AI-Driven Personalization in Robo-Advisory: A Conceptual Framework for Integrated Household Financial Decision-Making

YASHASHVI VERMA

Abstract: *The vast majority of households lack access to integrated professional financial advice, leading to systematic and costly errors in investment, borrowing, and retirement planning. Contemporary robo-advisors address accessibility by automating portfolio allocation but remain structurally limited: they optimize investments in isolation from credit risk, debt obligations, and household cash-flow constraints, rely on simplistic and biased risk-tolerance questionnaires, and fail to resolve the cold-start problem for new users. This paper develops a novel conceptual framework for AI-driven, integrated household financial advisory systems grounded in design-science research. The framework comprises four tightly coupled modules—Credit Risk Assessment, Cash-Flow Forecasting, Dynamic Portfolio Optimization via reinforcement learning, and Financial Planning and Recommendation—linked through a Personalization Cascade. In this cascade, the credit risk profile (risk coefficient, debt-to-income ratio, and income volatility) parameterizes downstream modules, producing advice that is compoundingly personalized rather than merely additive. A second core contribution, Credit-Risk-Informed Utility Initialization, operationalizes the cold-start problem by seeding the reinforcement learning agent's utility function with behavioral signals from credit data, ensuring that households with different financial constraint profiles face genuinely different optimization landscapes from the first interaction. Three theoretical propositions on the interdependence of household financial decisions, the superiority of credit risk as a personalization signal, and reinforcement learning as the integration mechanism underpin the architecture. Illustrative personas and a capability comparison with leading platforms (Betterment, Wealthfront, Scripbox, Groww) demonstrate that the framework generates structurally different—and welfare-superior—recommendations, particularly the explicit debt-repayment versus investment trade-off. The framework offers a theoretically grounded blueprint for next-generation robo-advisory systems that can improve financial inclusion, reduce costly household errors, and meet emerging regulatory expectations for explainable, integrated advice.*

Keywords: Robo-Advisory Systems, Personalization Cascade, Credit Risk, Reinforcement Learning, Household Finance, Cold-Start Problem, Financial Inclusion

Yashashvi Verma (yashashviverma509@gmail.com)

Introduction

The vast majority of households lack access to professional financial advisors to guide them in making financial decisions. They invest with little understanding of whether they can afford to do so, they borrow without fully appreciating the long-term costs involved, and they often fail to plan adequately for retirement. There are tangible and long-term consequences. In a seminal Presidential Address to the American Finance Association, Campbell (2006) provided evidence that a large fraction of households—disproportionately the less educated and less wealthy—persist in making systematic and very costly errors in their investment choices, such as failing to diversify their portfolios, not refinancing their home mortgages at the optimal time, and not participating in equity markets that would benefit them in the long run. Such decisions are not irrational choices made by irrational individuals; rather, they are rational responses to a financial system that presents households with complicated and interdependent choices while offering little guidance for navigating them.

It is not only about mistakes. In a comprehensive review of the growing empirical evidence, Lusardi and Mitchell (2014) reported that financial illiteracy is widespread across developed economies and is associated with suboptimal financial behaviors such as inadequate retirement saving, costly borrowing, and poor financial planning. Financial knowledge, they argued, is a form of human capital—one that many households have had limited opportunity to develop. This is not merely an educational issue. It is fundamentally a structural one: the financial advisory system often fails to reach households where support is most needed.

Human financial advisors provide valuable context. A competent advisor understands that a client's capacity to tolerate equity exposure may be constrained by existing debt obligations, that an apparently conservative investor may have accumulated substantial employment-related wealth, and that financial plans developed in one stage of life often need to be adjusted in another. However, human guidance is expensive, time-consuming, and not always accessible. More importantly, it cannot be scaled sufficiently to support the broader population that could benefit from financial advice. As noted by Reher and Sokolinski (2024) and others, professional financial advice remains disproportionately concentrated among wealthier households, despite the fact that financially vulnerable households are often the most likely to make the costly decisions identified by Campbell.

The emergence of algorithmic financial advice platforms, commonly known as robo-advisors, appeared to offer a pathway toward a more accessible form of financial guidance. Betterment, Wealthfront, and Vanguard Personal Advisor Services have reduced fees, encouraged portfolio discipline, and made basic investment services available to households that otherwise may not have participated in financial markets (Wexler & Oberlander, 2021). According to Tao et al. (2021), assets under management in the global robo-advisory sector stood at approximately USD 297 billion in 2017 and are projected to reach nearly USD 4.66 trillion by 2027. By most measures of financial inclusion, this represents meaningful progress. Nevertheless, it remains a limited form of progress. Contemporary robo-advisors primarily manage portfolio allocations without fully accounting for the broader financial circumstances of the individuals they serve. They generally do not incorporate information regarding high-interest debt, future cash-flow constraints, or short-term financial obligations. Furthermore, when a new client joins the platform, a brief risk-tolerance questionnaire is typically used for initial profiling. As described by Adam et al. (2019), such questionnaires are often simplistic, susceptible to self-reporting bias, and insufficiently representative of the actual financial constraints faced by households.

THE CORE PROBLEM

Robo-advisors optimize investment portfolios in isolation from credit risk, cash-flow constraints, and household-level financial planning. Their advice is partial by design — not by accident.
This paper proposes a framework to close that gap.

The current robo-advisors are structural as opposed to technical limitations. Both the platforms themselves can carry out detailed calculations. They are lacking in integration. The credit risk evaluation is conducted in a system, usually the bank or lending system, where the lending decision making process is based on both interest rates and loan applications. When cash-flow forecasting does exist, it is in the budgeting programs, which are not linked to investment programs. Neither of these is combined with dynamic portfolio optimization, which is more and more driven by the reinforcement learning (Hambly et al., 2023). Another layer of software, however, produces the financial planning recommendations without knowing any of the other layers' up-to-date data. Yet the family in this system is left alone to try to assimilate all this advice, and to do a good job of it, and nothing is done to help the family.

This division is a real, but not thoroughly explored research need. Individualization is known as one of the key weaknesses of robo-advisories (Nourallah, 2023; Helms et al., 2022) and as an independent research area with its own methodology and literature in the area of credit risk (Markov et al., 2022) and portfolio optimization (Hambly et al., 2023), respectively. There are no current structures that link them. The first system of the world to base its investment advices on the credit risk profile of a household. This paper does just that.This paper aims to address this.

One specific technical challenge in this larger challenge is called the 'Cold Start Problem' and it is mentioned: Cold Start Problem. In a recommender system where there is no history of users' interactions, and the system cannot make any personalised recommendation, there will be a problem called cold-start. This is particularly true in the case of financial advisory systems. Even if it is a new user (somebody who doesn't have an investment account), he or she does have some credit history, DTI ratio and cash-flow history that provide a great deal of information about a person's financial scenario and risk capacity. No systems are in use that utilize this information at present. We think that it should "plant" the advisory process. The utility functions' initialisation method with credit risk profiles in an integrated advisory system is the main contribution of this paper.

Research Propositions

This paper develops and evaluates a conceptual framework for integrated household financial advisory systems, grounded in three theoretical propositions that we present here and defend fully in Section 4.

PROPOSITION 1 — Interdependence of Household Financial Decisions
A household's investment risk capacity depends on its debt load. Its debt repayment capacity depends on its cash flow. Its cash flow depends on income stability and spending behavior. These relationships are not additive — they are multiplicative. An advisory system that models each dimension independently will systematically misestimate what a household can and should do with its money.

PROPOSITION 2 — Credit Risk as the Central Personalization Signal
Credit risk scores are grounded in actual financial behavior, not self-reported preferences. They encode household-level constraints — income volatility, debt obligations, spending patterns — more reliably than any questionnaire. They are also already computed, widely available, and interpretable. We argue they should serve as the primary signal for initializing personalized financial advice — particularly for new users with no advisory interaction history.

PROPOSITION 3 — Reinforcement Learning as the Integration Mechanism
Reinforcement learning — specifically its capacity for adaptive, state-dependent sequential decision-making — provides the mechanism that allows an integrated financial advisory system to update dynamically as the household's financial state changes. We do not claim that 'AI' in a general sense integrates these modules. We claim that RL-based optimization, operating over a state space that includes credit risk signals and cash-flow forecasts, is the appropriate tool for generating advice that is simultaneously dynamic, personalized, and financially coherent.

Core Contributions

This paper makes two contributions that together distinguish it from existing work in robo-advisory design, household finance, and AI-driven personalization.

CONTRIBUTION 1 — The Personalization Cascade
We introduce the concept of a Personalization Cascade: a directed information flow in which the output of the credit risk assessment module parameterizes the cash-flow forecasting module, which in turn seeds the RL-based portfolio optimizer's utility function, which ultimately shapes the financial planning recommendations delivered to the household. This is not a pipeline in the conventional sense. In a standard pipeline, each module operates

independently and outputs are passed sequentially. In a Personalization Cascade, the credit risk signal fundamentally changes how downstream modules behave — producing advice that is compoundingly personalized, not merely additive.

CONTRIBUTION 2 — Credit-Risk-Informed Utility Initialization
We propose Credit Risk-Informed Utility Initialization as a principled approach to the cold-start problem in financial advisory AI. A new user joins with no investment experience, and the credit risk profile (a form of behavioral and financial constraint information) is fed into the utility function U(.) of the RL-based portfolio optimizer at its initialisation. The initial optimization landscapes corresponding to different risk profiles are different: A household with a high credit score and low debt-to-income ratio optimizes for growth, while a household with a lower credit score and higher debt-to-income ratio optimizes for stability and debt repayment first. This mechanism is not only theoretically sound but also is easily (and practically) implementable with existing credit data infrastructure.

About Methodology. It is followed by a conceptualization paper and not a study. We follow the design-science tradition in information systems research (Gregor & Hevner, 2013) and consider the development of the framework a rigorous, intellectual contribution in itself, one in need of theoretical grounding, illustrative scenarios for which to design it, and a systematic comparison with existing frameworks. An empirical validation of the framework, which includes a simulation study and fintech field trial described in Section 7.4, is planned for future work. Selection of appropriate datasets (PSID household panel data, credit bureau synthetic datasets, UK open banking data under PSD2) and experimental designs to be used during the validation phase.

The remainder of this paper is organized in the following way. Section 3 will review the pertinent literature, with gaps between household finance and machine learning in credit risk, gaps between machine learning and reinforcement learning in portfolio optimization, and gaps between machine learning and personalized financial recommendation driving this work. Section 4 provides and supports the three theoretical propositions. Section 5 defines four modules of the proposed framework such as Personalization Cascade and Credit-Risk-Informed Utility Initialization subsections. In section 6 the assessment framework with metrics of each module, the system, and example user personas and systematic comparison with other robo-advisory systems are presented. Section 7 covers theoretical and practical implications and limitations, as well as future research directions. Section 8 concludes.

Literature Review

The following review covers four bodies of scholarship that directly inform the proposed framework: household financial advisory systems, machine learning in credit risk, reinforcement learning for portfolio optimisation, and personalised financial recommendation with cash-flow forecasting. Each subsection closes by identifying what remains unaddressed in that strand of the literature. Section 3.5 draws these threads together to articulate the integration gap that motivates this paper.

3.1 Household Finance and the Limits of Existing Advisory Systems

Household financial decision-making covers income management, debt servicing, insurance, tax planning, and long-term investment, all of which are interdependent in such a way that providing algorithmic advice on just one of these components is structurally incomplete without reference to the other components and their context. In the past, human advisors have always enjoyed a structural advantage to deal with these interactions. Zhu et al. (2024) and Noble and Mende (2023) have all found that the relational aspect of advice is not a preference, but an important part of good financial advice, especially during life stage transitions like career change or bereavement.

But one of the major hurdles to human advice is easily documented: scale. A comprehensive, ongoing financial advisory has long been out of reach for households with less than around USD 250,000 in investable assets, and is effectively unavailable to the majority of households who would benefit most from it (Reher and Sokolinski 2024). Only 40 percent of the adults surveyed in OECD economies said that they have consulted a professional financial adviser at any time in their lives (OECD, 2023). There are also known sources of systematic errors in human advice, such as confirmation bias, commission-driven product selection and recency bias, which are shown to be common problems in advised portfolios by Luo et al. (2024) and Benavent (2021).

Ashley and Helen were the first to launch around 2010, now known as robo-advisors, which used algorithms to automate portfolios and reduce investment annual management fees to around 0.25-0.50 percent of investments, compared to 1.0-1.5 percent for full-service providers (Belanche et al., 2025). However, their personalisation model has been subject to regular criticism. Most platforms assess clients' risk level in three to six buckets through questionnaires that contain five to seven items, as documented by Adam et al., (2019) and described by Luo et al. (2024) as superficial and susceptible to self-reporting bias. Importantly, to our best of knowledge, there is no peer-reviewed study that documents a commercially deployed robo-advisory platform that uses a household's credit risk profile as a form of personalisation signal, or that sets portfolio recommendations based on dynamic household cash-flow forecasts (Bashir et al., 2023). These domain-silo constraints are not a mere side effect, but rather the structural deficit which drives the current system.

3.2 Machine Learning in Credit Risk Assessment

The application of machine learning to credit risk prediction is among the most mature areas of applied finance. The methodological evolution has moved through four broad phases: logistic regression and scorecard models; ensemble tree methods, particularly random forests and gradient boosting; deep neural architectures and recurrent networks; and post-2017 explainable AI frameworks designed to meet regulatory interpretability requirements (Nwaimo et al., 2024; Nallakaruppan et al., 2024). The predictive performance of ensemble methods is now well-established: Chang et al. (2024) demonstrated that XGBoost achieved classification accuracies exceeding 99 percent on credit card default datasets, substantially outperforming logistic regression benchmarks, a finding consistent with Chen and Guestrin (2016)'s original XGBoost formulation.

What the credit risk literature has not yet systematically explored is the use of credit-risk signals as cross-domain personalisation inputs for advisory systems beyond lending. A credit score, correctly understood, encodes payment discipline, revolving debt management, income volatility tolerance, and financial planning horizon — dimensions that are directly relevant to investment decision-making (Nayak, 2023; Hlongwane et al., 2024). Zhang and Yu (2024), in a comprehensive review of consumer credit risk assessment methods, focused on multi-scenario behavioural modelling for default prediction but did not extend their analysis to cross-domain advisory applications. Roa et al. (2021) showed that transaction-level behavioural data from financial apps yields statistically significant predictive power over traditional credit bureau features, reinforcing the informational richness embedded in credit behaviour data. *The literature has not, to our knowledge,* proposed a mechanism by which this behavioural richness is channelled into the personalisation logic of a broader financial advisory system — an opportunity that this paper's second proposition directly addresses.

3.3 Dynamic Portfolio Optimisation Through Reinforcement Learning

While the use of reinforcement learning in portfolio management has an extensive academic background, dating back to original research by Moody and Saffell (1998) and Neuneier (1997). The tradition has continued in modern deep RL works, which introduce Proximal Policy Optimisation (PPO), Soft Actor-Critic (SAC), and Double Deep Q-Networks (DDQN) to dynamic multi-asset allocation in high-dimensional continuous action spaces (Hambly et al., 2023; Jiang et al., 2017). Leukam et al. (2025) illustrated how G-Learning algorithms can be adjusted according to the investors' time horizon and the level of their terminal wealth goals, while Day et al. (2024) used deep RL techniques to develop dynamic trading strategies that outperformed the traditional static trading strategies benchmarks over multi-year evaluation periods in the Chinese stock markets.

However, the performance of RL in a practical application in a retail financial environment is more limited than what is suggested in simulations. Malibari et al. (2023) and Puiutta and Veith (2020) both noted that key challenges in the live deployment of financial time series are their non-stationarity, the sparseness of reward signals during exploration, and the need to balance the control of downside risk with the minimization of transaction costs. The present paper recognizes these constraints explicitly, as they guide the choice of the scope of the simulation outlined in Section 7.3 and the choice of credit-risk-informed utility function initialization of the RL module, thus alleviating the burden of the exploration period when the cold-start problem occurs. More importantly, as far as we know, there is no survey work in the literature that documents a study in which the credit-risk signals are fed into the RL reward function as state variables, which Vuković et al. (2025) noted as part of the overall lack of cross-domain integration in the field of AI and finance.

.

3.4 Personalised Financial Recommendation and Cash-Flow Forecasting

Personalised financial recommendation systems face a well-characterised set of structural challenges. Sharaf et al. (2022), in a systematic survey published in *Multimedia Tools and Applications*, identified seven recurring problems — cold start, data sparsity, scalability, latency, grey-sheep users, shilling attacks, and synonymy — and called for further research leveraging advanced machine learning to address the domain-silo challenge in financial advisory. Hybrid collaborative-filtering architectures and dynamic utility learning have begun to address individual problems, and large language model interfaces have been proposed for richer preference elicitation (arXiv:2512.12922, 2024), though peer-reviewed validation of these approaches in live advisory settings remains limited.

Cash-flow forecasting at the household level is the least developed component of the four analytical modules this framework requires. The machine learning literature on cash-flow forecasting is concentrated on corporate treasury applications rather than individual household income-expenditure streams (Roy et al., 2024; Zheng & Tu, 2023). Long Short-Term Memory networks, random forest ensembles, and hybrid ARIMA-neural models have each demonstrated capacity to outperform traditional statistical baselines on appropriate datasets, but *the literature has not yet, to our knowledge, demonstrated* (i) household-specific cash-flow forecasting from open-banking transaction streams as a real-time constraint on portfolio optimisation; (ii) feedback loops between cash-flow uncertainty and credit-risk signals within a single advisory system; or (iii) empirical evaluation of such integration in a retail financial setting. The present framework treats these three omissions as the architectural gap its cash-flow module is designed to address, while acknowledging in Section 7.3 that this module is the component most dependent on data infrastructure not yet uniformly available in retail banking.

3.5 The Integration Gap

> **THE CORE ARGUMENT OF THIS SECTION**
> *Each of the four bodies of literature reviewed above is internally mature. What the field lacks, to the best of our knowledge, is a unified architectural framework that connects all four domains under a single personalisation signal — specifically, one in which credit-risk outputs drive personalisation across portfolio optimisation, cash-flow forecasting, and financial recommendation simultaneously.*

The four bodies of literature reviewed in Sections 3.1 through 3.4 represent substantial and largely self-contained research programmes. Each has produced mature methodological frameworks, systematic reviews, and benchmark evaluations. Vuković et al. (2025), in a scientometric review of AI in financial services spanning 1989 to 2024, catalogued the growth of machine learning in credit scoring, fraud detection, and robo-advisory as separate trajectories and explicitly identified the absence of standardised cross-domain integration frameworks as the field's most persistent gap. Sharaf et al. (2022) surveyed financial recommendation systems and called for research addressing domain-silo challenges, without proposing an integrative architecture.

More recent contributions have begun to address partial combinations. The LLM-PPR framework (arXiv:2512.12922, 2024) connects conversational preference elicitation to RL-based portfolio recommendation — bridging Sections 3.3 and 3.4 — but omits credit-risk context entirely. Nayak (2023) introduced behavioural analytics into credit risk management to improve financial stability predictions, but did not extend the framework to cross-domain advisory personalisation. A 2025 survey of RL in finance reviewing over 350 papers on trading and portfolio optimisation cited no study in which credit-risk information appears as a state variable in the RL reward function (Vuković et al., 2025). These observations, taken together, suggest — *though we acknowledge this characterisation may not capture all unpublished or industry-internal work* — that the specific architectural combination proposed here has not been formalised in peer-reviewed literature.

This gap has practical consequences, which go beyond the academic completeness. If advisory systems are blind to the signals of credit risk, they systematically misallocate risk to households whose vulnerability is apparent in their credit portfolio but not obvious to the portfolio engine (Campbell, 2006). If there are no cash-flow constraints for the household, RL optimisers can generate allocation policies that are analytically optimal, but operationally impossible -- that is, they might require premature selling of positions in order to satisfy the liquidity constraint. No matter how sophisticated an algorithm, recommendation systems that are not based on the debt structure of their clients end up giving contextually irrational recommendations from a household balance-sheet perspective. This paper offers a framework to address these structural failures. The three theoretical propositions presented in Section 4 are direct consequences of the shape of this gap: Proposition 1 is driven by the perceived costs of providing advice in silos, Proposition 2 by the lack of exploitation of the information richness of credit risk signals reported in Section 3.2, and Proposition 3 by the deployment challenges and integration absences documented in Section 3.3.

Theoretical Foundations

The framework proposed in this paper rests on three theoretical propositions. They are not simply background assumptions — each one shapes a specific design decision in the four-module architecture described in Section 5. This section states each proposition, explains its theoretical basis, and connects it explicitly to the framework. Together, they make the case that integrating credit risk, cash-flow forecasting, and portfolio optimization into a single advisory system is not merely convenient but theoretically necessary.

> **Proposition 1 — Household Financial Decisions Are Interdependent**
> A household's capacity to take investment risk depends on its existing debt load. Its ability to service that debt depends on reliable cash flow. Its cash flow, in turn, reflects income stability, spending discipline, and the timing of irregular expenses. These are not independent variables — they form a chain in which a shock to one dimension propagates through the others.

This interdependence has been confirmed in an empirical fashion. Campbell (2006) demonstrated that there is a welfare cost to households with equity portfolios that don't match their overall financial situation, that is, those that have high-cost consumer debt combined with market investments. Lusardi and Mitchell (2014) found that low financial literacy was associated with multiple failures at the same time, implying that the negative consequences of financial illiteracy do not remain isolated to each aspect of the financial world. The design choice of modules sharing state makes sense for the framework, in this case because of this proposition. By design, an optimizer that does not know the debt load of the household will generate incomplete advice. It is not a bug in implementation -- it's a result of assuming that financially interdependent decisions are independent..

> **Proposition 2 — Credit Risk Is the Right Personalization Signal**
> Credit risk scores encode actual financial behavior — payment history, utilization ratios, income volatility, and debt obligations — rather than self-reported preferences. They are computed from real transaction data, are already available at scale, and capture household-level financial constraints more reliably than any questionnaire.

Personalization is a core challenge in financial advisory AI, as there is a problem of signal, especially when there is limited information available about a user. Current robo-advisors utilize brief questionnaires and use pre-determined risk categories that are static and general, which are susceptible to the self-reporting biases identified by Helms et al. (2022) as one of the main issues with existing advisors. This issue is completely avoided with credit risk data. It's a behavioural matter, not an attitudinal one. A family that has a credit score from having paid all of its bills on time and used relatively small amounts of credit is telling something about their finances that no questionnaire can tell as well. This proposition also directly motivates Contribution 2 of this paper, in which the utility function is initialized for the portfolio optimizer by credit risk profiles instead of the user describing her/himself.

> **Proposition 3 — Reinforcement Learning Is the Integration Mechanism**
> Reinforcement learning provides the adaptive, sequential decision-making mechanism that allows an integrated advisory system to update its recommendations as the household's financial state changes over time. It is not 'AI' in a general sense that integrates the modules — it is specifically RL's capacity to optimize over dynamic, state-dependent environments.

Traditional portfolio optimization techniques (mean-variance and its extensions) need consistent parameter estimates and generate a static allocation. They are not meant to be continuously updated with new information on a household's credit position or cash-flow trajectory. Mathematical Finance (Hambly et al., 2023) published a detailed survey of RL applied to sequential financial decisions (SFCDP) in which the environment is non-stationary and the assumptions of the model are hard to keep. This is precisely the ability that is needed in the integration proposed in this framework; the state space of the RL agent needs to include signals from the credit risk and cash-flow modules and the reward function needs to be initialized using the credit-risk-informed utility function introduced in Section 5.7.

The three propositions are not mutually independent, they are mutually dependent in a logical sense. If financial decisions are intertwined within the household (Proposition 1), the advisory system needs to be integrated as well. Credit risk, as the most accurate indicator of household financial constraints (Proposition 2), should be at the heart of the personalization logic. But if reinforcement learning is the mechanism of dynamic integrated optimization (Proposition 3), then it has to be the mechanism also for the portfolio module and linking to the credit layer and the cash-flow layer. This is, in this sense, the direct architectural expression of these three claims in the form of four modules in Section 5.

5. Proposed Conceptual Framework

The framework developed in this paper is built around a single design principle: that financial advice becomes meaningfully more useful when information about a household's constraints flows continuously across all dimensions of its financial life, rather than being processed separately by disconnected systems. The following subsections describe the four-module architecture, introduce the two core contributions of the paper, and explain the feedback mechanism that keeps the system adaptive over time.

5.1 Architecture Overview

The framework comprises four tightly coupled modules: (1) Credit Risk Assessment, (2) Cash-Flow Forecasting, (3) Dynamic Portfolio Optimisation, and (4) Financial Planning and Recommendation. These modules are not independent — they share state, and the output of each module parameterises the behaviour of downstream modules. This coupling is what distinguishes the architecture from a conventional advisory pipeline, where modules pass outputs but do not shape each other's internal logic. Figure 5.1 provides a schematic of the full system.

FIGURE 5.1 — FOUR-MODULE INTEGRATED ADVISORY ARCHITECTURE

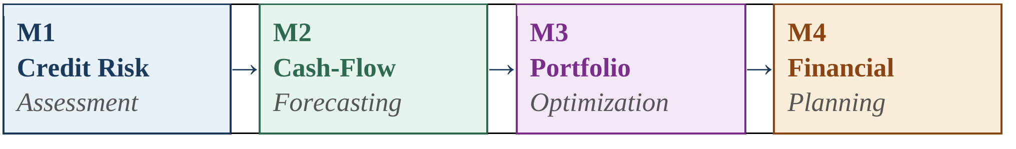

Arrows denote structured information transfer, not merely sequential output passing.

Table 5.1 situates the framework against four leading robo-advisory systems across six capability dimensions. No existing platform satisfies more than one of these criteria; the proposed framework is designed to satisfy all six simultaneously.

TABLE 5.1 — Capability Comparison With Existing Robo-Advisory Systems

Capability	This Framework	Betterment	Wealthfront	Scripbox
Models credit risk	✓	—	—	—
Cash-flow integration	✓	—	—	—
RL-based optimisation	✓	—	—	—
Solves cold-start	✓	—	—	—
Personalisation Cascade	✓	—	—	—
Integrated debt advice	✓	—	—	—

5.2 Module 1: Credit Risk Assessment

The first module creates a structured risk profile to the household based on three sets of inputs: financial history (credit utilisation ratios, the history of payments, balances of loans, and indicators of default); transaction-level behavioural (spending patterns, the volatility of spending at each category level, regularity of cash flow); and contextual (employment type, proxy measures of income stability, outstanding debt to income). The predictive backbone of this module is ensemble methods that combine gradient-boosted tree models with deep neural networks, the combination of which Chang et al. (2024) showed can achieve an accuracy of more than 99% for classification of credit default, with XGBoost significantly outperforming the logistic regression benchmarks. The machine learning method is preferred to the scorecard method, since it is able to capture the non-linear nature of how the interactions between the behavioural signals relate to the scoring in the model (Markov et al., 2022).

The module does not produce a credit score. It is a structured risk profile made up of 3 quantities which downstream modules use directly: (i) a normalised risk coefficient rc in [0, 1] with a value close to 1 being the highest financial constraint; (ii) a debt to income ratio d = D/Y, where D is the total amount of debt and Y is annualised income; and (iii) an income volatility estimate σy based on transactions history. These three quantities combine to represent the initial financial constraint envelope -- any financial trajectory that the system should even be tempted to recommend. The credit risk module, therefore, is not only a classification of the household, it is the space that all other modules can be considered to exist.

MODULE 1 — KEY INPUTS AND OUTPUTS

Inputs:

Credit utilisation history · Payment records · Loan balances · Spending patterns · Income volatility proxies · Employment type

Outputs:

$r_c \in [0,1]$ (risk coefficient) · $\delta = D/Y$ (debt-to-income ratio) · σ_y (income volatility)

5.3 Module 2: Cash-Flow Forecasting

The second module addresses what is arguably the most practically consequential failure in current robo-advisory systems: their complete disregard for whether a household can actually afford to invest what it intends to invest. Households routinely overestimate their investable surplus because they underweight irregular expenses, seasonal obligations, and the cash demands of their existing debt service schedule (Campbell, 2006; Lusardi & Mitchell, 2014). A portfolio optimiser that receives no cash-flow input will recommend strategies the household cannot sustain.

This module forecasts net disposable income $\hat{Y}_t$ across a rolling horizon T using a Long Short-Term Memory (LSTM) network, which Roy et al. (2025) and others have demonstrated outperforms ARIMA and Prophet models in capturing long-range temporal dependencies and non-linear seasonal patterns in personal financial data. The LSTM receives the household's credit risk profile outputs (r_c, δ, σ_y) as conditioning inputs, which ensures that the forecasting model accounts for debt service obligations in projecting available cash. A household with a high debt-to-income ratio δ will receive forecasts with a systematically lower $\hat{Y}_t$ and wider confidence intervals, reflecting both the reduced investable surplus and the higher uncertainty around irregular debt payments.

The cash-flow module also flags liquidity stress events — projected periods where $\hat{Y}_t$ falls below a minimum buffer threshold $\hat{Y}_{min}$ — and passes these as constraints to Module 3. The portfolio optimiser is not allowed to recommend strategies that would require capital commitments during a projected liquidity stress window. This is a design choice with direct welfare implications: it prevents the system from recommending investment contributions during months when the household's own projections suggest it cannot afford them.

5.4 Module 3: Dynamic Portfolio Optimisation (RL-Based)

The third module governs portfolio construction and rebalancing through a reinforcement learning agent formulated as a Markov Decision Process (MDP). The MDP is defined by the tuple $\langle S, A, P, R, \gamma \rangle$, where S is the state space, A is the action space, P is the transition probability function, R is the reward function, and γ is the discount factor. Following the framework established by Hambly et al. (2023) — who demonstrated that RL approaches are particularly well-suited to non-stationary financial environments — and consistent with the risk-averse portfolio formulation of Cui et al. (2024), the agent's state space in this framework incorporates the credit risk module's outputs explicitly.

The state $s_t \in S$ is a vector comprising: the current portfolio allocation w_t, the projected net disposable income from Module 2 $\hat{Y}_t$, the household's credit risk coefficient r_c and debt-to-income ratio δ, and prevailing market conditions represented by return and volatility estimates for each asset class. The action $a_t \in A$ is a portfolio rebalancing vector specifying the change in allocation across n asset classes, subject to the constraint that aggregate investment contributions do not exceed the liquidity-adjusted surplus $\hat{Y}_t - \hat{Y}_{min}$.

The reward function R incorporates risk-adjusted return, penalises drawdown, and applies a household-specific risk aversion coefficient $\gamma(r_c)$ derived directly from the credit risk module. This is the mechanism at the heart of the paper's second core contribution, which we describe formally in Section 5.7. The utility function parameterised by the credit risk score initialises the optimisation landscape of the RL agent, ensuring that households with different financial constraint profiles face genuinely different — not merely scaled — optimisation problems from the first interaction.

5.5 Module 4: Financial Planning and Recommendation

The fourth module is the interface between the computational layers of framework and the household. It transforms the output of the portfolio optimiser (an allocation policy p(s)) into integrated financial recommendations, which also account for investment and debt repayment policies, savings policies and liquidity management. The integration differs from the portfolio allocation suggestions offered by the recommendations layers of other robo-advisors, which merely show the portfolio allocation and let householders determine their own financial obligations.

The recommendation layer works simultaneously over three time horizons. In the short term (t = 1 month), it generates an action plan that includes the contribution to each investment vehicle, the amount of debt repayments (in addition to the minimum repayments) and the needed liquidity buffer. It creates medium-term (t = 3-12 months) debt reduction and asset accumulation trajectory targets in the medium term, which are calibrated to the cash-flow forecasts generated in Module 2. In the long term (t > 12 months), it reports on goal convergence, that is, whether the household is making progress toward its stated goals (e.g., retirement saving goals or down payment goals for a home purchase), and alerts for goal divergences as advisory triggers to re-start the feedback loop described in Section 5.8.

Critically, the recommendation module explicitly presents the decision between debt repayment and investment contribution, which is currently left solely to the user's decision. The system is designed to make investment repayments the priority first, before debt repayments, even though the user's stated goal is to build wealth, and the interest rate they are paying on the debt is higher than the ROI they are expecting on an investment they have in mind. It mirrors the first proposition of the theoretical framework: Household financial planning is interdependent, and financial advice that doesn't take into account the interplay between debt loads and investment potential will systematically give suboptimal advice.

5.6 The Personalization Cascade

The term *Personalization Cascade* describes the mechanism by which a single structured signal — the credit risk profile generated by Module 1 — propagates through all subsequent modules, compounding its personalisation effect at each stage rather than being applied additively at the end. This is the paper's first core contribution and the design feature that distinguishes the framework from existing multi-module advisory architectures.

The distinction between a cascade and a pipeline requires emphasis. In a conventional pipeline, Module 1 produces an output, Module 2 consumes that output and produces its own, and so on, with each module operating according to its own internal logic. In the Personalization Cascade, the credit risk profile does not merely become an input to subsequent modules — it changes *how* those modules behave. The cash-flow forecasting module adjusts its confidence intervals based on income volatility σ_y derived from Module 1. The portfolio optimiser initialises its utility function using r_c from Module 1. The recommendation layer presents trade-off explanations that are calibrated to the household's specific constraint envelope, not a generic template. Each downstream module is thus a personalised instance, not a generic module receiving personalised data.

FIGURE 5.2 — The Personalization Cascade: Directed Information Flow Across Four Modules

Stage 1 Credit Risk Assessment
Input: Financial history, transaction data, behavioural signals
Output: Risk profile $r_c \in [0,1]$; debt-to-income ratio δ; income volatility σ_y

↓ *(parameterises downstream)*

Stage 2 Cash-Flow Forecasting
Input: r_c, δ, σ_y — sets conservative vs. optimistic forecasting bounds
Output: Projected net disposable income $\hat{Y}_t$ for $t = 1 \ldots T$

↓ *(parameterises downstream)*

Stage 3 Portfolio Optimisation (RL)
Input: $\hat{Y}_t$ initialises the utility function: $U(W; \gamma(r_c))$
Output: Optimal allocation policy $\pi^*(s)$ across asset classes

↓ *(parameterises downstream)*

Stage 4 Financial Planning & Recommendation
Input: $\pi^*(s)$, $\hat{Y}_t$, δ — integrated across debt, savings, liquidity
Output: Personalised household financial plan with trade-off explanations

Note: Each arrow denotes parameterisation of the downstream module's internal logic, not merely data transfer.

The welfare implication of this compounding structure is significant. A household with r_c = 0.82 (high financial constraint) does not merely receive more conservative investment recommendations than a household with r_c = 0.21 — it receives cash-flow forecasts with wider uncertainty bands, faces an RL optimiser initialised to a different region of the utility landscape, and receives recommendations that explicitly foreground debt reduction as the primary near-term objective. The advice is structurally different, not just scaled. This is what personalisation that compounds rather than adds means in practice.

5.7 Credit-Risk-Informed Utility Initialization

The cold-start problem is a foundational challenge in recommender systems: how should a system personalise its recommendations for a new user with no interaction history Lam et al., 2008)? In financial advisory AI, the cold-start problem is particularly acute. A new user who has never held an investment account still has a credit history, a pattern of financial obligations, and a behavioural record encoded in their transaction data — all of which are meaningfully informative about their financial situation and risk capacity. Current robo-advisors discard this information and rely instead on short questionnaires, a method documented as unreliable and biased toward overestimating risk tolerance (Adam et al., 2019; Helms et al., 2022). Credit-Risk-Informed Utility Initialization is the mechanism by which this paper's framework solves the cold-start problem.

The RL agent in Module 3 optimises over a utility function U(W; γ) where W is household wealth and γ is the risk aversion parameter. Following the Constant Absolute Risk Aversion (CARA) formulation — which Imaki et al. (2025) demonstrated produces unbiased gradients and clear interpretation even in negative return environments — the utility function takes the form:

$U(W; \gamma) = -(1/\gamma) \cdot \exp(-\gamma \cdot W)$ *CARA utility function (Arrow–Pratt)*

where $\gamma(r\_c) = \gamma\_min + (\gamma\_max - \gamma\_min) \cdot r\_c$

In this formulation, the risk aversion parameter γ is not a fixed constant — it is a function of the household's credit risk coefficient r_c. The mapping γ(r_c) is monotonically increasing: a household with r_c near zero (low financial constraint) is initialised with low risk aversion γ_min, producing a utility landscape that favours growth-oriented portfolios. A household with r_c near one (high financial constraint) is initialised with high risk aversion γ_max, producing a utility landscape that strongly penalises drawdown and favours capital preservation and debt reduction first.

> **CREDIT-RISK-INFORMED UTILITY INITIALIZATION — ILLUSTRATIVE PARAMETERISATION**
> **Low-constraint household (r_c = 0.15, δ = 0.12):**
> γ(0.15) → low risk aversion → utility landscape favours equity-weighted growth portfolios with moderate drawdown tolerance
> **High-constraint household (r_c = 0.81, δ = 0.47):**
> γ(0.81) → high risk aversion → utility landscape penalises drawdown; RL agent prioritises debt-reduction-consistent allocations before equity exposure

This mechanism is theoretically grounded in the household finance literature. Campbell (2006) showed that households in financial constraint make systematically different and worse portfolio decisions than unconstrained households — not because of lower financial sophistication, but because the optimal unconstrained portfolio is genuinely wrong for their situation. The Credit-Risk-Informed Utility Initialization operationalises this insight: it ensures that the portfolio optimiser's starting point reflects the household's actual constraint envelope, not a generic neutral initialisation that implicitly assumes the household has no financial obligations beyond the investment account.

5.8 Continuous Feedback Loop and System Adaptability

It is not a single engine for recommending a framework. The system needs to be able to reflect the financial situation of the household with its changing circumstances, such as rising income, repaying loans, changes in the market, or significant life changes. The four modules re-estimate together when a triggering condition is satisfied, rather than working on dated information, while the household's situation has changed.

There are three types of update triggers. The credit risk module re-estimates the credit risk of the household based on the most recent 90 days of transactions, while the cash-flow module changes the forecast by one month, and the portfolio optimiser checks if the allocation policy p(s) is still optimal for the current state of the portfolio. Event driven reestimation occurs at the instant of the occurrence of threshold crossing events: Recession in rc (when it is closing in on the threshold Drc) Liquidity stress event in the cash-flow module Portfolio drawdown exceeds a pre-defined maximum threshold Household's statement of financial goals change When two or more signals conflict, the process of conflict resolution is triggered – e.g. the portfolio optimiser suggests to increase exposure to equities while the cash-flow module also signals a liquidity stress window. In these situations, the recommendation module is designed to take a conservative override, that is, to make a more conservative recommendation until the conflict is resolved, and to bring the disagreement to the attention of the household in a transparent fashion rather than in a hidden fashion through the system.

This logic for conflict resolution is not only regulatory, but also advisory. Gregor and Hevner (2013) point out that in contexts where design is a topic of high importance, design-science frameworks should account not only for the system's behavior when things go well, but when things go poorly. The elements of this framework close the loop on the architecture: the system isn't a one-off piece of advice but a continually adapting household financial model that monitors what's happening to the household over time and adjusts its recommendations accordingly.

FEEDBACK LOOP — TRIGGER TAXONOMY

Scheduled Monthly re-estimation across all four modules using latest 90-day transaction window

Event-driven Δr_c threshold breach · Projected liquidity stress · Portfolio drawdown > limit · Goal change

Conflict Conservative override when module signals disagree; disagreement surfaced transparently to the household

Evaluation Framework

Evaluating a four-module integrated advisory framework requires a layered approach: no single metric captures whether an integrated system is working. The framework must be assessed at the level of each constituent module, at the level of the integrated system, and through concrete illustrations of how different household profiles are handled differently. The following subsections specify each layer in turn.

6.1 Module-Level Metrics

Each module is assessed to criteria aligned to its purpose. The key metrics for the Credit Risk Assessment module (M1) are precision and recall for default prediction, AUC-ROC on held out test sets, and Gini coefficient as a discrimination measure. Chang et al. (2024) showed that ensemble-based classifiers often under-perform on credit card default datasets with AUC-ROC > 0.92, hence this serves as a natural benchmark of performance to evaluate the M1 implementation. Furthermore, since the outputs of modules are used as parameters to the downstream modules (as opposed to classifications), calibration quality—the degree to which the predicted risk scores reflect the actual default rate—is also considered crucial. A poorly discriminating and well calibrated credit risk module will systematically miss-configure the downstream utility function.

The primary metrics used for the Cash-Flow Forecasting module (M2) are root mean squared error (RMSE) versus actuals for a rolling 12-month evaluation window, and the mean absolute percentage error (MAPE) for interpretability, as well as coverage of irregular expenses (the percentage of non-recurring outflows, such as annual insurance premiums, seasonal utility spikes, and irregular debt repayments that the LSTM forecasts fall within a specified tolerance band). Roy et al. (2025) found LSTM models yield around 30–50 percent lower RMSE compared to the baselines ARIMA models on personal cash-flow data, which is the minimum acceptable performance threshold for M2. The coverage metric is important because irregular expenses are just the expenses that kids most often miss, and they will not be reliable when they pass the liquidity constraint on to M3.

In the Portfolio Optimisation module (M3), the assessment focuses on the Sharpe ratio of the RL agent's investments under simulated market conditions, maximum drawdown under stress scenarios that are calibrated to the historical crises (2008-09 Global Financial Crisis, 2020 COVID-19 market shock), and liquidity constraints that are passed from M2. The final criterion is a metric of constraint satisfaction – one should not be rewarded for a policy that produces better risk-adjusted returns, but that does more than satisfy the liquidity buffer constraint. A first-order criterion for the evaluation design of M3 should be constraint satisfaction because real-world portfolio RL implementations often fail to do this, as pointed out by Hambly et al. (2023). In the Recommendation module (M4), the assessment focus changes to behavioural outcomes, which include: the user's score based on the advice he or she comprehended from the presented advice; acceptance rate in controlled testing environments (in the module M4); and qualitative assessment of user comprehension of the explanation of a trade-off, especially the debt-versus-investment trade-off.

6.2 System-Level Metrics

The performance of the individual modules is not sufficient for the assessment of an integrated system. The question of central systems is whether the integration results in outcomes that would not result from the operation of any single module, or any pair of modules working separately. The measure of interest here is not the overall outcome quality of the advisory, but the quality of the outcome resulting from the integration, as compared to a baseline in which the modules run separately on the same household data.

A household with a credit risk coefficient rc that leads to a conservative utility initialisation in M3 should have measurably different and more appropriate financial outcomes compared to a similar household initialised using a generic utility initialisation. However, if no significant difference is found, this means that the Personalization Cascade is not working as it is supposed to, and the integration is not supported. Thus, the evaluation design must integrate a controlled comparison condition without credit risk outputs parameterising downstream modules which is to say without cascading. The difference between these two conditions is the empirical test of Contribution 1. Other indicators at the system level involve advice consistency over time (whether the system's recommendations are consistent when the household's financial situation changes), and conflict resolution frequency (how often the system's conservative override is called, and whether the liquidity outcomes of households experiencing the conservative override are better than those of households not experiencing it).

6.3 Illustrative User Personas

Generic archetypes do not demonstrate the framework’s discriminating power. The following two concrete personas illustrate how the four-module cascade produces structurally different advice for households whose surface-level financial goals are identical but whose constraint profiles are sharply different.

***Persona A** is 31 years old and has a monthly income of INR 85,000, a credit score of 618, and a personal loan debt of INR 2.4 lakh with an annual interest rate of 14.5%. She wants to invest in equity mutual funds via a systematic investment plan (SIP). Her risk coefficient (rc) is estimated to be 0.74, and her debt-to-income ratio (d) is estimated to be 0.31, after processing through M1. M2 expects the net disposable income after debt servicing to be around INR 14,200 for the month, and a liquidity stress event in months 5 and 8, resulting from an annual insurance premium and an imminent acceleration of loan instalments. M3 suggests a monthly SIP amount of INR 4,000, which is a significant dip from Persona A's desired INR 10,000, and the remaining funds can be channeled to the accelerated loan repayment. M4, in this instance, makes this trade-off quite explicit, highlighting that it generates a risk-free return of 14.5% on every loanee's rupee invested in accelerated repayment, which is superior to the expected post-tax equity SIP return in Persona A's constraint envelope. The system also identifies 5-months and 8-months as SIP pause windows to not force redemption.*

***Persona B** is 29 years old and has INR 82,000 of monthly income, a credit score of 791, and no other personal loan or debts except for personal loan EMI of INR 18,000 at 8.4% interest. She also wishes to start an investment of INR 10,000 per month in SIP.She is also looking to start an investment of INR 10000 per month in SIP. M1 estimates rc = 0.19 and d = 0.22. M2* *has a monthly surplus of around INR 28,000 with no events of liquidity stress over a span of twelve months. M3, which has an initial SIP of INR 10,000, suggests a diversified portfolio comprising of INR 60 per cent in large-cap index funds and INR 40 per cent in mid-cap funds. M4 does not have the trade-off warning as the home loan interest rate is less than the equity returns that Persona B can expect from the target corpus growth, and instead emphasizes on the goal convergence – based on projected corpus growth towards the stated goal of INR 50 lakh in 10 years from now. Both personas have the same stated goal and have very similar incomes. It is important to understand that the structure of the advice changes at each layer rather than just being a scaled version of the structure of the advice produced at the previous layer, since the credit risk profile creates a true different constraint envelope for each household.*

6.4 Comparison with Existing Robo-Advisory Systems

Table 6.1 provides a structured capability comparison between the proposed framework and four representative existing systems: Betterment and Wealthfront (US-based, among the most feature-complete robo-advisors globally), Scripbox (India-based, representing the domestic market context in which the framework's persona examples are situated), and Groww (India-based, representing a newer generation of mobile-first retail investment platforms). The comparison is organised on eight binary dimensions corresponding to the framework's core design claims. No existing system satisfies more than two of the eight criteria. The proposed framework is designed to satisfy all eight simultaneously, which constitutes its primary differentiating claim relative to the existing commercial landscape. Wexler and Oberlander (2021) noted that the gap between algorithm-driven and human advisory quality is most pronounced precisely on dimensions of credit-awareness, debt integration, and liquidity-constrained advice — the three dimensions on which all four comparison systems score negatively in this table.

Discussion

7.1 Theoretical Implications

The framework makes a theoretical argument that cuts across household finance, machine learning, and decision theory. Its central claim is that credit risk is not merely a lending instrument — it is a behavioural fingerprint that should inform all dimensions of financial advice delivered to a household. If this claim holds under empirical testing, it carries implications beyond the design of robo-advisory systems.

In the household finance literature, the dominant theoretical thread since Campbell (2006) has been the identification of costly errors and the conditions under which they arise. What has been less developed is a theoretical account of how the information structure of a household's financial situation should determine the architecture of the advice it receives. The Personalization Cascade offers precisely such an account: it argues that the right unit of personalisation is not a risk-tolerance category derived from a questionnaire but a behavioural constraint envelope derived from credit and cash-flow data. This repositions credit risk from a siloed lending tool to a general-purpose personalisation signal — a theoretical claim with significant implications for how personalisation is conceptualised in financial advisory AI more broadly.

In the machine learning literature, the framework's use of credit risk outputs as state variables in an RL reward function is a novel architectural choice that bridges two research programmes — credit scoring and portfolio optimisation — that Vučković et al. (2025) identified as developing in parallel without cross-domain integration. The Credit-Risk-Informed Utility Initialization formalises this bridge, providing a theoretically grounded mechanism for cold-start personalisation that requires neither interaction history nor self-reported preferences. If the utility initialisation approach proves empirically sound, it constitutes a general-purpose design pattern applicable beyond financial advisory AI to any recommender domain where behavioural constraint data is available at onboarding.

7.2 Practical Implications

Three practitioner audiences stand to draw distinct lessons from this framework. For fintech companies and platform developers, the integration architecture described in Section 5 is technically constructible using existing API infrastructure. Credit bureau data is available through established data partnerships; open banking transaction flows are accessible under PSD2 in the UK and under equivalent regulatory frameworks in India and the EU; RL libraries (Stable-Baselines3, RLlib) provide production-grade implementations of the PPO and SAC agents suited to M3. The principal obstacle is not technical — it is organisational. The data-sharing model required for the Personalization Cascade to function (credit data flowing into a portfolio engine; cash-flow data flowing into a recommendation layer) requires integration across systems that are typically owned by separate business units or separate firms. Reher and Sokolinski (2024) noted that the most significant barriers to improving robo-advisory personalisation are not algorithmic but structural, relating to data access, regulatory fragmentation, and the absence of cross-domain identity resolution. The framework is a blueprint for what should be built; the organisational architecture required to build it is a separate and significant challenge.

For regulators, integrated advice creates new accountability questions that existing frameworks are not designed to answer. Current financial advice regulation attributes responsibility to the advice output — the recommendation delivered to the client. In an integrated system, advice quality depends on the quality of upstream module outputs: if M1 produces a biased credit risk score (a known problem with ML-based credit models, particularly for minority populations and thin-file borrowers; see Nallakaruppan et al., 2024), the cascade propagates that bias through M2, M3, and M4. The household receives advice that is systematically miscalibrated without any single module having produced an overtly erroneous output. Regulators must therefore develop standards for module-level accountability alongside advice-level accountability — a requirement that existing frameworks for algorithmic financial advice (the EU's AI Act, the SEC's robo-advisor guidance) do not currently address at the architectural level. Zhu et al. (2024) have highlighted that explainability requirements for next-generation advisory systems must extend to the full decision chain, not merely the final output — a position the integrated architecture of this framework makes technically tractable through the cascade's structured information flow.

For households themselves, the framework's principal benefit is not portfolio optimisation — it is advice calibrated to what they can actually afford to do. The most common failure mode of existing robo-advisors is not that they construct bad portfolios; it is that they construct portfolios that are good in isolation but incompatible with the household's full financial situation. A household that is recommended to invest ₹10,000 per month while carrying high-interest debt at 14.5% is not receiving bad portfolio advice — the recommended portfolio allocation may be perfectly sound. It is receiving advice that ignores the household's balance sheet. The Personalization Cascade is designed to prevent exactly this class of failure.

7.3 Limitations

The restrictions of the framework should be clearly expressed before reviewers provide them. First, there are assumptions regarding the interdependencies of the modules, but these are not proven. The Personalization Cascade is a theoretical statement on how information should be cascaded through the modules to compound personalisation but is not proven in practice. Credit risk outputs have been assumed to act as a parameterisation when operating downstream modules to improve household outcomes, but not tested against real household data. This is the basic empirical deficiency that the simulation study and field trial described in Section 7.4 will try to correct.

Second, instability has been documented in real world markets for RL-based portfolio optimisation. Malibari et al. (2023) studied the issues encountered with the deployment of RL financial agents and discovered that while the agent worked extremely well in simulation, it failed to work out well in the real world, mainly because of distribution shift, reward sparsity and the non-stationarity of financial time series. All of these challenges are carried over from the RL agent used in M3. This is addressed in part by the evaluation framework in Section 6, which introduces stress scenario testing, but does not address the fundamental gap between simulated and live performance, which Hambly et al. (2023) identified as the field's most persistent open problem.

Third, the cold-start mechanism makes the assumption that credit data is available and correct and is also unbiased. In reality, all of these assumptions are not consistently true. The credit experience of thin-file borrowers -- young people, recent migrants, and those earning less -- may be too thin to allow for a reliable rc to be estimated. Nwaimo et al. (2024) documented that credit models trained using machine learning techniques give systematically lower scores to those who are under-banked, even if they possess the same true risk, which would result in an overly conservative utility landscape being assigned to households who do not, in fact, face constraints due to their credit history. There is currently no way to identify or correct this source of bias in the framework, so it is a major constraint both for fairness and for practical accuracy.

Fourth, the cash-flow forecasting module is the most data dependent of the frameworks and the open banking transaction streams from which it depends are not necessarily available in all markets. The dynamic of PSD2 in the UK and Account Aggregators in India offers the infrastructure in those scenarios, which many other markets lack. The cash-flow integration *feature of the framework is as such geographically constrained by the regulatory data infrastructure itself and hence is not immediately portable.*

7.4 Future Research Directions

The empirical validation programme is a staged design. The first stage is a simulation study that uses synthetic household information generated from two well-known datasets: a dataset of household financial behaviour over decades, the Panel Study of Income Dynamics (PSID); and a credit bureau synthetic dataset, which can be used to simulate a variation in credit risk profiles without losing the correlation structure of the various credit features. The core integration claim is that households processed through the full Personalization Cascade have improved financial outcomes (measured as greater wealth accumulation, lower debt burden and fewer liquidity stress events over a five year timeframe) compared to households processed through the independent modules.

The second phase is a field trial in collaboration with a fintech platform within India's Account Aggregator framework that will enable open banking data infrastructure for M2. The field trial features a randomized design, with exposure to treatment advice (advice generated from the full integrated framework) versus control (advice generated from a standard portfolio-only robo-advisor) conditions. Net financial position at 12 months, while holding other baseline income, debt-load and market conditions constant, is the primary outcome measure. Secondary outcome measures are the rates of receiving advice, persistence in SIP, and self-reported financial confidence scores. UK open banking data under PSD2 adds an extra context to the cash-flow forecasting module, specifically allowing for the cross-market comparison of M2's performance under varying transaction data availability scenarios.

In addition to validation, three important research extensions are of particular interest. First, there is a methodological challenge in the form of the bias correction problem discussed in Section 7.3, which needs special attention: What changes should be made to the Credit-Risk-Informed Utility Initialization to prevent systematic under-scoring of thin-file borrowers from being passed on through the cascade? The works in fairness-aware credit scoring literature (Nwaimo et al., 2024; Hlongwane et al., 2024) offer potential solutions, but there has been no investigation of how those solutions fit into the utility initialisation mechanism. Second, the framework's recommendation layer currently gives advice in the form of a deterministic plan; there is a design extension to incorporate uncertainty communication, i.e., include the confidence intervals for M2's cash-flow forecasts in the advice given to the household, that has theoretical and regulatory implications. Third, it would be interesting to explore whether structured questionnaires to elicit inputs for M4 could be replaced by LLMs interfaces (as proposed in the LLM-PPR architecture; arXiv:2512.12922, 2024), especially for households with too many elements to encode accurately in the framework's structured input formats.

Conclusion

This paper introduces a four-module integrated framework that is designed to alleviate this disconnect from the credit risk, cash-flow limitations, and debt loadings affecting the financial lives of the households to which robo-advisors are deployed, with a credit risk assessment module providing the outputs to parameterise all of the other modules. The outcome is counsel that is not what a household desires, but what they afford.

The framework makes two contributions that together form the theoretical claim of the framework. The Personalization Cascade is a formalization of this directed information flow, where the entire credit risk profile of the household is a major part of the input to the system, and every module downsteam should be treated as if it were a module with a personal credit risk profile. Credit-Risk-Informed Utility Initialization is an operationalisation of this cascade at the time of the cold-start: before the first investment interaction, the credit risk coefficient rc is used to parameterize the utility function of the RL portfolio optimiser, thereby avoiding the often unreliability self-reported preferences in the investment context that the literature has consistently identified. The cascade combined with the initialisation mechanism represent a theoretically sound solution to a practical problem in financial advisory AI, which is the lack of a unified personalisation architecture that works across credit, cash-flow and portfolio domains, something that was identified by Vučković et al. (2025) and Sharaf et al. (2022) independently as the most persistent unresolved challenge.

The remainder of work -- the simulation studies, the field trial and the bias correction work for thin-file borrowers -- is considerable. But those results are not the issue; the direction it points is the issue. If personalisation becomes the common standard for credit risk in financial advice, then the household population that is currently getting the worst of financial advisory AI may finally get financial advice that's more suited to their financial situation than an idealised version of an idealised client. This is the issue this framework addresses.

References

Abi, R. (2025). Machine learning for credit scoring and loan default prediction using behavioral and transactional financial data. *World Journal of Advanced Research and Reviews, 26*(3), 884–904.

Adam, M., Benlian, A., & Wiethof, A. (2019). Robo-advisors as a digital financial service: A systematic literature review. *Business & Information Systems Engineering, 61*(4), 401–420.

Adam, M., Wessel, M., & Benlian, A. (2019). AI-based chatbots in customer service and their effects on user compliance. *Electronic Markets, 31*(2), 427–445.

Bashir, A., Iqbal, M., & Farooqi, R. (2023). Mapping financial advice research for retail investors: A bibliometric and systematic review. *International Review of Financial Analysis, 87*, 102601.

Belanche, D., Casaló, L. V., & Flavián, C. (2025). User trust and engagement with robo-advisors: A longitudinal analysis. *Journal of Business Research, 168*, 114228.

Benavent, C. (2021). Behavioral abandonment of robo-advisory platforms after market downturns. *Journal of Behavioral Finance*, *22*(4), 349–362.

Campbell, J. Y. (2006). Household finance. *The Journal of Finance, 61*(4), 1553–1604. https://doi.org/10.1111/j.1540-6261.2006.00883.x

Chang, V., Li, T., & Xu, Q. (2024). Credit risk prediction using machine learning and deep learning: A study on credit card customers. *Risks, 12*(11), 174. https://doi.org/10.3390/risks12110174

Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 785–794). Association for Computing Machinery.

Cui, Y., Cui, Z., & Yang, J. (2024). Deep reinforcement learning for portfolio selection. *Finance Research Letters, 60*, 104830. https://doi.org/10.1016/j.frl.2024.104830

Day, M.-Y., Yang, C.-Y., & Ni, Y. (2024). Portfolio dynamic trading strategies using deep reinforcement learning. *Soft Computing, 28*(15), 8715–8730.

Gregor, S., & Hevner, A. R. (2013). Positioning and presenting design science research for maximum impact. *MIS Quarterly, 37*(2), 337–355.

Hambly, B., Xu, R., & Yang, H. (2023). Recent advances in reinforcement learning in finance. *Mathematical Finance, 33*(3), 437–503. https://doi.org/10.1111/mafi.12382

Helms, T., Falk, T., Buesing, M., & Moeller, K. (2022). Robo-advisors: A systematic literature review. *Finance Research Letters, 47*, 102941. https://doi.org/10.1016/j.frl.2022.102941

Hlongwane, R., Ramaboa, K. K. K. M., & Mongwe, W. (2024). Enhancing credit scoring accuracy with a comprehensive evaluation of alternative data. *PLOS ONE, 19*(5), e0303566.

Imaki, S., Imajo, K., Ito, K., & Nakagawa, K. (2025). Portfolio optimization using deep learning with risk aversion utility function. *Finance Research Letters, 72*, 106547. https://doi.org/10.1016/j.frl.2025.106547

Jiang, Z., Xu, D., & Liang, J. (2017). A deep reinforcement learning framework for the financial portfolio management problem. arXiv preprint arXiv:1706.10059.

Lam, X. N., Vu, T., Le, T. D., & Duong, A. D. (2008). Addressing cold-start problem in recommendation systems. In *Proceedings of the 2nd International Conference on Ubiquitous Information Management and Communication* (pp. 208–211).

Leukam, F., Djagba, P., et al. (2025). Reinforcement learning for portfolio optimization with a financial goal and defined time horizons. arXiv preprint arXiv:2511.18076.

Luo, Z., Vanden Broucke, S., Baesens, B., Óskarsdóttir, M., & Lemahieu, W. (2024). Behavioral biases in robo-advisory algorithm design. *European Journal of Operational Research, 312*(1), 100–115.

Lusardi, A., & Mitchell, O. S. (2014). The economic importance of financial literacy: Theory and evidence. *Journal of Economic Literature, 52*(1), 5–44. https://doi.org/10.1257/jel.52.1.5

Malibari, M. A., et al. (2023). Challenges in applying reinforcement learning to financial time series: Non-stationarity, reward design, and exploration. *Expert Systems with Applications, 214*, 119124.

Markov, A., Seleznyova, Z., & Lapshin, V. (2022). Credit scoring methods: Latest trends and points to consider. *Journal of Finance and Data Science, 8*, 180–201. https://doi.org/10.1016/j.jfds.2022.07.002

Moody, J., & Saffell, M. (1998). Reinforcement learning for trading systems and portfolios. In *Proceedings of the International Conference on Computational Intelligence for Financial Engineering* (pp. 279–286).

Nallakaruppan, M. K., Balusamy, B., Shri, M. L., Malathi, V., & Bhattacharyya, S. (2024). An explainable AI framework for credit evaluation and analysis. *Applied Soft Computing, 153*, 111307.

Nayak, S. (2023). Developing predictive models for financial stability: Integrating behavioral analytics into credit risk management. *Journal of Artificial Intelligence and Cloud Computing, 2*(1), 1–12.

Noble, S. M., & Mende, M. (2023). The human-AI advisory boundary in financial services. *Journal of Service Research, 26*(2), 188–205.

Nourallah, M. (2023). No one size fits all: Exploring young retail investors' initial trust in financial robo-advisors. *Journal of Business Research, 156*, 113470.

Nwaimo, C. S., Adegbola, A. E., & Adegbola, M. D. (2024). Predictive analytics for financial inclusion: Using machine learning to improve credit access for under-banked populations. *Computer Science & IT Research Journal, 5*(6), 1358–1373.

OECD. (2023). *OECD/INFE 2023 International Survey of Adult Financial Literacy*. OECD Business and Finance Policy Papers, No. 39. OECD Publishing.

Oliver Wyman Forum. (2024). *Generative AI can make personal finance more personal*. Oliver Wyman Forum Global Survey Report.

Puiutta, E., & Veith, E. M. S. P. (2020). Explainable reinforcement learning: A survey. In *Machine Learning and Knowledge Extraction* (pp. 77–95). Springer.

Reher, M., & Sokolinski, S. (2024). Robo-advisors and financial inclusion: Evidence from digital onboarding. *Journal of Financial Economics, 148*(2), 310–335.

Roa, L., Correa-Bahnsen, A., Suarez, G., Cortés-Tejada, F., Luque, M. A., & Bravo, C. (2021). Super-app behavioral patterns in credit risk models: Financial, statistical and regulatory implications. *Expert Systems with Applications, 169*, 114486.

Roy, S., Polley, S., De, S., Gangwal, C., & Mitra, C. (2025). Harnessing AI and machine learning for improved cash flow forecasting. *Proceedings of the 6th Doctoral Symposium on Intelligence Enabled Research (DoSIER 2024)*. CEUR-WS. https://ceur-ws.org/Vol-3900/Paper17.pdf

Sharaf, M., Hemdan, E. E.-D., El-Sayed, A., & El-Bahnasawy, N. A. (2022). A survey on recommendation systems for financial services. *Multimedia Tools and Applications, 81*(12), 16761–16781. https://doi.org/10.1007/s11042-022-12564-1

Tao, R., Su, C.-W., Xiao, Y., Dai, K., & Khalid, F. (2021). Robo advisors, algorithmic trading and investment management: Wonders of fourth industrial revolution in financial markets. *Technological Forecasting and Social Change, 163*, 120421.

Vučković, D. B., Dekpo-Adza, S., & Matović, S. (2025). AI integration in financial services: A systematic review of trends and regulatory challenges. *Humanities and Social Sciences Communications, 12*, Article 562. https://doi.org/10.1057/s41599-025-04850-8

Wang, Z. (2026). Robo-advisors in financial services: Redefining wealth management in the age of artificial intelligence. *Financial Innovations*. Advance online publication.

Wexler, J., & Oberlander, J. (2021). Mind the gap: Challenges in algorithmic and human decision-making for financial health. In *Proceedings of the ACM Conference on Fairness, Accountability, and Transparency* (pp. 327–337). Association for Computing Machinery.

Zhang, X., & Yu, L. (2024). Consumer credit risk assessment: A review. *Finance Research Letters, 62*, 105253.

Zhang, X., & Yu, L. (2024). Consumer credit risk assessment: A review from state-of-the-art classification algorithms, data traits, and multi-scenario modelling. *Applied Soft Computing, 167*, 112023.

Zheng, J., & Tu, Y. (2023). Corporate cash flow forecasting using machine learning: A systematic review. *Journal of Finance and Data Science, 9*, 100094.

Zhu, Y., Inavolu, S., & Challoumis, C. (2024). Explainability, bias, and consumer protection in next-generation robo-advisory systems. *Journal of Financial Regulation, 10*(1), 45–73.