Interpretable Simulation-Optimization for Dynamic Balance Sheet Management

Pedro Júdice[1], Luís Pinto[2*], José Luís Santos[2]

[1]Instituto Universitário de Lisboa (ISCTE-IUL), Business Research Unit (BRU-IUL), Avenida das Forças Armadas, Lisbon, 1649-026, Portugal.
[2]University of Coimbra, CMUC, Department of Mathematics, Apartado 3008, EC Santa Cruz, Coimbra, 3001-501, Portugal.

*Corresponding author(s). E-mail(s): luisp@mat.uc.pt;
Contributing authors: Pedro.Judice@iscte-iul.pt; zeluis@mat.uc.pt;

Abstract

The management of bank balance sheets is a major determinant of banks' long-term performance and sustainability. Therefore, it is one of the most critical problems in bank management. Using autoregressive processes for stochastic deposit flows and credit losses, and vector autoregressive processes for interest rates, we develop a simulation-optimization method that devises balance sheet policy functions based on state variables such as interest rates or leverage levels. Performance analysis on an independent testing set shows that the algorithm outperforms other established methodologies, delivering an increase in the return-to-risk ratio of up to 68% when compared to optimizing time-independent static allocations and an increase greater than 593% when compared to equal-weight and 60/40 policies. In addition to performance, we emphasize interpretability, developing an algorithm that allows the analysis of the parametric allocations and interpretation, and can thus be used in practice to support banks' senior management dynamic decisions.

Keywords: balance sheet management, economic environment simulation, dynamic model, interpretability, bi-objective Pareto optimization, policy function.

1 Introduction

1.1 Motivation

The management of bank balance sheets is one of the most important practical problems faced today. According to statista.com, the total value of worldwide banking assets is 180 trillion dollars, making managing those assets one of the most critical problems for financial institutions. The task is not easy, as banks have to carefully balance the return and the risk of failure due to liquidity or solvency-driven events. Also, banks should optimize the balance sheet at each point in time as a function of the economic environment. If a bank takes on too much risk, it will most likely fail when the first financial crisis hits the door. If the bank takes little or no risk, it cannot generate enough return to pay its employees and costs and is also destined to fail. The sweet spot is between taking no risk and excessive risks.

The risks for the bank can be manifold. A bank is exposed to runs from creditors (liquidity risk), such as those witnessed in 2023 in the Silicon Valley Bank. Interest rate rises penalize fixed-rate assets (interest rate risk), as recorded during the Savings and Loan Crisis in the early eighties. Bear markets penalize stock market investments (market risk). Spikes in defaults force banks to absorb high losses (credit risk); the Subprime Crisis is a good illustration of credit risk. Risk-taking also entails many options, as bank managers have to carefully select the bank strategy while also taking advantage of the economic environment and the different prospective returns on each asset.

The balance sheet management problem is in its essence a dynamic problem, in that allocations should adjust to the economic environment, for example interest rates. Also, the dynamics for the risk factors for banks should include long-term features of these dynamics, such as the autocorrelations present in the stochastic processes for default risk, liquidity risk and interest rate risk, as well as the interdependency of wholesale and retail interest rates. Due to the number of state variables involved, the purpose of our research is to establish a simulation-optimization method to determine balance sheet policy functions on environment variables, assuming stochastic deposit flows, stochastic credit losses, and stochastic interest rates.

Our research also focuses on interpretability, and we look to develop a dynamic balance sheet policy model that weighs performance and complexity. The lack of interpretability is a major barrier hindering the adoption of intelligent decision support systems in areas such as banking management (Doumpos, Zopounidis, Gounopoulos, Platanakis, & Zhang, 2023; Kaya, 2019). Without the possibility of expert validation, artificial intelligence generally lacks credibility and trust. Human oversight of the decision-making processes is critical, at least at the command level, to risk-manage the outputs given by the system. Interpretable models are thus a big help for their effective administration, as a model validator will understand the model and gain confidence in the output decision of the system, even in unforeseen circumstances. An interpretable decision model is also easier to communicate to senior management than a black-box model, making the decision-making process more secure.

Finally, an optimal strategy should be tested on a separate testing dataset, validating the results that have been used in the training dataset and give confidence on the methodology that is proposed.

1.2 Related literature

In this section we survey the literature on bank balance sheet management and related literature. Bank balance sheet management is related to portfolio theory (Markowitz (1952)) and asset-liability management models from insurance and pension funds, although it has fundamental differences. First, in classical portfolio theory and in pension fund insurance ALM, assets are marked-to-market; in contrast, the bulk of bank assets are recorded at book value, making earnings less volatile in this setting. Second, classical portfolio theory and pension fund and insurance ALM focus on the market risk of assets, whereas in banking credit risk plays a much more important role. Third, the nature of liabilities is very different. In classical portfolio theory, usually one does not consider a liability structure; in pension fund and insurance ALM, liabilities are stable and long-dated. In contrast, in bank management liabilities are short-term and volatile; funding liquidity risk (the risk of bank runs) plays a much more important role. Whereas in pension fund and insurance ALM one has stable liabilities to invest in mostly liquid assets, in bank management liabilities are short-dated and volatile and finance long-term illiquid assets (such as mortgages and corporate loans). Fourth, whereas in classical portfolio theory and pension fund ALM one allows for a full rebalancing of the portfolio (even if such rebalancing entails transaction costs), in bank balance sheet management the allocations to loans have to evolve in a smooth way, since a bank cannot immediately dispose of such assets in an immediate way.

For the sake of completeness, we briefly review some of the literature in portfolio theory, pension fund and insurance ALM, and bank balance sheet management. We will focus more on the latter, since it is the subject of the research we present. It will be impossible to make a thorough review on all of these subjects, so we highlight some of the landmark papers.

Portfolio theory has its roots in Markowitz (1952), who laid the framework to be used in several settings. Another landmark model is the consumption-investment model for and individual investor by

Merton (1969). Campbell and Viceira (2002) contains several examples of applications of portfolio theory to the individual investor case. In the pension fund and ALM case, Ziemba and Mulvey (1998) is a comprehensive collection of research papers in this field. Some landmark models include G.C. Boender (1997), Cariño et al. (1994), Consigli and Dempster (1998), Hibiki (2006), Kouwenberg (2001), Mulvey and Vladimirou (1989), Yu, Tsai, and Huang (2010), and Zenios (1995).

As we mentioned before, there are fundamental differences between the bank balance sheet management problem and the pension fund ALM problem. As a consequence, bank balance sheet models have their own literature that we now describe in more detail. Again, we will not be able to list all the papers in this very rich area, but we will highlight some of the output in this area. We distribute our review between static and dynamic models. Static models assume constant policies within a certain time horizon, whereas dynamic policies change as a function of the evolving state variables.

We start with static models. Kosmidou and Zopounidis (2004) used goal programming to determine optimal strategies, while assessing the sensitivity of the solution to different interest rate scenarios. Schmaltz, Pokutta, Heidorn, and Andrae (2014) optimize a balance sheet assuming compliance with Basel III capital requirements. Júdice and Zhu (2021) optimize balance sheets using duality, and making the link between optimal balance sheet and shadow prices of interest and credit risk. The approach of static balance sheets is further developed for more general risk measures in Maier-Paape, Júdice, Platen, and Zhu (2023). Júdice, Pinto, and Santos (2021) conducted static optimization for a long-term horizon. We also highlight Yan, Zhang, and Wang (2020), who used robust optimization for credit portfolios, and Sirignano, Tsoukalas, and Giesecke (2016), who conducted the optimization for large-scale credit portfolios.

Concerning dynamic optimization, i.e., optimizing balance sheet strategies that change over time, we start by referring the landmark model of Kusy and Ziemba (1986), which is a stochastic programming approach, using a three-period, two-point distribution of the joint evolution of interest rates and liquidity flows, although it does not consider credit risk. Oğuzsoy and Güven (1997) is also a stochastic programming model, assuming a random variable with three outcomes for outstanding deposits, although it does not specify interest rate and credit risk processes. Hałaj (2016) posits a dynamic model with liquidity risk, but assumes that interest rates are deterministic, and implements the model over a two-period horizon. We also highlight Mukuddem-Petersen and Petersen (2006), who use continuous dynamic programming, which comes with the cost that the setting has to be simplified. Specifically, the authors do not account for stochastic liquidity flows (liquidity risk) and assume independent lognormal returns for loans, whereas the returns for loans tend to be autocorrelated.

As a summary, most of the research on dynamic balance sheet strategies falls into the realm of stochastic or dynamic programming. To devise optimal balance sheet policies, the methods typically sacrifice the distributional assumptions of the risk factors, due to the increase in complexity from dynamic programs with high dimensions (the curse of dimensionality). Also, except for Mukuddem-Petersen and Petersen (2006), the dynamic policy functions are more difficult to interpret, since they do not give a parametric description between the state variables and the optimal strategy.

1.3 Contributions

In this section, we highlight the contributions of our research and how they improve the state of the art in bank balance sheet management.

First, using simulation optimization, we devise an interpretable policy function in a setting that encompasses interest rate, credit, and liquidity risk. Mukuddem-Petersen and Petersen (2006) has devised an intepretable policy function in a simplified setting that overlooks liquidity risk and assumes lognormal returns for loans. Our approach relies on improved distributional processes for credit risk, taking into account autocorrelations, the dependencies between retail and market interest rates via an autoregressive process, and a stochastic process for liquidity risk that includes autocorrelation. Compared to the stochastic programming approaches, such as Kusy and Ziemba (1986) or Oğuzsoy and Güven (1997), our parametric approach allows us to identify the relationship between the policy values and the state variables. Having a parsimonious relation between the state variables and the policy function mitigates the risk of overfitting, as validated in our tests.

Second, as mentioned before, we emphasize the difference of the bank balance sheet setting and the classical portfolio, pension fund, and insurance ALM. We highlighted the differences in the previous section. As a consequence, our model differs significantly from these settings, since it has the features that are inherent to bank management.

Third, compared to previous dynamic models in bank ALM, we use improved distributional assumptions. As we highlighted in the literature review, settings such as Kusy and Ziemba (1986) and Oğuzsoy and Güven (1997) use point distributions. Hałaj (2016) does not consider stochastic interest rates and his model is developed over two periods. Mukuddem-Petersen and Petersen (2006) uses a continuous time setting, so has to engage in simplification such as assuming that deposit funds are fully invested in cash, the optimization focusing on the securities part and not on the loan allocation, and not specifying dynamics for the deposit outflows or interest rates. Hałaj (2016) uses a two-period model and does not assume stochasticicty in the interest rates. The dynamic model we present is based on stochastic processes for credit losses, market and deposit interest rates, and liquidity outflows. All these processes account for autocorrelation.

Fourth, compared to previous dynamic approaches we cited in the literature, we test the results in a separate dataset, allowing us to validate the approach. To the best of our knowledge, out-of-sample testing has been conducted for static models (Coelho, Santos, and Júdice (2023)), but not for dynamic balance sheet models. This validation is important, as overfitting in the training phase often leads to poor results in separate testing datasets (Bailey, Borwein, López de Prado, and Zhu (2014)). We use a testing dataset to evaluate the method and test it against several established allocation methodologies, and confirm the superior performance of the model, of 42% against a static optimization method and more than 1200% against equal-weight and 60/40 heuristic allocations, if the bank starts with a leverage ratio of 5%. If the bank starts with a higher leverage ratio of 15%, the dynamic optimization model outperforms the static optimization model by 68% and the tested heuristic models by more than 593%. As we will see in Sections 3.5 and 3.6, these striking differences stem from the fact that, in a leveraged setting, the amount of capital to absorb losses is low, thus penalizing strategies that overweigh risky assets.

1.4 Organization

We organize the paper as follows. Section 2 describes the proposed methodology; we present the dynamic balance sheet model and briefly overview the economic risk factors generator and the balance sheet model. Section 3 is devoted to numerical experiments, including hyperparameter optimization and training of the tuned dynamic allocation model. This section also presents the simulated trajectories that comprise our dataset. Still in Section 3, we examine the performance of our model through a series of computational tests and compare the model results to alternative policies using a testing dataset. Section 4 concludes.

2 Proposed methodology

As illustrated in Fig. 1, we formulate a bi-objective policy function optimization model that makes the trade-off between return and risk. The optimization framework relies on a bank balance sheet simulator, which computes the return and risk metrics for a given dynamic balance sheet model, over a multi-annual period. We assume that the bank can allocate assets across cash, loans, (10-year) bonds, and stocks, i.e., $\alpha = (\alpha_{cash}, \alpha_{loans}, \alpha_{bonds}, \alpha_{stocks})$, where α is a policy function of state variables. The bank balance sheet simulator is fed by a scenario generator for the relevant risk factors and formed by a suitable set of equations and regulatory restrictions governing the bank's dynamics.

In the following sections, we go through the four modules of our methodology: scenario generator that describes the economic environment, bank balance sheet simulator, dynamic bank policy model, and return and risk metrics.

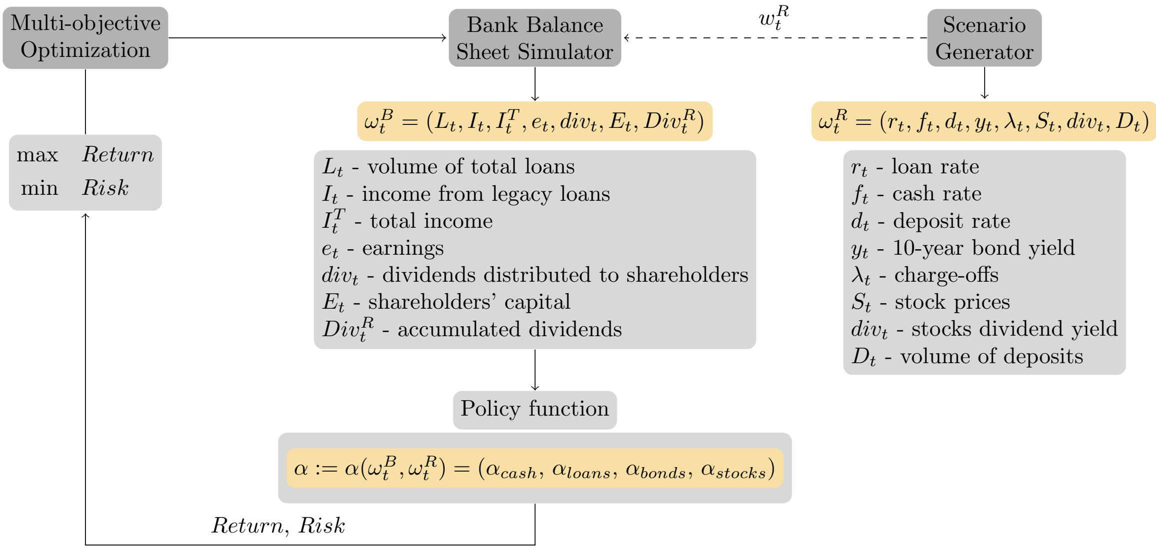

Fig. 1 Schematic representation of the proposed balance sheet model. The goal is to find the policy function $\alpha := \alpha(\omega_t^B, \omega_t^R)$ that maximizes return while minimizing risk. A bank balance sheet simulator, guided by an economic scenario generator and the policy function α, determines the return and risk metrics. The bank allocates assets to cash, loans, bonds, and stocks. The random variables w_t^R and w_t^B encapsulate the risk factors and the bank balance sheet evolution.

2.1 Risk factor scenario generator

The scenario generation for the exogenous risk factors (interest rates, stock prices, deposit volume dynamics, and credit losses) that describe the economic environment is based on the work of Costa, Faias, Júdice, and Mota (2020), who adapted the model presented by Birge and Júdice (2013) to include stochastic deposit volumes.

Risk factors	Parameter	Description	Model
Liquidity risk	D_t	volume of deposits	$D_{t+1} = c + aD_t + bm_t + \epsilon_t$
Interest rate risk	r_t	Interest rate on new loans	$X_{t+1}^* = AX_t^* + bm_t + \epsilon_{t+1}$,
	f_t	interest rate on cash	with $X_t^* = (r_t^*, f_t^*, d_t^*, y_t^*)$
	d_t	deposit rate	
	y_t	10-year bond yield	
Credit risk	λ_t	charge-offs/credit losses	$\lambda_{t+1}^* = c + a\lambda_t^* + bm_t + \epsilon_{t+1}$
Market risk	S_t	stock prices	$S_{t+1}^* = c + a\delta_t^* + \epsilon_{t+1}^S$
	div_t	stock dividends yields	$\delta_{t+1}^* = c + a\delta_t^* + \epsilon_{t+1}^\delta$

Table 1 Risk factors variables and respective models: deposit volumes, interest rates, credit losses, and stock prices. Here, c, b, a, and A represent parameters, m_t momentum terms, and ϵ_t residuals, that are specific to each stochastic process. We set $r_t^* = g(r_t) - \widehat{g(r_t)}$, where the hat denotes the long-term mean and $g(x) = \ln(x)$, for $0 < x < 1$, and $g(x) = x - 1$, for $x \geq 1$. We define similarly f_t^*, d_t^*, and y_t^*. We also set $\lambda_t^* = N^{-1}(\lambda_t)$, with N^{-1} the inverse of the standard normal cumulative distribution, $S_t^* = \ln((S_t + div_t)/S_{t-1})$, and $\delta_t^* = \ln(div_t/S_t)$, with div_t the dividends from stock market investments at time t. The interest rate model was calibrated using US public data from FRED (Federal Reserve Economic Data), from 1971 to 2016. The data for credit risk is also taken from the FRED database, starts at 1985 and ends at 2016. We used stock market data from the Robert Shiller Irrational Exuberance Database, from 1946 to 2016.

The generation of scenarios uses vector autoregressive processes, which confer realism to the framework and will be utilized for the long-term assessment of state-dependent bank policies. The risk factors, the corresponding models, and the variables used throughout this work are listed in Table 1. We now describe the risk factors in more detail. The scenario generator consists of four modules: liquidity risk, interest rate risk, credit risk, and market risk for stock positions.

The liquidity risk module is taken from Costa et al. (2020), who calibrate a panel data model to a sample of different banks. The model has an autoregressive component, and also includes a momentum variable m_t that accounts for the persistence in increases and decreases in deposits. The interest rate risk model is based on the work of Birge and Júdice (2013) and was further improved by Costa et al. (2020). It consists of a vector autoregressive process that captures the dependencies between market and retail rates, and also has a momentum component. The credit risk model was also devised by Birge and Júdice (2013), and accounts for autocorrelation and momentum in credit losses. The stock price model is based on a vector autoregressive process for stock prices and dividends, inspired on the work by Campbell and Viceira (2002). The reader can refer to the articles above for the rationale and tests associated with these models.

We simulate K trajectories for each of the eight risk factors described in Table 1. Let us further assume that the initial period for the bank is $t_0 = 1$ and that the bank has to manage its allocation decisions within the time frame $t = t_0, \ldots, T$. The risk factors are encapsulated in a random variable $\omega^R \in \mathbb{R}^{8\times K\times T}$, representing the economic environment that will feed the bank balance sheet simulator. By

$$\omega_t^R = (D_t,\, r_t,\, f_t,\, d_t,\, y_t,\, \lambda_t,\, S_t,\, div_t) \tag{1}$$

we denote the realization of the risk factors trajectories ω^R at time t.

2.2 Balance sheet simulator

The simulation model for the balance sheet generalizes the balance sheet simulators of Birge and Júdice (2014) and Júdice et al. (2021) to state-dependent policies, and is summarized in Table 2. The state

State variables	Parameter	Description	Equation
Loans	L_t	volume of total loans	$L_{t+1} = (L_t + \overline{\alpha}_{loans})(1 - p - \lambda_{t+1})$
	I_t	income obtained from legacy loans	$I_{t+1} = (I_t + r_t\overline{\alpha}_{loans})(1 - p - \lambda_{t+1})$
Income	I_t^T	total income from balance sheet at time t	$I_{t+1}^T = I_t + r_t\overline{\alpha}_{loans} + \overline{\alpha}_{cash} f_t - d_t D_t - \lambda_{t+1}(L_t + L_t^{new}) + \overline{\alpha}_{stocks} S_{t+1}^* + \overline{\alpha}_{bonds}(y_t + Dur(y_t) m_{t+1}^y)$
	e_t	earnings at time t	$e_{t+1} = I_{t+1}^T - c_{t+1}$
	Div_t	dividends distributed to bank shareholders at time t	$Div_{t+1} = \max(R_p e_{t+1}, 0)$
Earnings	E_t	amount of shareholder capital	$E_{t+1} = \max(E_t + e_{t+1} - Div_{t+1}, 0)$
	Div_t^R	accumulated dividends	$Div_{t+1}^R = (1 + f_t) Div_t^R + Div_{t+1}$

Table 2 Bank state variables and equations used to simulate the performance of a bank's policy α. We also define the equation $\overline{\alpha}_{cash} + (L_t + \overline{\alpha}_{loans}) + \overline{\alpha}_{bonds} + \overline{\alpha}_{stocks} = E_t + D_t$, which states that assets equal liabilities. In the total loans L_t and income I_t equations, $p = 0.1$ is the amortization ratio. In the total income I_t^T equation, we have $Dur(y_t) = 1/y_t(y_t + 1)^{10} - 1/y_t$, which is the modified duration for bonds, and in the earnings e_t equation, we have $c_{t+1} = c(E_t + D_t)$, which computes operating costs with $c = 0.015$. In the distributed dividends Div_t equation, $R_p = 0.5$ is the dividend payout ratio.

variables are loans (with dynamics for volume of total loans L_t and income from legacy loans I_t), income (with dynamics for total income from balance sheet I_t^T, earnings e_t, and dividends distributed to shareholders Div_t), and accumulated wealth (with dynamics for shareholders' capital E_t and accumulated dividends Div_t^R). We introduce a state-dependent policy function $\alpha = (\alpha_{cash}, \alpha_{loans}, \alpha_{bonds}, \alpha_{stocks})$ that determines the proportions of funds allocated to the different asset classes.

Given the policy function $\alpha = (\alpha_{cash}, \alpha_{loans}, \alpha_{bonds}, \alpha_{stocks})$, with values between 0 and 1, we allocate the proportions to absolute values (dollar or euro) based on the available funding; we define $\overline{\alpha}_{cash} = \alpha_{cash}(E_t + D_t)$, $\overline{\alpha}_{bonds} = \alpha_{bonds}(E_t + D_t)$, $\overline{\alpha}_{stocks} = \alpha_{stocks}(E_t + D_t)$, and $\overline{\alpha}_{loans} = \max(\alpha_{loans}(E_t + D_t) - L_t, 0)$. The most important equation is the one for E_t,

$$E_{t+1} = \max(E_t + e_{t+1} - Div_{t+1},\, 0), \tag{2}$$

which describes the capital position of the bank. Here, we set the initial capital at $E_0 = 0.05D_0$, with D_0 the initial volume of deposits. We also assume that a bank needs to be compliant with a Tier 1 capital ratio limit of at least 10%. The balance sheet evolution is encapsulated again by a random variable or a path ω^B, where ω_t^B is the realization of the balance sheet at time t under path ω^B. Specifically,

$$\omega_t^B = \left(L_t,\, I_t,\, I_t^T,\, e_t,\, div_t,\, E_t,\, Div_t^R\right). \tag{3}$$

For a given time t, the policy function α will depend solely on the information available at time t, i.e., it will disregard the information prior to time t. Therefore, it is a Markov decision. Specifically,

$$\alpha = \alpha(\omega_t^R, \omega_t^B), \tag{4}$$

i.e., it will depend on the realizations of the risk factors ω^R and the balance sheet state variables ω^B at time t. The model used to parameterize the policy function $\alpha = \alpha(\omega_t^R, \omega_t^B)$ is described in the next section.

2.3 Dynamic policy model

We now address the bank policy function that depends on the environment variables. As we will see later in the paper, in order to solve this problem we will need to resort to sophisticated optimization algorithms. Given the high dimensionality of the problem, both in the risk factors ω_t^R and balance sheet state variables ω_t^B, it is undesirable to posit the policy function α as a function of all the environment variables (ω_t^R, ω_t^B), as this may lead to overfitting and a lack of interpretability. Therefore, we need to restrict the number of variables to be included in the policy function. We propose the following six variables to reflect the state of the bank balance sheet: loan rate (x_1), bond yield (x_2), cash rate (x_3), charge-offs (x_4), a leverage ratio E_t/D_t, which we call (x_5), and the loan allocation as a percentage of the total funding $L_t/(E_t + D_t)$, also denoted (x_6).

In principle, one would anticipate that many of these variables would be highly relevant for balance sheet decisions. For example, the interest rate environment heavily influences banks' net interest margin. Therefore, one would anticipate that the interest rates on loans, bonds, or cash can affect the policy decision. One could also expect that the credit risk environment, dictated by the charge-off rate, could be an important variable when considering how much to allocate to credit. Other balance sheet variables, such as the loan allocation or the bank's leverage, should also be relevant. In particular, banks with lower leverage should be able to take on more risk. In any case, the variable selection algorithm we present below will dictate the most relevant variables for the problem.

As illustrated in Fig. 2, the proposed policy function model can be represented as a neural network architecture. A critical aspect of adopting intelligent decision support models is the possibility of analyzing results and the interpretation by financial decision-makers. To meet this demand, we drive the investigation toward parsimonious neural networks. The neural network is shallow, with one input layer (blue neurons) and one activation layer with activation function σ (orange neurons). A normalization step of the output of the activation layer gives the allocation vector α (represented in yellow).

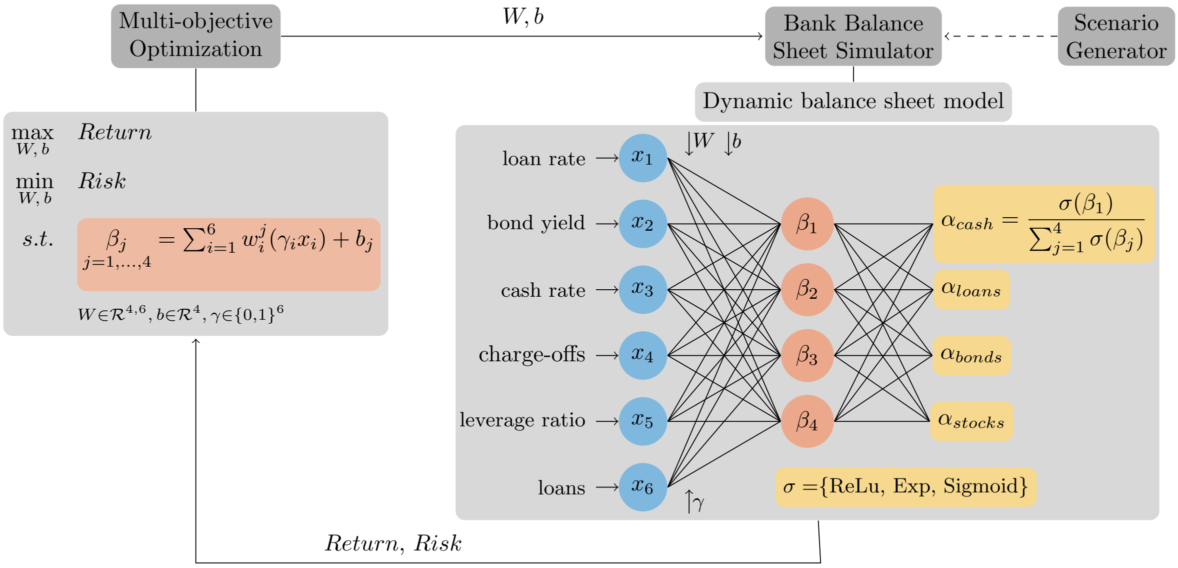

Fig. 2 Detailed representation of the dynamic policy model $\alpha = (\alpha_{cash}, \alpha_{loans}, \alpha_{bonds}, \alpha_{stocks})$. The weights W and the biases b are the learnable parameters of the model. The tunable hyperparameters are γ and σ. The binary hyperparameter γ, with elements in $\{0,1\}^6$, determines which variables $\{x_i\}_{i=1}^6$ are included in the model. For example, $\gamma = (0,0,1,0,1,0)$ selects the variables x_3 (cash rate) and x_5 (leverage ratio). The hyperparameter σ represents three possible activation functions, $\mathrm{ReLU}(x) = \max(0,x)$, $\mathrm{Exp}(x) = e^x$, and $\mathrm{Sigmoid}(x) = (1+e^{-x})^{-1}$.

Let us see how the neural network policy function works. In blue, represented by $\{x_i\}_{i=1}^6$, we find the initial layer formed by six neurons associated with the six variables considered. All neurons in the input layer send weighted-adjusted information to the four neurons β_j, $j = 1,\dots,4$, in the activation layer (orange circles). That is, for the neuron β_j, as shown in orange on the left gray block of Fig. 2, we get

$$\beta_j = \sum_{i=1}^{6} w_i^j(\gamma_i x_i) + b_j, \quad j = 1,\dots,4, \tag{5}$$

where $w_i^j \in \mathbb{R}$ are the weights, $b_j \in \mathbb{R}$ the bias, and $\gamma_i \in \{0,1\}$ a binary hyperparameter that indicates whether or not to include the variable x_i in the model. For instance, $\gamma = (0,0,1,0,1,0)$ selects the variables x_3 (cash rate) and x_5 (leverage ratio). The hyperparameter γ allows for 63 distinct model configurations, representing all possible combinations with a number of variables between one and six. The weight w_i^j value conveys information about the influence of the input variable x_i on neuron β_j, with a value near zero suggesting that this variable x_i is not as important as a variable with a larger weight. The bias b is an additional parameter that adds flexibility to the neural network allowing, e.g., to change the location of the activation function by shifting the argument, similarly to the role played by the intercept in a linear regression model.

Last, to get normalized allocations that sum to one, we divide the value of the output of the activation layer by the sum of the value of the four activation neurons. For instance, as shown in yellow in Fig. 2, for the allocation to cash we get

$$\alpha_{cash} = \frac{\sigma(\beta_1)}{\sum_{j=1}^{4} \sigma(\beta_j)}, \tag{6}$$

with β_j, $j = 1,\dots,4$, given by (5) and σ the hyperparameter representing one of three possible activation functions:

$$\mathrm{ReLu}(x) = \max(0,x), \quad \mathrm{Exp}(x) = e^x, \quad \text{and} \quad \mathrm{Sigmoid(x)} = \frac{1}{1+e^{-x}}.$$

The activation function σ adds non-linearity to the model.

2.4 Risk and return metrics

Given a trajectory ω^R for the risk factors in Table 1, one can simulate the evolution of the balance sheet given the equations described in Table 2 and a policy function α defined by the dynamic model proposed in the previous section. The return of the policy α on a trajectory ω^R is given by the bank's capital plus accumulated dividends at time T, that is,

$$Ret(\omega^R) = \left(\frac{E_T + Div_T^R}{E_0}\right)^{\frac{1}{T}} - 1, \tag{7}$$

where T the number of years in the simulation. The return for a given policy function α takes the expectation of the random variable $Ret(\omega^R)$ over its possible realizations, so that

$$Return = \mathbb{E}(Ret(\omega^R)). \tag{8}$$

Given a policy function α, we also need to evaluate its risk. First, we use a random variable $f(\omega_R)$ which indicates if the bank enters into a failure or default under the risk factor trajectory ω_R and the policy α. This variable is defined by $f(\omega^R) = 1$ if the bank runs out of liquid assets or reaches the tier 1 limit ratio, and $f(\omega^R) = 0$ otherwise. Risk for a policy function α is defined as the probability of the bank failing, i.e., averaging the random variable f over the possible realizations of ω^R under the policy function α:

$$Risk = \mathbb{E}(f(\omega^R)). \tag{9}$$

Let us justify the choice for this risk indicator in comparison to alternatives. First, it has a clear advantage that it can simultaneously encapsulate events of bank failure due to solvency shocks (such as a sudden an increase in interest rates or spikes in defaults), and liquidity shocks (due to the absence of liquid assets to withstand deposit withdrawals).

Second, it is a long-term measure. For the horizon we consider, which is typically a period of 10 years or more, we are interested in the long-term sustainability of the bank, and not so much on the year-on-year changes. Therefore, we chose this measure in comparison to other measures such as the yearly standard deviation of returns. The long-term nature of the risk indicator also means that it will include events such as the bank not generating enough income to cover its operating costs. Over the course of one year, this may not be so problematic. However, this problem compounds over the course of many years, as a sequence of many years of negative return on equity increases the likelihood of the failure of the bank.

Third and finally, similar measures have been proposed in the context of pension fund ALM. For example, C. Boender, van Aalst, and Heemskerk (1998) reports the use of the probability of underfunding as a risk measure in this context.

These risk and return metrics give rise to a non-continuous and non-differentiable multi-objective optimization problem, and we solve it using the concept of non-dominated Pareto solutions and an efficient multi-objective optimization algorithm, the NSGAII (Non-dominant Sorting Genetic Algorithm II) (Deb, Pratap, Agarwal, & Meyarivan, 2002).

3 Numerical experiments

In this section, we assess the effectiveness of our dynamic bank policy model using a training-validation-testing scheme. It consists of four stages: first, use the scenario generator in Table 1 to generate a comprehensive dataset simulating diverse economic scenarios; second, tune the hyperparameters γ (variable selection) and σ (activation function); third, train the tuned model followed by results analysis and interpretation; and finally, test and compare the model's performance against other established methodologies.

3.1 Dataset generation

To build a robust dataset representing diverse economic environments, we fed the scenario generator with four initial points that cover several decades (1985, 1995, 2005, and 2015), generating 128.000 risk factor trajectories with a time frame of 10 years, 32.000 from each initial point. To further reduce the dependency on the initial conditions and increase variety in the trajectories, we let the simulations run for 10 years, considering only the 10 years afterward. Some of the generated trajectories shown in Fig. 3 illustrate the dataset diversity. For subsequent analysis, we randomly split the dataset into 50% for training, 25% for validation, and 25% for testing.

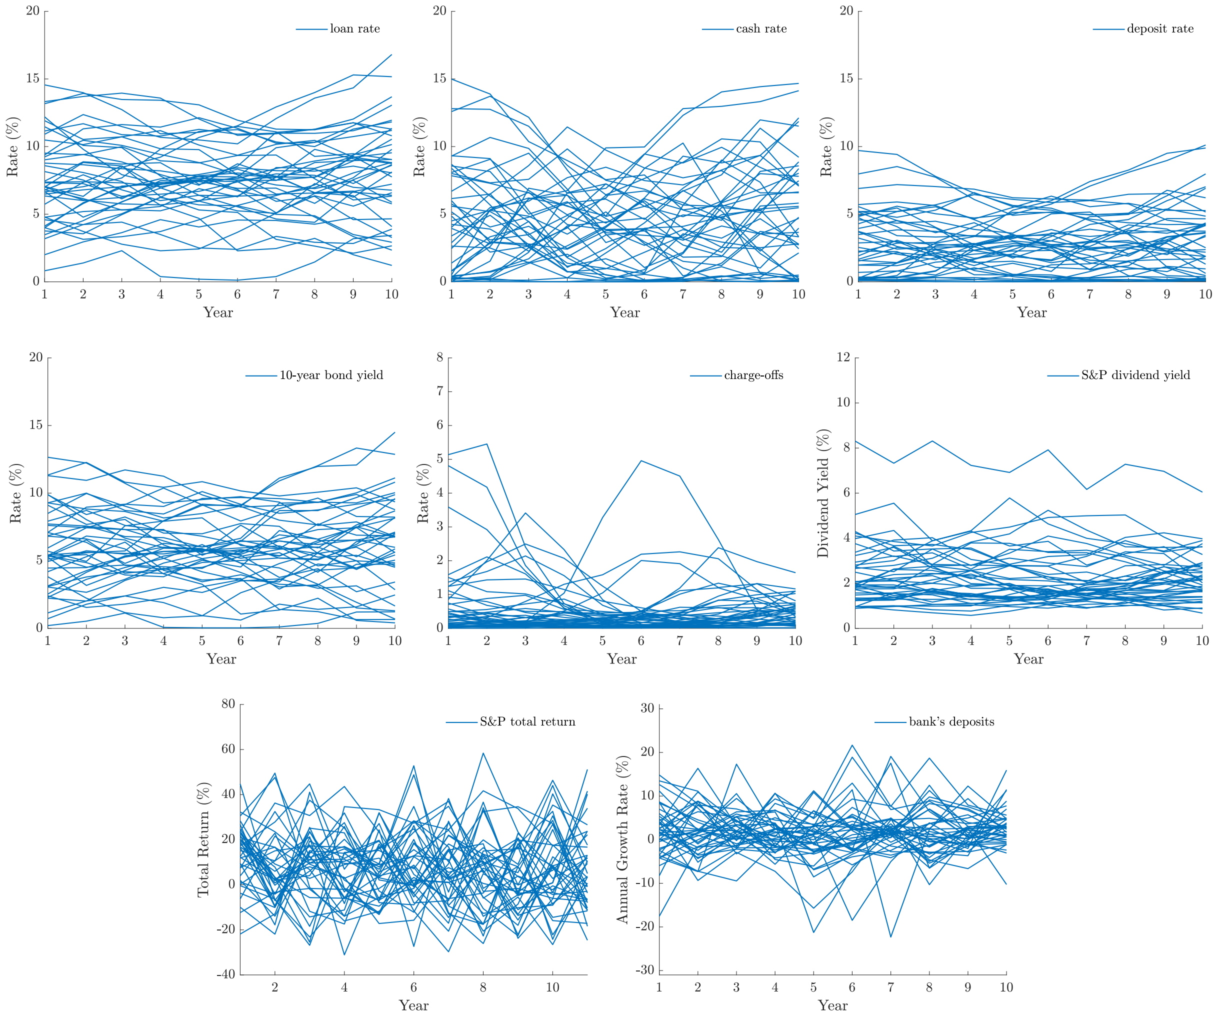

Fig. 3 Some examples of the interest rate (loan rate, cash rate, deposit rate and 10-year bond yield), charge-offs, S&P dividend yield and total return, and bank's deposits (annual growth rate) data used for the numerical experiments. The full dataset consists of 128.000 risk factor trajectories with a time frame of 10 years, split into 50% for training, 25% for validation, and 25% for testing.

3.2 Hyperparameter optimization

We resort to grid search for optimizing the hyperparameters γ (variable selection) and σ (activation function). This exhaustive method evaluates all possible hyperparameter combinations, enabling us to find the one that yields the best model performance.

Given 63 configurations for γ and three for σ, the total number of configurations explored by grid search is 189. Each grid search iteration involves two stages: training the model defined by the corresponding hyperparameter combination in the training set and assessing the model's performance in the validation set. Considering the inherent trade-off between return and risk, we rank the models based on their average return-to-risk ratio, thus associating each model with a single performance measure. We assume that a higher return-to-risk ratio indicates a more favorable model, as it implies a higher return for the same level of risk. Algorithm 1 describes the main steps of the grid search optimization.

Algorithm 1 Grid search algorithm for hyperparameter optimization.

---

find optimal hyperparameter combination *in:* training set, validation set *out:* γ^*, σ^*
// search all 189 hyperparameter combinations.
for $i = 1$ to 189 **do**
  Define $model_i$ associated with the ith hyperparameter combination.
  Train $model_i$ in the training set using the concept of Pareto front and the multi-objective optimization algorithm NSGAII. We obtain a set of Pareto points where each point corresponds to a given policy.
  // search all N^i Pareto front points/policies of $model_i$.
  for $n = 1$ to N^i **do**
    Apply $policy_n^i$ (nth policy of the Pareto front of $model_i$) to the validation set.
    Calculate the return and risk of $policy_n^i$ in the validation set.
  end for
  Calculate the average return-to-risk ratio of $model_i$ in the validation set.
end for
Find the $model_i$ with the highest average return-to-risk ratio.
Return the associated hyperparameter combination (γ^*, σ^*).

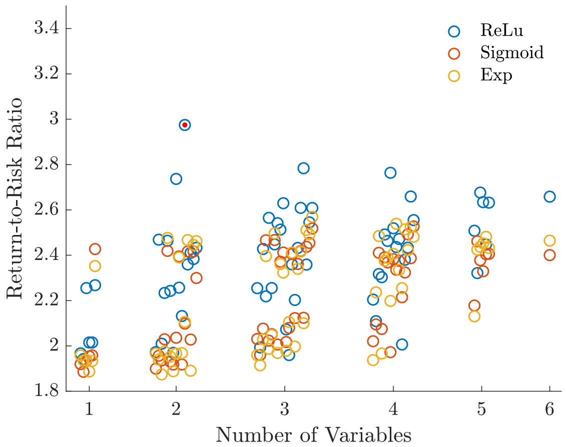

Fig. 4 Average return-to-risk ratio (higher is better) obtained during hyperparameters tuning (variable selection and activation function). According to the return-to-risk ratio results, the ReLu activation with the variables cash rate and leverage ratio (highlighted with a red dot) emerged as the best combination.

As Fig. 4 shows, the ReLu activation with the variables cash rate and leverage ratio (highlighted with a red dot) yields the best performance. This finding led us to choose this model configuration for further analysis. The phenomena of underfitting and overfitting explains why a two-variable model performs better than those with fewer or more variables. Beyond performance, this two-variable configuration leads to a readily interpretable model due to its parsimonious number of variables. This interpretability will allow valuable insights into the model's decision-making process.

As shown in Fig. 5, after hyperparameter tuning, the policy function depends only on the two selected variables: cash rate and the bank's leverage ratio. The reduced number of variables makes analyzing their impact on the policy easier. For example, ignoring the normalization factor, the allocation to stocks, α_{stocks}, is obtained as follows:

$$\alpha_{stocks} = \max(0, \beta_4) = \max(0, w_1^4 x_1 + w_2^4 x_2 + b_4),$$

with x_1 as the cash rates and x_2 as the bank's leverage ratio. The weights reflect the variables' importance. If, for example, w_1^4 is zero and w_2^4 is positive, the allocation to stocks is unaffected by cash rates and increases with the bank's leverage ratio. The bias b_4 shifts the activation function horizontally, controlling where it is triggered. If b_4 is positive (negative), the ReLu function shifts to the left (right), expanding (contracting) the range of values $w_2^4 x_2$ that trigger the allocation to stocks.

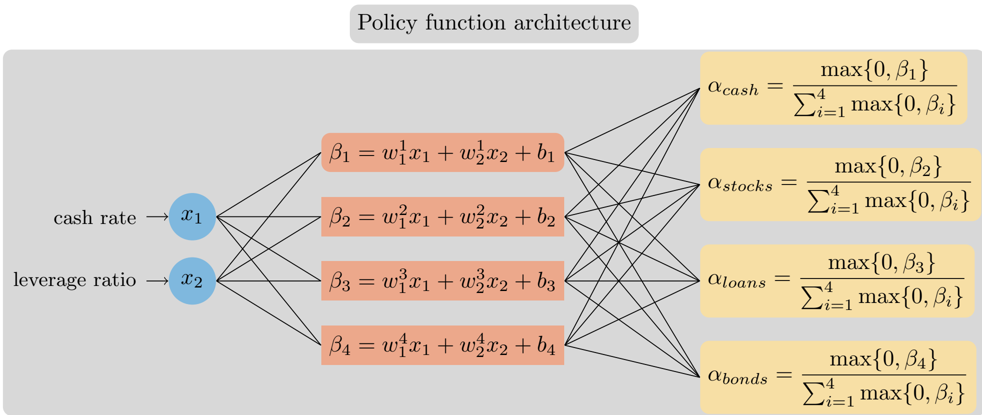

Fig. 5 Final architecture of the policy function after hyperparameter tuning. This tuned model relies on the variables cash rate and bank's leverage, denoted by x_1 and x_2, and on the ReLu activation function, $\text{ReLu}(x) = \max\{0, x\}$.

3.3 Model training and results analysis

This section focuses on training and analyzing the tuned dynamic policy model. We establish logical relationships between the variables (cash rate and leverage ratio) and the allocations from the policy function, by looking in detail at the values of the neural network weights w and biases b. For training, we use the full training dataset and the concept of Pareto front. Let N_P be the number of points in the optimal Pareto front. The parameters ω and b will vary within the Pareto front, so we will have N_P different realizations within the front. The allocation α dictated by the policy function will also vary for each time t and scenario ω^R.

In Fig. 6, on the top left, we depict the risk and the return associated with the Pareto front. In Fig. 6, on the top right, we give the Pareto allocations obtained by averaging over the paths and over the years. As we can readily observe, the Pareto front is increasing in return and risk. As expected, for a higher level of return and risk, there is a higher allocation to stocks and a lower allocation to cash. Also, notice that, for all policies in the front, the balance sheet allocates a significant amount to liquid assets (cash and stocks) to hedge severe decreases in deposit volumes, thus mitigating liquidity risk. As we observed before, running out of liquid assets is one event that dictates the bank's failure.

Averaging over the trajectories, we show, on the bottom left of Fig. 6, the average annual asset allocations for the minimum risk policy function. The minimum risk policy holds significant cash and other assets necessary to generate a return to compensate for operating costs and avoid default. Similarly, on the bottom right of Fig. 6, we show the average time evolution of the maximum return policy. We observe that this policy will allocate significantly to riskier assets such as loans or stocks. The varying

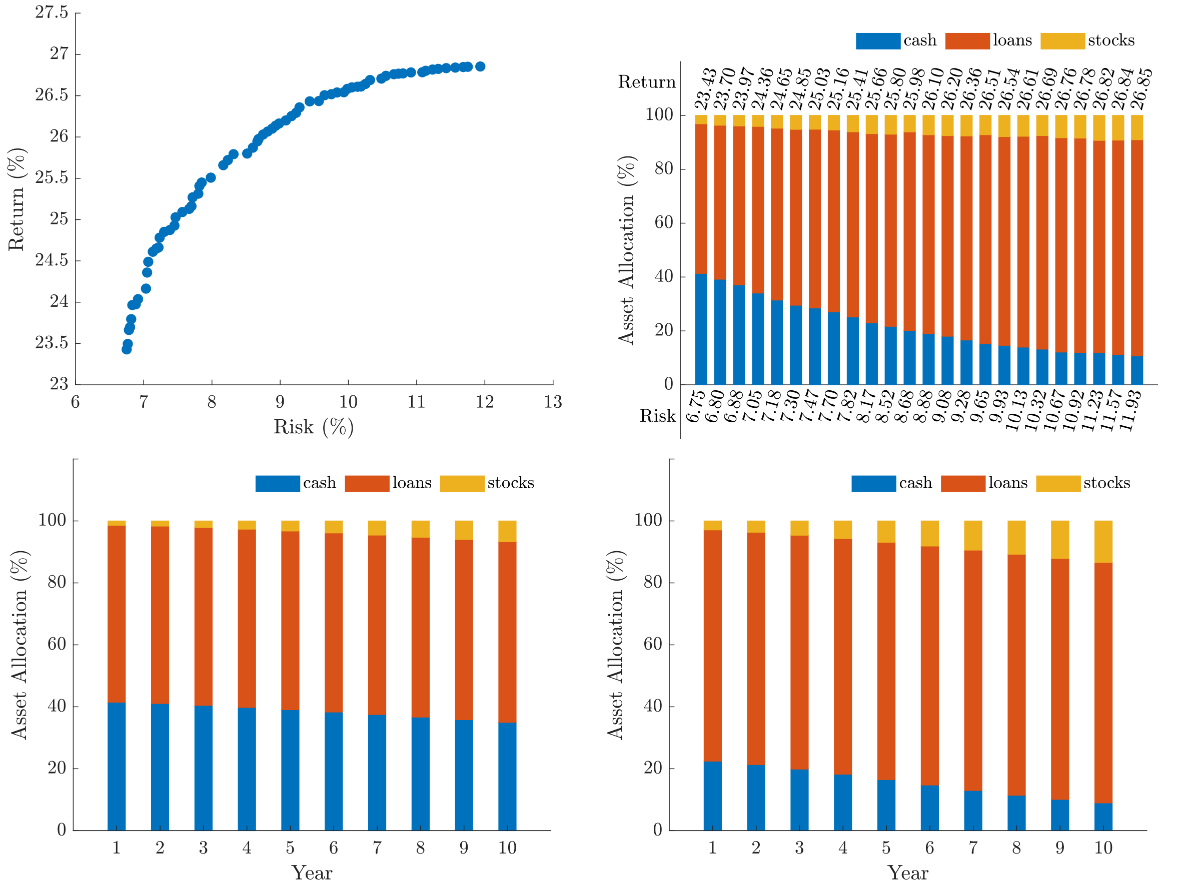

Fig. 6 Top row, on the left: training Pareto front obtained with the tuned dynamic model. Top row, on the right: Pareto optimal asset allocation. The allocations are obtained by averaging the dynamic allocations over the paths and over the years. Bottom row, from left to right: average asset allocation obtained with the minimum risk and maximum return Pareto solutions. The average is over the total number of paths. The horizontal axis denotes the year in the simulation horizon of 10 years.

average proportions show the dynamic nature of the policy algorithm. This adaptability allows the model to construct an optimized strategy over time, gradually increasing exposure to riskier assets (stocks) as returns increase and the bank's capital position improves.

In Fig. 7 and 8, we show, for each Pareto point, the weights w and the biases b associated with the four neurons $\beta_i = w_1^i x_1 + w_2^i x_2 + b_i$, $i = 1, \ldots, 4$, in the activation layer, where x_1 is the cash rate and x_2 the bank's leverage ratio. We recall that β_i, $i = 1, \ldots, 4$, are associated to the allocation to cash, loans, bonds, and stocks, respectively (see Fig. 5). The values for w and b are highly interpretable. Let us start by looking at the graph of β_1, the neuron associated with the cash position (Fig. 7, on the left). This neuron exhibits a positive sensitivity w_1^1 to cash rates and a negative sensitivity w_2^1 to leverage ratio. This result is highly intuitive. When cash rates increase, there is an incentive to increase the cash allocation. The fact that the sensitivity is negative to the capital position in the bank stems from the fact that, when capital is higher, the bank invests more in stocks, as we will see below. The bias b_1 is always positive, causing the ReLu activation function, $\alpha_{cash} = \max\{0,\, w_1^1 x_1 + w_2^1 x_2 + b_1\}$, to shift to the left and increasing the neuron output.

Looking into the neuron β_2 for the loan position (Fig. 7, on the right), we observe that b_2 is positive, shifting the activation function, $\alpha_{loans} = \max\{0,\, w_1^2 x_1 + w_2^2 x_2 + b_2\}$, to the left and increasing the neuron output. It is also relatively high, suggesting that loan allocation is significant independently of the economic environment. This finding is intuitive when we consider that loans, here classified at book value, tend to be less risky than bonds and stocks (whose market price variations impact the bank's capital

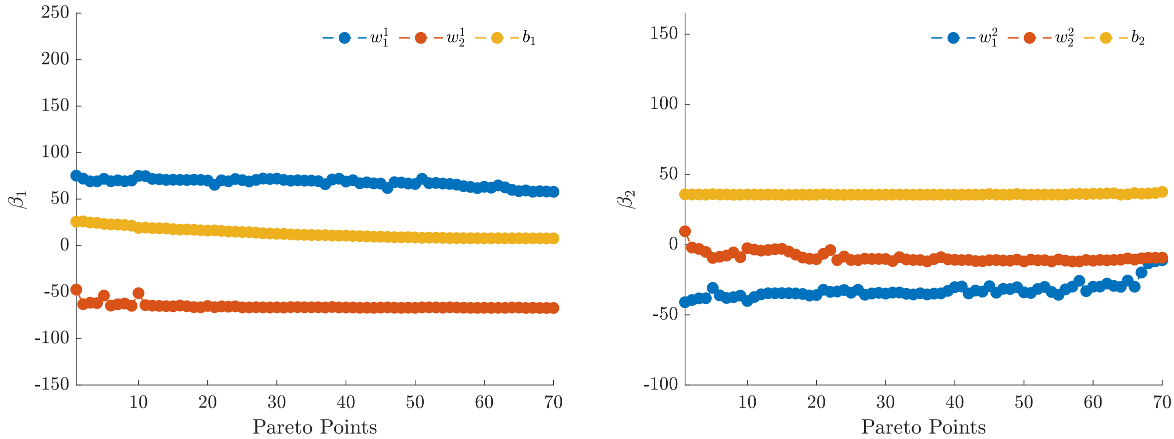

Fig. 7 Left to right and top to bottom: for each one of the 70 Pareto points, we show the weights w and biases b associated with the neurons β_1 and β_2. Each neuron is defined by the formula $\beta_i = w_1^i x_1 + w_2^i x_2 + b_i$, $i = 1, \ldots, 4$, where x_1 is the cash rate and x_2 is leverage ratio. These neurons are directly related to the allocations $\alpha_{cash} = \max\{0, \beta_1\}/\widehat{\beta}$ and $\alpha_{loans} = \max\{0, \beta_2\}/\widehat{\beta}$, with $\widehat{\beta} = \sum_{i=1}^{4} \max\{0, \beta_i\}$.

position) and deliver higher returns than cash. We also observe a negative sensitivity to the cash rate w_1^2, revealing an incentive to increase loan allocation when interest rates are lower. In a low-interest-rate scenario, the narrow net interest margin between cash rates and rates paid on deposits generates limited income. This income is far from fulfilling the bank's operating costs, and the dynamic model looks to generate more return by increasing allocation to loans (and also to stocks, as we will see next).

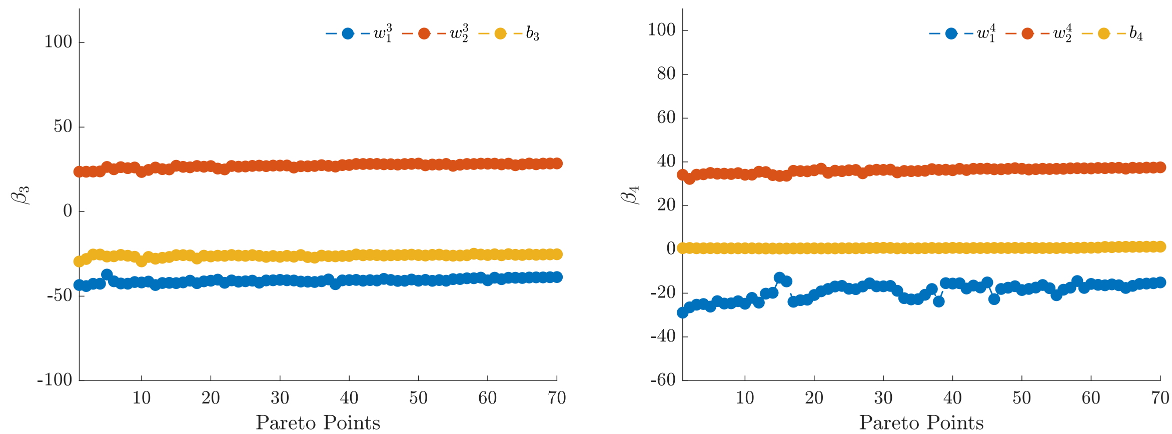

Fig. 8 Left to right and top to bottom: for each one of the 70 Pareto points, we show the weights w and biases b associated with the neurons β_3 and β_4. Each neuron is defined by the formula $\beta_i = w_1^i x_1 + w_2^i x_2 + b_i$, $i = 1, \ldots, 4$, where x_1 is the cash rate and x_2 is leverage ratio. These neurons are directly related to the allocations $\alpha_{bonds} = \max\{0, \beta_3\}/\widehat{\beta}$ and $\alpha_{stocks} = \max\{0, \beta_4\}/\widehat{\beta}$, with $\widehat{\beta} = \sum_{i=1}^{4} \max\{0, \beta_i\}$.

The neuron β_3 (Fig. 8, on the left) for bonds is conditioned by a very negative bias, which in practice will mean that the allocation will be zero in most cases, as we will see in the following sections. Unlike loans, the fluctuations in bond prices impact the bank's capital position and thus generate more risk, while not generating enough return for the price fluctuations. When generating return through liquid assets, the optimizer prefers the investment in stocks to bonds.

Finally, we analyze the neuron β_4 (Fig. 8, on the right) concerning the allocation to stocks. The model shows a negative sensitivity to interest rates, indicating the model's preference to increase the allocation to stocks and decrease the cash allocation under the low-interest rate regimes. Under low interest rates, to avert default, the bank needs to increase the exposure to stocks to generate enough return to compensate for operating costs. Also, the allocation to stocks is highly dependent on the leverage ratio. Banks with higher leverage can allocate more to riskier asset classes, such as stocks.

In summary, the dynamic model interpretability analysis reveals the following balance sheet strategy:

- Consistently high allocation to loans, the asset with the better risk-return trade-off.
- Increased cash allocation in periods of high interest rates and increased exposure to loans and stocks in periods of low interest rates.
- Increased exposure to stocks when the bank capital buffers are high.

In Figure 9 we illustrate the dynamic nature of the model by showing two examples of how the model reacts under two different economic scenarios, by comparing the behavior under these two scenarios. The dotted line corresponds to an economic scenario of decreasing interest rates and higher leverage ratios. Under this scenario, the bank will have a higher exposure to stocks and loans, and a lower allocation to cash. The full line corresponds to a scenario of increasing interest rates and low leverage ratios. We observe the opposite behavior of the model in this case: lower allocations to stocks and loans, and higher allocations to cash.

In the upcoming section *3.5 Model testing*, we test the dynamic model under diverse economic scenarios and evaluate the effectiveness of this dynamic and intuitive strategy.

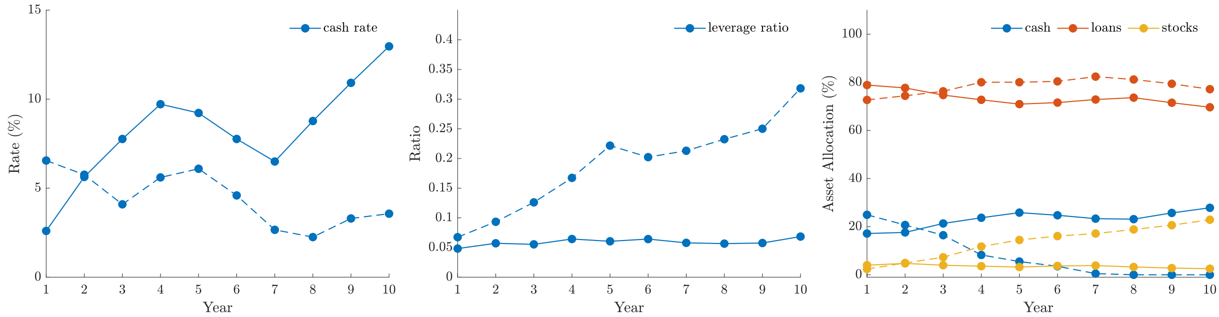

Fig. 9 Examples of the reaction of the model to two different economic samples/trajectories. The full line corresponds to a trajectory of increasing interest rates and lower leverage levels, whereas the dotted line corresponds to a trajectory of decreasing interest rates and higher leverage levels. As we can observe, the allocations determined by the model vary as a function of the economic environment.

3.4 Static model

As stated in the introduction, most of the recent models for bank asset management use static allocations (Brito & Júdice, 2022; Júdice et al., 2021; Júdice & Zhu, 2021; Schmaltz et al., 2014; Sirignano et al., 2016; Yan et al., 2020), so we should compare the proposed dynamic model to a static model to highlight the relevance of dynamic policies. In our framework, we can obtain a static model ignoring the input of economic variables and defining the constant policy function

$$(\alpha_{cash},\, \alpha_{loans},\, \alpha_{bonds},\, \alpha_{stocks}) = \frac{1}{\widehat{\theta}}(\theta_1, \theta_2, \theta_3, \theta_4), \quad \theta_i \in \mathbb{R}_0^+,\, i = 1, \ldots, 4, \tag{10}$$

with $\widehat{\theta} = \sum_{i=1}^{4} \theta_i$. The optimal static allocations are obtained by formulating the bi-objective optimization problem of finding $\theta_i \in \mathbb{R}_0^+$, $i = 1, \ldots, 4$, that maximizes *Return* and minimizes *Risk*.

In Fig. 10, on the left, we depict the risk and return associated with the Pareto front of the static model. Fig. 10, on the right, shows the Pareto static allocations, revealing that higher returns and risks are associated with increased loan investments and decreased cash positions. Comparing these results

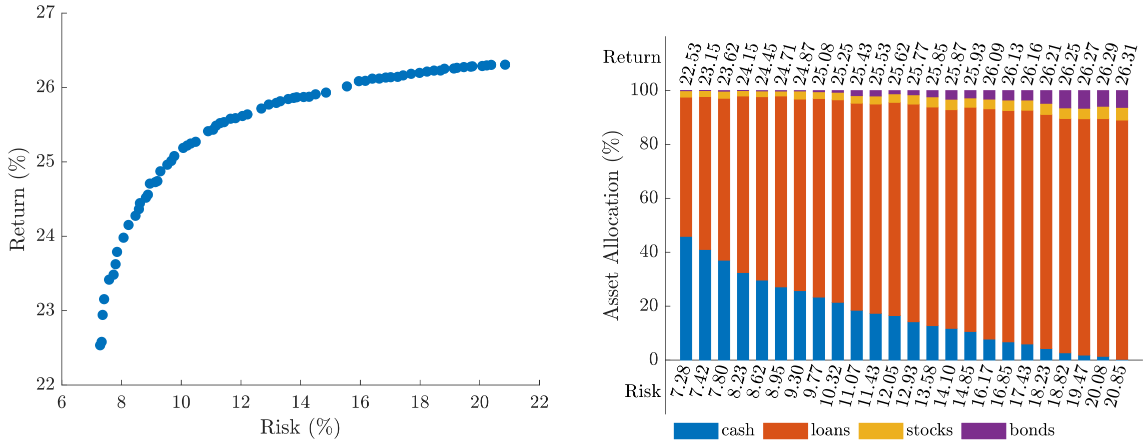

Fig. 10 On the left: training Pareto front obtained with the static model. On the right: Pareto optimal asset allocation for the static model, i.e., with constant allocations.

with Fig. 6, we observe that the dynamic model has superior adaptability and flexibility, achieving higher returns at lower risks with an allocation pattern that includes greater allocation to stocks and cash, particularly at higher returns. Overall, we also observe that both models favor loan allocation, owing to the asset's low volatility, here classified at book value. In contrast, bonds and stocks are considerably more volatile, as the fluctuations in the valuations have a direct impact on capital, explaining why we see low bond allocations that are even zero in the dynamic model. As illustrated in the following section, the dynamic model allocates to bonds in some particularly favorable economic conditions, like a high leverage ratio.

In the following section, we assess the performance of our dynamic model and compare it to alternative asset allocation policy functions, namely the static model and classical heuristic strategies. We conduct the experiments on the testing dataset, which we kept separate from the tuning and training phases to ensure an unbiased comparison.

3.5 Model testing

To benchmark our dynamic model against the static model, we apply the dynamic and static models calibrated to the training data to the separate test data. We emphasize that having a superior performance in the training dataset does not guarantee a superior performance in the separate testing data set. As it is generally known, overfitting in training often leads to poor performance out-of-sample, as evidenced by Bailey et al. (2014).

We also evaluate the performance of the dynamic model against two traditional heuristic allocation strategies, namely, the equal-weight policy and the 60/40 policy, which allocates 60% to higher-risk assets and 40% to lower-risk assets. In the case of investment portfolio optimization, DeMiguel, Garlappi, and Uppal (2007) have shown that the equal-weight strategy outperforms several optimized allocations out-of-sample, so it is always a helpful benchmark. In our case, the 1/N allocation policy assumes a 25% asset allocation to each asset class. In the case of pension funds' portfolios, Chaves, Hsu, Li, and Shakernia (2011) have shown that a 60/40 portfolio that includes 60% stocks and 40% bonds performs at least as well as equal-weight and risk parity portfolios. In our case, the 60/40 policy assumes a 20% allocation to cash, 30% to loans, 20% to bonds, and 30% to stocks.

The results in Fig. 11, on the left, show that the proposed dynamic model significantly outperforms the static and heuristic policies. The dynamic model presents a 2.65 return-to-risk ratio, outperforming the static model by 42% and the equal-weight and 60/40 strategies by more than 1000% and 5000%, respectively. The heuristic strategies, such as the equal-weight policy, are not particularly suitable for the amount of leverage we consider in the problem, as we consider an initial leverage ratio of 5% (considering the ratio of initial shareholders' capital to deposits): for such a leverage ratio, these strategies allocate

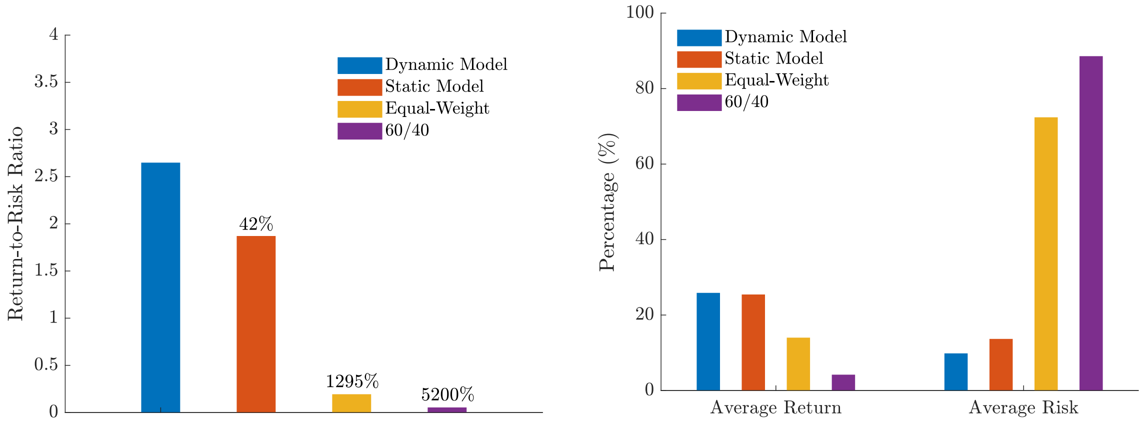

Fig. 11 Comparison, on the testing dataset, between the proposed dynamic model, the static model, and the heuristic policies (equal-weight and 60/40). On the left: return-to-risk ratio (higher is better). The numbers above the bars report the outperformance of the dynamic model compared to the static, equal-weight, and 60/40 allocations. On the right: average return (higher is better) and risk (lower is better). Here, the initial capital base is 5% of the initial deposit volume.

excessively to risky assets, exacerbating the risk under this leveraged environment, and significantly deteriorating the return to risk ratios.

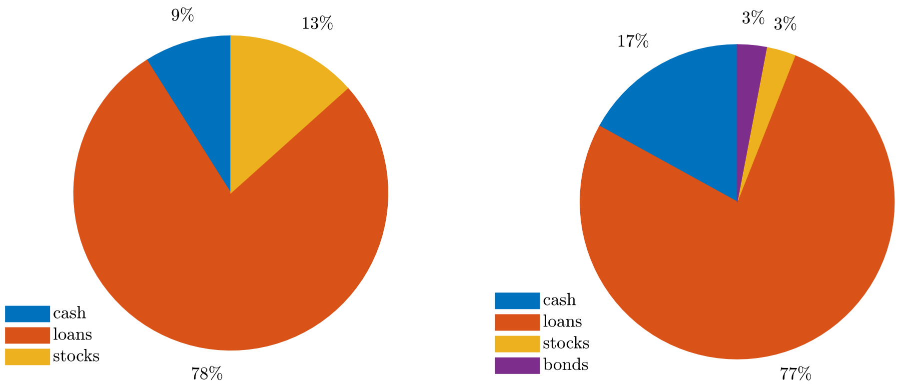

Fig. 12 From left to right: comparison, on the testing dataset, between the proposed dynamic model's average asset allocation at the end of the 10-year simulation period and the static model's asset allocation.

Fig. 11, on the right, reveals that when compared to the static model, the dynamic model has a higher return (26% vs. 25%) at significantly lower risk (10% vs. 14%), explaining the dynamic model's superior return-to-risk ratio. Compared to the equal-weight and 60/40 strategies, the dynamic model delivers a considerably higher return at significantly lower risk.

In Fig. 12, we compare the asset allocation of the static model (on the right) to the average asset allocation of the dynamic model (on the left) in the final year of simulation (year 10). This figure highlights the dynamic model's ability to adapt to the economic environment, as it can capitalize on timely stock allocations to deliver the highest return.

3.6 Sensitivity to the leverage level

In Fig. 13, we analyze the impact of increasing the bank's initial capital from 5% to 15% of initial deposits. As seen in Fig. 13, on the left, the dynamic model significantly outperforms both the static and heuristic models. The dynamic model presents a 4.02 return-to-risk ratio, outperforming the static model by 68% and the equal-weight and 60/40 strategies by more than 500%. Fig. 13, on the right, reveals that a higher return at lower risk is behind the dynamic model's superior performance. As expected, comparing the results to the ones in Fig. 11, lower leverage leads to lower risk across all models.

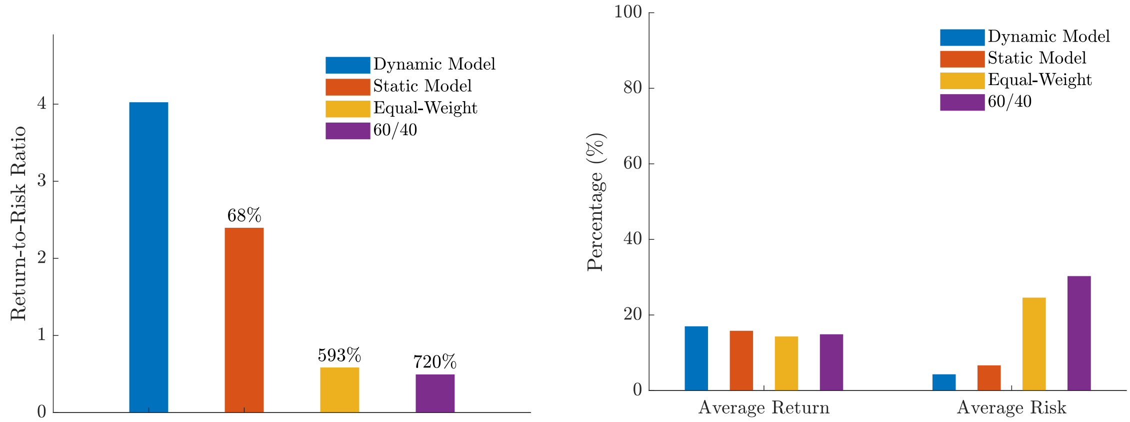

Fig. 13 Comparison, on the testing dataset, between the proposed dynamic model, the static model, and the heuristic strategies (equal-weight and 60/40). On the left: return-to-risk ratio (higher is better). The numbers above the bars report the outperformance of the dynamic model compared to the static, equal-weight, and 60/40 policies. On the right: average return (higher is better) and risk (lower is better). Here, the initial capital base is 15% of the initial deposit volume against a baseline of 5%.

Fig. 14 shows the dynamic model's average asset allocation. Compared to the baseline scenario (Fig. 12), the data reveals that the dynamic model effectively exploits the lower risk of bankruptcy associated with lower leverage. We see a strategic increase in stock allocation (from 13% to 21%) and the introduction of a small allocation to bonds. We highlight that the static model allocation remains constant independently of the economic environment. The lower risk of bankruptcy also explains why the 60/40 strategy, more exposed to risky assets (loans and stocks), presents a significantly better return under this low leverage scenario (Fig. 13, on the right).

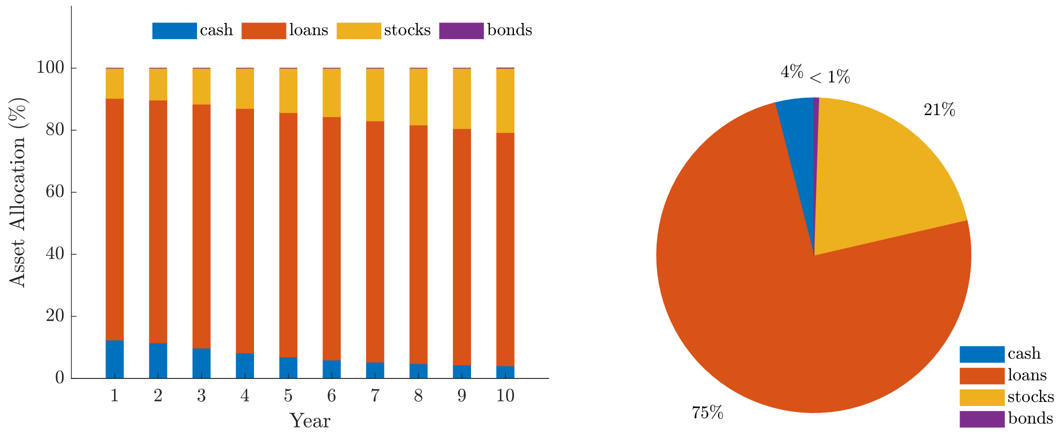

Fig. 14 Yearly dynamic model's average asset allocation (on the left) and at the end of the 10-year simulation period (on the right). Here, the initial capital base is 15% of the initial deposit volume against a baseline of 5%.

3.7 Sensitivity to credit losses

Next, we analyze the impact of contrasting credit loss environments on the model's performance. We create a high and a low credit losses scenario by splitting the testing dataset in half based on the average value of the charge-off rates. We present the results in Fig. 15. In both scenarios, the dynamic model significantly outperforms the static and heuristic models in the metric return-to-risk ratio (Fig. 15, on the left). The dynamic model outperforms the static model by more than 45% and heuristic models by more than 800%. Fig. 15, on the right, reveals that the dynamic model delivers the highest return at the lowest risk. As expected, the models present lower risk and a better return-to-risk ratio in the lower credit losses scenario (Fig. 15, bottom left), particularly the dynamic and static models, which are more exposed to loans.

These results demonstrate the dynamic model's ability to adapt to changing economic conditions, a capability not shared by static models. Fig. 16 further illustrates this adaptability. In the high-credit loss scenario (Fig. 16, on the left), the dynamic model recognizes the increased risk and takes mitigation actions, reducing exposure to stocks and loans while increasing cash allocation. On the other hand, in the favorable low-credit loss scenario (Fig. 16, on the right), the dynamic model increases stock, loan, and bond allocation. This responsiveness to changing economic environments is one of the dynamic model's fundamental advantages over static models.

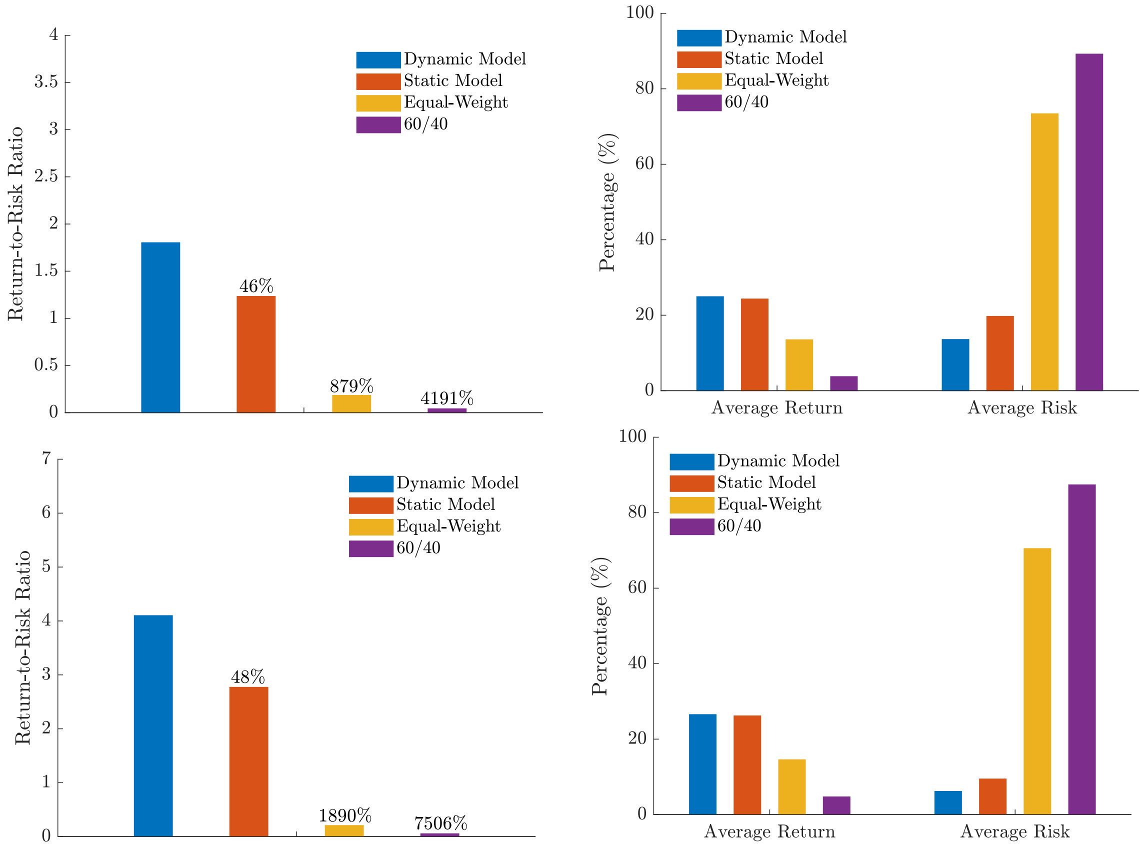

Fig. 15 Comparison, on the testing dataset, between the proposed dynamic model, the static model, and the heuristic strategies (equal-weight and 60/40). On the left: return-to-risk ratio (higher is better). The numbers above the bars report the outperformance of the dynamic model compared to the static, equal-weight, and 60/40 policies. On the right: average return (higher is better) and risk (lower is better). Top row: high-credit loss environment. Bottom row: low-credit loss environment. The testing dataset is split into two based on the average charge-off rate.

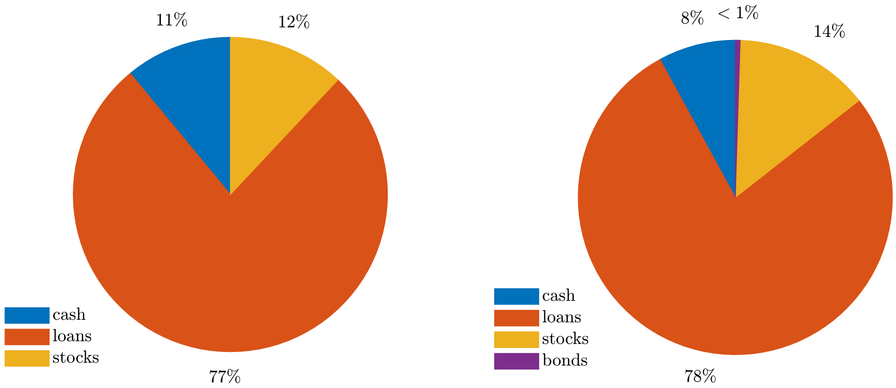

Fig. 16 Comparison, on the testing dataset, of the dynamic model's average allocations at the end of the 10-year simulation period in two distinct scenarios. On the left, high-credit loss environment. On the right, low-credit loss environment. The testing dataset is split into two based on the average charge-off rate.

3.8 Sensitivity to interest rates

In our last example, we analyze the impact of contrasting interest rate environments on the model's performance. We create a high and a low interest rates scenario by splitting the testing dataset in two based on the average value of the interest rates: 10-year bond yield, deposit rate, loan rate, and cash rate. The high interest rate scenario includes half of the testing scenarios with the highest average interest rates, whereas the low interest rate scenario includes half of the testing scenarios with lowest average interest rates.

We present the results in Fig. 17. First, we observe that the return-to-risk ratio of the dynamic strategy is favored by high interest rates. In the high-interest rate regime, this strategy has a return-to-risk ratio of 6.75, compared with a return to risk ratio of 1.56 in the low-interest rate regime. The increase in profitability is expected, in line with the findings by Claessens, Coleman, and Donnelly (2018), who document that prolonged low interest rates have a negative impact on bank profitability. In both scenarios, the dynamic model significantly outperforms the static and heuristic models in the return-to-risk ratio (Fig. 17, on the left). The increase in performance is very marked for the high-interest rate environment, where the dynamic strategy yields a 106% increase in performance compared to the optimal static policy and more than 3360% compared to the selected heuristics.

Fig. 18 further emphasizes the dynamic model's adaptability to changing economic conditions. In the low-interest-rate scenario (Fig. 18, left), the model strategically reduces the cash allocation while increasing exposure to loans. Under low interest rates, net interest margins are much lower, so the optimal dynamic strategy suggests that the bank should decrease its cash buffer to invest in assets with higher return. Conversely, in the high-interest rate scenario (Fig. 18, on the right), the dynamic model increases the cash allocation and decreases the allocation to loans. The increased allocation to stocks stems from the higher accumulation of capital under the high-interest rate regime.

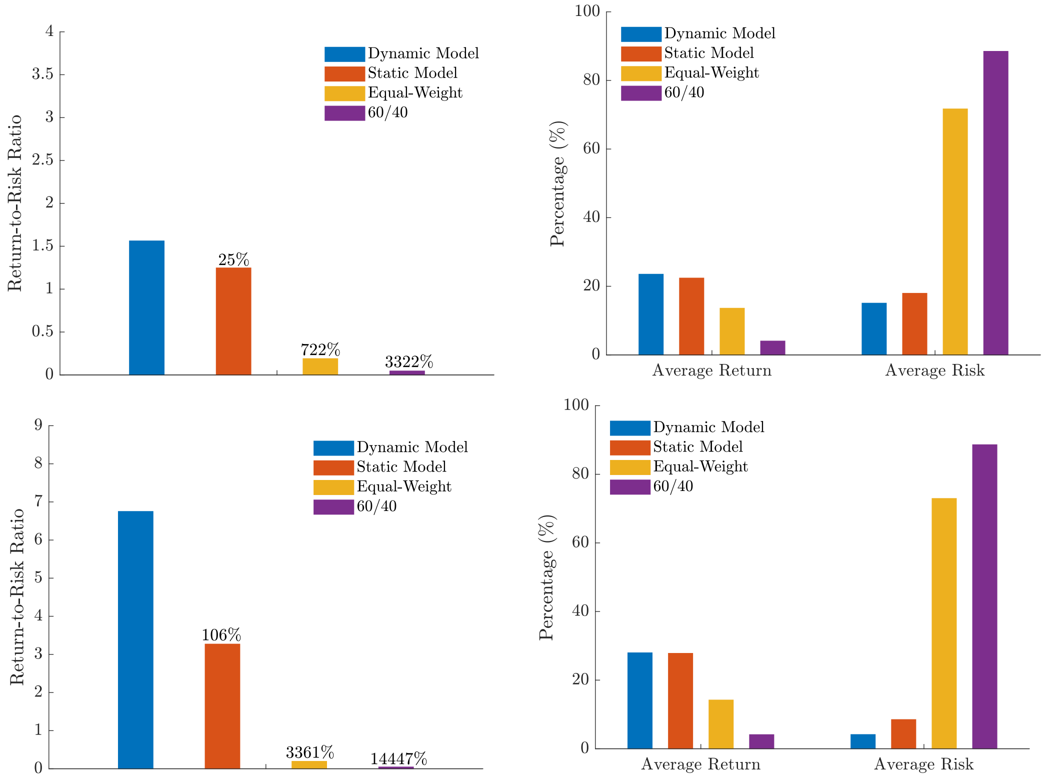

Fig. 17 Comparison, on the testing dataset, between the proposed dynamic model, the static model, and the heuristic strategies (equal-weight and 60/40). On the left: return-to-risk ratio (higher is better). The numbers above the bars report the outperformance of the dynamic model compared to the static, equal-weight, and 60/40 allocations. On the right: average return (higher is better) and risk (lower is better). Top row: low-interest rate environment (10-year bond yield, deposit rate, loan rate, and cash rate). Bottom row: high-interest rate environment. The testing dataset is split into two based on the average interest rates.

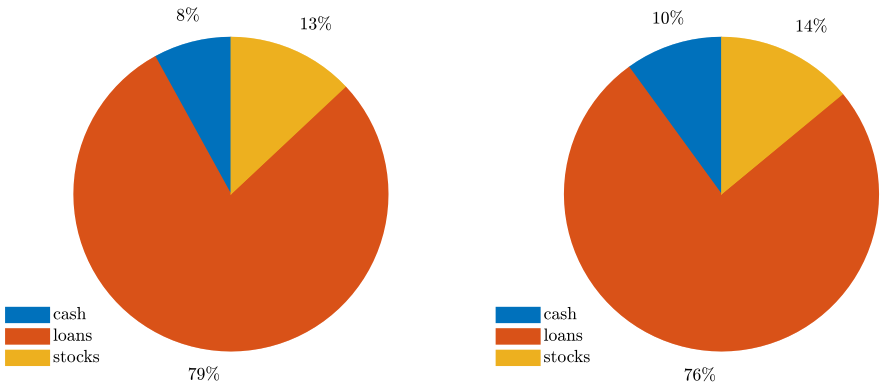

Fig. 18 Comparison, on the testing dataset, of the dynamic model's average allocations at the end of the 10-year simulation period in two distinct scenarios. On the left: low-interest rate environment (10-year bond yield, deposit rate, loan rate, and cash rate). On the right: high-interest rate environment. The testing dataset is split into two based on the average interest rates.

4 Conclusions

This research develops a dynamic balance sheet management methodology for a bank that determines the balance sheet policy as a function of environment variables. The methodology determines Pareto frontiers for the optimal policy functions, based on the NSGAII, a multi-objective algorithm that does not require derivatives, due to the non-continuity and non-differentiability of our problem.

We proceeded to develop a hyperparameter optimization algorithm based on scores of return and risk. The algorithm created a parsimonious policy function for the asset allocation as a function of the cash rates and bank's leverage ratio.

The model delivers a highly interpretable model, with highly intuitive sensitivities of the allocations with respect to the input variables. Also, environments of low interest rates, which are less profitable, steer the allocation towards riskier assets such as loans and stocks so that the bank can generate enough return to avoid default.

We also used a testing set with different paths to assess the validity of our results. Our testing results confirm that the dynamic strategy has better adaptability to the economic environment, effectively managing risk while adjusting asset allocations to capitalize on favorable economic conditions. This adaptability translates into superior performance, achieving considerably higher return-to-risk ratios than static strategies (42% for a low leverage ratio and 68% for a high leverage ratio) and the heuristics 1/N (more than 593%) and 60/40 (more than 720%), and thus is suitable for use in practice.

Dynamic balance sheet modeling is a very rich field, as it can generate considerable gains for banks as a function of the economic variables. We hope this research will help develop dynamic balance sheet management models that can be used in practice.

Acknowledgements. Partially supported by the Centre for Mathematics of the University of Coimbra (funded by the Portuguese Government through FCT/MCTES, DOI 10.54499/UIDB/00324/2020), by the project Complex Risk Management in the Big Data Regime (PTDC/MAT-AP/ 1286/2021), and by Fundação para a Ciência e a Tecnologia, grant UIDB/00315/2020 (DOI: 10.54499/UIDB/00315/2020).

Declarations

- **Conflict of interest** The authors have no conflict of interest to declare that are relevant to the content of this article.

References

Bailey, D.H., Borwein, J.M., López de Prado, M., Zhu, Q.J. (2014). Pseudo-mathematics and financial charlatanism: The effects of backtest overfitting on out-of-sample performance. *Notices of the American Mathematical Society*, *61*(5), 458, https://doi.org/10.1090/noti1105

Birge, J.R., & Júdice, P. (2013). Long-term bank balance sheet management: Estimation and simulation of risk-factors. *Journal of Banking & Finance*, *37*(12), 4711–4720, https://doi.org/10.1016/j.jbankfin.2013.07.040

Birge, J.R., & Júdice, P. (2014). Assessing the long-term effects of bank policies. *SSRN Electronic Journal*, , https://doi.org/10.2139/ssrn.2442194

Boender, C., van Aalst, P., Heemskerk, F. (1998). Modelling and management of assets and liabilities of pension plans in the netherlands. *World-wide asset and liability modelling* (pp. 561–580). Cambridge University Press.

Boender, G.C. (1997). A hybrid simulation/optimisation scenario model for asset/liability management. *European Journal of Operational Research*, *99*(1), 126–135, https://doi.org/10.1016/s0377-2217(96)00387-6

Brito, R.P., & Júdice, P. (2022, March). Efficient credit portfolios under ifrs 9. *International Transactions in Operational Research*, *30*(5), 2453–2484, https://doi.org/10.1111/itor.13137

Campbell, J.Y., & Viceira, L.M. (2002). *Strategic asset allocation: portfolio choice for long-term investors.* Clarendon Lectures in Economic.

Cariño, D.R., Kent, T., Myers, D.H., Stacy, C., Sylvanus, M., Turner, A.L., ... Ziemba, W.T. (1994). The russell-yasuda kasai model: An asset/liability model for a japanese insurance company using multistage stochastic programming. *Interfaces*, *24*(1), 29–49, https://doi.org/10.1287/inte.24.1.29

Chaves, D., Hsu, J., Li, F., Shakernia, O. (2011, February). Risk parity portfolio vs. other asset allocation heuristic portfolios. *The Journal of Investing*, *20*(1), 108–118, https://doi.org/10.3905/joi.2011.20.1.108

Claessens, S., Coleman, N., Donnelly, M. (2018). "low-for-long" interest rates and banks' interest margins and profitability: Cross-country evidence. *Journal of Financial Intermediation*, *35*, 1-16, https://doi.org/https://doi.org/10.1016/j.jfi.2017.05.004

Coelho, C., Santos, J.L., Júdice, P. (2023, November). The performance of bank portfolio optimization. *International Transactions in Operational Research*, *31*(3), 1458–1485, https://doi.org/10.1111/itor.13395

Consigli, G., & Dempster, M. (1998). Dynamic stochastic programmingfor asset-liability management. *Annals of Operations Research*, *81*, 131–162, https://doi.org/10.1023/a:1018992620909

Costa, S., Faias, M., Júdice, P., Mota, P. (2020, oct). Panel data modeling of bank deposits. *Annals of Finance*, *17*(2), 247–264, https://doi.org/10.1007/s10436-020-00373-1

Deb, K., Pratap, A., Agarwal, S., Meyarivan, T. (2002). A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, *6*(2), 182–197, https://doi.org/10.1109/4235.996017

DeMiguel, V., Garlappi, L., Uppal, R. (2007). Optimal versus naive diversification: How inefficient is the 1/N portfolio strategy? *Review of Financial Studies*, *22*(5), 1915–1953, https://doi.org/10.1093/rfs/hhm075

Doumpos, M., Zopounidis, C., Gounopoulos, D., Platanakis, E., Zhang, W. (2023, April). Operational research and artificial intelligence methods in banking. *European Journal of Operational Research*, *306*(1), 1–16, https://doi.org/10.1016/j.ejor.2022.04.027

Hałaj, G. (2016, nov). Dynamic balance sheet model with liquidity risk. *International Journal of Theoretical and Applied Finance*, *19*(07), 1650052, https://doi.org/10.1142/s0219024916500527

Hibiki, N. (2006). Multi-period stochastic optimization models for dynamic asset allocation. *Journal of Banking & Finance*, *30*(2), 365–390, https://doi.org/10.1016/j.jbankfin.2005.04.027

Júdice, P., Pinto, L., Santos, J.L. (2021). Bank strategic asset allocation under a unified risk measure. *Expert Systems with Applications*, *185*, 115574, https://doi.org/10.1016/j.eswa.2021.115574

Júdice, P., & Zhu, Q.J. (2021, dec). Bank balance sheet risk allocation. *Journal of Banking & Finance*, *133*, 106257, https://doi.org/10.1016/j.jbankfin.2021.106257

Kaya, O. (2019). *Artificial intelligence in banking* (Tech. Rep.). Deutsche Bank Research.

Kosmidou, K., & Zopounidis, C. (2004). Combining goal programming model with simulation analysis for bank asset liability management. *INFOR: Information Systems and Operational Research*, *42*(3), 175–187, https://doi.org/10.1080/03155986.2004.11732701

Kouwenberg, R. (2001). Scenario generation and stochastic programming models for asset liability management. *European Journal of Operational Research*, *134*(2), 279–292, https://doi.org/10.1016/s0377-2217(00)00261-7

Kusy, M.I., & Ziemba, W.T. (1986). A bank asset and liability management model. *Operations Research*, *34*, 356–376, https://doi.org/10.1287/opre.34.3.356

Maier-Paape, S., Júdice, P., Platen, A., Zhu, Q.J. (2023). *Scalar and vector risk in the general framework of portfolio theory: A convex analysis approach.* Springer International Publishing.

Markowitz, H. (1952). Portfolio Selection. *Journal of Finance*, *7*(1), 77-91, https://doi.org/j.1540-6261.1952.tb01525.

Merton, R.C. (1969). Lifetime portfolio selection under uncertainty: The continuous-time case. *The Review of Economics and Statistics*, *51*(3), 247, https://doi.org/10.2307/1926560

Mukuddem-Petersen, J., & Petersen, M.A. (2006, aug). Bank management via stochastic optimal control. *Automatica*, *42*(8), 1395–1406, https://doi.org/10.1016/j.automatica.2006.03.012

Mulvey, J.M., & Vladimirou, H. (1989). Stochastic network optimization models for investment planning. *Annals of Operations Research*, *20*(1), 187–217, https://doi.org/10.1007/bf02216929

Oğuzsoy, C.B., & Güven, S. (1997). Bank asset and liability management under uncertainty. *European Journal of Operational Research*, *102*(3), 575–600, https://doi.org/10.1016/s0377-2217(96)00241-x

Schmaltz, C., Pokutta, S., Heidorn, T., Andrae, S. (2014). How to make regulators and shareholders happy under Basel III. *Journal of Banking & Finance*, *46*, 311–325, https://doi.org/10.1016/j.jbankfin.2014.05.031

Sirignano, J.A., Tsoukalas, G., Giesecke, K. (2016). Large-scale loan portfolio selection. *Operations Research*, *64*(6), 1239–1255, https://doi.org/10.1287/opre.2016.1537

Yan, D., Zhang, X., Wang, M. (2020). A robust bank asset allocation model integrating credit-rating migration risk and capital adequacy ratio regulations. *Annals of Operations Research*, *299*(1-2), 659–710, https://doi.org/10.1007/s10479-020-03571-2

Yu, T.-Y., Tsai, C., Huang, H.-T. (2010). Applying simulation optimization to the asset allocation of a property–casualty insurer. *European Journal of Operational Research*, *207*(1), 499–507, https://doi.org/10.1016/j.ejor.2010.04.012

Zenios, S.A. (1995). Asset/liability management under uncertainty for fixed-income securities. *Annals of Operations Research*, *59*(1), 77–97, https://doi.org/10.1007/bf02031744

Ziemba, W., & Mulvey, J. (1998). *Worldwide asset and liability modeling.* Cambridge University Press.