publication/2023/may/model-risk-management-principles-for-banks-ss]

Bansak, Kirk and Elisabeth Paulson (2024): Public attitudes on performance for algorithmic and human decision-makers, published in: www.pubmed.ncbi.nlm.nih.gov, URL: [https://pmc.ncbi.nlm.nih.gov/articles/PMC11631221/].

Barocas, Solon and Andrew D. Selbst (2016): Big Data's Disparate Impact, published in: www.cs.yale.edu, URL: [https://www.cs.yale.edu/homes/jf/BarocasSelbst.pdf].

CFPB (2015): Data Point: Credit Invisibles, published in: www.consumerfinance.gov, URL: [https://files.consumerfinance.gov/f/201505_cfpb_data-point-credit-invisibles.pdf].

Crisanto, Juan Carlos, Katharina Kienecker, Jermy Prenio and Eileen Tan (2020): From data reporting to data-sharing: how far can suptech and other innovations challenge the status quo of regulatory reporting?, FSI Insights on policy implementation, No. 29, published in: www.bis.org, URL: [https://www.bis.org/fsi/publ/insights29.pdf Bank for International Settlements].

Demirgüç-Kunt, Asli, Leora Klapper, Dorothe Singer and Saniya Ansar (2021): The Global Findex Database 2021 Financial Inclusion, Digital Payments, and Resilience in the Age of COVID-19, published in: www.worldbank.org, URL: [https://documents1.worldbank.org/curated/en/099818107072234182/pdf/IDU-06a834fe908933040670a6560f-44e3f4d35b7.pdf].

European Commission (2023): Framework for financial data access, URL: [https://finance.ec.europa.eu/digital-finance/framework-financial-data-access_en].

European Commission, (2025). AI Act, Article 10: Data and data governance, URL: [https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-10].

European Banking Authority (2024). Special topic – Artificial intelligence (adoption trends, methods, sourcing patterns), URL: [https://www.eba.europa.eu/publications-and-media/publications/special-topic-artificial-intelligence].

European Union (2015): Directive (EU) 2015/2366 of the European Parliament and of the Council of 25 November 2015 on payment services in the internal market, published in: www.eur-lex.europa.eu, URL: [https://eur-lex.europa.eu/eli/dir/2015/2366/oj/eng].

European Union (2024): Regulation (EU) 2024/1689 of the European Parliament and the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations, published in: www.eur-lex.europa.eu, URL: [https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=OJ%3AL_202401689].

Federal Reserve (2025): Alternative Data: Expanding Access to Credit (Consumer & Community Context), URL: [https://www.federalreserve.gov/publications/files/consumer-community-context-20251017.pdf].

Frank, Roland (2025): Fintech 2040 – Trajectories for the Evolution of the Fintech Ecosystem, published in: riverty.com, URL: [https://www.riverty.com/49a2b2/globalassets/media-images/insights/fintech-2040/riverty_fintech2040.pdf].

Financial Stability Board (2024). The Financial Stability Implications of Artificial Intelligence, URL: [https://www.fsb.org/uploads/P14112024.pdf].

FinRegLab (2023): Explainability & Fairness in Machine Learning for Credit Underwriting: Policy Analysis, URL: [https://finreglab.org/wp-content/uploads/2023/12/FinRegLab_2023-12-07_Research-Report_Explainability-and-Fairness-in-Machine-Learning-for-Credit-Undewriting_Policy-Analysis.pdf].

Gambacorta, Leonardo, Yiping Huang, Han Qiu and Jingyi Wang (2019). How do machine learning and non-traditional data affect credit scoring?, BIS Working Papers No. 834, URL: [https://www.bis.org/publ/work834.pdf].

Government of India (2025): Account Aggregator Framework, published in: www.financialservices.gov.in, URL: [https://financialservices.gov.in/beta/en/account-aggregator-framework].

Krivorotov, G. (2023). Machine learning-based profit modeling for credit card underwriting, Journal of Banking & Finance, Volume 149, URL: [https://www.sciencedirect.com/science/article/abs/pii/S0378426623000213].

Monopolkommission (2018): Algorithms and collusion: Excerpt from Chapter I of the XXII Biennial Report of the Monopolies Commission, URL: [https://www.monopolkommission.de/images/HG22/Main_Report_XXII_Algorithms_and_Collusion.pdf].

NIST (2023): Artificial Intelligence Risk Management Framework, published in: www.nvlpubs.nist.gov, URL: [https://nvlpubs.nist.gov/nistpubs/ai/nist.ai.100-1.pdf].

OECD (2017): Algorithms and Collusion – Background Note by the Secretariat (DAF/COMP(2017)4). OECD Competition Committee, URL: [https://one.oecd.org/document/DAF/COMP(2017)4/en/pdf].

OECD (2022): Data Shaping Firms and Markets, OECD Digital Economy Papers, No. 344, URL: [https://www.oecd.org/content/dam/oecd/en/publications/reports/2022/12/data-shaping-firms-and-markets_b0cbadeb/7b1a2d70-en.pdf].

Pies, Ingo (2009): Pies, Ingo (2009): Das ordonomische Forschungsprogramm, Discussion Paper, No. 2009-7, Martin-Luther-Universität Halle-Wittenberg, published in: econstor.eu, URL: [https://www.econstor.eu/bitstream/10419/170317/1/dp2009-07.pdf].

Riverty (2025): Fintech 2040: Where money becomes invisible, URL: [https://www.riverty.com/en/business/insights/fintech-2040/].

Song, Fenghua and Anjan Thakor (2022): Ethics, capital and talent competition in banking, Journal of Financial Intermediation, Volume 52, URL: [https://www.sciencedirect.com/science/article/abs/pii/S104295732200016X].

Stiglitz, Joseph E. and Andrew Weiss (1981): Credit Rationing in Markets with Imperfect Information, The American Economic Review, Volume 71, Issue 3, URL: [https://pages.ucsd.edu/~aronatas/project/academic/Stiglitz%20credit.pdf].

The Berlin Group (2025): PSD2 Access to Bank Accounts, published in: www.berlin-group.org, URL: [https://www.berlin-group.org/psd2-access-to-bank-accounts].

World Bank (2024). The Use of Alternative Data in Credit Risk Assessment: The Opportunities, Risks, and Challenges, URL: [https://documents1.worldbank.org/curated/en/099031325132018527/pdf/P179614-3e01b947-cbae-41e4-85dd-2905b6187932.pdf].

Zytek, Alexandra, Sara Pido, Sarah Alnegheimish, Laure Berti-Equille, and Kalyan Veeramachaneni (2024): Explingo: Explaining AI Predictions using Large Language Models, published in: arXiv.com, URL: [https://arxiv.org/pdf/2412.05145].

About the Autors:

Anita Klingel is Head of AI Science and Transformation at the IPAI Foundation in Germany. She is an expert in AI governance & testing, having published several books and papers on those subjects.

Prior to her role at IPAI, she led the AI team at the German state's in-house consultancy, where she advised various levels of the public sector on the implementation of AI systems. In 2023, she was recognized by Handelsblatt as a "Young Leader in GovTech."

Jan Richard is a Product Marketing leader at Riverty and an eCommerce expert with over 15 years of experience in order-to-cash and post-purchase processes. He operates at the intersection of payments, consumer finance, and digital business models, helping merchants turn financial infrastructure into a strategic growth lever.

Jan convenes executive roundtables with global eCommerce leaders, advancing dialogue on payment innovation and operational excellence. He is known for translating complex payment and accounting structures into scalable, customer-centric solutions.

Florian Haubold is a marketing leader at Riverty, responsible for brand, communications, and go-to-market initiatives across Europe. He focuses on positioning financial solutions in embedded finance, digital payments, and AI-driven credit.

Florian leads cross-functional marketing programs and thought-leadership initiatives such as Fintech 2040, translating complex financial and technological topics into clear market narratives. His background spans more than 20 years in financial services, digital business models, and growth marketing across European markets.

About Riverty:

Riverty, the Fintech company of **Bertelsmann**, supports thousands of merchants and over 28 million consumers by processing more than 80 million transactions monthly. Offering flexible payments, debt collection, and smart accounting solutions, Riverty empowers businesses and consumers with cutting-edge financial services. With a dedicated team of over 4,000 employees across 11 countries in Europe and North America, Riverty is a leader in delivering comprehensive financial solutions.

www.riverty.com

This publication is a thought leadership contribution. It does not constitute product communication, financial promotion, or a statement of operational AI deployment.

riverty.com

RIVERTY