Optimizing Digital Credit, Risk Modelling, and Inclusive Finance Under Uncertainty

Timothy M. Bisakaya
The Co-operative University of Kenya

ROADMAP

Structure of the Talk

01	**Introduction**	*The digital credit paradigm, the taxonomy of uncertainty, and the inclusion–risk trilemma*
02	**System Architecture**	*Data pipelines, alternative-data features, graph risk, and real-time decisioning infrastructure*
03	**The Framework**	*Selection bias, explainability, drift, survival models, decision theory, bandits, CVaR, fairness, ECL*
04	**Methodology & Approach**	*Model-class selection, stress testing, and adversarial robustness in practice*
05	**Conceptual Findings**	*Applied outcomes from an emerging-market deployment*
06	**Conclusions & Future Directions**	*Synthesis and open research frontiers*
07	**References**	*Primary sources underpinning the mathematical framework*

The Digital Credit Paradigm Shift

MOTIVATION

From collateralized banking to instant, uncollateralized micro-credit

- **Inputs:** shift from structured, static, near-Gaussian data to high-frequency, non-Gaussian, high-dimensional behavioural signals.
- **Decisioning:** micro-second, zero human-in-the-loop scoring under ultra-thin margins — precision in default prediction is not optional.
- **Drivers:** mobile-money logs, GSM call-detail records (CDRs), e-commerce transaction streams.
- **Tension:** instant liquidity demand versus systemic over-indebtedness risk.

SIGNAL → DECISION PIPELINE



< 200 ms decision latency at the point of the last stage

Order-of-magnitude scale: millions of thin-file applicants; loan tenors of 7–30 days; feature vectors refreshed hourly.

Taxonomy of Uncertainty

DEFINITION

σ

Aleatoric Uncertainty

Irreducible noise

$$\sigma^2_\text{aleatoric} = \text{Var}(\text{ income} \mid \text{borrower})$$

- Intrinsic randomness in borrower income and cash-flow volatility
- Does not shrink with more data on the same applicant
- Practitioner response: capital buffers & macro-hedging

∇

Epistemic Uncertainty

Model / data deficiency

$$\sigma^2_\text{epistemic} = \text{Var}(\text{ model posterior})$$

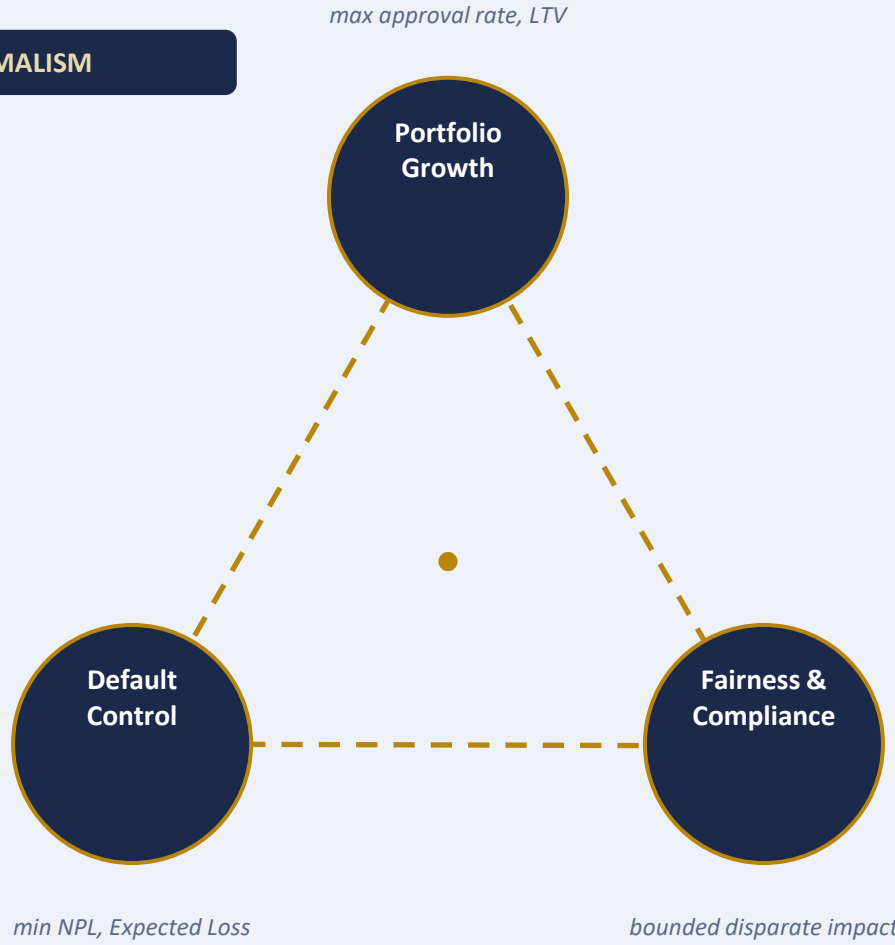
- Lack of historical labels for thin-file / first-time applicants
- Shrinks as data accumulates — the object of active learning
- Practitioner response: bandits, conservative initial limits

Total predictive variance decomposes as:

$$\text{Var}[y] = \text{E}[\sigma^2_\text{epistemic}] + \text{E}[\sigma^2_\text{aleatoric}]$$

The Inclusion–Risk Trilemma

FORMALISM



Optimizing any two vertices tends to degrade the third

- A pure profit objective over-serves high-signal segments and under-serves thin-file, protected-class applicants.
- A pure risk-minimization objective collapses approval rates — the classic "conservative underwriting trap."
- A pure fairness constraint, applied naively, can raise portfolio-level Expected Loss beyond risk appetite.

$$\max_{\theta} E[\text{Profit}(\theta)]$$
$$\text{s.t. } EL(\theta) \leq \text{cap}, \ |\text{ParityGap}(\theta)| \leq \varepsilon$$

→ requires joint, constrained, multi-objective optimization — not sequential heuristics.

High-Dimensional Feature Extraction

DATA PIPELINE

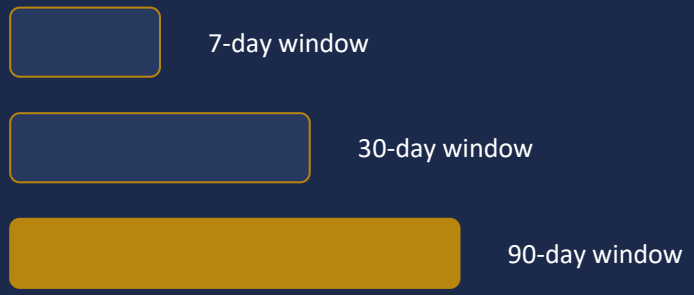
From raw digital footprint to predictive feature

- **Sources:** mobile-money transfers, airtime top-ups, app-usage metadata, geolocation stability.
- **Sparsity:** $p \gg n$ in early-tenure applicants — more candidate features than observations.
- **Sampling:** irregular, multi-rate time series across channels; no common clock.

Velocity feature:

$$\text{velocity} = \text{outflow_7d} / \text{outflow_30d}$$

AGGREGATION WINDOWS



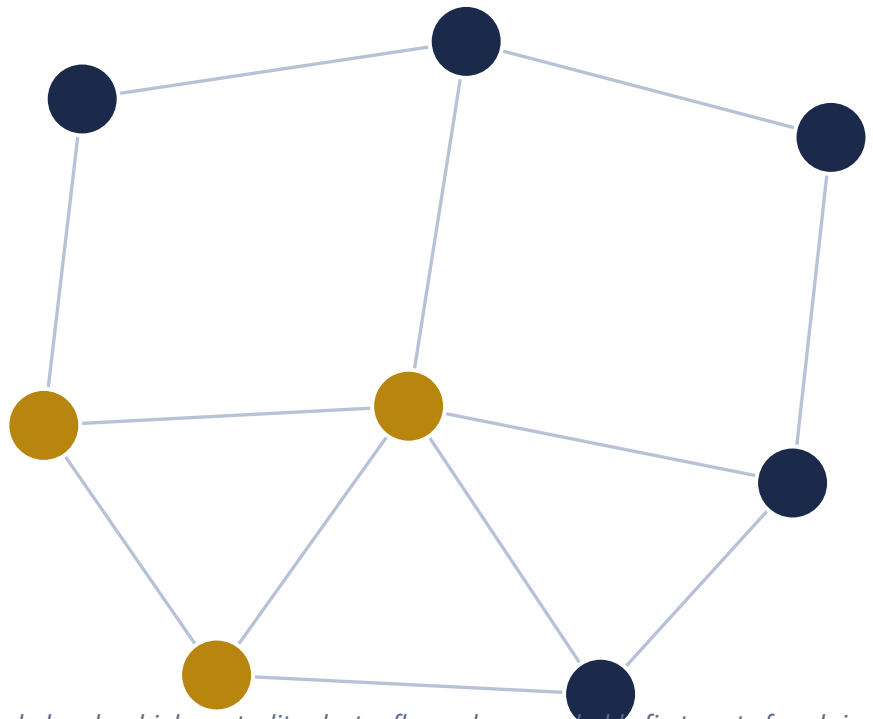
Feature families engineered per window:

- Velocity ratios across windows (short vs. long horizon)
- Entropy of the transaction-partner network
- Recency, frequency, monetary (RFM) decompositions

5

Graph Analytics & Network-Based Risk

GRAPH THEORY



Shaded nodes: high-centrality cluster flagged as a probable first-party fraud ring

Financial ecosystems as directed, weighted graphs

$$G = (V, E, W)$$

$$h_v^{(k)} = AGG\{ h_u^{(k-1)} : u \in N(v) \}$$

Graph Neural Networks propagate node embeddings over *k* hops of the transaction / social graph.

- PageRank-type centrality proxies economic stability
- Detects contagion & syndicate risk invisible to per-node features

MLOps Architecture for Real-Time Underwriting

SYSTEMS DESIGN



Feature Store

Dual online (low-latency, e.g. Redis) & offline (batch, e.g. Snowflake / BigQuery) stores with point-in-time correctness.



Inference Engine

Micro-service architecture (gRPC, containerized REST) executing end-to-end in < 200 ms.



Monitoring Layer

Automated drift detection, data-quality checks, latency SLAs, alerting.

The Reject Inference Problem

SELECTION BIAS

Models trained only on accepted applicants inherit sample selection bias

$$Y \text{ observed} \Leftrightarrow S = 1$$
$$\text{target unobserved: } P(Y=1 \mid X, S=0)$$

Y = default outcome, *S* = historical accept/reject decision.

- The counterfactual default probability of a rejected applicant is never realised.
- Naive re-scoring on accepts-only drifts the model toward hyper-conservative underwriting.

Correction techniques

1

Heckman Correction

Two-stage parametric model of the selection process, using an exclusion restriction.

2

Inverse Propensity Weighting

Re-weight accepted observations by 1 / P(S=1 | X) to approximate the full population.

3

Semi-Supervised Pseudo-Labeling

Boundary-adjacent exploration assigns soft labels to rejected applicants.

Explainable AI: Shapley Additive Explanations

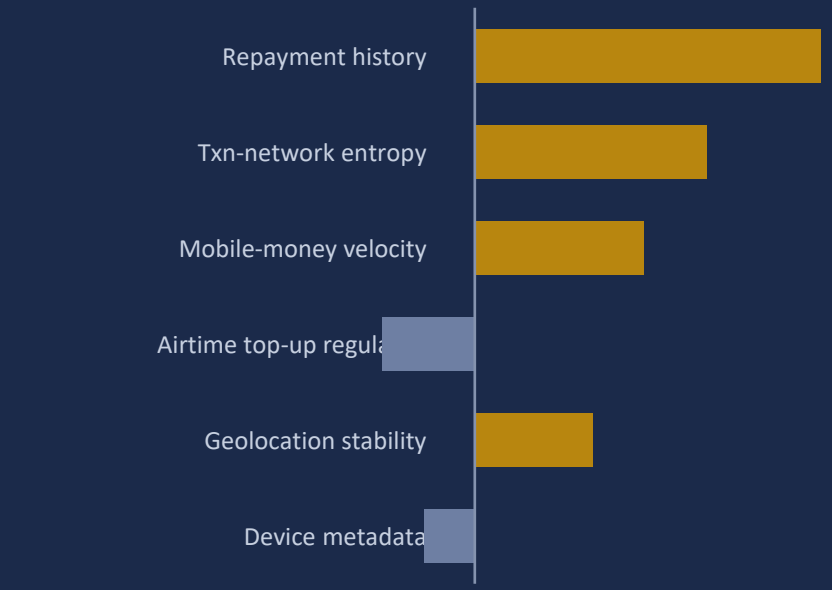
GAME THEORY

Local feature attribution via cooperative game theory

$$\phi_i(x) = \sum_{S \subseteq F \setminus \{i\}} \left[|S|!(|F|-|S|-1)!/|F|! \right] \cdot \left[f(S \cup \{i\}) - f(S) \right]$$

Interpretation

- $\phi_i(x)$ is the marginal contribution of feature i, averaged over all coalitions S
- Satisfies efficiency, symmetry, dummy, and additivity axioms



Non-Stationary Environments & Concept Drift

DISTRIBUTION SHIFT

Covariate Shift

$$P(X) \text{ changes, } P(Y|X) \text{ stable}$$

Applicant population composition shifts; the risk relationship itself does not.

Concept Drift

$$P(Y|X) \text{ changes over time}$$

Economic shocks (inflation, climate events) invalidate the risk relationship itself.

$$\text{PSI} = \Sigma(\%\text{Actual} - \%\text{Expected}) \times \ln(\%\text{Actual} / \%\text{Expected})$$

PSI > 0.20 conventionally triggers mandatory model review

Illustrative drift signal over time



Remediation: continuous-learning pipelines, adaptive windowing, drift-triggered retraining.

Time-to-Default: Survival Analysis

HAZARD MODELS

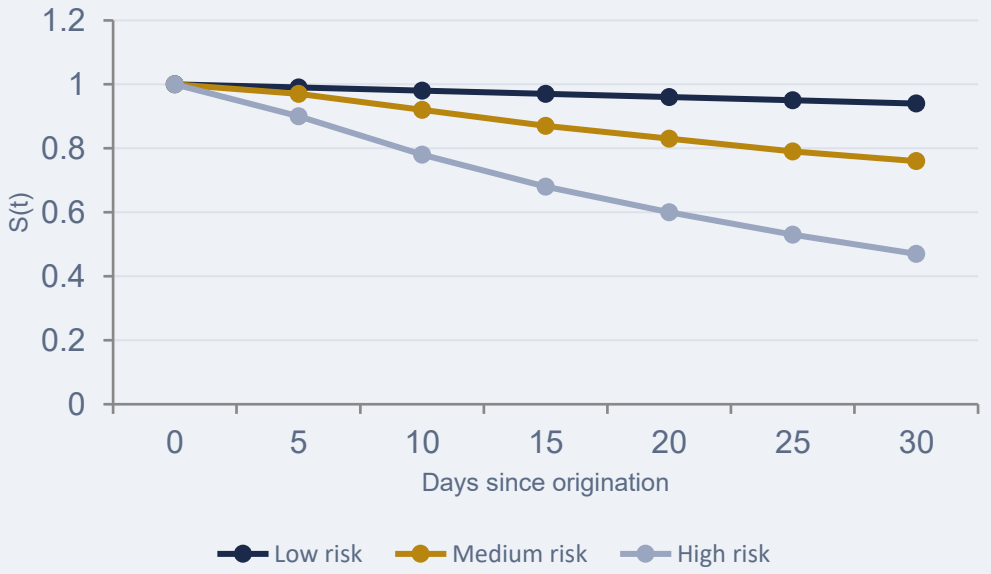
From binary classification to continuous-time hazard rates

$$h(t|X) = h_0(t) \cdot \exp(\beta^T X)$$

Cox Proportional Hazards

- h(t|X) = instantaneous default hazard at tenor t given features X
- Random Survival Forests & neural survival models capture non-linear, time-varying covariates
- Essential for micro-loans spanning 7–30 day tenors with heterogeneous risk trajectories

Illustrative survival curves by risk tier



Enables dynamic timing of optimal collection interventions per risk tier.

From Probability to Policy

DECISION THEORY

A risk score P(Default) is not an economic action — it must be mapped to one.

$$E[\text{Profit} \mid X, L] = L \cdot [(1 - P(D|X,L)) \cdot r - P(D|X,L) \cdot LGD] - C_{acquisition}$$

L	r	LGD	$P(D\|X,L)$
Loan amount	Interest rate / fee	Loss given default	Predicted default probability

Action space

- **Approve / Reject**
- **Credit Limit Assignment**
- **Pricing Tier Selection**

Exploration vs. Exploitation: Contextual Bandits

ONLINE LEARNING

Safely testing credit policy on marginal populations

$$a_t = \text{argmax}_a \left(\theta_a^T x_t + \alpha \sqrt{x_t^T A_a^{-1} x_t} \right)$$

- LinUCB: balances immediate expected revenue against reduction of parameter uncertainty
- Automates controlled credit-line increases and sub-prime experimentation
- Explicit exploration budget bounds portfolio risk exposure

EXPLOIT vs. EXPLORE



Illustrative traffic allocation on a borderline-reject segment

α controls the exploration bonus width — a wider confidence bound trades short-run yield for faster resolution of epistemic uncertainty.

Portfolio-Level Risk & Robust Optimization

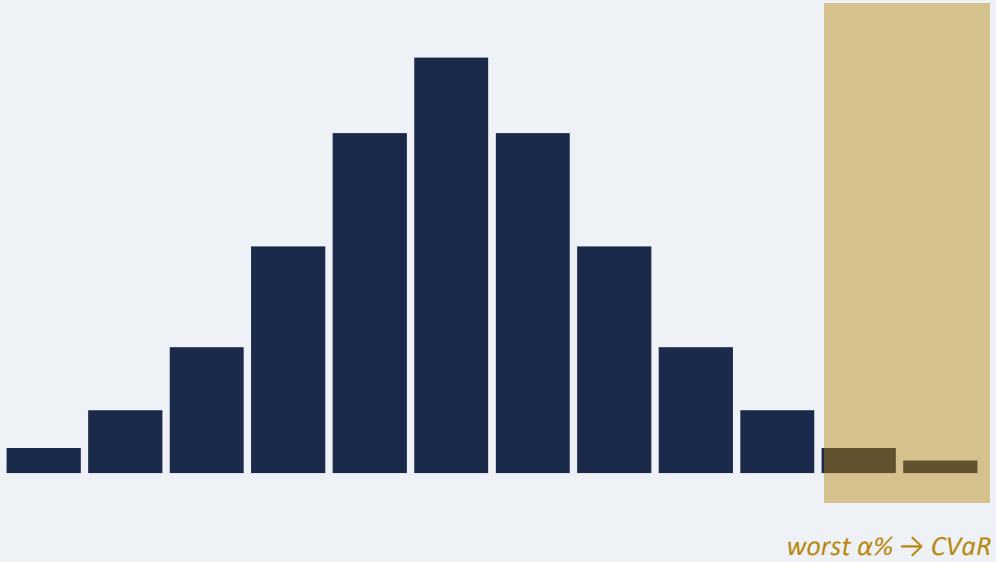
CVaR

Individual scoring ignores correlated portfolio-wide default risk

$$\min_{w \in \Omega} \mathrm{CVaR}_{\alpha}(w) \quad \text{s.t. } E[R(w)] \geq R_{\text{target}}$$

- CVaR_α(w) captures the average loss in the worst α-tail of portfolio outcomes
- Macro-covariate correlation modeled via Gaussian / t-Copulas
- Protects the balance sheet against regional downturns & regulatory shocks

Illustrative loss distribution & CVaR tail



VaR_α marks the tail cutoff; CVaR_α averages losses beyond it — a coherent, sub-additive risk measure unlike VaR alone.

Algorithmic Fairness: Competing Definitions

FAIRNESS CRITERIA

Demographic Parity

$$P(\hat{Y}=1|A=0) = P(\hat{Y}=1|A=1)$$

Equal approval rates across protected groups, regardless of true outcome.

Equalized Odds

$$P(\hat{Y}=1|A=0,Y=y) = P(\hat{Y}=1|A=1,Y=y)$$

Equal true- and false-positive rates across groups, for each outcome y.

These criteria are, in general, mutually incompatible — the choice must be explicit and documented.

Mitigation approaches

Pre-processing

Disparate-impact remover, re-weighing

In-processing

Fairness penalty term added to the loss function

Post-processing

Group-specific threshold adjustment

Capital Allocation: ECL under IFRS 9

PROVISIONING

Integrating ML risk outputs directly into regulatory financial reporting

$$\text{ECL} = \text{PD} \times \text{LGD} \times \text{EAD}$$

PD

Probability of Default

Lifetime PD calibrated via survival models under forward-looking macro scenarios.

LGD

Loss Given Default

Modeled dynamically from recovery costs and collection-channel effectiveness.

EAD

Exposure at Default

Dynamic, utilization-dependent exposure at the moment of default.

Model Selection, Robustness & Adversarial Testing

M

Model Class Selection

- Structural scorecards (logistic / GAM): monotone, regulator-legible
- Ensembles (XGBoost, LightGBM): benchmark predictive performance
- Selection weighs AUC-ROC, Gini, and Expected Calibration Error jointly

Δ

Stress-Testing Protocol

- Structural equation models isolate causal policy effects from shocks
- Monte Carlo simulation of severe liquidity / velocity contractions
- Buffers and limits contract automatically ahead of default clusters

A

Adversarial Robustness

- Causal feature selection favors invariant, non-gameable drivers
- Adversarial training optimizes against worst-case perturbations
- Guards against strategic feature manipulation (Goodhart's Law)

A validated model is not a deployed model — each column above is a required gate before production release.

Case Study: Emerging-Market Deployment

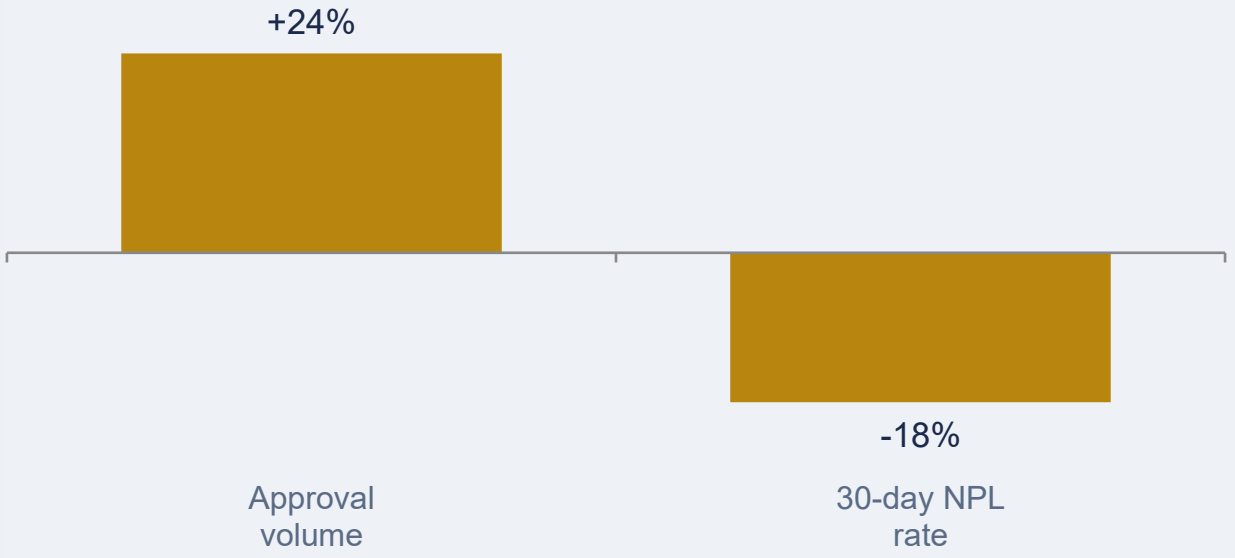
Intervention

- Replaced static decision trees with dynamic XGBoost + SHAP interpretability layer
- Implemented LinUCB contextual bandits on borderline-reject population
- Deployed automated, drift-triggered retraining cycles

Governance outcome

Fully compliant, automated adverse-action notice generation at scale.

Quantifiable outcomes



+24% approval expansion **–18%** 30-day NPL reduction

CONCLUSIONS & FUTURE DIRECTIONS

Rigor Meets Operation

Optimizing digital credit requires bridging mathematical rigor — uncertainty quantification, robust optimization, causal inference — with pragmatic operational systems: MLOps, governance, and compliance.

FRONTIERS OF RESEARCH & PRACTICE

Causal Machine Learning

Shifting from predictive correlation to causal effect estimation for limit-setting and collections.

Privacy-Preserving ML

Federated learning for cross-institution risk modeling without centralized data pooling.

Generative Synthetic Data

Simulating macro-economic stress on thin-file datasets via conditional GANs / diffusion models.

True financial inclusion is not merely giving access to credit — it is pricing and managing risk so sustainably that both institution and borrower thrive.

Thank you — Questions & Discussion

REFERENCES

Primary Sources

1. Heckman, J. J. (1979). Sample Selection Bias as a Specification Error. Econometrica, 47(1), 153–161.

2. Cox, D. R. (1972). Regression Models and Life-Tables. Journal of the Royal Statistical Society, Series B, 34(2), 187–220.

3. Lundberg, S. M., & Lee, S.-I. (2017). A Unified Approach to Interpreting Model Predictions. Advances in Neural Information Processing Systems (NeurIPS) 30.

4. Li, L., Chu, W., Langford, J., & Schapire, R. E. (2010). A Contextual-Bandit Approach to Personalized News Article Recommendation. Proceedings of the 19th International Conference on World Wide Web (WWW).

5. Rockafellar, R. T., & Uryasev, S. (2000). Optimization of Conditional Value-at-Risk. Journal of Risk, 2(3), 21–41.

6. Hardt, M., Price, E., & Srebro, N. (2016). Equality of Opportunity in Supervised Learning. Advances in Neural Information Processing Systems (NeurIPS) 29.

7. Kipf, T. N., & Welling, M. (2017). Semi-Supervised Classification with Graph Convolutional Networks. International Conference on Learning Representations (ICLR).

8. Page, L., Brin, S., Motwani, R., & Winograd, T. (1999). The PageRank Citation Ranking: Bringing Order to the Web. Stanford InfoLab Technical Report.

9. International Accounting Standards Board (IASB). (2014). IFRS 9 Financial Instruments — Expected Credit Loss Impairment Model.

10. Goodhart, C. A. E. (1975). Problems of Monetary Management: The U.K. Experience. Papers in Monetary Economics, Reserve Bank of Australia.