P2P Lending & Robo-Advisors: The Fairness-Accuracy Tradeoff

Algorithms that promise financial inclusion may be encoding the very biases they claim to overcome

Digital Finance

Why Would an Algorithm Designed to Help You Get a Loan Reject You for Where You Live?

The Paradox

P2P lending platforms promise to democratize credit – no more stuffy bank managers making subjective decisions. Algorithms assess everyone equally. Or do they?

What makes algorithmic lending attractive:

- No human bias (no racism, sexism, or mood-dependent decisions)
- Thousands of data points analyzed in milliseconds
- Access for "thin-file" borrowers excluded by traditional banks
- Lower costs passed to borrowers as lower interest rates

What algorithmic lending cannot escape:

- Training data reflects historical lending decisions – which were biased
- ZIP codes correlate with race. Job title correlates with gender
- An algorithm that maximizes accuracy will exploit every correlation – including the discriminatory ones
- "The algorithm is only as fair as the data it learned from"

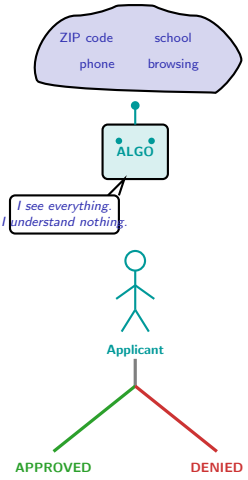

P2P platforms replaced the biased loan officer with an algorithm – but the algorithm learned from the biased loan officer's decisions.

Have You Ever Been Judged by Data You Didn't Choose to Share?

Reflection Prompt

Think about the digital footprint you leave every day. Your phone logs where you go. Your browser records what you read. Your payment history shows what you buy, when, and where.

Now imagine a lender using ALL of this to decide whether you deserve a loan.

Most of us would say: "I have nothing to hide." But consider these scenarios:

- You searched for "bankruptcy advice" to help a friend – now YOUR risk score drops
- You live in a zip code with high default rates, even though YOU have never missed a payment
- You changed jobs three times in two years – the algorithm sees instability, you see career growth
- You browse on an older phone model – the algorithm correlates this with lower income

Each data point is accurate. The INFERENCE drawn from it may not be.

This is the core tension: more data improves accuracy for the population, but may harm INDIVIDUALS who do not fit the pattern.

Alternative data promises inclusion for the 1.7 billion unbanked adults – but the same data can be used to exclude in new, invisible ways.

What Separates a Credit Score from an Algorithmic Verdict?

Dimension	Traditional FICO	ML Credit Scoring	Fair-ML Scoring
Input data	5 factors (payment, debt, length, mix, new)	Hundreds of features	Constrained features
Model type	Linear, static	Non-linear, dynamic	Non-linear + fairness constraints
Explainability	High (weights known)	Low (black box)	Medium (constrained)
Regulation	Established (ECOA, FCRA)	Evolving (EU AI Act)	Emerging best practice
Accuracy	Moderate	Highest	Slightly reduced
Fairness	Unintentional bias	May amplify bias	Designed to mitigate
Coverage	Excludes thin-file	Includes thin-file	Includes thin-file

The critical distinction is the OBJECTIVE FUNCTION

- Traditional FICO: minimize prediction error (accuracy only)
- ML Black-Box: maximize AUC – accuracy at any cost, including fairness
- Fair-ML: maximize AUC SUBJECT TO fairness constraints (demographic parity, equalized odds)

Why it matters for P2P lending

P2P platforms that use unconstrained ML may offer better rates to some borrowers while systematically disadvantaging others – not because of creditworthiness, but because of correlated attributes.

The objective function is a MORAL choice disguised as a TECHNICAL one – what you optimize for determines who gets credit.

Follow One Loan Application Through Three Different Scoring Models

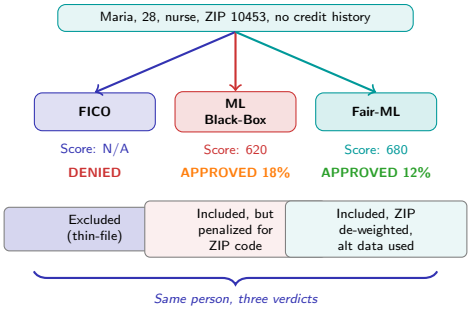

Three models, three outcomes

- **FICO:** Maria has no credit history. Traditional scoring cannot evaluate her. She is excluded entirely.
- **ML Black-Box:** The model finds signal – but her ZIP code (South Bronx) correlates with race. She is assigned high risk partly because of where she lives, not how she behaves.
- **Fair-ML:** The model constrains protected-correlated features. Maria's rent payments and employment stability drive the score.

Key insight: The "best" model depends on what you mean by "best." Accuracy? Fairness? Inclusion? You cannot maximize all three simultaneously.

Maria is illustrative. But the pattern is real: 26 million Americans are 'credit invisible' and ML models score them very differently depending on the objective function.

How Do You Build a Model That Is Both Accurate and Fair?

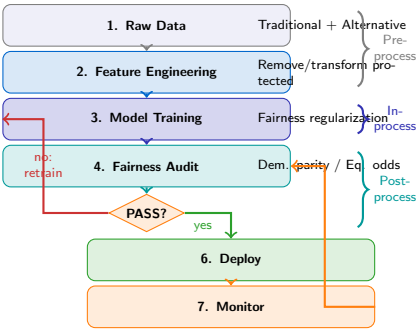

Three intervention points

- **Pre-processing:** Remove or transform features that correlate with protected attributes. Risk: removing a feature does not remove the information if other features are correlated.
- **In-processing:** Add fairness constraints to the optimization objective. The model trades accuracy for fairness during training. This is where the "dial" is set.
- **Post-processing:** Adjust model outputs after scoring to equalize approval rates across groups. Simple, but may violate individual fairness.

No approach is perfect. Each creates a different tradeoff. The choice is institutional, not technical.

Pre-processing removes features, in-processing constrains optimization, post-processing adjusts outputs – each creates a different fairness-accuracy tradeoff.

What Happens When the Fair Model Costs More and the Accurate Model Discriminates?

The Double Bind

P2P platforms face a genuine dilemma. If they use the most accurate model, they risk:

- Regulatory penalties: EU AI Act classifies credit scoring as "high-risk AI." Fines up to 3% of global turnover
- Reputational damage: investigative journalism exposes algorithmic discrimination (ProPublica, The Markup)
- Legal liability: US ECOA prohibits discrimination. "The algorithm did it" is not a legal defense

If they use the fairest model, they risk:

- Higher default rates: fairness constraints may approve genuinely riskier borrowers
- Competitive disadvantage: higher defaults mean higher rates, losing low-risk borrowers to competitors
- Adverse selection: good borrowers leave, bad borrowers stay – the lemons problem returns

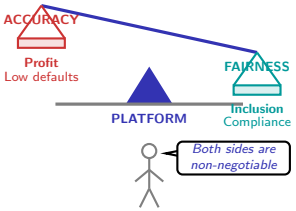

The platform that gets the tradeoff wrong loses either its license or its customers.

Regulators demand fairness. Markets demand accuracy. No P2P platform has fully solved this tension – it may be inherently unsolvable.

Where Is Algorithmic Lending Actually Expanding – and Where Is It Contracting?

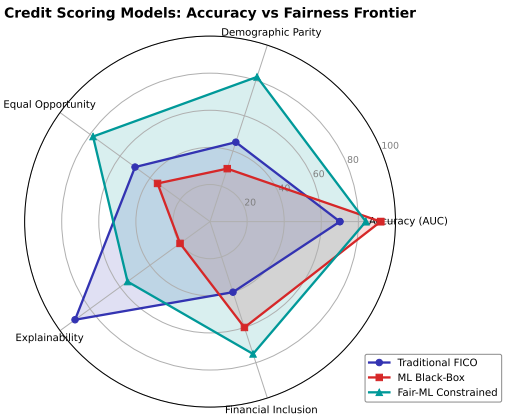

Illustrative comparison based on published research

No model dominates on all dimensions

- FICO excels at explainability but fails on inclusion – 1.7B adults have no credit history
- ML Black-Box achieves highest accuracy but scores worst on fairness
- Fair-ML offers the best balance but sacrifices peak accuracy – the 8-point AUC gap costs real money

Expanding:

- Emerging markets (M-Shwari, CASHe) – mobile data fills the thin-file gap
- SME lending (Funding Circle, Kabbage) – alternative data from operations
- BNPL (Klarna, Affirm) – real-time decisioning at point of sale

Contracting:

- Mortgage lending – regulatory scrutiny too high for black-box models
- Markets where EU AI Act enforcement has begun

The radar chart reveals the fundamental insight: expanding one dimension contracts another. No model dominates on all axes.

Who Wins and Who Loses When Algorithms Replace Loan Officers?

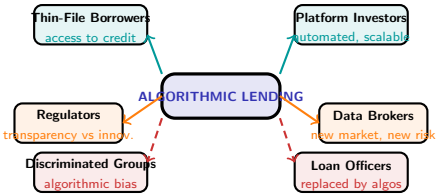

Winners

+ **Thin-file borrowers:** Previously excluded, now scored using alternative data. Immigrants and gig workers gain access.
+ **Platform investors:** Automated underwriting scales without hiring. Consistent decisions reduce operational risk.

Losers

– **Discriminated groups:** Communities redlined by banks may be redlined by algorithms – invisibly.
– **Loan officers:** Automated decisioning eliminates the traditional credit analyst role.

Mixed impact

∼ **Regulators:** Every decision is logged, but the model itself is a black box.
∼ **Data brokers:** New market for alternative data, but GDPR and EU AI Act constrain collection.

Algorithmic lending does not eliminate discrimination – it transforms it from visible (a loan officer's prejudice) to invisible (a model's feature weights).

The Fairness Dial: Every Model Makes a Choice – Where Do You Set It?

The P2P Lending Fairness Scorecard
When evaluating any algorithmic lending platform, ask these five questions.

1. **What is the objective function?** Does it optimize for accuracy alone, or include fairness constraints? If the platform cannot answer this, that IS the answer.
2. **What data does the model use?** Are protected-correlated features (ZIP code, employer, school) included? How are they handled?
3. **Can the model explain its decisions?** EU AI Act requires explanations for high-risk AI. Can a rejected applicant understand WHY?
4. **How is fairness measured?** Demographic parity? Equalized odds? Individual fairness? The choice of metric determines what "fair" means.
5. **Who audits the model?** Internal team? External auditor? No one? Independent audits are the only credible check.

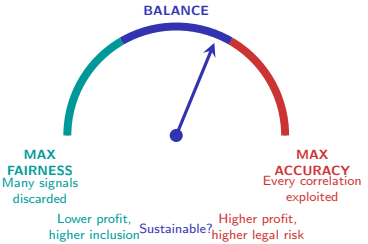

Every lending platform sets this dial – most do not tell you where.

There is no objectively 'correct' position on the dial – it is a value judgment. The five questions above help you assess where any platform has placed it.

Your Challenge: Audit a Credit Scoring Model for Bias

Mini-Challenge (15 minutes)

You are hired as a fairness auditor for a European P2P lending platform that uses ML credit scoring. The platform's model has the highest accuracy in its market – but a journalist has published an article showing that approval rates for applicants in minority neighborhoods are 40% lower than for identical applicants in majority neighborhoods.

Your deliverable: A one-page audit plan answering each of the following:

1. **Data audit:** Which features correlate with protected attributes? List at least three features that could serve as proxies for race, gender, or age.
2. **Metric selection:** Would you measure fairness using demographic parity, equalized odds, or individual fairness? Justify your choice.
3. **Intervention recommendation:** Would you use pre-processing, in-processing, or post-processing? What are the risks of your choice?
4. **Stakeholder impact:** Who benefits and who is harmed by your proposed intervention? Consider borrowers, investors, the platform, and regulators.
5. **Regulatory compliance:** How does your recommendation align with the EU AI Act's requirements for high-risk AI systems?

Conclude with a one-sentence recommendation: Should the platform accept a reduction in accuracy to improve fairness, and if so, how much?

The best way to understand the fairness-accuracy tradeoff is to sit on the uncomfortable side of the table – where every decision has consequences for real people.