

Ai-Enabled Family-Linked Credit Intelligence For Global Digital Banking Platforms: A Strategy, Analytics, And Agentic Governance Framework

Sriram Polavakalipalayam Ramakrishnan^{1,*}

¹Thayer School of Engineering at Dartmouth and Tuck School of Business at Dartmouth,
Dartmouth College, Hanover, NH, USA, ORCID: 0009-0002-2319-3194
Corresponding Author: Sriram Polavakalipalayam Ramakrishnan^{1,*}

Abstract— This academic research paper develops the Family-Linked Credit Intelligence (FLCI) framework, a conceptual contribution at the intersection of financial technology, household economics, and responsible AI. Motivated by the goal of strengthening family financial bonds through shared, consented, and transparent financial visibility, FLCI examines how federally insured credit unions operating under NCUA supervision can better serve internationally mobile members whose cross-border financial histories are often invisible to domestic bureaus. The framework integrates consented, verified immediate-family signals into an explainable analytics stack designed to align with Regulation P, Regulation B, SR 11-7-informed model risk governance, and the voluntary NIST AI Risk Management Framework. The work is purely conceptual and mathematically grounded; no member-level or proprietary data is used, and the study is conducted as an independent research contribution intended to complement, not substitute for, regulator-led supervisory judgment. We specify cloud-native MLOps with feature stores, champion/challenger governance, drift and fairness monitoring, and graph-based fraud detection. A non-authoritative agentic layer generates audit-ready artifacts through tool-using LLMs under human approval. Under stated assumptions, illustrative scenario analysis suggests potential membership expansion of 5 to 14% and potential incremental annual revenue of roughly \$15 to 38 million, alongside a research-level hypothesis that transparent family-shared financial visibility may foster stronger household financial bonds, subject to pilot validation.

Keywords— agentic AI, AI underwriting, Azure cloud, credit unions, cross-border finance, family linkage, fraud graphs, IFRS 9, metaheuristics, model risk management, reinforcement learning.

Disclaimer. This paper is an independent academic research contribution grounded in U.S. regulatory guidance, international supervisory frameworks, public technical standards, and publicly available industry data, including NCUA publications. It is undertaken in a scholarly spirit of collaboration with the supervisory community, not in opposition to it, and does not represent the systems, strategies, or operations of any specific institution. The framework recognizes and respects the primary authority of U.S. financial regulators, including the NCUA, and is intended to inform research discourse on how technology can responsibly serve family-scale financial well-being. All architectures are proposed reference models intended solely for academic and research purposes.

I. INTRODUCTION

Internationally mobile members, including expatriate professionals, globally distributed households, and cross-border earners, frequently present fragmented credit histories alongside multi-currency cashflows and assets held across multiple jurisdictions. Underwriting decisions that depend solely on a domestic bureau produce inconsistent outcomes, high manual-review volumes, and delayed digital experiences. Cross-border data compounds these difficulties by introducing missingness and asynchrony: income evidence arrives in varying formats, asset documentation differs by jurisdiction, and foreign cashflows may be intermittent. Without explicit data-product contracts and

lineage controls, institutions face operational brittleness and audit friction during examinations .

Fraud has professionalized: synthetic identities, mule networks, and coordinated rings exploit onboarding, account takeovers, and real-time payments. Such behaviour manifests as relational patterns (shared devices, reused addresses, repeated beneficiary changes), making graph analytics central to modern defences .

This paper develops FLCI as an academic research contribution aimed at federally insured credit unions that serve internationally mobile populations. We note with appreciation the U.S. credit-union sector's long history of member-focused service under NCUA supervision, which provides the regulatory foundation that makes responsible family-linked analytics possible. Using Q4 2023 NCUA industry data as a fixed baseline for this conceptual analysis, the U.S. credit union sector reached 139.3 million members and \$2.26 trillion in assets , demonstrating both scale and stability; within this mature sector, thin-file and cross-border member experiences represent an opportunity for further service innovation. A professional reassigned internationally, for example, may carry years of foreign payment history not yet captured in a U.S. bureau. FLCI explores how, with full NCUA-supervised governance, member-scoped and consent-driven family signals may help families navigate these transitions while preserving the regulatory protections that U.S. members depend on.

FLCI does not rely on any universal-family-dataset assumption. Visibility is confined to members' own records and evidence from legitimate workflows; derived indicators replace raw attributes wherever possible, governed by purpose limitation, minimization, and revocation . Cloud-hosted platforms require due diligence and exit planning . Explainability and fairness are governance requirements ; security controls span ISMS, secure SDLC, and ASVS testing .

II. MOTIVATION AND OBJECTIVES

Thin-file and cross-border members can experience friction from limited domestic bureau history, even when they have stable income and meaningful assets abroad, leading to higher manual-review rates and delayed disbursements. Bounded household context reduces that uncertainty when consented, minimized, and explainable ; bias in credit-adjacent domains compounds across protected classes . Cross-border wire and real-time payment volumes create attack surface for coordinated rings; graph anomaly signals support step-up and investigation routing, provided governance prevents spurious soft-link action . Legacy core constraints motivate modular architectures with API gateways, incremental rollout, and tokenized access .

Objectives. (1) Define a member-scoped household-aware risk model feasible under privacy and supervisory expectations. (2) Preserve interpretability through explicit mathematical formulations. (3) Implement consent-driven data procurement with revocation and audit. (4) Specify a proposed cloud reference architecture for ingestion, graph computation, model lifecycle management, and observability. (5) Incorporate RL and metaheuristics as governance-first recommendation methods. (6) Define an agentic layer that generates audit-ready artifacts while keeping humans in control. (7) Quantify the illustrative institutional growth and revenue impact of family linkage for a credit union of comparable scale. (8) Contribute to the academic discourse on how consent-driven, family-scale financial transparency may help strengthen household financial bonds, support shared financial planning, and promote inclusive access to credit for thin-file family members.

Fair-Lending Safeguard. Family-linked variables are used only under explicit member consent and may not be used to infer prohibited-basis characteristics or to disadvantage an applicant based on the status of relatives or associates. Adverse-action reasons remain applicant-specific and are generated solely from the applicant's own credit profile, consistent with ECOA and NCUA fair-lending guidance.

Novel Technical Contributions. (C1) Policy-bound household graph edges, each persisting only while accompanied by link confidence q_{ij} (Eq. [eq:linkconf]), immutable consent provenance, and a jurisdictional-scope tag governing downstream use. (C2) Dual-track relationship scoring combining q_{ij} with a normalization-bounded Household Stability Index (Eq. [eq:hsi]) and a closed-form Family-Linked Risk band (Eq. [eq:flr]). (C3) Composite probability of default (Eq. [eq:pd]) extended by a provably bounded family-linked adjustment (Eq. [eq:pdadj]) whose factor lies in $(0.75, 1]$. (C4) Hard/soft fraud-graph weighting (Eq. [eq:edgeweight]) with a hard-link corroboration rule and a contestability-aware edge lifecycle. (C5) Capacity-aware RL triage with bounded action space, offline training, and human approval gating. (C6) Non-authoritative agentic governance producing schema-locked, RAG-grounded, cryptographically signed artifacts. (C7) Illustrative closed-form growth model (Eq. [eq:linkpool]–[eq:revenue]) linking governed linkage to potential membership and revenue under stated assumptions and NCUA-supervised constraints.

III. METHODS

A. Relationship Verification and Data Boundary

FLCI recognizes immediate-family linkages only when: (a) the relative is also a member and the relationship is captured through account or membership relations, (b) the member provides evidence for a legitimate service request, or (c) the member authorizes ingestion through lawful, transparent mechanisms. The feature layer stores only relationship type, verification status, and timestamps; all underlying artifacts remain access-controlled. Derived stability indicators replace raw paystubs or statements wherever possible.

B. Consent, Tokenization, and Audit

Consent is captured per purpose and logged immutably. Scoped tokens enforce least-privilege access to approved data products; upon revocation, tokens are invalidated and downstream removal workflows fire. A jurisdiction-aware policy engine enforces purpose limitation, regional consent semantics, and cross-border use restriction at both the feature store and at inference time, distinct from data-residency controls. Audit readiness is maintained through consent, access, lineage, and decision logs that support examiner review.

C. Data Products, Feature Store, and Quality Gates

Each FLCI data product carries an explicit contract specifying schema, meaning, freshness, and acceptable missingness. Quality gates enforce schema checks, distribution anomaly detection, and reconciliation. A feature store supports both online serving and offline training with time travel, backfills, and leakage controls that restrict training to information available at decision time.

D. Mathematical Formulation

1) Link Confidence

For member i , relative j , relationship type r , and binary evidence signals $s_1, \dots, s_M \in \{0, 1\}$:

$$q_{ijr} = \sigma(\eta_0 + \sum_{m=1}^M \lambda_m (2s_m - 1)), \quad \sigma(x) = \frac{1}{1+e^{-x}} \quad (1)$$

where η_0 is a prior bias and $\lambda_m > 0$ are evidence weights. When the relationship type is clear from context we write $q_{ij} \equiv q_{ijr}$ for brevity.

2) Household Stability Index (HSI)

Let $\mathbf{S}_i = [S_i^{(1)}, \dots, S_i^{(K)}]^\top$ denote min-max normalized stability bands (income band, tenure band, asset band, etc.) for member i , each scaled to $[0,1]$. The HSI is:

$$HSI_i = \mathbf{w}^\top \mathbf{S}_i = \sum_{k=1}^K w_k S_i^{(k)}, \quad w_k \geq 0, \quad \sum_k w_k = 1 \quad (2)$$

The normalization and weight constraints guarantee $HSI_i \in [0,1]$.

3) Family-Linked Risk Band (FLR)

For the verified immediate-family set $\mathcal{F}_i$ of applicant i :

$$FLR_i = \frac{\sum_{j \in \mathcal{F}_i} q_{ij} \cdot \omega_{r(i,j)} \cdot \mathbf{1}[\tilde{y}_j = 1]}{\varepsilon + \sum_{j \in \mathcal{F}_i} q_{ij} \cdot \omega_{r(i,j)}}, \quad \varepsilon > 0 \quad (3)$$

where $\omega_{r(i,j)}$ is a relationship-type weight, $\tilde{y}_j \in \{0,1\}$ is the good-standing indicator for relative j (1 = satisfactory repayment), and ε prevents division by zero.

4) Probability of Default (PD)

The composite model integrates traditional features $\mathbf{X}_i$, household features $\mathbf{H}_i$, cross-border features $\mathbf{C}_i$, graph features $\mathbf{G}_i$, and the family-linked risk indicator. The final scoring layer is logistic for interpretability:

$$P(D_i = 1) = \sigma(\beta_0 + \beta_x^\top \mathbf{X}_i + \beta_h^\top \mathbf{H}_i + \beta_c^\top \mathbf{C}_i + \beta_g^\top \mathbf{G}_i + \beta_f FLR_i) \quad (4)$$

The expected sign of β_f is negative ($\beta_f < 0$), so that a higher family-linked risk band (stronger family repayment) reduces PD. The sign is validated during champion/challenger calibration.

Calibration is monitored via Brier score and reliability curves; discrimination via ROC/AUC; stability via PSI and population drift. SHAP attributions are stored alongside model version and feature hash.

E. Fraud Graph Analytics

Graph edges include hard links (verified membership relationships, joint accounts, authorized users) and soft links (shared devices, overlapping addresses, repeated beneficiary patterns). Let q_{uv} denote the hard-link confidence (from Eq. [eq:linkconf]) and $s_{uv} \in [0,1]$ the behavioural soft-link score. The composite edge weight is a convex combination controlled by mixing parameter $\alpha \in [0,1]$:

$$w_{uv} = \alpha \cdot q_{uv} + (1 - \alpha) \cdot s_{uv} \quad (5)$$

yielding $w_{uv} \in [0,1]$ since both components are bounded.

Three complementary anomaly dimensions are computed:

$$A_u^{\text{deg}} = \sum_{v \in \mathcal{N}(u)} w_{uv} \quad (6)$$

$$A_u^{\text{vel}} = N_u(t - \Delta t, t) / \Delta t \quad (7)$$

$$r_u = (1 - \rho) b_u + \rho \sum_v P_{uv} r_v \quad (8)$$

where $\mathcal{N}(u)$ is the neighbour set, $N_u(t - \Delta t, t)$ counts new edges in the window, b_u is the intrinsic risk prior, P_{uv} is the row-normalised transition weight, and $\rho \in (0,1)$ is the damping factor. With P row-stochastic and $\rho < 1$, ([eq:propagation]) admits a unique fixed point by Perron–Frobenius.

The composite Graph Risk Score (GRS) is:

$$GRS_u = \theta_1 A_u^{\text{deg}} + \theta_2 A_u^{\text{vel}} + \theta_3 r_u \quad (9)$$

Governance policy prohibits adverse action from soft links alone; corroboration by at least one hard signal is required before escalation, and any decision touching a soft link

generates a contestable artifact entitled to member challenge through a documented appeal channel. GRS-based alerts are triaged into step-up authentication, investigator queue, or watchlist monitoring, as shown in Fig. 4.

F. RL for Capacity-Aware Triage

RL recommends bounded operating points subject to human approval. The MDP defines state s_t as backlog depth, staffing capacity, drift, volatility, and alert mix; action a_t adjusts thresholds within governance-approved bounds $[a^{min}, a^{max}]$; reward R_t penalizes policy violations, latency breaches, and fairness regressions while rewarding timely resolution and false-positive reduction. Training is offline; deployment begins in shadow mode with rollback criteria. An intermediate step uses offline contextual bandits for threshold recommendation under fixed policies. The triage loop is illustrated in Fig. 5.

G. Metaheuristics for Multi-Objective Optimization

Metaheuristics search for Pareto-optimal candidates across competing objectives:

$$\min_{\theta} [\mathcal{L}, \text{FraudLeak}, \Delta_{\text{Fair}}, \text{Latency}, \text{InvCap}] \quad (10)$$

where $\mathcal{L}$ is the credit loss function, FraudLeak is the missed-fraud rate, Δ_{Fair} is the fairness gap, Latency is decisioning time, and InvCap is investigator capacity utilization, all subject to regulatory and operational constraints.

NSGA-II generates the Pareto front; local search refines selected regions. All outputs require governance sign-off.

H. Agentic Governance Layer

A non-authoritative agentic layer generates governance artifacts: consent and provenance summaries, explanation drafts, policy checklists, pilot reports, and investigation packets. Tool-using LLMs operate through schema-locked interfaces with least-privileged scopes; ReAct-style orchestration structures generation, and retrieval-augmented generation restricts narratives to controlled policy text and model cards. Outputs are cryptographically signed and persisted to an immutable decision log. The agentic architecture is depicted in Fig. 3.

I. Prototype Workflow

Algorithm [alg:proto] specifies the prototype skeleton. The agentic layer generates evidence and artifacts; humans retain exclusive authority over credit and fraud decisions.

Algorithm 1: FLCI Agentic Prototype Workflow

Input: Application ID a , Member ID m , Purpose purpose

Output: decision, audit_log

1. $(f, \text{prov}) \leftarrow \text{Ingest}(a)$
2. $\text{scope} \leftarrow \text{VerifyConsent}(a, \text{purpose})$
3. $(\hat{p}_c, \hat{p}_h, \text{cal}) \leftarrow \text{ScoreModels}(f)$
4. $(\text{attr}, \text{stab}) \leftarrow \text{Explain}(\hat{p}_c, f)$
5. $(\text{GRS}, \text{corr}) \leftarrow \text{FraudGraphCheck}(m)$
6. $\text{checklist} \leftarrow \text{RunPolicyChecklist}()$
7. $\text{packet} \leftarrow \text{GenerateAnalystPacket}()$
8. $\text{decision} \leftarrow \text{HumanApprove}(\text{packet})$
9. **return** (decision, audit_log)

J. Azure Reference Architecture

The proposed reference architecture (Fig. 1) follows a zone-separated design using commercially available cloud services. Digital channels connect through API Management. Consent is enforced via a dedicated service with immutable logging. Ingestion lands into a zone-separated data lake. Databricks computes derived features and graph anomalies. Azure ML manages the model registry, deployment slots, and monitoring endpoints. Observability is built on centralized monitoring, incident playbooks, kill-switch logic, and automated rollback. This architecture is a conceptual reference model; it does not depict any institution's production systems.

Fig. 1. FLCI reference architecture. Digital channels connect through API Management. Data flows through zone-separated storage to feature engineering, model serving, and the non-authoritative agentic governance layer. All decisions require human approval.




Figure 1. FLCI reference architecture. Digital channels connect through API Management. Data flows through zone-separated storage to feature engineering, model serving, and the non-authoritative agentic governance layer. All decisions require human approval.

Fig. 2. Governance gate flowchart. Each decision point enforces a control boundary; failures trigger escalation or fallback rather than pass-through.




Figure 2. Governance gate flowchart. Each decision point enforces a control boundary; failures trigger escalation or fallback rather than pass-through.

Fig. 3. Agentic governance prototype. The LLM agent invokes schema-locked tools to generate artifacts (analyst packets, governance checklists, signed audit entries) but holds no decision authority. All outputs require human approval.




Figure 3. Agentic governance prototype. The LLM agent invokes schema-locked tools to generate artifacts (analyst packets, governance checklists, signed audit entries) but holds no decision authority. All outputs require human approval.

Fig. 4. Graph-based fraud detection pipeline. Transaction events are decomposed into hard and soft link graphs, scored across degree, velocity, and propagation dimensions, and triaged by composite GRS.




Figure 4. Graph-based fraud detection pipeline. Transaction events are decomposed into hard and soft link graphs, scored across degree, velocity, and propagation dimensions, and triaged by composite GRS.

Fig. 5. RL-based capacity-aware triage loop. The offline-trained agent observes operational state and recommends bounded threshold adjustments. Human approval is required before any production change.




Figure 5. RL-based capacity-aware triage loop. The offline-trained agent observes operational state and recommends bounded threshold adjustments. Human approval is required before any production change.

IV. EXPECTED OUTCOMES AND PILOT PLAN

This paper is conceptual and reports no empirical results. Expected pilot hypotheses include reduced manual review through bounded stability indicators, improved fraud resilience through graph anomaly signals, stronger supervisory defensibility through consent logs, and improved operational stability through monitoring and rollback.

TABLE I: PILOT REPORTING METRICS (ILLUSTRATIVE)

Category	Metric	Baseline	Target
Calibration	Brier Score	0.18	≤ 0.15
Discrimination	AUC-ROC	0.72	≥ 0.78
Stability	PSI	N/A (new)	< 0.10
Fairness	Δ_{DP}	0.08	≤ 0.05
Fairness	Δ_{EO}	0.06	≤ 0.04
Operational	p95 Latency	2.1s	$\leq 1.5s$
Operational	Exception Rate	12%	$\leq 8\%$
Fraud	Alert FP Rate*	35%	$\leq 25\%$

Fraction of alerts subsequently cleared as non-fraud.

TABLE II: GOVERNANCE MAPPING TO SUPERVISORY THEMES

Theme	FLCI Control	Reference
Privacy	Consent, minimization, tokenization	Reg. P [1], Safeguards [10]
Model Risk	Champion/challenger, calibration, drift	SR 11-7 [4]
Fairness	Δ_{DP} , Δ_{EO} monitoring	ECOA [3], AI RMF [21]
Third-Party	Due diligence, exit planning	NCUA [11], Basel [12]
Resilience	Kill-switch, rollback, incident response	DORA [13], CSF [20]
Security	ISMS, ASVS, pen testing	ISO 27001 [18], OWASP [19]

TABLE III: AGENTIC GOVERNANCE RISKS AND CONTROLS

Risk	Impact	Control
Prompt injection	Unauthorized tool calls	Schema-locked tools, input validation
Data leakage	PII exposure	RAG grounding, output filtering
Hallucination	Incorrect claims	Retrieval-only narratives, human review
Tool misuse	Excessive access	Least-privilege tokens, audit trail
Version drift	Inconsistent behavior	Prompt version control, regression tests

TABLE IV: STAGED ROLLOUT PLAN

Stage	Scope	Duration	Gate
1	Shadow mode	8 to 12 wk	Calibration $\leq$ baseline
2	5% pilot	8 to 12 wk	AUC $\geq$ target; PSI < 0.10
3	25% expanded	12 to 16 wk	All KPIs within bounds
4	Full production	Ongoing	Continuous monitoring

A pilot should publish its acceptance criteria ex ante, specifying the maximum permitted fairness-gap delta, maximum latency increase, minimum calibration improvement, and explicit guardrails for soft-link usage. Agentic artifacts are assessed by governance throughput metrics and security metrics. RL and metaheuristic recommendations

undergo offline backtesting and limited-scope shadow mode before any policy change . Stress testing covers currency volatility spikes, employment interruption proxies, and coordinated fraud bursts; the objective is to verify that the system degrades to baseline decisioning rather than producing unexplainable outputs .

V. INSTITUTIONAL BENEFITS AND ECONOMIC GROWTH IMPACT

Credit unions differ from commercial banks structurally: members are simultaneously depositors, borrowers, and cooperative owners, so growth compounds across all three dimensions. NCUA industry data confirms the potential: the US credit union sector added approximately 4.0 million new members in 2023 , insured shares and deposits grew 2.2% year-over-year through Q4 2023 , and the aggregate net worth ratio stood at 10.70%, above the 7% well-capitalized threshold . FLCI's family-linkage mechanism converts latent household relationships into active memberships.

K. Family-Linked Growth Model

Let N_0 denote the current membership base. We define the eligible linkage pool as

$$L = N_0 \cdot \varphi \cdot \kappa \quad (11)$$

where $\varphi \in (0,1)$ is the fraction of members whose immediate family includes at least one eligible non-member, and κ is the average number of such eligible relatives per linked member.

First-order enrollment from direct family outreach yields

$$\Delta N_1 = L \cdot \gamma \quad (12)$$

with conversion rate γ .

For a mid-sized credit union with $N_0 = 250,000$ members (representative of the NCUA's upper-quartile asset tier), illustrative estimates suggest $\varphi \approx 0.30$ and $\kappa \approx 1.5$. A conservative $\gamma = 0.10$ yields $\Delta N_1 \approx 11,250$ new members over a three-year rollout horizon.

L. Word-of-Mouth Amplification

Newly enrolled family members constitute a secondary referral channel. Each satisfied member generates r additional referrals per period, with a decay factor $\delta < 1$ accounting for diminishing reach. The cumulative word-of-mouth gain over T periods, with $r\delta < 1$ ensuring geometric-series convergence, is

$$\Delta N_{\text{wom}} = \Delta N_1 \cdot \frac{r\delta(1-(r\delta)^T)}{1-r\delta} \quad (13)$$

With $r = 0.3$ and $\delta = 0.7$ over $T = 3$ years, the secondary wave adds roughly 26% atop first-order conversions, giving $\Delta N_{\text{total}} = \Delta N_1 + \Delta N_{\text{wom}}$.

M. Revenue and Capital Impact

The incremental revenue attributable to family-linked growth is

$$\Delta R = \Delta N_{\text{total}} \cdot (\bar{d} \cdot \eta + \bar{f}) \quad (14)$$

where $\bar{d}$ is the average deposit per member, η is the net interest margin, and $\bar{f}$ is the average annual non-interest income per member.

Drawing on NCUA industry aggregate data for upper-quartile credit unions : $\bar{d} \approx \$30,000$ (average member shares at upper-quartile asset-tier institutions), $\eta = 2.5\%$ (a conservative scenario assumption), and $\bar{f} \approx \$200$ (non-interest income per member). Table 5 projects the illustrative impact across three scenarios.

TABLE V: ILLUSTRATIVE PROJECTED INSTITUTIONAL IMPACT OF FLCI (THREE-YEAR HORIZON)

Metric	Conserv.	Moderate	Optimistic
Conversion rate γ	0.10	0.15	0.25
Direct enrollments	~11,000	~17,000	~28,000
Referral additions	~3,000	~4,000	~7,000
Total member gain	~14,000	~21,000	~35,000
Membership growth (%)	~5	~8	~14
New deposits (\$M)	~420	~630	~1,050
Incremental NII (\$M)	~11	~16	~26
Fee uplift (\$M)	~3	~4	~7
Total revenue gain (\$M)	~14	~20	~33

On the capital side, additional deposits increase total assets. For illustration, using the Q4 2023 sector aggregate net worth ratio of 10.70% as a baseline for a well-capitalized credit union with total assets $A_0 \approx \$7.5\text{B}$, the net worth ratio after absorbing ΔA in new deposits is $\text{NWR}_0 \cdot A_0 / (A_0 + \Delta A)$. Under the conservative-to-moderate scenarios in Table 5 ($\Delta A \approx \$420\text{M}$ to $\$630\text{M}$), this yields a reduction of approximately 55 to 80 basis points, well above the 7% well-capitalized threshold and manageable through retained earnings and routine asset-liability adjustments.

N. Credit Enhancement for Family Members

Beyond institutional growth, FLCI may deliver benefit to the family members themselves. A thin-file dependent, such as the spouse of an international professional reassigned from Geneva to Nairobi, faces punitive rates or outright denial when applying for consumer credit in a new jurisdiction. Under FLCI, the family-linked PD adjustment is

$$\text{PD}_j^{\text{adj}} = \text{PD}_j^{\text{base}} \cdot (1 - \lambda \cdot \text{HSI}_i \cdot q_{ij}) \quad (15)$$

where λ is a bounded credit enhancement coefficient ($0 < \lambda \leq 0.25$), calibrated under SR 11-7 governance, $\text{HSI}_i \in [0,1]$ is the household stability index (Eq. [eq:hsil]), and $q_{ij} \in (0,1)$ is the link confidence (Eq. [eq:linkconf]). Since $\lambda \cdot \text{HSI}_i \cdot q_{ij} < 0.25$, the adjustment factor remains in $(0.75,1]$, guaranteeing $\text{PD}_j^{\text{adj}} \in [0,1]$.

Under risk-based pricing, the rate premium is approximately proportional to PD. A thin-file applicant facing a 600 to 800 bps risk premium, with mid-to-strong family parameters ($\lambda \in [0.10,0.25]$, $\text{HSI}_i \in [0.5,0.9]$, $q_{ij} \in [0.5,0.9]$), may produce a PD adjustment factor in the illustrative range of 3 to 20%, which may translate into an indicative 50 to 150 bps pricing improvement depending on pilot calibration, product type, and policy constraints. The downstream effect is a virtuous cycle: family members who receive fair pricing become long-tenure depositors and advocates, reinforcing both Eq. [eq:linkpool] and Eq. [eq:wom].

O. Strengthening Family Financial Bonds: A Research Hypothesis

A central motivation of this research is the hypothesis that transparent, consent-driven, family-shared financial visibility may support stronger household financial bonds. Prior household-economics scholarship has documented associations between financial communication and household outcomes; for instance, Dew, Britt, and Huston found that financial disagreements were a relatively strong predictor of household strain in their sample of 4,574 couples. Motivated by this prior literature—while making no causal claims—FLCI posits that shared financial dashboards, co-signed savings goals, and joint

repayment cadences may help family members coordinate toward shared financial objectives. Formally, let D_{base} denote a baseline household-coordination-friction indicator over a fixed horizon; a conceptual FLCI-influenced coordination indicator can be written as

$$D_{\text{rel}} = D_{\text{base}} \cdot (1 - \mu \cdot \text{HSI}_i \cdot q_{ij} \cdot T_{\text{share}}) \quad (16)$$

where $\mu \in (0, 0.5]$ is a bounded coordination-benefit coefficient, $T_{\text{share}} \in [0, 1]$ is the joint-visibility indicator, and HSI_i, q_{ij} are reused from Eq. [eq:hsi] and Eq. [eq:linkconf]. Each multiplicand lies in $[0, 1]$, so $D_{\text{rel}} \in [0, D_{\text{base}}]$.

This formulation is strictly a research hypothesis to motivate future empirical inquiry, not a validated causal claim; any empirical investigation of behavioral effects would require separate IRB-governed research protocols. The enduring academic contribution here is the framing itself: that credit-union technology, designed with consent and transparency at its core, may serve as an instrument for strengthening family financial bonds alongside its traditional risk and inclusion objectives.

P. Cybersecurity, Privacy, and NCUA Supervisory Ecosystem

FLCI is designed to map onto key NCUA supervisory themes relevant to privacy, cybersecurity, credit risk, BSA/AML, and third-party oversight. FICUs enumerated in the NCUA Credit Union Locator, insured to \$250,000 by the NCUSIF, reporting via CUOnline Call Reports and Financial Performance Reports, benefit from consumer-facing transparency through MyCreditUnion.gov and supervisory enforcement via NCUA Administrative Orders. The threat model aligns with the NCUA Cybersecurity Resources program (ACET, ISE, incident reporting), NIST CSF, ISO/IEC 27001, and OWASP ASVS, addressing prompt injection, model exfiltration, soft-link spoofing, token replay, and federated drift via schema-locked tools, signed audit logs, tokenized API gateways, and kill-switch rollback. Privacy beyond Regulation P and the FTC Safeguards Rule enforces purpose limitation, derivation-over-rawness, scoped tokenization, and revocation cascades, consistent with DORA and EBA GL 2025/02. NCUA's 2026 Supervisory Priorities (credit risk, cybersecurity, BSA/AML, information security) map directly onto Eq. [eq:pd], the controls above, the graph anomaly pipeline (Sec. 3.5), and the Azure architecture (Sec. 3.10). FLCI fairness monitoring ($\Delta_{\text{DP}}, \Delta_{\text{EO}}, \text{SHAP review}$) is designed to remain consistent with ECOA and the voluntary NIST AI RMF, and is intended to align with the NCUA's ongoing supervisory stewardship of U.S. credit unions, which practitioners are encouraged to follow as supervisory practice continues to evolve.

VI. DISCUSSION AND CONCLUSIONS

FLCI's academic defensibility rests on bounded scope: no universal family dataset exists; only member-scoped links from legitimate workflows and explicit consent are used, with derived indicators replacing raw attributes. Graph analytics may support fraud resilience while governance prevents spurious soft-link action via corroboration, step-up authentication, and contestable artifacts. RL and metaheuristics are proposed only under bounded action spaces, offline validation, and logged human approval; the agentic layer remains non-authoritative, tool-locked, and fully logged. Growth in Section 5 depends on outreach, onboarding ease, and cross-jurisdictional eligibility; the sector share-to-loan ratio near 1.17 (implied by the 85.2% loan-to-share ratio reported in) suggests asset-liability buffer for absorbing new deposits. **Limitations.** The framework is conceptual and strategy-focused; any pilot implementation must validate predictive value, fairness, and agentic throughput through staged shadow-mode rollout with examiner-facing documentation under NCUA supervision. The behavioral hypothesis on strengthening family financial bonds is offered as a research framing only and would require separate IRB-governed studies. **Conclusion.** This academic research paper contributes a

reference-level design for member-scoped, consented family signals; explicit interpretable formulations; graph fraud analytics; governed RL and metaheuristic optimization; a non-authoritative agentic prototype; a research hypothesis on strengthening family financial bonds; and design alignment with major NCUA supervisory priorities relevant to privacy, cybersecurity, credit risk, and governance . Under stated assumptions, illustrative scenario analysis suggests potential membership expansion of 5 to 14% and potential incremental annual revenue of roughly \$15 to 38 million, with possible credit-access improvement for thin-file family members pending pilot validation. The authors wish to underscore that this framework is intended as a collaborative research complement to—never a substitute for—the supervisory judgment of the NCUA and other U.S. financial regulators. Future work includes controlled pilots, fairness validation, adversarial testing, IRB-governed behavioral studies, and examiner-informed governance verification .

VII. REFERENCES

- [1] National Credit Union Administration, "Privacy of consumer financial information (regulation p)," Federal Consumer Financial Protection Guide, 2023, <https://ncua.gov/regulation-supervision/manuals-guides/federal-consumer-financial-protection-guide>.
- [2] IFRS Foundation, "IFRS 9 financial instruments," International Financial Reporting Standards, 2014, <https://www.ifrs.org/issued-standards/list-of-standards/ifrs-9-financial-instruments/>.
- [3] CFPB, "Regulation B (equal credit opportunity act): 12 CFR part 1002," Consumer Financial Protection Bureau, 2023, <https://www.consumerfinance.gov/rules-policy/regulations/1002/>.
- [4] Federal Reserve, "SR 11-7: Guidance on model risk management," Supervisory and Regulation Letters, 2011, <https://www.federalreserve.gov/supervisionreg/srletters/sr1107.htm>.
- [5] Microsoft, "Azure well-architected framework," Microsoft Learn, 2024, <https://learn.microsoft.com/azure/well-architected/>.
- [6] Databricks, "Delta lake documentation," Delta Lake Project, 2024, <https://docs.delta.io/>.
- [7] L. Akoglu, H. Tong, and D. Koutra, "Graph based anomaly detection and description: A survey," Data Mining and Knowledge Discovery, vol. 29, no. 3, pp. 626–688, 2015.
- [8] Z. Chen et al., "Parallel path detection for fraudulent accounts in banks based on graph analysis," PeerJ Computer Science, vol. 9, p. e1268, 2023.
- [9] National Credit Union Administration, "Quarterly credit union data summary, q4 2023," NCUA, 2024, <https://ncua.gov/files/publications/analysis/quarterly-data-summary-2023-Q4.pdf>.
- [10] FTC, "Standards for safeguarding customer information (safeguards rule)," Federal Trade Commission, 2023, <https://www.ftc.gov/business-guidance/resources/ftc-safeguards-rule-what-your-business-needs-know>.
- [11] National Credit Union Administration, "Evaluating third party relationships (07-cu-13)," Letter to Credit Unions, 2007, <https://ncua.gov/regulation-supervision/letters-credit-unions-other-guidance/evaluating-third-party-relationships>.
- [12] Basel Committee on Banking Supervision, "Principles for the sound management of third-party risk," BIS, Dec. 2025, <https://www.bis.org/bcbs/publ/d605.htm>.
- [13] European Commission, "Digital operational resilience act (DORA): Regulation (EU) 2022/2554," European Commission Finance: Digital Operational Resilience Regulation, 2022, <https://finance.ec.europa.eu/regulation-and-supervision/financial-services->

legislation/implementing-and-delegated-acts/digital-operational-resilience-regulation_en.

- [14] European Banking Authority, "Guidelines amending guidelines EBA/GL/2019/04 on ICT and security risk management (EBA/GL/2025/02)," EBA Guidelines, applicable from 20 May 2025, 2025, <https://www.eba.europa.eu/regulation-and-policy/internal-governance/guidelines-on-ict-and-security-risk-management>.
- [15] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2017, pp. 4765–4774.
- [16] M. Hardt, E. Price, and N. Srebro, "Equality of opportunity in supervised learning," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2016, pp. 3315–3323.
- [17] S. Barocas, M. Hardt, and A. Narayanan, *Fairness and Machine Learning: Limitations and Opportunities*. MIT Press, 2023, <https://fairmlbook.org/>.
- [18] ISO/IEC, "ISO/IEC 27001: Information security management systems," International Organization for Standardization, 2022, <https://www.iso.org/isoiec-27001-information-security.html>.
- [19] OWASP, "Application security verification standard (ASVS)," OWASP Foundation, 2023, <https://owasp.org/www-project-application-security-verification-standard/>.
- [20] National Institute of Standards and Technology, "Cybersecurity framework (CSF) 2.0," National Institute of Standards and Technology, Tech. Rep., 2024, <https://www.nist.gov/cyberframework>.
- [21] NIST, "AI risk management framework (AI RMF 1.0)," National Institute of Standards and Technology, Tech. Rep., 2023, <https://www.nist.gov/itl/ai-risk-management-framework>.
- [22] Microsoft Corporation, "Azure API management documentation," Microsoft Learn, 2024, <https://learn.microsoft.com/azure/api-management/>.
- [23] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*, 2nd ed. MIT Press, 2018.
- [24] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan, "A fast and elitist multiobjective genetic algorithm: NSGA-II," *IEEE Transactions on Evolutionary Computation*, vol. 6, no. 2, pp. 182–197, 2002.
- [25] T. Schick, J. Dwivedi-Yu, R. Dessì, R. Raileanu, M. Lomeli, L. Zettlemoyer, N. Cancedda, and T. Scialom, "Toolformer: Language models can teach themselves to use tools," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2023, <https://arxiv.org/abs/2302.04761>.
- [26] Microsoft Databricks, "Azure databricks documentation," Microsoft Learn, 2024, <https://learn.microsoft.com/azure/databricks/>.
- [27] S. Yao et al., "ReAct: Synergizing reasoning and acting in language models," *arXiv preprint arXiv:2210.03629*, 2022.
- [28] Microsoft Azure, "Azure machine learning documentation," Microsoft Learn, 2024, <https://learn.microsoft.com/azure/machine-learning/>.
- [29] National Credit Union Administration, "Credit union system performance data and industry trends," NCUA Analysis Center, 2025, <https://www.ncua.gov/analysis/credit-union-corporate-call-report-data>.
- [30] J. Dew, S. Britt, and S. Huston, "Examining the relationship between financial issues and divorce," *Family Relations*, vol. 61, no. 4, pp. 615–628, 2012, <https://scholarsarchive.byu.edu/facpub/4526/>.
- [31] National Credit Union Administration, "MyCreditUnion.gov: Consumer financial protection resource," NCUA Consumer Education, 2025, <https://mycreditunion.gov/>.
- [32] National Credit Union Administration, "NCUA credit union locator: Federally insured credit unions," NCUA Mapping Service, 2025, <https://mapping.ncua.gov/ResearchCreditUnion>.

- [33] National Credit Union Administration, "National credit union share insurance fund (NCUSIF)," NCUA Support Services, 2025, <https://ncua.gov/support-services/share-insurance-fund>.
- [34] National Credit Union Administration, "CUonline: Profile and call report submission service," NCUA Regulatory Reporting, 2025, <https://ncua.gov/regulation-supervision/regulatory-reporting/cuonline>.
- [35] National Credit Union Administration, "Financial performance reports (FPR) for federally insured credit unions," NCUA Analysis, 2025, <https://ncua.gov/analysis/credit-union-corporate-call-report-data/financial-performance-reports>.
- [36] National Credit Union Administration, "NCUA administrative orders and enforcement actions," NCUA Enforcement, 2025, <https://ncua.gov/regulation-supervision/enforcement-actions/administrative-orders>.
- [37] National Credit Union Administration, "NCUA cybersecurity resources program," NCUA, 2025, <https://ncua.gov/cybersecurity>.
- [38] National Credit Union Administration, "NCUA's 2026 supervisory priorities," Letter to Credit Unions 26-CU-01, January 2026, 2026, <https://ncua.gov/regulation-supervision/letters-credit-unions-other-guidance/ncuas-2026-supervisory-priorities>.
- [39] National Credit Union Administration, "Removal of disparate impact from supervisory practice," Letter to Credit Unions, September 2025, 2025, <https://ncua.gov/regulation-supervision/letters-credit-unions-other-guidance/removal-disparate-impact>.
- [40] National Credit Union Administration, "Elimination of reputation risk from supervisory examinations," Letter to Credit Unions, September 2025, 2025, <https://ncua.gov/regulation-supervision/letters-credit-unions-other-guidance/elimination-reputation-risk>.