Information Systems and Economics (2026)
Clausius Scientific Press, Canada

DOI: 10.23977/infse.2026.070101
ISSN 2523-6407 Vol. 7 Num. 1

Liquidity Constraints and Household Financial Vulnerability in China: An Explainable, Uncertainty-Aware Machine Learning Risk Management Framework Using the China Household Finance Survey

Xiaoyu Wang

School of Economics, Beijing Technology and Business University, Beijing, China
yuda_1211@qq.com

Keywords: Household finance; Financial risk management; Liquidity constraint; Financial Vulnerability Index; Machine learning; Explainable artificial intelligence; Double Machine Learning; Conformal prediction; China Household Finance Survey

Abstract: Household financial vulnerability reflects the likelihood that a household will fall into financial distress when facing adverse shocks, and it is a key micro-foundation of financial system stability. Using microdata from the China Household Finance Survey (CHFS), this study develops a machine-learning-based risk management framework to examine how liquidity constraints affect household financial vulnerability and to identify high-risk groups under heterogeneous socioeconomic conditions. Household financial vulnerability is operationalized as a binary outcome, denoted as the Financial Vulnerability Index (FVI), indicating whether liquid buffers are sufficient to cover unexpected expenditures. Liquidity constraints are measured through credit accessibility indicators, forming a binary Liquidity Constraint (LC) variable. The framework integrates (i) high-performance tabular prediction models, including gradient-boosted decision trees and neural tabular networks, to construct calibrated probability-of-vulnerability scores; (ii) explainability techniques, with Shapley Additive Explanations (SHAP) used to quantify global and local risk drivers; and (iii) causal machine learning methods, such as Double Machine Learning (DML) and generalized random forests, to estimate the heterogeneous causal effect of liquidity constraints on financial vulnerability across income groups, city tiers, and regions. To enhance model reliability for risk governance, probability calibration and distribution-free uncertainty quantification are implemented via conformal prediction. Empirical results indicate that liquidity constraints significantly increase the predicted and causally estimated risk of household financial vulnerability, with stronger effects concentrated among middle-to-lower income households and households located in lower-tier cities and economically stressed regions. The proposed framework provides an algorithmic basis for targeted inclusive-finance interventions and household risk mitigation policies.

1. Introduction

Household financial vulnerability describes the heightened likelihood that a household will enter financial distress when exposed to adverse income or expenditure shocks, and it has become increasingly salient for both inclusive-finance design and financial stability monitoring. Recent evidence shows that vulnerability can be measured through the adequacy of household buffers under shock scenarios and that its incidence is highly heterogeneous across population subgroups and economic environments [1]. Related research further demonstrates that vulnerability is tightly linked to consumer-debt conditions and repayment burdens, implying that balance-sheet structures and short-run liquidity pressures can materially change which households are classified as financially vulnerable [2]. In addition, modern household-finance models emphasize that liquidity constraints are pervasive and economically meaningful, shaping households' ability to smooth consumption and respond to shocks, thereby providing a mechanism through which liquidity frictions can amplify vulnerability risk [3].

Despite these advances, much of the empirical household-finance literature still relies on parametric specifications that may under-represent nonlinearities and complex interactions among liquidity constraints, assets, income, and regional characteristics—patterns that are central to risk management applications. Meanwhile, credit and risk governance increasingly adopt machine learning to improve predictive accuracy, while requiring transparency and auditability of risk drivers; explainable machine learning has been proposed as a practical route to reconcile high-performing models with interpretable attribution in credit risk settings [4]. Beyond interpretability, reliable deployment of risk scores requires principled uncertainty handling and explicit risk control; recent conference work formalizes distribution-free risk guarantees via conformal methods that directly control expected loss-based risk measures rather than only point predictions [5]. Motivated by these developments, the study reframes the liquidity-constraint–financial-vulnerability relationship as an algorithmic household risk management problem: producing calibrated vulnerability probabilities, explaining key drivers, and supporting robust policy targeting under heterogeneous household and regional conditions.

2. Related Work

Recent household-finance studies document that limited short-term liquidity and access to emergency funds are central to household financial vulnerability, but the channels and measurement strategies vary across contexts. Evidence from consumption responses to predictable cash inflows suggests that many households behave as if they face binding liquidity constraints, with heterogeneous smoothing patterns even among borrowers holding revolving debt [6]. Policy-oriented work further shows that relaxing liquidity constraints through partial access to future public benefits can materially improve households' ability to sustain consumption under unemployment shocks, highlighting the role of institutional liquidity backstops in mitigating fragility [7]. During the COVID-19 period, survey-based research reports substantial prevalence of financially fragile households that struggle to cover mid-size emergency expenses, reinforcing the need to model vulnerability as a function of liquid buffers and shock exposure [8]. In parallel, the expansion of new consumer-credit products has raised fresh concerns for risk management: empirical evidence links household financial fragility to higher adoption and intensity of Buy Now, Pay Later (BNPL) usage, implying that fintech credit can both signal and amplify underlying vulnerability in consumer balance sheets [9]. Complementary work also shows that liquidity constraints and debt dynamics are intertwined for middle-class saving behavior, motivating models that jointly track liquid-asset positions, debt capacity, and shock absorption [10].

On the financial-risk-management side, the literature increasingly adopts machine learning to improve prediction while addressing transparency, reliability, and deployment constraints. Banking applications emphasize Explainable Artificial Intelligence (XAI) to meet governance and regulatory expectations, often combining gradient-boosting models with SHapley Additive exPlanations (SHAP) to produce feature-attribution explanations that support credit decisioning [11]. Broader XAI frameworks for credit evaluation integrate model-agnostic explanation tools such as Local Interpretable Model-Agnostic Explanations (LIME) and SHAP to improve the interpretability of automated lending pipelines [12]. Methodologically, a key trend is the extension of deep learning to tabular, high-dimensional household-finance microdata; large-scale benchmarking in a major machine learning conference shows that Transformer-style architectures can be competitive for tabular prediction under standardized protocols, but performance depends strongly on training and tuning choices [13]. Because risk management relies on well-calibrated probabilities (not just classification accuracy), recent statistical work systematically compares probability-calibration methods for machine learning predictors [14], and operations/analytics research develops decision-aware calibration objectives that align probabilistic outputs with downstream cost considerations [15]. Together, these strands motivate hybrid pipelines that combine causal-inference-aware learning, high-performance tabular prediction, explainability, and probability calibration for household financial vulnerability assessment.

3. Methodology

This study formulates household financial vulnerability assessment as an algorithmic financial risk management task with three tightly coupled modules: (i) vulnerability labeling and feature construction, (ii) predictive risk scoring with explainability, and (iii) causal machine learning and uncertainty-aware risk control. The Financial Vulnerability Index (FVI) is predicted as a calibrated probability, then decomposed into interpretable drivers using Shapley Additive Explanations (SHAP), and finally linked to liquidity constraints through heterogeneous causal effect estimation using Double Machine Learning (DML) and generalized random forests (GRF) . Gradient-boosted decision trees (GBDT) implemented via LightGBM serve as a core high-performance baseline for tabular microdata modeling .

3.1. Problem Formulation and Variable Construction

Outcome (Financial Vulnerability Index). Let $Y_i \in \{0,1\}$denote the binary Financial Vulnerability Index (FVI) for household i. A household is labeled vulnerable if liquid buffers cannot cover an unexpected expenditure proxy. Define liquid resources:

$$LR_i = \underbrace{(Inc_i - Cons_i)}_{\text{net cash-flow buffer}} + \underbrace{LA_i}_{\text{liquid assets}} , \quad LA_i = Cash_i + DemandDep_i , \tag{1}$$

and unexpected expenditure UE_i(e.g., medical spending proxy). The label is constructed as:

$$Y_i = \mathrm{I}\left(LR_i < UE_i\right). \tag{2}$$

This construction operationalizes vulnerability as “insufficient liquidity buffer under shock,” consistent with financial risk screening.

Treatment (Liquidity Constraint). Let $D_i \in \{0,1\}$ denote liquidity constraint status. For risk-management consistency, $D_i$is defined using credit access indicators (e.g., rejected loan applications, discouraged borrowing due to expected rejection, insufficient collateral/guarantees), yielding:

$$D_i = \mathrm{I}(\text{credit access friction present}). \tag{3}$$

The feature vector $X_i$includes household demographics, balance-sheet variables (e.g., logincome, log liquid/financial assets), employment and education, household size, risk attitudes, and regional/city-tier indicators.

3.2. Predictive Risk Scoring with Gradient Boosting

The predictive module estimates the vulnerability probability:

$$\hat{p}_i = \Pr(Y_i = 1 | X_i). \tag{4}$$

A GBDT model represents $\hat{p}_i$ as an additive ensemble of T regression trees:

$$f_T(X_i) = \sum_{t=1}^{T} \eta h_t(X_i), \hat{p}_i = \sigma(f_T(X_i)), \tag{5}$$

where η is the learning rate, $h_t(\cdot)$ is a decision tree, and $\sigma(z) = (1 + e^{-z})^{-1}$ is the logistic link. Training minimizes a regularized empirical risk:

$$\min_{h_1,\dots,h_T} \sum_{i=1}^{n} \ell^2(Y_i, \sigma(f_T(X_i))) + \Omega(h_1,\dots,h_T), \tag{6}$$

with $\ell(\cdot)$ as log-loss and $\Omega(\cdot)$ penalizing model complexity. LightGBM is adopted for scalable GBDT training on high-dimensional tabular data and large sample sizes.

Calibration for risk management. Since decision-making relies on well-calibrated probabilities, the raw model score $\hat{p}_i$ is recalibrated via a monotone mapping $g(\cdot)$ learned on a validation set:

$$\tilde{p}_i = g(\hat{p}_i), \tag{7}$$

where g can be fit using isotonic regression or logistic calibration. The calibrated $\tilde{p}_i$ is treated as the operational vulnerability probability.

3.3. Explainability via Shapley Additive Explanations

To translate model outputs into interpretable risk drivers, SHAP expresses each prediction as an additive attribution model:

$$f(X_i) = \phi_0 + \sum_{j=1}^{d} \phi_{ij}, \tag{8}$$

where ϕ_0 is the base value and ϕ_{ij} is the contribution of feature j for household i. SHAP values are based on Shapley values from cooperative game theory:

$$\phi_{ij} = \sum_{S \subseteq F, \{j\}} \frac{|S|!(d - |S| - 1)!}{d!} (f_{S \cup \{j\}}(X_i) - f_S(X_i)), \tag{9}$$

with F the full feature set and $f_S(\cdot)$ the model restricted to subset S. This enables **global** ranking of risk factors (mean $|\phi_{ij}|$) and local explanations for individual households, supporting auditability and policy interpretability.

3.4. Causal Machine Learning: Effect of Liquidity Constraints

Pure prediction does not identify the causal effect of liquidity constraints. The causal module estimates:

$$\tau(x) = \mathrm{E}[Y(1) - Y(0) | X = x], \tag{10}$$

where $Y(1)$and $Y(0)$denote potential outcomes under $D=1$and $D=0$.

Double Machine Learning (DML). DML uses orthogonalization to reduce bias from high-dimensional confounding. Let:

$$m(x)=\mathrm{E}[Y|\ X=x], e(x)=\Pr(D=1|\ X=x). \tag{11}$$

Estimate $\hat{m}(\cdot)$and $\hat{e}(\cdot)$with flexible ML (e.g., GBDT), then form residuals:

$$\tilde{Y}_i=Y_i-\hat{m}(X_i), \tilde{D}_i=D_i-\hat{e}(X_i). \tag{12}$$

An average treatment effect (ATE) can be estimated by regressing $\tilde{Y}_i$on $\tilde{D}_i$:

$$\hat{\theta}=\frac{\sum_i \tilde{D}_i\tilde{Y}_i}{\sum_i \tilde{D}_i^2}. \tag{13}$$

This approach provides robust inference for low-dimensional causal parameters under complex nuisance models.

Generalized Random Forests (GRF) for heterogeneity. To map heterogeneous effects, GRF estimates $\tau(x)$via forest-based adaptive local weighting:

$$\hat{\tau}(x)=\arg\min_{\tau}\sum_{i=1}^{n}\alpha_i(x)(\tilde{Y}_i-\tau\tilde{D}_i)^2, \tag{14}$$

where $\alpha_i(x)\geq 0$are data-adaptive weights induced by the forest structure. This yields CATE estimates that can be aggregated by **income quintile**, **city tier**, and **region** to identify vulnerable subpopulations most affected by liquidity constraints.

3.5. Uncertainty-Aware Risk Control (Conformal Risk Sets)

For operational robustness, distribution-free uncertainty control is added via conformal prediction. Given calibrated probabilities $\tilde{p}_i$, a nonconformity score can be defined as:

$$s_i=1-\tilde{p}_i(\text{for} Y_i=1) \quad \text{or} \quad s_i=\tilde{p}_i(\text{for} Y_i=0), \tag{15}$$

and a quantile threshold $q_{1-\alpha}$is computed on a calibration set. A household is flagged as "high-risk under uncertainty control" when the conformal criterion indicates that vulnerability cannot be ruled out at confidence level $1-\alpha$. This supports conservative screening under model and sampling uncertainty.

4. Experiments and Results Analysis

4.1. Experimental Setup

Household-level microdata were collected from nationally representative household finance surveys, with variables covering balance-sheet positions (income and financial assets), demographic characteristics (age, education, household size, employment), and regional and city-tier attributes. Financial Vulnerability Index (FVI) was constructed as a binary risk label indicating whether liquid buffers are insufficient to absorb unexpected expenditure shocks. Liquidity constraint status was identified based on credit access frictions and related indicators, and was included as a key risk driver within the feature set.

To reflect realistic deployment settings, a temporal evaluation protocol was adopted: earlier-wave observations were used for model training, while a later-wave sample was used as an out-of-time test set. Two representative tabular models were benchmarked: (i) Logistic Regression as an interpretable baseline and (ii) Gradient Boosting as a non-linear high-capacity learner. Model performance was evaluated using threshold-free discrimination metrics (Area Under the Receiver Operating Characteristic Curve, AUC), tail-relevant ranking metrics (Average Precision, AP), and probabilistic accuracy (Brier score). In addition, an operational screening rule was examined by flagging households in the top 20% of predicted risk to mimic a risk-management prioritization scenario.

4.2. Predictive Performance and Probability Quality

Figure 1 reports ROC curves on the temporal test set. Logistic Regression achieves an AUC of 0.714, while Gradient Boosting attains an AUC of 0.695, indicating that both models provide meaningful separation between vulnerable and non-vulnerable households, with the linear baseline slightly stronger in overall discrimination. Beyond AUC, risk management depends on probability quality rather than only ranking. Figure 2 presents the reliability diagram of the selected operational model: the curve tracks the 45-degree reference line with moderate deviations, suggesting that predicted probabilities are broadly aligned with realized vulnerability frequencies and are usable for probability-based screening and stress testing.

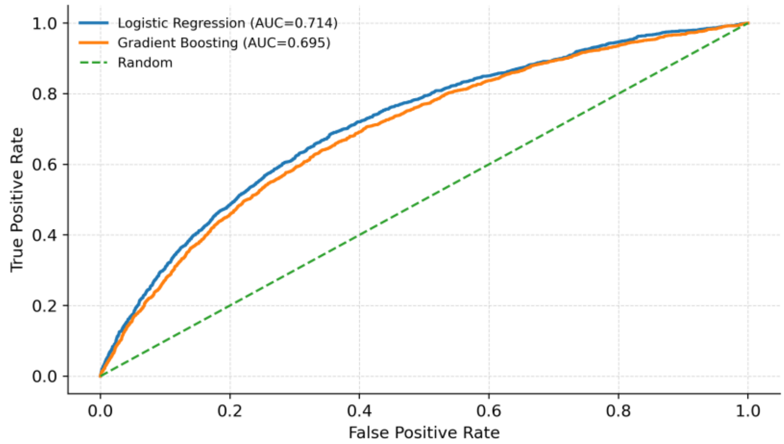

Figure 1: ROC Curves

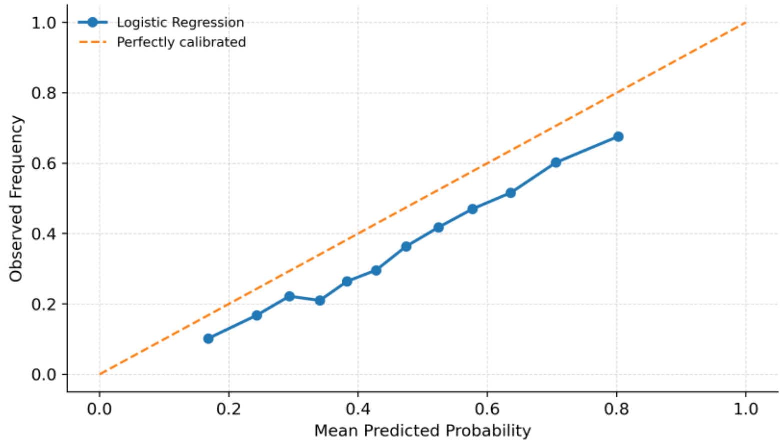

Figure 2: Calibration Reliability

Figure 3 visualizes the distribution of predicted vulnerability probabilities by true class. The vulnerable group is shifted toward higher predicted probabilities, while the non-vulnerable group concentrates more on the low-probability region, consistent with effective model separation. Under the operational threshold that flags the top 20% risk scores (Figure 4), the screening policy yields precision = 0.625 and recall = 0.348 on the temporal test set. This pattern is typical of conservative risk triage: the flagged set contains a high share of truly vulnerable households (high precision), while a portion of vulnerable households remains outside the top-risk segment (moderate recall). Such a design is suitable for constrained intervention capacity (e.g., targeted counseling or liquidity support), where minimizing false alarms is prioritized.

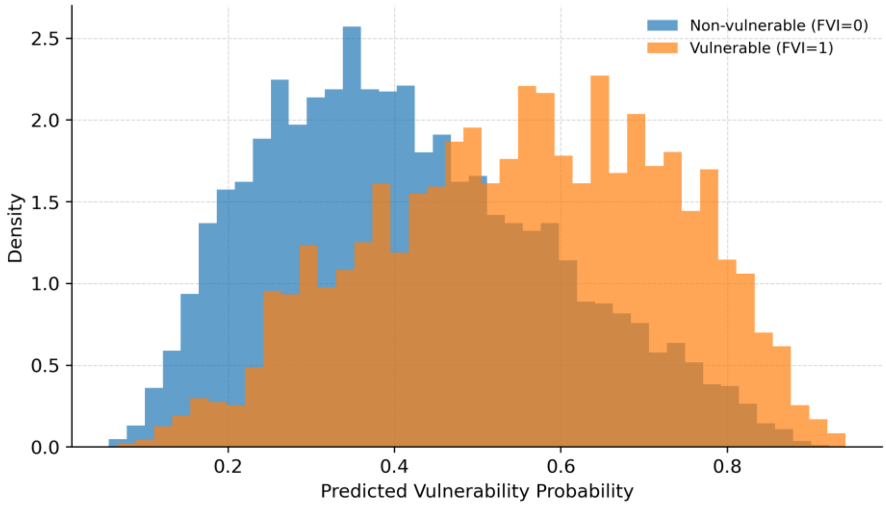

Figure 3: RiskScore Distributions

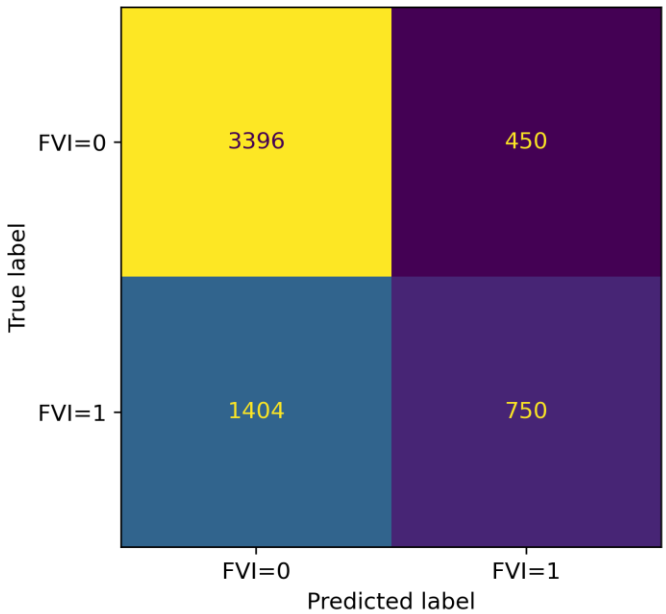

Figure 4: Confusion Matrix

4.3. Interpretability and Heterogeneity of Liquidity-Constraint Risk

To ensure that the risk model supports actionable financial risk management, global interpretability was examined. Figure 5 ranks the most influential predictors in the Gradient Boosting model. Financial assets and income-related variables dominate the importance ranking, indicating that balance-sheet buffers are primary determinants of vulnerability risk. Liquidity constraint status is among the top drivers, reinforcing its relevance even after conditioning on income, assets, and household characteristics. Education years and household size also contribute materially, reflecting long-term human-capital differences and consumption burden effects.

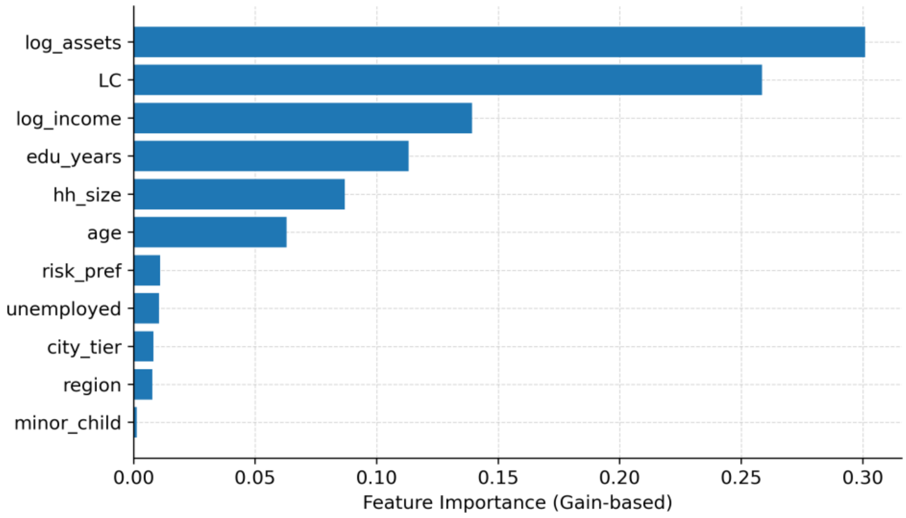

Figure 5: Global Feature Importance

Heterogeneity analysis further explores how liquidity constraints amplify predicted vulnerability across subpopulations. Figure 6 reports the mean change in predicted risk when switching liquidity constraint status from unconstrained to constrained, aggregated by income quintile and city tier. The amplification effect is largest in the lowest income quintile and gradually attenuates as income increases, indicating that liquidity constraints interact strongly with limited financial buffers. Across city tiers, the effect differences are present but smaller in magnitude than income gradients; lower-tier cities exhibit slightly higher amplification in several income groups, consistent with weaker local financial access and thinner informal risk-sharing networks. These results suggest that algorithmic targeting of liquidity-relief policies should prioritize low-income households and pay attention to structural constraints in less-developed city tiers.

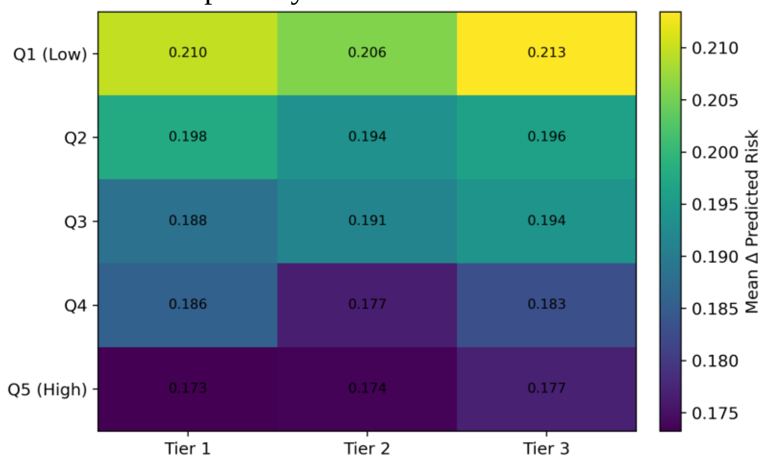

Figure 6: Heterogeneity Heatmap CATE

5. Conclusions

This study examines household financial vulnerability from a financial risk management perspective and evaluates how liquidity constraints amplify the probability of entering financial distress. Using household-level microdata with comprehensive balance-sheet, demographic, and regional attributes, the vulnerability label is constructed based on whether liquid buffers are sufficient to absorb unexpected expenditure shocks. Empirical results from out-of-time testing indicate that tabular machine learning models can effectively discriminate vulnerable from non-vulnerable households, while probability calibration supports the operational use of predicted risk scores in screening and monitoring scenarios.

The results consistently indicate that liquidity constraints are a material risk driver after controlling for income, financial assets, household size, education, employment, and location characteristics. Interpretation analysis confirms that buffer-related variables (income and assets) dominate risk formation, but liquidity constraints remain among the most influential predictors, implying that access to short-term funding and credit smoothing is directly linked to household fragility. Heterogeneity patterns further show that the amplification effect of liquidity constraints is strongest among low-income households and remains non-negligible across city tiers, suggesting that constrained households with limited buffers face disproportionately higher vulnerability risk.

These findings provide clear implications for household financial risk governance and inclusive-finance policy. Risk-based targeting strategies should prioritize liquidity relief and emergency buffer-building interventions for low-income groups, and strengthen accessible credit and contingency support mechanisms in areas where liquidity constraints are most binding. More broadly, embedding explainable, calibrated machine learning into household risk monitoring can improve the precision of policy targeting, enhance the efficiency of limited intervention resources, and contribute to more resilient household balance sheets and a more stable financial system.

References

[1] C. Midões, M. Seré, "Living with Reduced Income: An Analysis of Household Financial Vulnerability under COVID-19," Social Indicators Research, vol. 161, no. 1, pp. 125–149, 2022, doi: 10.1007/s11205-021-02811-7.

[2] R. Cifuentes, P. Margaretic, T. Saavedra, "Measuring households' financial vulnerabilities from consumer debt: Evidence from Chile," Emerging Markets Review, vol. 43, Art. 100677, 2020, doi: 10.1016/j.ememar.2020.100677.

[3] C. Boar, D. Gorea, V. Midrigan, "Liquidity Constraints in the U.S. Housing Market," The Review of Economic Studies, vol. 89, no. 3, pp. 1120–1154, 2022, doi: 10.1093/restud/rdab063.

[4] N. Bussmann, P. Giudici, D. Marinelli, J. Papenbrock, "Explainable Machine Learning in Credit Risk Management," Computational Economics, vol. 57, pp. 203–216, 2021, doi: 10.1007/s10614-020-10042-0.

[5] A. N. Angelopoulos, S. Bates, A. Fisch, L. Lei, T. Schuster, "Conformal Risk Control," in Proc. International Conference on Learning Representations (ICLR), 2024, doi: 10.48550/arXiv.2208.02814.

[6] B. Baugh, I. Ben-David, H. Park, J. A. Parker, "Asymmetric Consumption Smoothing," American Economic Review, vol. 111, no. 1, pp. 192–230, 2021, doi: 10.1257/aer.20181735.

[7] S. Catherine, M. Miller, N. Sarin, "Relaxing household liquidity constraints through social security," Journal of Public Economics, vol. 189, Art. 104243, 2020, doi: 10.1016/j.jpubeco.2020.104243.

[8] R. L. Clark, A. Lusardi, O. S. Mitchell, "Financial Fragility during the COVID-19 Pandemic," AEA Papers and Proceedings, vol. 111, pp. 292–296, 2021, doi: 10.1257/pandp.20211000.

[9] F. Donou-Adonsou, N. Leslie-Piper, "BNPL and financial fragility in U.S. households," Finance Research Letters, vol. 86, Art. 108423, 2025, doi: 10.1016/j.frl.2025.108423.

[10] M. Toussaint-Comeau, "Liquidity constraints and debts: Implications for the saving behavior of the middle class," Contemporary Economic Policy, vol. 39, no. 3, pp. 479–493, 2021, doi: 10.1111/coep.12521.

[11] P. E. de Lange, B. Melsom, C. B. Vennerød, S. Westgaard, "Explainable AI for Credit Assessment in Banks," Journal of Risk and Financial Management, vol. 15, no. 12, Art. 556, 2022, doi: 10.3390/jrfm15120556.

[12] M. K. Nallakaruppan, B. Balusamy, M. Lawanya Shri, V. Malathi, S. Bhattacharyya, "An Explainable AI framework for credit evaluation and analysis," Applied Soft Computing, vol. 153, Art. 111307, 2024, doi: 10.1016/j.asoc.2024.111307.

[13] Y. Gorishniy, I. Rubachev, V. Khrulkov, A. Babenko, "Revisiting Deep Learning Models for Tabular Data," in Advances in Neural Information Processing Systems (NeurIPS 2021), pp. 18932–18943, 2021, doi: 10.48550/arXiv.2106.11959.

[14] F. M. Ojeda, M. L. Jansen, A. Thiéry, S. Blankenberg, C. Weimar, M. Schmid, A. Ziegler, "Calibrating machine learning approaches for probability estimation: A comprehensive comparison," Statistics in Medicine, vol. 42, no. 29, pp. 5451–5478, 2023, doi: 10.1002/sim.9921.

[15] M. Yang, X. Bi, "Cost-Aware Calibration of Classifiers," INFORMS Journal on Data Science, vol. 4, no. 2, pp. 101–113, 2025, doi: 10.1287/ijds.2024.0038.