



## OPEN ACCESS

## EDITED BY

Peter Schwendner,  
Zurich University of Applied Sciences,  
Switzerland

## REVIEWED BY

Jan-Alexander Posth,  
Zurich University of Applied Sciences,  
Switzerland  
Bertrand Kian Hassani,  
University College London, United Kingdom

## \*CORRESPONDENCE

Josef Teichmann  
✉ josef.teichmann@math.ethz.ch

RECEIVED 01 March 2023

ACCEPTED 28 April 2023

PUBLISHED 22 May 2023

## CITATION

Krabichler T and Teichmann J (2023) A case study for unlocking the potential of deep learning in asset-liability-management. *Front. Artif. Intell.* 6:1177702. doi: 10.3389/frai.2023.1177702

## COPYRIGHT

© 2023 Krabichler and Teichmann. This is an open-access article distributed under the terms of the [Creative Commons Attribution License \(CC BY\)](https://creativecommons.org/licenses/by/4.0/). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.

# A case study for unlocking the potential of deep learning in asset-liability-management

Thomas Krabichler<sup>1</sup> and Josef Teichmann<sup>2\*</sup>

<sup>1</sup>Centre for Banking and Finance, Eastern Switzerland University of Applied Sciences, St. Gallen, Switzerland, <sup>2</sup>Stochastic Finance Group, ETH Zurich, Zurich, Switzerland

The extensive application of deep learning in the field of quantitative risk management is still a relatively recent phenomenon. This article presents the key notions of Deep Asset-Liability-Management (“Deep ALM”) for a technological transformation in the management of assets and liabilities along a whole term structure. The approach has a profound impact on a wide range of applications such as optimal decision making for treasurers, optimal procurement of commodities or the optimization of hydroelectric power plants. As a by-product, intriguing aspects of goal-based investing or Asset-Liability-Management (ALM) in abstract terms concerning urgent challenges of our society are expected alongside. We illustrate the potential of the approach in a stylized case.

## KEYWORDS

asset-liability-management, deep hedging, machine learning in finance, portfolio management, reinforcement learning

## 1. Introduction

Mismatches between opposing fixed and floating prices along a whole term structure are omnipresent in various business fields and raise substantial business risks. Already humble approaches lead to analytically intractable mathematical models entailing high-dimensional allocation problems with constraints in the presence of frictions. This impediment often leads to over-simplified modeling (which still needs sophisticated technology), uncovered risks or overseen opportunities. Spectacular successes of deep learning techniques in language processing, image recognition, learning games from *tabula rasa*, or risk management, just to name a few, are stimulating our imagination that *asset-liability-management* (ALM) might enter a new era. This novel approach, briefly denoted “Deep ALM”, raises a huge field of mathematical research questions such as feasibility, robustness, inherent model risk and inclusions of all economic aspects. If we answer these questions, this could possibly change the ALM practice entirely.

ALM is a research topic that may be traced back to the seventies of the previous century when Martin L. Leibowitz and others developed cash-flow matching schemes in the context of the so-called *dedicated portfolio theory*; see [Fabozzi et al. \(1992\)](#). In the eighties and nineties, ALM was an intensively studied research topic in the mathematical community. This unanimously important branch of quantitative risk management has been dominated by techniques from *stochastic control theory*; e.g., see [Consigli and Dempster \(1998\)](#) by Giorgio Consigli and Michael A. H. Dempster or [Frauendorfer and Schürle \(2003\)](#) by Karl Frauendorfer and Michael Schürle for articles utilizing *dynamic programming*. Further aspects of the historical evolution are treated in [Ryan \(2013\)](#) by Ronald J. Ryan. However, already over-simplified modeling involves a high complexity both from the analytical and the technological viewpoint. A cornerstone in this regard is [Kusy and Ziemba \(1983\)](#) by Martin I. Kusy and William T. Ziemba. They consider a rather advanced and flexible ALM

model by specifying an economic objective (e.g., the maximization of the present value of future profits), a series of constraints (e.g., regulatory requirements and liquidity assurance) and penalty costs for constraint violations.

Prevalent approaches applied in the industry such as solely immunizing oneself against parallel shifts of the yield curve have remained suboptimal until this day; e.g., see [Spillmann et al. \(2019\)](#) by Martin Spillmann, Karsten Döhnert, and Roger Rissi for a recent exposition. Particularly, the empirical approach with static replicating portfolios described in the Section 4.2 of [Spillmann et al. \(2019\)](#) is often applied in mid-sized retail banks. Although the impact of stress scenarios is investigated in the light of regulatory requirements, the holistic investment and risk management process of many treasury departments remains insufficient. More sophisticated approaches, such as for instance *option-adjusted spread models* (e.g., see [Bardenhewer, 2007](#) by Martin M. Bardenhewer), have not become established in the financial industry. Therefore, despite many achievements, it is fair to say that ALM has not entirely fulfilled its high expectations.

The fresh approach on *deep hedging* by Hans Bühler, Lukas Gonon, Josef Teichmann, and Ben Wood (see [Bühler et al., 2019](#)) might pave a way for a new era in quantitative risk management. It is clear that classical yet analytically intractable problems from ALM can be tackled with techniques inspired by deep reinforcement learning in case of an underlying Markovian structure; see [Goodfellow et al. \(2016\)](#) by Ian Goodfellow, Yoshua Bengio, and Aaron Courville for a comprehensive overview of deep learning and [Sutton and Barto \(1998\)](#) by Richard S. Sutton and Andrew G. Barto for an introduction to reinforcement learning in particular. Reinforcement learning has found many applications in games (e.g., see [Silver et al., 2018](#) by a research team of Google's DeepMind), robotics (e.g., see [Kober et al., 2013](#) by Jens Kober, J. Andrew Bagnell, and Jan Peters) and autonomous vehicles (e.g., see [Kiran et al., 2020](#) by a joint publication of Valeo and academic partners). Recent monographs on machine learning for financial applications are [López de Prado \(2018\)](#) by Marcos López de Prado and [Dixon et al. \(2020\)](#) by Matthew F. Dixon, Igor Halperin, and Paul Bilokon. Regarding applications of neural networks for hedging and pricing purposes, some 200 research articles have accumulated over the last 30 years; [Ruf and Wang \(2019\)](#) by Johannes Ruf and Weiguan Wang provides a comprehensive literature review. Only a handful of these articles exploit the reinforcement learning paradigm. In contrast to this obvious applications of deep reinforcement learning to ALM, *deep hedging approaches to ALM* are simpler but still provide the solution relevant for business decisions. Notice, however, that neither Markovian assumptions are needed, nor value functions or dynamic programming principles: Deep ALM will simply provide an artificial asset-liability-manager who precisely solves the business problem (and not more) in a convincing way, i.e., provides ALM strategies along pre-defined future scenarios and stress scenarios. The fundamental ideas will be further elaborated on in the sections below. To the best of our knowledge, we are the first to tackle ALM systematically by deep hedging approaches.

There are many fundamental problems arising in mathematical finance that cannot be treated with explicit analytic formulae. Exemplarily, explicit pricing formulae for American options in

continuous time are yet unknown even in the simplest settings such as the Black-Scholes-framework. This is incredible in that most exchange traded options are of American style. Optimal stopping problems in genuine stochastic models involve so-called parabolic integro-differential inequalities, which can be numerically approximated by classical methods up to three dimensions; e.g., see the Section 10.7 in [Hilber et al. \(2013\)](#). The curse of dimensionality impedes the feasibility of classical finite element methods for applications involving more than three risk drivers. With today's computational power of a multi-core laptop, the runtime in six dimensions would last for many centuries. A popular circumvention is the regression based least squares Monte-Carlo method proposed by Francis A. Longstaff and Eduardo S. Schwartz; see [Longstaff and Schwartz \(2001\)](#). Some statisticians even perceive machine learning as a (non-linear) generalization of such regression techniques. Machine Learning therefore just introduces another regression basis to the Longstaff-Schwartz-algorithm. As an example a spectacular success was achieved by Sebastian Becker, Patrick Cheridito, and Arnulf Jentzen when they utilized deep neural networks to price American options in as much as 500 dimensions below 10 min; see [Becker et al. \(2019\)](#) for further details. In Deep ALM we even leave away the usual (and classically very useful) value function approach, which is also applied in the Longstaff-Schwartz-algorithm, but we go right away for an optimal trading strategy which is a priori completely untrained at each point in time (and not related to each other over time). As a consequence one is forced to abandon the fundamental desire of directly solving a very complex mathematical equation and must approve of the paradigm shift to the deep learning principle: "What would a clever, generic, well-experienced and non-forgetful artificial financial agent with a decent risk appetite do?" This change of mindset opens the room for many possibilities not only in quantitative finance. Computationally very intensive techniques such as nested Monte-Carlo become obsolete. We aim at utilizing these advances to tackle problems that were formulated back in the eighties, e.g., by Martin I. Kusy and William T. Ziemba (see [Kusy and Ziemba, 1983](#)), and have remained unsolved for the desired degree of complexity ever since. Thus, the technological impact of Deep ALM for the financial industry might be considerable.

The Markov assumption, roughly speaking that future states only depend on the current state and not the past, usually leads to comparatively tractable financial models and respective algorithms. Real world financial time series often feature a *leverage effect* (i.e., a relationship between the spot and the volatility) and a *clustering effect* (i.e., a persistence of low and high volatility regimes). Unless volatility becomes a directly traded product itself (which generally is not the case) and is perceived as an integrable part of the state, an effective hedging strategy can hardly be found analytically. This can be analyzed within the concept of *platonic financial markets* that was introduced in [Cuchiero et al. \(2020\)](#) by Christa Cuchiero, Irene Klein, and Josef Teichmann. Generally observable events and decision processes are adapted to a strict subfiltration of the whole market filtration that is generated by the obscure market dynamics. This concept goes in a similar direction as *hidden Markov models*; e.g., see the textbook [Elliott et al. \(1995\)](#) by Robert J. Elliott, Lakhdar Aggoun, and John B. Moore for a reference. Regarding Deep ALM, this inherent

intractability of platonic financial markets is not an impediment at all. Certainly, the replication portfolios along a trained deep neural network only depend on generally observable states. However, the financial market dynamics do not have to satisfy any Markovianity restrictions with respect to the observable states and/or traded products. The deep neural network jointly learns the pricing and the adaptation of hedging strategies in the presence of *market incompleteness* and even *model uncertainty*; e.g., see [El Karoui and Quenez \(1995\)](#) by Nicole El Karoui and Marie-Claire Quenez and [Avellaneda et al. \(1995\)](#) by Marco Avellaneda, Arnon Levy, and Antonio Parás for further details. Smaller filtrations or uncertainty can be coined in the framework of partial observability, which can be embedded into the Markovian framework at the price of pushing dimension toward infinity. Deep ALM approaches circumvent this step but rather directly write strategies in a flexible way as trainable functions of the observations. Besides these described aspects of partial observability, also frictions, like transaction costs or price impacts, can be easily treated within Deep ALM. Hence, Deep ALM unifies key elements of mathematical finance in a fundamental use case and it allows for the first time a tractable framework with fully realistic assumptions.

An essential prerequisite for the viability of Deep ALM in a treasury department is to come up with a sufficiently rich idea of the macro-economic environment and bank-specific quantities such as market risk factors, future deposit evolutions, credit rate evolutions and migrations, stress scenarios and all the parameterizations thereof. In the end, the current state of a bank must be observable. In contrast to dynamic programming, machine learning solutions are straightforward to implement directly on a cash-flow basis without any simplifications (e.g., a divisibility into simpler subproblems becomes redundant), and their quality can be assessed instantaneously along an arbitrarily large set of validation scenarios. Moreover, the paradigm of letting a smart artificial financial agent (similar to an artificial chess player) with a lot of experience choose the better option than human competitors allows to tame the curse of dimensionality, which is the pivotal bottleneck of classical methods. Further aspects such as intricate price dynamics, illiquidity, price impacts, storage cost, transaction cost and uncertainty can easily be incorporated without further ado; all these aspects are hardly treated in existing frameworks. Despite the higher degree of reality, Deep ALM is computationally less intensive than traditional methods. Risk aversion or risk appetite respectively can be controlled directly by choosing an appropriate objective associated with positive as well as negative rewards in the learning algorithm. Regulatory constraints can be enforced through adequately chosen penalties. Hence, Deep ALM offers a powerful and high-dimensional framework that supports well-balanced risk-taking and unprecedented risk-adjusted pricing. It surpasses prevalent replication strategies by far. Moreover, it allows to make the whole financial system more resilient through effective regulatory measures.

The article is structured as follows. In the Section 2, we describe a prevalent replication methodology for retail banks. Furthermore, we specify similar use cases from other industries. Subsequently, we explain the key notion of Deep ALM in the Section 3. A hint about the tremendous potential of Deep ALM is provided in the Section 4 in the context of a prototype. While accounting for

the regulatory liquidity constraints, a deep neural network adapts a non-trivial dynamic replication strategy for a runoff portfolio that outperforms static benchmark strategies conclusively. The last section elaborates on implications of Deep ALM and future research. Common misunderstandings from the computational and regulatory viewpoint are going to be clarified.

## 2. A static replication scheme

### 2.1. The treasury case of a retail bank

The business model of a retail bank relies on making profits from *maturity transformations*. To this end, a substantial part of their uncertain liability term structure is re-invested. The liability side mainly consists of term and non-maturing deposits. Major positions on the asset side are corporate credits and mortgages. Residual components serve as liquidity buffers, fluctuation reserves and general risk-bearing capital. The inherent optimization problem within this simplified business situation is already involved. Whereas most cash-flows originating from the liability side are stochastic, the bank is obliged to ensure a well-balanced liquidity management. Regarding short- and long-term liquidity, regulations require them to hold a *liquidity coverage ratio* (LCR) and a *net stable funding ratio* (NSFR) beyond 100%. In contrast, market pressure forces them to tighten the liquidity buffers in order to reach a decent return-on-equity. *Asset-liability-committees* (ALCO) steer the balance sheet management through risk-adjusted pricing and hedging instruments. Due to the high complexity, many banks predominantly define and monitor static replication schemes. It is worth noticing that ALCOs nowadays often meet only once a month and supervise simplistic measures such as the modified duration and aspects of hedge accounting on loosely time-bucketed liabilities. Many opportunities are overseen, and many risks remain unhedged; e.g., see the introductory chapter of [Spillmann et al. \(2019\)](#). Crucial business decisions are inconsistent over time and rather reactionary than pro-active.

[Table 1](#) illustrates the replication scheme of a generic retail bank. It is inspired by Section 4.2 in [Spillmann et al. \(2019\)](#). The notional allocated to fixed income currently amounts to EUR 100 m, which are rolled-over monthly subject to some static weights. The first line corresponds to a liquidity buffer that is held in non-interest bearing legal tender.

Exemplarily, 15% or EUR 15 m of the currently allocated capital are invested in fixed income instruments with a maturity of 1y. This portion is again split into equal portions, such that roughly the same volume matures each month. In order to end up with a closed cycle, this refers to as maintaining 12 tranches with a volume of EUR 1.25 m ( $= 1/12 \times 15m$ ). This redemption amount is re-invested again over a horizon of 1y at the latest market rates. The same roll-over procedure is conducted for all other maturities. This leaves us with 202 rolling tranches in total. Under the current premises, 6 tranches with a total volume of EUR 24.125 m mature each month and are re-invested consistently. The portfolio weights are only revised from time to time. They are typically the result of empirical and historical considerations. On the one hand, it is lucrative to increase the portion in longer terms as long-term yields

TABLE 1 Static investment scheme of a fixed income portfolio.

Term in Months	Weight in %	Volume in mEUR	Number of tranches	Volume in mEUR per tranche
0	10	10		
1	15	15	1	15.000
3	15	15	3	5.000
6	15	15	6	2.500
12	15	15	12	1.250
60	15	15	60	0.250
120	15	15	120	0.125
	100	100	202	24.125

are often but not always higher than short-term yields. On the other hand, if only a small portion is released each month, this leaves the bank with a significant interest and liquidity risk. This static scheme is typically accompanied with selected hedging instruments such that the sensitivities with respect to parallel shifts of the yield curve or the so-called *modified durations* for the asset and liability side roughly coincide.

Certainly, dynamic strategies (including a restructuring of all tranches) are superior to their static counterparts regarding profit potential and resilience of the enterprise. However, dealing directly with 202 tranches is often not feasible computationally. If we only add a simple dynamic feature and allow for arbitrary weights in the monthly re-investment process, this already introduces a considerable additional complexity from the analytical viewpoint. Even if one reduces the dimensionality of the yield curve dynamics, one faces an intricate nested optimization exercise over all re-investment instances. Utilizing traditional methods such as dynamic programming and accounting for a reasonable set of constraints is far from being trivial. Moreover, due to the necessary model simplifications, it is uncertain whether this additional complexity really pays off. It is therefore not surprising that many retail banks have been implementing a rather static replication scheme as described above. In the near future, this is likely to change once the novel and computationally much less intensive Deep ALM with a huge potential will be deployed.

## 2.2. Similar use cases

The notion of replication schemes may be applied analogously to other use cases. In the following, we describe three of them.

### 2.2.1. Actuarial perspective

The situation described above can be transferred one-to-one to the situation of insurance companies. The aspects that change are completely different regulatory circumstances, the asset illiquidity of hedging instruments (the *retrocession*) and the stochastic cash-flow generation as well as separate extreme risks on the liability side. Life insurance companies and pension funds must account for the exact product terms and the mortality characteristics of their policyholders. This includes, for instance, longevity and

pandemic risks. Non-life insurance companies model the annual loss distribution and its unwinding over consecutive years.

### 2.2.2. Procurement

The producing industry is faced with a long and complex process chain. Several raw commodities need to be bought, delivered, stored, processed, and designed into a consumer product. In order to control the inherent business risks and ensure profitability, a maturity transformation in the procurement process becomes inevitable. Even though the exact amount of the required materials is yet unknown, the company's exposure toward adverse price movements needs to be hedged accordingly. Despite the seemingly large daily trading volumes, any order comes with a price impact. Controlling the inventory is also worthwhile due to funding and storage cost. Therefore, slightly more involved than the treasury case described above, we have an ALM optimization in the presence of illiquidity, storage cost, transaction cost and uncertainty.

### 2.2.3. Hydroelectric power plants

The technological advances regarding renewable energy and the deregulation are having a huge impact on electricity markets. Wind and solar energy make the electricity prices weather-dependent and more volatile. Still, some segments of future markets have remained comparatively illiquid. A peculiarity of the electricity market are its intricate price dynamics (including negative prices) and a wide range of more or less liquidly traded future contracts over varying delivery periods (e.g., intraday, day ahead, weekend, week, months, quarters, years). This makes the estimate of *hourly price forward curves* (HPFC) that account for seasonality patterns (intraday, weekday, months), trends, holiday calendars and aggregated price information inferred from various future contracts (base, peak and off-peak) a very challenging exercise. Electricity producers may incur significant losses due to adverse contracts with locked-in tariffs and mis-timed hydropower production in reservoirs. This environment forces the suppliers to conduct a rigorous ALM, which comprises, for instance, a competitive pricing of fixed-for-floating contracts and daily optimized power production plans. The situation of delivering competitive yet still profitable margins in the presence of uncertainty is a situation banks have

been accustomed to for decades (with a different commodity and different rate dynamics). Deep ALM can be tailored to this very situation and provides exactly the technological solution to this delicate business problem. All one needs to establish is an appropriate scenario generator for the future markets and the HPFC. Many business-specific peculiarities can hardly be reflected by traditional optimization techniques. As an illustration, a turbine cannot be turned on and off on an hourly basis. This simple-looking constraint renders the dynamic programming principle impossible. However, this constraint can be accounted for in Deep ALM without further ado by incorporating suitable penalties.

### 3. Deep ALM

Beyond professional judgment, practical problems arising in ALM are often tackled either by linear programming or dynamic programming. The latter requires that the problem is divisible into simpler subproblems. For non-linear and high-dimensional cases, one typically exploits Monte-Carlo techniques in order to derive or validate strategies empirically along a handful of criteria. As a matter of fact, the level of sophistication for replicating strategies remains fairly limited due to the high intricacy and the lack of an adequate constructional approach. Deep ALM makes all these impediments obsolete. One directly implements the arbitrarily complex rule book of the use case and lets a very smart and non-forgetful artificial financial agent gain the experience of many thousand years in a couple of hours. By incentivizing the desired behavior and stimulating a swift learning process, one reaches superhuman level in due course.

We consider a balance sheet evolution of a retail bank over some time grid  $\{0, 1, 2, \dots, N\}$  in hours, days, weeks or months. The roll-forward is an alternating process between interventions and stochastic updates over time; (see Figure 1). At each time instance  $t$ , an artificial asset-liability-manager, subsequently denoted the “artificial financial agent” (AFA), assesses the re-investment of matured products as well as the restructuring of the current investment portfolio. This involves various asset classes, not just fixed income instruments. Correspondingly, the AFA is active in both the *primary market* of newly issued instruments and the *secondary market* of previously issued and circulating instruments. There may be further eligible transformations of this initial state with transaction cost, e.g., granting further credits or building additional liquidity buffers. The balance sheet components  $A_t$  (assets),  $L_t$  (liabilities) and  $E_t$  (equity) come with the superscripts “pre” and “post”, which refer to as before and after the balance sheet restructuring respectively. Preceding interventions, an auxiliary step for the calculation of taxes may be necessary. Subsequently, the economic scenario generator performs a stochastic roll-forward of the balance sheet. All accounts are updated to the latest circumstances and macro-economic factors, e.g., certain products pay off cash amounts such as coupons or dividends, certain products need to be written off due to the bankruptcy of the referenced entity, clients withdraw certain portions of their deposits, etc. It needs to be noted that the described situation is kept simplistic for illustrative purposes. It is the main target of future research work to reach an acceptable level of sophistication in order

to make the technology eligible for real world ALM. A first step toward that goal is Englisch et al. (2023).

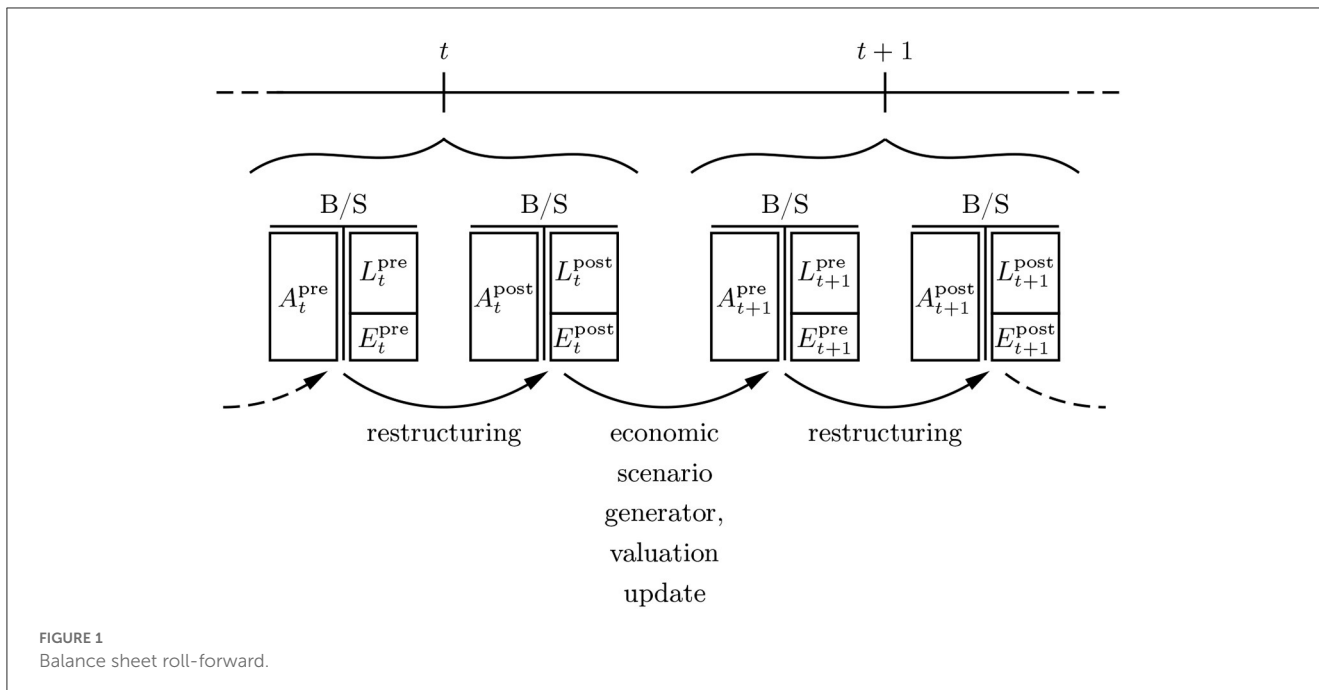
The superiority of the dynamic replication scheme requires that the trained AFA can perfectly trade-off benefits from a restructured portfolio and the incurred transaction cost of the restructuring. Furthermore, the trained AFA must impeccably weigh up different objectives such as complying with regulatory constraints, earning a risk-adjusted spread, maintaining an adequate level of liquidity, keeping sufficient margins for adverse scenarios, etc. The fundamental idea is to parameterize the decision process at each time instance  $t$  through a deep neural network (DNN) and to incrementally improve its performance through techniques inspired by machine learning. By concatenating these deep neural networks along the time axis, one ends up at a holistic dynamic replication strategy that easily deals with non-stationary environments. The learning process incentivizes the maximization of an adequately chosen reward function by incrementally updating the weights of the feedforward networks through *backpropagation*; (see Figure 2). This idea allows to tame the curse of dimensionality. These new possibilities of Deep ALM allow for a more frequent supervision while accounting for arbitrarily many side constraints and complicated market frictions such as illiquidity and offsetting effects of transaction cost in hedges. Thus, the full power of dynamic strategies can be deployed.

Traditional numerical schemes describe an algorithm either in order to solve an equation or in order to optimize a certain functional. Typically, algorithms are only accepted by the mathematical community if they come with two supplements, namely a proof that the recipe terminates with a certain accuracy at the desired level, on the one hand, and a statement about the rate-of-convergence, on the other hand. In the context of Deep ALM, we coin the notion of *convincing strategies* because we are likely not able to prove that the randomly initialized deep neural network will end up at some optimal replication strategy. In fact, we intend to derive non-trivial but still realistic dynamic replication strategies within a reasonable time that significantly outperform those strategies currently used in the financial industry. From the mathematical viewpoint, Deep ALM entails a delicate paradigm shift. The results will be derived empirically based on scenarios representing the experience of many thousand years. By repeating the learning process a few times, we will corroborate the robustness of the approach. Furthermore, we will validate and challenge the performance of the deep neural network on unseen testing data. If the gradient descent algorithm can only further reduce the training loss with some overfitting strategy at the cost of the validation loss, we have ended up at a local optimizer. Even though we will not be able to prove global optimality, the strategy might still beat conventional models as well as seasoned professionals and, thus, be more than convincing.

## 4. Empirical study on a stylized case

### 4.1. The runoff case

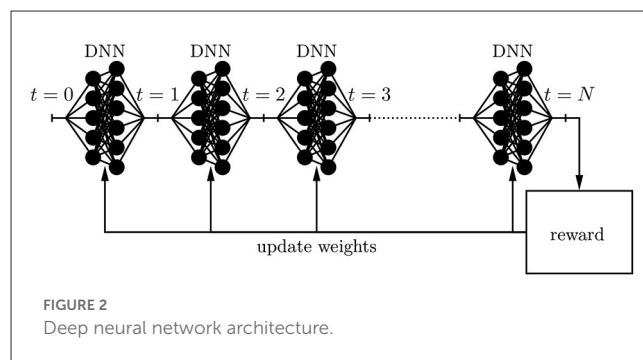
We aim at optimizing the replication strategy when unwinding a runoff portfolio over 10 years in the context of an equity index and a large fixed income universe. At every time instance, roughly



200 bonds are active. This prototype demonstrates the feasibility of Deep ALM in a stylized case. In the end, every balance sheet roll-forward is a superposition of runoff portfolios. Furthermore, we give an indication about the potential of dynamic replication strategies.

The liabilities are backed by a static legacy investment portfolio. Roughly two thirds of the assets are invested in the fixed income market and 10% are invested in equities. The remainder is held in cash. Every month, six new bonds are issued in the primary market with the face amount of 100 monetary units and the maturities 1, 3, 6m, 1, 5, and 10y. The first three series only pay a coupon at the final redemption date, the other three pay semi-annual coupons. At issuance, all bonds trade at par. The legacy investment portfolio has attributed equal portions to all of the six bond series and has conducted an equally distributed roll-over strategy based on historical data from the European Central Bank (ECB). Future yield curve scenarios are simulated consistent with a principal component analysis. For simplicity, we assume that there is no secondary market. Hence, bonds cannot be sold on in the market and must be held to maturity. Furthermore, we preclude credit defaults. Regarding equities, we generated scenarios based on a discretized geometric Brownian motion.<sup>1</sup> Changing the position in equities involves proportional transaction cost. Due to legal requirements, at least 10% of the asset portfolio must be held in cash; this is just an arbitrary choice in order to incorporate a regulatory constraint. Any discrepancy is penalized monthly with a high interest rate. The initial balance hosts assets worth 100 monetary units. The term structure of liabilities with the same

<sup>1</sup> We make this popular choice for reasons of simplicity and reproducibility. For well-founded considerations in the financial industry, more complicated models or more realistic scenarios would be advisable. Still, the applicability of the approach would not be affected.



face amount is spread across the time grid according to some beta distribution.

## 4.2. The implementation strategy

A well-established approach to tackle quantitative finance problems with deep reinforcement learning is to follow the standard routine inspired by games. A game is a stochastic environment in which a player can act according to a set of rules and can influence the course of events from start to finish to a certain degree. The game terminates after a prespecified number of rounds. A rule book includes the model drivers and their interdependence, the constraints and the penalties for constraint violations, the control variables and the objective function. Notably, we actually do not solve the game for all initial states nor do we assume any Markovian structure. We implemented the “game” according to the rule book specified below in a Python script such that the random outcome of any strategy could be simulated arbitrarily often. To this end, we defined an *economic*

scenario generator that hosted the yield curve dynamics and spot prices for equities. Having this environment in place, one can test the performance of any naive or traditionally derived replication portfolio. The step of translating the rule book into model dynamics is crucial for understanding the game, for ensuring a well-posed optimization problem and for streamlining the logic of the subsequent tensorization. The evolution of  $\mathbb{R}^N$ -valued quantities can be inferred from historical time series with a principal component analysis as proposed by Andres Hernandez in Hernandez (2016). Subsequently, the game is embedded in a deep neural network graph architecture using TensorFlow such that an AFA can run many games in parallel and adaptively improve its performance. The graph constitutes of many placeholders and formalizes the logic of the game. The first decision process of the AFA is managed through the randomly initialized weights of a neural network with a handful of hidden layers. Its input layer is a parameterization of the initial state, concerning both the balance sheet structure and the market quotes, and the output layer characterizes the eligible balance sheet transformation. The balance sheet roll-forward is accomplished by connecting the output of the neural network with a whole bunch of efficient tensor operations reflecting the balance sheet restructuring, new inputs for the non-anticipative scenario updates and the input layer of the next decision process. Iterating this process yields the desired deep neural network.

### 4.3. Term structures

We are given a discrete time grid  $\mathbb{T}_0 = \{0, 1, 2, \dots, N\}$  in months, where  $N = 120$  refers to as a planning horizon of 10 years. The initial term structure of liabilities with the face amount of 100 monetary units is spread across the time grid according to a beta distribution with the parameters  $a = 1.5$  and  $b = 2.5$ , i.e., if  $F_{a,b}$  denotes the corresponding cumulative distribution function, then the cash-flow

$$L^{(t)} := \left( F_{a,b}(t/N) - F_{a,b}((t-1)/N) \right) \times 100$$

becomes due in month  $t \in \mathbb{T} := \mathbb{T}_0 \setminus \{0\}$ . We briefly write  $L := (L^{(1)}, L^{(2)}, \dots, L^{(N)})$  for the collection of all payables. At some fixed time instance, any active bond is characterized through its future cash-flows  $B = (B^{(1)}, B^{(2)}, \dots, B^{(N)}) \in \mathbb{R}^N$ . If  $Y = (Y^{(1)}, Y^{(2)}, \dots, Y^{(N)}) \in \mathbb{R}^N$  denotes the current yield curve with the yield  $Y^{(T)}$  for the maturity  $T/12$  in years and  $D = (D^{(1)}, D^{(2)}, \dots, D^{(N)}) \in \mathbb{R}^N$  with  $D^{(T)} = e^{-T/12 \cdot Y^{(T)}}$  for all  $T \in \mathbb{T}$  the consistent discount factors, then the value of  $B$  is simply  $V(B) = \langle D, B \rangle$ , where  $\langle \cdot, \cdot \rangle$  stands for the standard inner product. The coupons are chosen such that each bond trades at par. The cash-flow generation at issuance is straightforward. It corresponds to solving a linear equation. Exemplarily, a bond with a maturity of 5y and semi-annual coupons carries an annualized coupon rate

$$c = \frac{12}{6} \cdot \frac{1 - D^{(60)}}{\sum_{k=1}^{10} D^{(6 \cdot k)}}$$

Therefore, it holds

$$B^{(6 \cdot k)} = \frac{c}{2} \cdot 100 \quad \text{for all } k \in \{1, 2, \dots, 9\}, \quad B^{(60)} = \left( 1 + \frac{c}{2} \right) \cdot 100,$$

and all other components are zero. When going from one time instance to the next,  $B^{(1)}$  is paid off and one simply applies to the scheme  $B$  the linear update operator

$$U = \begin{pmatrix} 0 & 1 & & 0 \\ & 0 & 1 & \\ & & \ddots & \ddots \\ & & & 0 & 1 \\ 0 & & & & 0 \end{pmatrix} \in \mathbb{R}^{N \times N}.$$

A  $k$ -fold application of the update operator is denoted by  $U^k$ . We utilize data from the European Central Bank (ECB) in order to model the historical and future yield curves. ECB publishes for each working day the Svensson parameters  $\beta_0, \beta_1, \beta_2, \beta_3, \tau_1$ , and  $\tau_2$  of its yield curve. The corresponding yield curve is the vector  $Y = (Y^{(1)}, Y^{(2)}, \dots, Y^{(N)}) \in \mathbb{R}^N$  with

$$Y^{(T)} = \beta_0 + \beta_1 \frac{1 - e^{-T/(12\tau_1)}}{T/(12\tau_1)} + \beta_2 \frac{1 - e^{-T/(12\tau_1)}}{T/(12\tau_1) - e^{-T/(12\tau_1)}} + \beta_3 \frac{1 - e^{-T/(12\tau_2)}}{T/(12\tau_2) - e^{-T/(12\tau_2)}}$$

We fix some historical time series in  $\mathbb{R}^N$  and conduct a principal component analysis. To this end, we assume stationary daily yield curve increments  $\Delta X$  with expected value  $\mathbb{E}[\Delta X] \equiv \mu \in \mathbb{R}^N$  and covariance matrix  $cov(\Delta X) \equiv Q \in \mathbb{R}^{N \times N}$ . We spectrally decompose  $Q = \Lambda \Lambda^T$  into the normalized eigenvectors  $\Lambda \in \mathbb{R}^{N \times N}$  and the ordered eigenvalues  $L = \text{diag}\{\lambda^{(1)}, \lambda^{(2)}, \dots, \lambda^{(N)}\}$  with  $\lambda^{(1)} \geq \lambda^{(2)} \geq \dots \geq \lambda^{(N)}$ . The last historical yield curve is the day when we acquire the runoff portfolio. The stochastic yield curve increments from 1 month to the next are consistently sampled according to

$$\Delta Y = 22 \cdot \mu + \sum_{k=1}^n Z_k \Lambda_{\cdot k},$$

where we choose  $n = 3$  and  $Z_k \sim \mathcal{N}(0, \sigma_k^2)$  for  $\sigma_k = \sqrt{22 \cdot \lambda^{(k)}}$ . The factor 22 accounts for the number of trading days per month. If  $i$  denotes the maturity in number of months, the legacy bond portfolio has been investing each month a face amount of  $15/i$  monetary units in that particular bond series for  $i \in I := \{1, 3, 6, 12, 60, 120\}$ . This is consistent with the replication scheme from Table 1.

### 4.4. The equity market

We assume that the stochastic process  $S = (S_t)_{t \in \mathbb{T}_0}$  follows a discretized geometric Brownian motion, i.e., it holds for all  $t \in \mathbb{T}$

$$\log \frac{S_t}{S_{t-1}} \sim \mathcal{N} \left( \left( m - \frac{1}{2} s^2 \right) \cdot \frac{1}{12}, \frac{s^2}{12} \right)$$

with initial level  $S_0 = 100$ , drift  $m = 5\%$  and volatility  $s = 18\%$ . Changing the position in equities involves proportional transaction cost  $\kappa = 0.50\%$ .

### 4.5. The balance sheet

We proceed iteratively. Let  $Y_t \in \mathbb{R}^N$  denote the current yield curve for  $t \in \mathbb{T}_0 \setminus \{N\}$  and  $D_t$  consistent discount factors. Before restructuring, the left hand side of the balance sheet with a total present value of  $A_t^{\text{pre}}$  consists of three additive components: the held cash  $C_t^{\text{pre}}$ , the legacy bond portfolio with the value  $V_t^{\text{pre}}$ , and the equity position worth  $G_t^{\text{pre}} = \Delta_t^{\text{pre}} \cdot S_t$ . The right hand side of the balance sheet consists of the liabilities

$$L_t^{\text{pre}} = \langle D_t, U^t L \rangle$$

and the residue  $E_t^{\text{pre}} = A_t^{\text{pre}} - L_t^{\text{pre}}$ .

### 4.6. The deep neural network architecture

For any  $t \in \mathbb{T}_0 \setminus \{N\}$ , we consider a feedforward neural network

$$F_t = (\phi \circ W_t^{(2)}) \circ (\phi \circ W_t^{(1)}) \circ (\phi \circ W_t^{(0)})$$

with some affine functions

$$W_t^{(0)} : \mathbb{R}^{10} \rightarrow \mathbb{R}^{15}, W_t^{(1)} : \mathbb{R}^{15} \rightarrow \mathbb{R}^{15}, W_t^{(2)} : \mathbb{R}^{15} \rightarrow \mathbb{R}^7.$$

and the ReLU activation function  $\phi(x) = \max\{x, 0\}$ . The input layer consists of the *leverage ratio*  $E_t^{\text{pre}}/A_t^{\text{pre}}$ , the *liquidity ratio*  $C_t^{\text{pre}}/A_t^{\text{pre}}$ , the *risk portion*  $G_t^{\text{pre}}/A_t^{\text{pre}}$ , the holding in equities  $\Delta_t^{\text{pre}}$  and the yields  $Y_t^{(i)}$  for  $i \in I$  of the latest bond series. The output layer reveals the outcome of the restructuring at the time instance  $t$ . The first six components represent the holdings  $h_t^{(i)}$  for  $i \in I$  in the latest bond series, the last component represents the holding  $\Delta_t^{\text{post}}$  in equities. Note that the activation function implies long-only investments.

### 4.7. The restructuring

The restructuring implicates the updated balance sheet items

$$\begin{aligned} C_t^{\text{post}} &= C_t^{\text{pre}} - 100 \cdot \sum_{i \in I} h_t^{(i)} - \left( (\Delta_t^{\text{post}} - \Delta_t^{\text{pre}}) \right. \\ &\quad \left. + \kappa \cdot |\Delta_t^{\text{post}} - \Delta_t^{\text{pre}}| \right) \cdot S_t, \\ V_t^{\text{post}} &= V_t^{\text{pre}} + 100 \cdot \sum_{i \in I} h_t^{(i)}, \\ G_t^{\text{post}} &= \Delta_t^{\text{post}} \cdot S_t, \\ A_t^{\text{post}} &= C_t^{\text{post}} + V_t^{\text{post}} + G_t^{\text{post}}, \\ L_t^{\text{post}} &= L_t^{\text{pre}}, \\ E_t^{\text{post}} &= A_t^{\text{post}} - L_t^{\text{post}}. \end{aligned}$$

### 4.8. The roll-forward

Let  $B_t^{\text{post}} \in \mathbb{R}^N$  denote the aggregated future cash-flows of the whole fixed income portfolio after restructuring at time  $t$ . Furthermore, let  $\pi^{(1)} : \mathbb{R}^N \rightarrow \mathbb{R}$  denote the projection onto the first component. If one does not adhere to the regulatory constraint of holding at least 10% of the assets in cash, one is penalized with liquidity cost of  $p = 24\%$  per annum on the discrepancy. Once the yield curve has been updated stochastically to the state  $Y_{t+1}$  with consistent discount factors  $D_{t+1}$  and the equity index has been updated stochastically to the state  $S_{t+1}$ , this leaves us with the roll-forward

$$\begin{aligned} C_{t+1}^{\text{pre}} &= C_t^{\text{post}} + \pi^{(1)}(B_t^{\text{post}}) - \pi^{(1)}(U^t L) \\ &\quad - \frac{p}{12} \cdot \max\{10\% \cdot A_t^{\text{post}} - C_t^{\text{post}}, 0\}, \\ V_{t+1}^{\text{pre}} &= \langle D_{t+1}, UB_t^{\text{post}} \rangle, \\ G_{t+1}^{\text{pre}} &= \Delta_t^{\text{post}} \cdot S_{t+1} = \Delta_{t+1}^{\text{pre}} \cdot S_{t+1}, \\ A_{t+1}^{\text{pre}} &= C_{t+1}^{\text{pre}} + V_{t+1}^{\text{pre}} + G_{t+1}^{\text{pre}}, \\ L_{t+1}^{\text{pre}} &= \langle D_{t+1}, U^{t+1} L \rangle, \\ E_{t+1}^{\text{pre}} &= A_{t+1}^{\text{pre}} - L_{t+1}^{\text{pre}}. \end{aligned}$$

### 4.9. Objective

For different maturities  $T \in \mathbb{T}$ , a common approach in mathematical finance is to maximize  $\mathbb{E}[u(E_T^{\text{post}}/E_0^{\text{pre}})]$ , where

$$u(x) = \begin{cases} \frac{x^{1-\gamma} - 1}{1-\gamma} & , \text{ if } \gamma \neq 1 \\ \log x & , \text{ if } \gamma = 1 \end{cases}$$

denotes the iso-elastic utility function with constant relative risk aversion  $\gamma \geq 0$ . This is not directly applicable in our case since we cannot prevent  $E_T^{\text{post}}$  from becoming negative. Instead, we aim at maximizing

$$\mathbb{E}\left[ u\left( \left( \varepsilon + \phi(E_T^{\text{post}}) \right) / E_0^{\text{pre}} \right) \right]$$

for a small constant  $0 < \varepsilon \ll 1$ . Provided that the final equity distribution is positive, the case  $\gamma = 1$  corresponds to the quest for the *growth optimal portfolio* that maximizes the expected log-return; see [Platen and Heath \(2006\)](#). An intriguing alternative is quadratic hedging, where one aims at minimizing

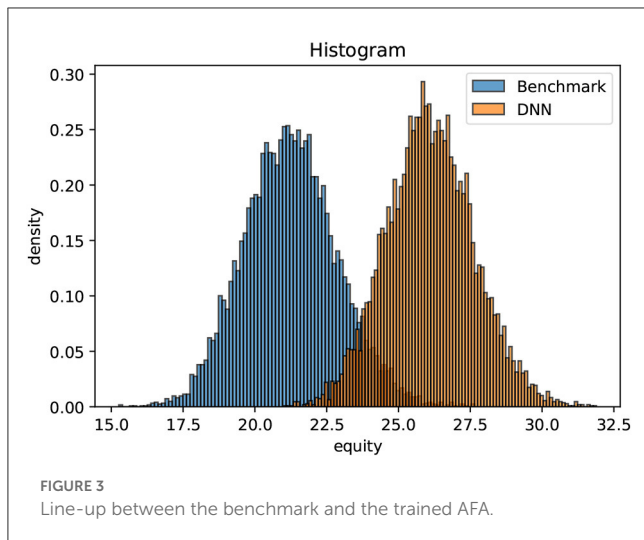
$$\mathbb{E}\left[ \left( E_T^{\text{post}} - (1+r)^{T/12} E_0^{\text{pre}} \right)^2 \right]$$

for some annualized target return-on-equity  $r > 0$ . Either case leads to convincing strategies.

### 4.10. Results

The optimization is only non-trivial for non-negative yield curves. If the mid-term yields are consistently negative (as they have been in the EUR area for quite some years), holding cash





is superior to the considered fixed income investments (within the simplified premises specified above). As starting point, we choose the ECB yield curve as per December 31, 2007. In that case, the initial value of the liabilities amount to  $L_0^{\text{pre}} = 86.0$  monetary units. As outlined above, the exercise of optimizing the equity distribution after the settlement of all liabilities is extremely challenging with conventional methods. The investment universe consists at each time instance of 204 active assets, namely the fixed income universe, an equity index and legal tender. We propose a simple benchmark strategy that allocates the available cash beyond 10% of  $A_t^{\text{pre}}$  to bonds with a maturity of  $1m$ , provided that the coupons are positive. Equities are divested linearly over time. Moreover, we train an AFA on 10 000 unwinding scenarios following the ideas of Deep ALM. An until only recently inaccessible dynamic strategy outperforms the benchmark strategy conclusively after a short learning process of roughly 10 min. Using another 10 000 previously unseen validation scenarios, the enclosed chart illustrates the final equity distributions after 10 years for a classical static strategy in blue and for the dynamic strategy in orange. The dynamic strategy does not involve extreme risk taking. It simply unveils hidden opportunities. The systematic excess return-on-equity with respect to the benchmark is beyond 2% per annum; (see Figure 3).

## 5. Outlook and further research

### 5.1. Potential of Deep ALM

The real potential can hardly be foreseen. One achievement will lead to another and trigger further accomplishments. Despite the optimism, the technique is not mature enough and requires an adequate level of research. Market participants across all use cases, who we spoke with, like the intriguing idea of Deep ALM. However, only a few dare being the first movers in this precompetitive phase. ALM is at the heart of risk management for any enterprise in the financial industry. In the light of the new opportunities, it is not acceptable how simplistic and ineffective prevalent replication schemes are to this day. All of us are directly affected by unhedged

risks of bank deposits, insurance policies and pension funds. This can be illustrated exemplarily with Finland whose banking sector caused a devastating financial crisis in the early 1990s. In the aftermath, official unemployment rates escalated to roughly 18%. Deep ALM offers a powerful framework that supports well-balanced risk-taking and unprecedented risk-adjusted pricing. Thus, its rigorous application may prevent a financial crisis sooner or later. Likewise, it is becoming increasingly important to spend resources wisely. Deep ALM promotes this attitude with significant efficiency gains in the procurement of commodities and the energy production; see also Curin et al. (2021) for a real world application. These are promising prospects for Deep ALM indeed.

### 5.2. Necessity of further research

Some people argue that machine learning will facilitate the enterprise risk management, make it more effective and make some of the current tasks even redundant. We agree that the new opportunities will allow for more consistent and sound risk taking. However, we are convinced that legitimate applications of Deep ALM will entail a lot of additional work since the previously inaccessible analyses do not come for free. Moreover, they will raise a huge portion of new economic challenges and obligations. Thus, it is essential that the research community prepares a solid ground for this technological transformation.

### 5.3. Feasibility

Optimizing a whole term structure of assets and liabilities by utilizing techniques inspired by deep reinforcement learning is a novel approach and its feasibility needs to be demonstrated. In the light of recent advances, we are optimistic that this goal can be achieved in due course. Nonetheless, this challenge opens a great field of research questions from the mathematical, computational and economic viewpoints. Regarding feasibility, it is an important concern that we do not just want to feed an utterly complex mathematical system into a powerful supercomputer and see whether we get something out. In contrast, we aim at formulating and optimizing dynamic replication schemes that meet the requirement of industrial applicability, on the one hand, and operability on a moderately enhanced computer, on the other hand. Particularly, this involves an assessment whether the behavior of the trained AFA is interpretable, whether it follows sound economic reasoning and whether the non-anticipative dynamic replication strategies fulfill the basic requirement of practical viability. Solving a high-dimensional utility maximization problem that entails unrealistic leverage or arbitrary rebalancing frequencies is futile for real world applications. For this purpose, Deep ALM must trade off many objectives such as cost and benefit of liquidity, drawdown constraints, moderate leverage ratios, profit expectation, regulatory interventions, risk appetite, short sale restrictions, transaction cost, etc. Most aspects can be accounted for in the model by means of penalties that incentivize the AFA's behavior accordingly. Others such as drawdown constraints and profit expectation typically enter the target function of the optimization. Finding the right balance

between all the different goals without adversely affecting the robustness of the learning process entails some engineering work. However, once a convincing and stable economic model choice has been established, this has a tremendous impact on the economic research field. Previously inaccessible price tags to crucial economic notions such as, for instance, cost of liquidity for a retail bank can be evaluated quantitatively. All one needs to do is to line up performance measures of trained AFAs when short sale restrictions of the cash position are enforced and dropped respectively.

## 5.4. Goal-based investing

A closely related subject to Deep ALM is goal-based investing; e.g., see [Browne \(1999\)](#) by Sid Browne for a neat mathematical treatment. One aims at maximizing the probability to reach a certain investment goal at a fixed maturity. From the theoretical viewpoint with continuous rebalancing, this is equivalent to replicating a volume-adjusted set of digital options (unless the target return is smaller or equal than the risk-free rate). Despite being very intuitive, this result is only of limited use for practical applications. Discrete rebalancing and the payoff-discontinuity at the strike lead, as one gets closer and closer to maturity, to an inherent shortfall risk way beyond the invested capital. It needs to be investigated whether Browne's setting can be modified accordingly for real world applications. This involves more realistic dynamics of the economic scenario generator, exogenous income (as inspired by Section 7 of [Browne, 1999](#)), maximal drawdown constraints and a backtesting with historical time series. An attempt in this direction can be found in [Krabichler and Wunsch \(2021\)](#).

## 5.5. Generalization

If we take one step further and look at ALM as an abstract framework, Deep ALM may contribute crucially to urgent challenges of our society such as combating the climate change or dealing with pandemics. It trains a player to reach a certain goal while minimizing the involved cost. Regarding the climate change, there is a complex trade-off in implementing protection measures, incentivizing behavioral changes, transforming the energy mix, etc. Regarding COVID-19, rigorous emergency policies were enforced throughout the world in order to protect individual and public health. These measures involved extreme cost such as temporary collapses of certain industries, closed schools, restricted mobility, etc. These are again high-dimensional problems, which cannot be tackled with classical methods. Deep ALM is a first step toward these more difficult challenges.

## 5.6. Impediments

A challenge will be to overcome the common misunderstanding that deep learning requires a lot of historical data (that banks or commodity companies usually do not have), that deep learning is an incomprehensible black box, and that regulators will never approve of Deep ALM. Firstly, the above approach does not

necessarily require any historical data. All one needs are model dynamics of ordinary scenarios and stress scenarios. The training data can be fully built upon these premises synthetically. Still, for the calibration of selected parameters, the availability of historical data can of course be useful. Secondly, the performance of deep neural networks can be validated without further ado along an arbitrarily large set of scenarios. To our mind, this is one of the most striking features of Deep ALM (which is not satisfied by traditional ALM models). Traditional hybrid pricing and risk management models entailed an intensive validation process. This intensive work becomes obsolete and can be invested into stronger risk management platforms instead. Thirdly and lastly, as long as one provides the regulators with convincing arguments regarding the accuracy and resilience of the approach (e.g., in terms of a strong validation and model risk management framework), they will never object to use deep learning. As a rule, regulators emphasize being model agnostic.

## 5.7. Regulatory perspective

The complex regulatory system in place is supposed to make the financial system less fragile. However, designing and implementing effective policies for the financial industry, which do not promote wrong incentives and which are not (too) detrimental to economic growth, is very demanding. This is highlighted exemplarily in the IMF working paper ([Sharma et al., 2003](#)). An illustration that certain measures do not necessarily encourage the desired behavior can be found in [Grossmann et al. \(2016\)](#) by Martin Grossmann, Markus Lang, and Helmut Dietl. Initially triggered by the bankruptcy of Herstatt Bank (1974), the Basel Committee on Banking Supervision has established a framework that specifies capital requirements for banks. It was revised heavily twice since its publication in 1988. The current framework, also known as "Basel III" (see [Basel Committee on Banking Supervision, 2011](#)), was devised in the aftermath of the 2007/2008 credit crunch. Amongst others, it proposed requirements for the LCR and the NSFR. The impact of these liquidity rules was empirically assessed in [Dietrich et al. \(2014\)](#) by Andreas Dietrich, Kurt Hess, and Gabrielle Wanzanried. Due to limited historical data and the intricate complexity between the plethora of influencing factors, the true impact of the new regime can hardly be isolated. Deep ALM opens an extremely appealing new way of appraising the effectiveness of regulatory measures. To this end, one investigates how a very smart and experienced AFA changes her behavior when certain regulatory measures are enabled or disabled respectively. This approach allows to truly identify the impact of supervisory interventions and the compatibility amongst each other. These quantitative studies will be a promising supplement of Deep ALM.

## 5.8. Model risk management

The supervisory letter SR 11-7 of the U.S. Federal Reserve System ("Fed") has become a standard for considerations of *model*

*risk management*; see Board of Governors of the Federal Reserve System (2011). Controlling the inherent model risk and limitations will be a key challenge before Deep ALM may be exploited productively. While the performance of proposed strategies can be checked instantaneously for thousands of scenarios within the model, an assessment of the discrepancy between the Deep ALM model and the real world is a crucial research topic. Once a robust Deep ALM learning environment has been established, these aspects can be explored by means of various experiments. For instance, one can analyse the sensitivities of the acquired replication strategies with respect to the model assumptions or the performance of those in the presence of parameter uncertainties. These uncertainties may concern both the assumed states (e.g., the term structure of deposit outflows) and the concealed scenario generation for both the training and the validation (e.g., the drift assumptions of equities). Some people propose to utilize *generative adversarial networks* (GANs) in order to identify both the nature and the impact of adverse scenarios, on the one hand, and to improve the AFA through additional training iterations even further, on the other hand; e.g., see Flaig and Junike (2022).

## 5.9. Risk policy perspective

Another intriguing question is the quest for locally stable strategic asset allocations. Typically, the trained neural network will reshuffle the balance sheet structure right from the beginning. Characterizing *no-action regions* corresponds to locally optimal initial states. Concerning the model risk management, revealing the link between different model assumptions and the internal risk policies, on the one hand, and the no-action regions, on the other hand, will be a powerful tool for risk committees.

## References

- Avellaneda, M., Levy, A., and Parás, A. (1995). Pricing and hedging derivative securities in markets with uncertain volatilities. *Appl. Math. Finance*. 2, 73–88. doi: 10.1080/13504869500000005
- Bardenhewer, M. M. (2007). “Modeling non-maturing products,” in *Liquidity Risk Measurement and Management*, eds L. Matz and P. Neu (John Wiley & Sons). doi: 10.1002/9781118390399.ch10
- Basel Committee on Banking Supervision (2011). *Basel III: A Global Regulatory Framework for More Resilient Banks and Banking Systems (revised version)*. Available online at: <https://www.bis.org/publ/bcbs189.htm>
- Becker, S., Cheridito, P., and Jentzen, A. (2019). Deep optimal stopping. *J. Mach. Learn. Res.* 20, 1–25. Available online at: <http://jmlr.org/papers/v20/18-232.html>
- Board of Governors of the Federal Reserve System (2011). *SR 11-7: Guidance on Model Risk Management*. Available online at: <https://www.federalreserve.gov/supervisionreg/srletters/sr1107.htm>
- Browne, S. (1999). Reaching goals by a deadline: digital options and continuous-time active portfolio management. *Adv. Appl. Prob.* 31, 551–577. doi: 10.1239/aap/1029955147
- Bühler, H., Gonon, L., Teichmann, J., and Wood, B. (2019). Deep hedging. *Quant. Fin.* 19, 1271–1291. doi: 10.1080/14697688.2019.1571683
- Consigli, G., and Dempster, M. A. H. (1998). Dynamic stochastic programming for asset-liability management. *Ann. Oper. Res.* 81, 131–162. doi: 10.2139/ssrn.34780
- Cuchiero, C., Klein, I., and Teichmann, J. (2020). A fundamental theorem of asset pricing for continuous time large financial markets in a two filtration setting. *Theory Prob. Appl.* 65, 388–404. doi: 10.1137/S0040585X97T990022
- Curin, N., Kettler, M., Kleisinger-Yu, X., Komaric, V., Krabichler, T., Teichmann, J., et al. (2021). A deep learning model for gas storage optimization. *Decis. Econ. Fin.* 44, 1021–1037. doi: 10.1007/s10203-021-00363-6
- Dietrich, A., Hess, K., and Wanzenried, G. (2014). The good and bad news about the new liquidity rules of Basel III in Western European countries. *J. Bank. Fin.* 44, 13–25. doi: 10.1016/j.jbankfin.2014.03.041
- Dixon, M. F., Halperin, I., and Bilokon, P. (2020). *Machine Learning in Finance*. Berlin; Heidelberg: Springer Verlag. doi: 10.1007/978-3-030-41068-1
- El Karoui, N., and Quenez, M.-C. (1995). Dynamic programming and pricing of contingent claims in an incomplete market. *SIAM J. Control Opt.* 33, 29–66. doi: 10.1137/S0363012992232579
- Elliott, R. J., Aggoun, L., and Moore, J. B. (1995). *Hidden Markov Models*. Berlin; Heidelberg: Springer Verlag
- Englich, H., Krabichler, T., Müller, K. J., and Schwarz, M. (2023). Deep treasury management for banks. *Front. Artif. Intell.* 6, 1120297. doi: 10.3389/frai.2023.1120297
- Fabozzi, F. J., Leibowitz, M. L., and Sharpe, W. F. (1992). *Investing: The Collected Works of Martin L. Leibowitz*. Probus Professional Pub. Available online at: [https://books.google.ch/books/about/Investing.html?id=K0AUQAAMAAJ&redir\\_esc=y](https://books.google.ch/books/about/Investing.html?id=K0AUQAAMAAJ&redir_esc=y)
- Flaig, S., and Junike, G. (2022). Scenario generation for market risk models using generative neural networks. *arXiv preprint arXiv:2109.10072*. doi: 10.3390/risks10110199
- Fraudendorfer, K., and Schürle, M. (2003). Management of non-maturing deposits by multistage stochastic programming.

## Data availability statement

Publicly available datasets were analyzed in this study. This data can be found here: [https://www.ecb.europa.eu/stats/financial\\_markets\\_and\\_interest\\_rates/euro\\_area\\_yield\\_curves/html/index\\_en.html](https://www.ecb.europa.eu/stats/financial_markets_and_interest_rates/euro_area_yield_curves/html/index_en.html).

## Author contributions

Both authors listed have made a substantial, direct, and intellectual contribution to the work and approved it for publication.

## Funding

Open access funding by ETH Zurich.

## Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

## Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

- Eur. J. Oper. Res.* 151, 602–616. doi: 10.1016/S0377-2217(02)00626-4
- Goodfellow, I., Bengio, Y., and Courville, A. (2016). *Deep Learning*. MIT Press. Available online at: <https://mitpress.mit.edu/9780262035613/deep-learning/>
- Grossmann, M., Lang, M., and Dietl, H. (2016). Why taxing executives' bonuses can Foster risk-taking behavior. *J. Instit. Theoret. Econ.* 172, 645–664. doi: 10.1628/093245616X14689190842778
- Hernandez, A. (2016). *Model Calibration with Neural Networks*. Available online at: <https://ssrn.com/abstract=2812140>
- Hilber, N., Reichmann, O., Schwab, C., and Winter, C. (2013). *Computational Methods for Quantitative Finance*. Berlin; Heidelberg: Springer Verlag. doi: 10.1007/978-3-642-35401-4
- Kiran, B. R., Sobh, I., Talpaert, V., Mannion, P., Al Sallab, A. A., Yogamani, S., et al. (2020). Deep reinforcement learning for autonomous driving: a survey. *arXiv preprint arXiv:2002.00444*. Available online at: <https://ieeexplore.ieee.org/document/9351818>
- Kober, J., Bagnell, J. A., and Peters, J. (2013). Reinforcement learning in robotics: a survey. *Int. J. Robot. Res.* 32, 1238–1274. doi: 10.1177/0278364913495721
- Krabichler, T., and Wunsch, M. (2021). Hedging goals. *arXiv preprint arXiv:2105.07915*.
- Kusy, M. I., and Ziemba, W. T. (1983). *A Bank Asset and Liability Management Model*. IIASA Collaborative Paper. Available online at: <http://pure.iiasa.ac.at/id/eprint/2323/>
- Longstaff, F. A., and Schwartz, E. S. (2001). Valuing American options by simulation: a simple least-squares approach. *Rev. Fin. Stud.* 14, 113–147. doi: 10.1093/rfs/14.1.113
- López de Prado, M. (2018). *Advances in Financial Machine Learning*. John Wiley & Sons. doi: 10.2139/ssrn.3270269
- Platen, E., and Heath, D. (2006). *A Benchmark Approach to Quantitative Finance*. Berlin; Heidelberg: Springer Verlag. doi: 10.1007/978-3-540-47856-0
- Ruf, J., and Wang, W. (2019). Neural networks for option pricing and hedging: a literature review. *arXiv preprint arXiv:1911.05620*. doi: 10.2139/ssrn.3486363
- Ryan, R. J. (2013). *The Evolution of Asset/Liability Management*. CFA Institute Research Foundation Literature Review. Available online at: <https://ssrn.com/abstract=2571462>
- Sharma, S., Chami, R., and Khan, M. (2003). Emerging issues in banking regulation. *IMF Working Papers, WP/03/101*. doi: 10.5089/9781451852530.001
- Silver, D., Hubert, T., Schrittwieser, J., Antonoglou, I., Lai, M., Guez, A., et al. (2018). A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play. *Science* 362, 1140–1144. doi: 10.1126/science.aar6404
- Spillmann, M., Döhnert, K., and Rissi, R. (2019). *Asset Liability Management (ALM) in Banken*. Springer Gabler. doi: 10.1007/978-3-658-25202-1
- Sutton, R. S., and Barto, A. G. (1998). *Reinforcement Learning: An Introduction*. MIT Press. doi: 10.1109/TNN.1998.712192