

# Deep Reinforcement Learning for Optimal Investment and Saving Strategy Selection in Heterogeneous Profiles: Intelligent Agents working towards retirement

Ozhamaratli, Fatih  
ucabfoz@ucl.ac.uk

Barucca, Paolo  
p.barucca@ucl.ac.uk

June 12, 2022

## Abstract

The transition from defined benefit to defined contribution pension plans shifts the responsibility for saving toward retirement from governments and institutions to the individuals. Determining optimal saving and investment strategy for individuals is paramount for stable financial stance and for avoiding poverty during work-life and retirement, and it is a particularly challenging task in a world where form of employment and income trajectory experienced by different occupation groups are highly diversified. We introduce a model in which agents learn optimal portfolio allocation and saving strategies that are suitable for their heterogeneous profiles. We use deep reinforcement learning to train agents. The environment is calibrated with occupation and age dependent income evolution dynamics. The research focuses on heterogeneous income trajectories dependent on agent profiles and incorporates the behavioural parameterisation of agents. The model provides a flexible methodology to estimate lifetime consumption and investment choices for heterogeneous profiles under varying scenarios.

# Contents

<b>1</b>	<b>Introduction</b>	<b>2</b>
<b>2</b>	<b>Model</b>	<b>3</b>
2.1	Agent and Environment Cycle . . . . .	4
2.2	Behavioural Parameters of Agents . . . . .	6
2.3	Deep Policy Network for Optimal Saving, Investment and Liquidity . . . . .	7
<b>3</b>	<b>The Environment</b>	<b>8</b>
3.1	The Graph and Synthetic Population . . . . .	8
3.2	Simulation Processes . . . . .	9
3.3	Scaling . . . . .	9
<b>4</b>	<b>Results</b>	<b>10</b>
4.1	Labour, Income, Consumption and Wealth . . . . .	10
4.2	Saving Profiles . . . . .	10
4.3	Portfolio Allocation . . . . .	12
4.4	Effects of Behavioural Parameters . . . . .	15
<b>5</b>	<b>Conclusion</b>	<b>17</b>
<b>6</b>	<b>Future Work</b>	<b>17</b>
<b>7</b>	<b>Appendix</b>	<b>20</b>
7.1	Model Card . . . . .	20
7.2	Nomenclature . . . . .	21
7.3	Neural Architecture . . . . .	21
7.4	Graph Plot . . . . .	22
7.5	Cross-sectional analysis . . . . .	23
7.6	Tables . . . . .	25
7.7	Raw Plots and Behavioural Parameter Effect Plot . . . . .	28

## 1 Introduction

Retirement financing has been experiencing a clear transition trend from defined benefit (DB) schemes to defined contribution (DC) schemes, as reported by OECD [1], [2], [3]. DB schemes require scheme sponsors as ultimate guarantors of the schemes, and many DB schemes' deficit require their sponsors to bail out. Employers are preferring DC schemes, because the risk and responsibility of managing funds, longevity risk, and market risks are transferred to contributors in DC schemes. Furthermore, the contribution rates in DC schemes in the UK are on average significantly less, 5.1%, in comparison to DB average contributions of 28.5% [4]. The effects of shocks during the accumulation phase are critical, some people were raiding retirement accounts amid COVID-19. Pensioners in under-pensioned groups [5] were draining pots and faced significant wage shocks, and this will have an effect on lower future cumulative pot and earnings, exceptional government policies were critical to alleviate the effects of COVID-19 on pension savings and wages but a significant shock with effects to the labour market could not be avoided. It has become apparent how different professions can be affected differently by economic shocks, bringing the attention to the role of profile heterogeneity also in the context of pension management. For instance, the rise of the gig economy and irregular workforce participation modes enable more flexible work-life conditions, but introduce larger variations to income trajectories due to the lack of guaranteed income streams.

Previous research has addressed the income distribution and its relationship with age [6], that can be used to investigate the effects of demographic shifts and aging population on income. The increasing heterogeneity of career paths and income trajectories requires addressing the questions of how much to save in a more granular way, as well as how to allocate the savings between spendable liquid investments and illiquid retirement investments. In the literature the life-cycle models of income, consumption, and portfolio allocation have been analysed with various perspectives. Samuelson approached lifetime portfolio selection [7] as dynamic stochastic programming in discrete time and solved the many-period generalisation corresponding to lifetime planning of consumption and investment decisions. Merton formulated the continuous-time version [8] of the same approach for portfolio selection under uncertainty. Later he extended these results [9] to more general utility functions, price behaviour assumptions, and for income generated also from non- capital gains sources. A comprehensive study [10] proposes a life-cycle model of consumption and portfolio choice, as an intertemporal portfolio optimisation problem where labor income is assumed to be a risk-free asset, and calibrates it with real-world data by solving numerically, where asset choices are risky and risk-free. Coco et al. [10] present a model where risky income is invested in either risky asset or riskless asset, both are liquid and can be used for consumption, they model the income process explicitly and analytically; they solve the optimal portfolio allocation problem at a given age by numerical solution of their model with backward induction. A following study by Campanale et. al. [11] present a model, which includes explicit formulation of income process; but differs from previous research by introducing liquidity friction to risky asset, by charging an excess cost if consumption is financed through the risky asset. The model must be solved numerically via special algorithm, and the solution is described by authors as slow and difficulty, due to three continuous state variables, two continuous controls, and fixed transaction cost breaking the concavity of the objective function. The Campanale model assumes that a person has the freedom to switch between liquid and illiquid asset types, which is not the case with locked pension savings. Campanale et al. uses dynamic programming to optimise the Epstein-Zinn[12] preference utility of a household, given specific labour income process consisting of deterministic  $G(t)$  of third-order polynomial, and idiosyncratic shock. In the Campanale et al. model, the most important calibration challenge is the transaction cost, which also includes psychological and nonmonetary costs.

Further studies focus on liquid and non-liquid retirement savings accounts where liquidity is constrained by introducing cost to liquidate retirement savings[13] and [11]. Previous research

fails to address the heterogeneity of contributor profiles and falls short of addressing the idiosyncratic challenges of avoiding consumption crisis during unemployment periods and saving an adequate pension pot for retirement.

Advances in agent-based modelling of complex financial systems, increased computational power, and advances in techniques for optimising agent behaviour in complex environments are enablers and drivers of the research. Our research uses a novel methodology to address complex problems of financial systems, one such effort in the recent literature is the model called AI Economist[14], which uses AI-assisted deep reinforcement learning and implements an agent-based model to address the needs of socioeconomic challenges introduced by designing and testing economic policies, where modules called social planners are trained to discover tax policies in dynamic economies that can effectively trade off economic equality and productivity. A two-level deep reinforcement learning approach is applied to learn dynamic tax policies, based on economic simulations in which both agents and a government learn and adapt.

We introduce a simple model of contributor agents, who decide how much to save and how to allocate the savings, this decision is effected by simplified behavioural dynamics and information flow in the peer network. Agents decide and optimise their allocation strategy using a deep neural network trained with reinforcement learning. We introduce a simple simulation environment for the agents; which encapsulates employment and income dynamics. Our research represents a bridge between agent-based modelling of the pension system and deep reinforcement learning for finance.

We provide results from trained agents reflecting behaviour in par with the literature on aggregate level and agent-specific optimal behaviour with higher granularity for heterogeneous profiles. The model is dynamic, scalable, and can be calibrated to different scenarios. The results show that the balance between near-term consumption safety and retirement savings can be achieved by profile-specific allocation strategies.

Our research differs from previous studies also in terms of extensive behavioural modeling and parameterisation of the agents. It captures effect of information transmission and emphasizes consumption sensitivity against negative shocks, and covers the utility perception.

In addition, our model makes it possible for contributors to account for occupation-specific dynamics of life-time income trajectories, which in turn makes it possible to prepare against profile-specific income shocks by allocating savings to cash buffer at the right time frames of their lives.

Our research represents a significant first step to model pension finances in an agent-based model with deep reinforcement learning which permits model configurations with increased complexity and realism, in our paper we presented a simple two asset version with simple environment dynamics.

## 2 Model

We introduce a simple model where the agents interact with the simulation environment and optimise their savings behaviour. During each cycle, agents observe the environment in which they are situated; they choose to allocate their income between consumption, liquid and illiquid assets.

Each agent has a heterogeneous profile reflecting the occupation and demographic characteristics; these characteristics are determinants of the unique income and consumption trajectories of each agent. Agents also have characteristic behavioural parameters such as shock sensitivity, consumption utility, and peer-influence factor; which effect the way agents perceive the world and assign value to their stances. In particular, agents are bootstrapped in a social graph which is used for the transmission of information, such as employment status.

Each month, agents receive their income according to their employment. Simulated employment and market dynamics, such as asset return rates, are exogenous and provided by the

modeler according to empirical observations. The employment dynamics is dependent on heterogeneous profiles (occupation and demography) and includes the new employment of unemployed agents.

The agent first decides how much to save and how much to consume, and secondly the agent allocates the saved amount among a liquid asset and illiquid asset towards pension savings, each with different return rates. In order to make this financial decision, the agent’s profile, income, behavioural parameters, and peer information observed from the own social network are given as an input to a deep policy network.

Deep reinforcement learning and parallel simulation of nearly 30000 agents in 100M timesteps is used for training the deep policy network. The policy network learns an optimal saving and investment strategy for pension savings and avoiding a consumption crisis during unemployment due to insufficient liquid savings.

## 2.1 Agent and Environment Cycle

In order for the simulation to be integrated with existing frameworks, the AEC (Agent Environment Cycle)[15] is followed to also provide a standardised GYM-Like API. The simulations are vectorised and run in parallel. For the purpose of this research, the simulations are conducted in parallel utilising 32 processors, where each processor runs a cohort of more than a thousand agents. For each time step, all of the agents observe and act simultaneously.

Agents observe the environment; these observations include information regarding the market, graph, and agent’s own state including occupation, age, income, and wealth.

The agent action  $a_{i,t}$  is shaped by policy  $\pi_i$ , during learning the reward  $r_{i,t}$  for the agent is the sum of total discounted utility and penalty for consumption crisis, which denotes the situation where the agent cannot finance its consumption  $c_{i,t}$  governed by consumption dynamic  $C$ . Agent behaviour is shaped by influences from peers, individuality factor, consumption utility, and shock response characteristics.

The agent policies are modelled with a deep neural network, which takes as input agent specific observations and hidden-state:

$$a_{i,t} \sim \pi(o_{i,t}^{network}, o_{i,t}^{agent}, o_{i,t}^{market}, h_{i,t}; \theta) \quad (1)$$

The parameter variable  $\theta$  is not agent specific, but common for all contributor agents and the hidden state is updated during action inference of policy network.

- $o_{i,t}^{network}$ : Observation of the network
- $o_{i,t}^{agent}$ : Observation of own behavioural factors, income, and resources.
- $o_{i,t}^{market}$ : Observation of the market, such as
- $h_{i,t}$ : Hidden state. The updating of hidden state can be interpreted as agents updating their Risk Profile given observations and previous state. And in the future the hidden-state can be used as Risk Profile embedding.

The action space is as follows:

- $a_{i,t}^{save}$ : Decides to save  $x\%$  (and consuming  $(100 - x)\%$ )
- $a_{i,t}^{liquid}$ : Decide to allocate  $y\%$  to liquid asset  $x\%$  (and allocating the illiquid asset  $(100 - x)\%$ ).
- Saving and liquidity percentages are discretised into bins such as  $[0, 0.25, 0.5, 0.75, 1]$  in the model.

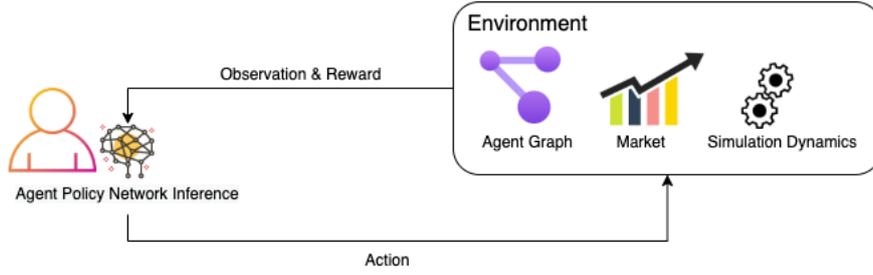


Figure 1: Agent and Environment

The full list of variables can be found in the Appendix.

The utility of an agent is expressed as the sum of the utility of consumption and of savings.

$$u_{i,t} = u_{i,t}^{consumption} + u_{i,t}^{savings} \quad (2)$$

Both the consumption and the savings utility functions are considered linear in this study. Constant relative risk aversion (CRRA) [16] is defined as :

$$crra(l, \eta) = \begin{cases} \frac{l^{1-\eta}-1}{1-\eta} & \eta \geq 0, \eta \neq 1 \\ \ln(l) & \eta = 1 \end{cases} \quad (3)$$

Individual consumption is calculated by multiplying the consumed amount with agent specific consumption utility factor,

$$u_{i,t}^{consumption} = \sum_{t=0 \dots t} [crra(C_{i,t}, q_{i,t})] \quad (4)$$

The utility of the saved amount equals the monetary value of the total savings.

$$u_{i,t}^{savings} = x_{i,t} \quad (5)$$

If the agent is unemployed or allocated insufficient funds to fulfil minimum consumption required by the modeller, then the liquid funds are used to finance consumption. If the funds are insufficient, a consumption crisis occurs, which impacts reward negatively with a consumption crisis penalty. If the agent consumes a lesser percentage than it is required to finance at least minimum consumption limit, then the invalid action penalty occurs. Consumption crisis is penalised if liquid savings cannot finance the decided consumption at time step  $t$ :

$$\mathbb{P}_{i,t} = \begin{cases} \psi & \text{if } C_{i,t} > x_{i,t}^{liquid} \\ 0 & \text{if } C_{i,t} \leq x_{i,t}^{liquid} \end{cases}$$

Invalid Actions, such as insufficient consumption rate allocation, are penalised if the decided consumption is less than minimum consumption limit despite having excess earnings that could be used for financing the minimum consumption amount:

$$\mathbb{A}\mathbb{P}_{i,t} = \begin{cases} (m - C_{i,t}) * \zeta & \text{if } m > C_{i,t} \text{ and } \mathbb{E}_{i,t} > C_{i,t} \\ 0 & \text{else} \end{cases}$$

The results of the actions are evaluated as utility difference of the new state of the agents, which constitutes the raw rewards of the actions.

$$\Delta_{i,t} = u_{i,t} - u_{i,t-1} \quad (6)$$

The consumption crisis and invalid action penalties are used to augment the utility difference to constitute the reward.

$$r_{i,t} = \Delta_{i,t} * f(\Delta_{i,t}, \kappa_{i,t}) - \mathbb{P}_{i,t} - \mathbb{A}\mathbb{P}_{i,t} \quad (7)$$

The negative utility difference is augmented with an agent’s shock perception modifier, in order to amplify the negative shocks according to the agent’s behavioural parameter.

$$f(\Delta, \kappa) = \begin{cases} 1 & \text{if } \Delta \geq 0 \\ e^\kappa & \text{if } \Delta < 0 \end{cases}$$

$$r_{i,0} = u_{i,0}^{savings} \quad (8)$$

The actions are percentage choices between consumption and savings, and investment choices between pension orientated non-liquid funds and liquid funds that can be used at any time to finance consumption, these funds have a vital function especially during the times of unemployment.

The AI Economist’s approach[14] stated as a multi-agent Markov Game can be defined by the tuple  $(S, A, r, \mathcal{T}, \gamma, o, \mathcal{J})$ , for the State Space, Action Space, Reward, State Transition Distribution, Utility Discount Rate, observation, and Agent Indices. The agents observe  $o_{i,t}$ , which is by the agent  $i$  observable fraction of the state  $s_{i,t}$ . The agent learns a policy,  $\pi(a_{i,t}|o_{i,t}, h_{i,t}; \theta_i)$ , and executes an action  $a_{i,t}$  according to the policy. The state of the environment is updated according to  $\mathcal{T}(s_{t+1}|s_t, a_t)$ . Agents maximise their  $\gamma$  discounted expected return, depending on the hidden state  $h_{i,t}$ , and the policy parameter  $\theta_i$ . Following the same reinforcement learning approach as AI Economist, the agent maximises its expected reward, depending on the actions  $a_{-i}$  and behavioural policies  $\pi_{-i}$  of the other agents, and environment transitions:

$$\max_{\theta_i} \mathbf{E}_{a_i \sim \pi_i, a_{-i} \sim \pi_{-i}, s' \sim \mathcal{T}} \left[ \sum_t \gamma^t r_{i,t} \right] \quad (9)$$

During action execution, the rewards and utilities are calculated and the tuples of observation, hidden state, and action vectors, as well as rewards are recorded to replay buffer, for future training. The actions are sometimes random for greedy exploration of the domain, the epsilon-greedy parameter is provided by the modeller. At regular intervals, the replay buffer is used for training the deep policy network by utilising the rewards with policy gradients. The model is trained with A2C, PPO and Policy Gradient methods. A2C training results in slightly better results, which are presented in the article.

## 2.2 Behavioural Parameters of Agents

For modeling behaviour, we base our parameterisation on the approach in [17], where the authors investigated the applicability of the Theoretical Domains Framework outside clinical uses for cross-disciplinary implementation and other research on behaviour change, and provided a simplified version containing 14 domains and 84 component constructs. Theoretical Domains Framework is wide and reports on pension behaviour tend to focus on few factors; for the scope of our research we chose three factors:

- Consumption Utility: How do they value current consumption? An agent specific consumption utility multiplier factor
- Shock Response Characteristics: How do they respond to the shock? A factor reflecting how sharp do agents react to the shock and how drastic are they decreasing their consumption.

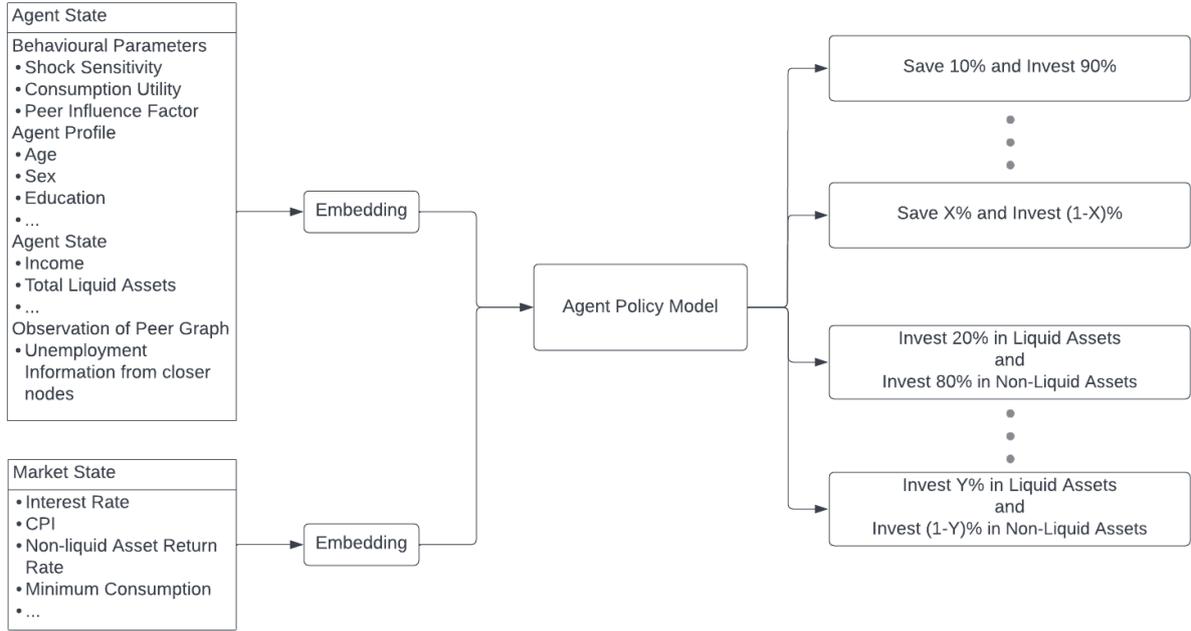


Figure 2: Policy Model

- **Individuality Factor**: How are they being affected by each others beliefs and decisions.

Some agents are optimistic and underestimate the severity of the shocks, and some agents are pessimist and overestimate the effect of the shocks. The shock sensitivity factor  $\kappa_{i,t}$  is a multiplier of the perceived shock effect, which is normalised for agents of the same occupation. It can be assigned from a normal distribution, can be controlled for experimentation, or fed from empirical report. The agents are affected by the peers and the shocks experienced by the peers (if  $z_{i,t-1} = 0$ ). The observations are informative for the closer peers on the graph, and becomes less informative for other agents with weaker connection on the graph. The shocks that effect the peers are also weighted with the shock-sensitiveness parameter. The behaviour parameters that are introduced in this section are fixed during the entire simulation, but later research can adapt the model to allow the parameters to change during lifetime.

### 2.3 Deep Policy Network for Optimal Saving, Investment and Liquidity

Agent observations are expressed as a single vector that comprises the concatenation of agent, market and graph vectors. The observation vector is passed through the deep neural network towards the LSTM, which updates the agent's hidden state and outputs a vector for next layer, which is softmaxed to output a vector representing the action probabilities. A single policy network is trained for all actions: The action can be as follows "(C25', 'L75')", where "C25" means consume 25% and save 75%; "L75" means allocate 75% of your savings to liquid asset, and 25% of your saving to nonliquid asset.

The hidden states from the model can be thought of as Risk Profile Embedding, which is updated by observations and processing the Agent Profile with the observed environment and shocks via a Deep Neural Architecture that can be found on Fig. 2. Reinforcement Learning is used for adjusting the allocation profile according to the Risk Profile Embedding also expressed as a hidden state. At each time step, the agent decides to allocate the income among consumption, savings, and investment classes. This is accomplished by a deep neural network constituted of several layers of a Feedforward Neural Network with nonlinearity as well as a recurrent unit of LSTM, which is responsible for acting as the memory of agents, and the

hidden states can be used as agent representations during further studies. The details of the neural architecture can be found in the Appendix.

There is a single action space unifying the choices of consumption and liquidity preference, which means that there are not two different networks for different decisions but one unified network which represents the collection of actions such as "(C25, L75)". An alternative choice was continuous action space of two different actions expressed as a vector of length two, which is expected to have more capability to capture relationship between similar amounts, especially if the vast amounts of simulation data are not available. Furthermore, the dimension of the action space is greatly reduced if the number of discrete consumption and investment type rate choices are high. Setting the reward function for the agents is the trickiest part of the training process, different reward function structures tend to give spurious and unintended conclusions especially due to the existence of both positive-effecting factors and negative-effecting factors, which makes the hyper-parameter tuning for the penalties paramount; and requires several trials to find the hyper-parameter enabling successful optimisation of the agents. Failing to tune the penalties results in unintended shortcuts that obstruct the main goal of optimising agent behaviour in an understandable and meaningful way.

### 3 The Environment

At each time step, first the environment operations are executed. Agent environment operations are executed as follows: first, the market dynamics are executed which ensures that assets are gaining value according to the calculated interest rates determined by the modeler. Secondly, essential population dynamics is executed such as ageing of agents, executing death operations according to the age-specific death probability. The Retirement Process checks if any new agents are required to retire due to age. If an agent retires, their retirement pension is calculated as a rate of their previous consumption at employment, according to the recommended guidelines of the OECD. If an agent is retired, then the agent collects pension from illiquid pension fund that they contributed during employment life. The agents that are not retired are processed to determine stochastically if they will lose their employment and, if so, for how long they will stay unemployed according to the unemployment duration distribution dependent on the occupation and age. Unemployed agents are assigned new income at their new jobs according to the income distribution depending on occupation and age. These distributions are fed as quantiled distribution tables to simulation. The employed agents receive their salary at each month according to their predetermined income.

The agents decide how to allocate their income between consuming and saving, and decide to allocate the saved amount in liquid and riskless assets, or illiquid and low-risk assets. The decision is shaped by learnt policy, and observations, which includes the market dynamics, information regarding actions and behaviour of peers, and considering the agent's own profile.

#### 3.1 The Graph and Synthetic Population

A synthetic representative population is used for the initialising agent population, information such as age, income, profession, education level, and other relevant background information are included.

We assume the employee network consists of three communities divided by income level as low, medium, high, and the three communities have significant intercommunity interaction but limited intracommunity interaction. The graph choice is based on the idea that geographical and social networks are also characterised by the socioeconomic clusters, and the choice of three communities with income levels is the simplification of the socioeconomic network. The synthetic database is generated according to the basic insights from the surveys. Later investigation could incorporate survey data to bootstrap the population and investigate geographical graph,

potential social network data, and known network structures to model connection between agents.

Observation of graphs can be done in several ways; a simplistic way to do is modelling information transmission between each agent and its vicinity; for the first iteration of the model this means that the neighbours and their neighbours, and transmission of employment information. A more advanced graph observation might be modelled as transmission of not just employment information, but also incorporating additional information such as occupation and the income or consumption data; moreover, the near-neighbour graph can be represented with state-of-the-art graph embedding methodology.

$$o_{a,t}^{network} = \sum_{a,b} \left[ \sum_{b,c} A_{b,c} \delta(\mathbb{E}_{c,t}, 0) \right] + \sum_{a,b} A_{a,b} \delta(\mathbb{E}_{b,t}, 0) \quad (10)$$

For the purpose of experimentation and investigation of the model, a synthetic but representative population can provide both fidelity and flexibility in a controlled environment. As a design choice for the synthetic population network, we include three clusters, which can be thought of as three neighbourhoods; these neighbourhoods possess nodes with three different income groups: high, medium, and low income. Each node is connected to its own neighbourhood node, and the neighbourhoods are connected to each other with specified weights. Agents are bootstrapped with one of the general occupation groups, occupation-specific incomes, employment status, and ages derived from USA Census Data[18]. Census data are used to generate the synthetic agent population.

### 3.2 Simulation Processes

The simulation is initialised by bootstrapping the agent population and processes. During each time step, the simulation dynamics such as getting income, getting employed if vacant, are applied first, then the agent decides to allocate income for the consumption or saving, and decides to save by investing in liquid assets, which can be liquidised easily during unemployment, or illiquid assets which are towards a future retirement, but usually have better return. Agents are bound by constraints such as the need to consume a minimum amount determined in light of government statistics[19] that determines a minimum consumption per individual.

The occupation-specific income for new employment is determined according to the summary tables from USA Census Data. The tables reflect the quantile breakdown, and the agents are probabilistically assigned to one of the income quantiles.

The unemployment events and employment processes are explicitly modelled and calibrated with the US Census Data[18]. The probability of unemployment and the duration of unemployment are determined according to the summary tables of the US Census.

Retirement age and retirement income can be accounted for in the system. For the sake of simplicity, initial simulations neglect the retirement period, by only focusing contribution period; but the system is later extended to cover the retirement period. Retirement income is defined as a fraction of the last income, fractional retirement income is recommended by international institutions, and this methodology is often also used in the literature[10].

The agent death probabilities are modelled using the Actuarial Life Table [20] in order to make the model comparable with existing models in the literature.

### 3.3 Scaling

The agent observations are continuously scaled and standardised, with an online methodology. This is due to the fact that the training dataset is generated continuously during simulation and the distribution of the observed dataset is not known in advance at the start of the simulation, but it can be learnt to an extent after several epochs, and these learnt

scales can be utilised in following training and inference as well. The relevant agent variables ("OCC\_CODE", "income", "consumption\_utility\_factor", "shock\_sensitivity\_factor", "individuality\_factor", "non\_liquid\_asset", "liquid\_asset") are transformed to a vector by concatenating categorical one hot vectors with the values of the continuous variables; here the standardisation of these categorical variables is challenging due to the variability of the aspects such as accumulated liquid assets. The huge value differences have the potential to introduce instability during the training of machine learning models. The market state captures important variables such as interest rates given to different asset classes as a dictionary, the market dictionary is transformed to a vector as well. The perception of graph by the agent is represented as a vector, for the sake of simplicity in this first version of the model the graph is represented by single-valued vector, which is calculated by second-order graph traversal of neighbouring agents in the network, the income information from neighbouring agents is incorporated according to the edge weights, as well as the individuality factor of the effected agent.

## 4 Results

We look at longitudinal trajectory plots and strategy breakdown per total asset size, which provide granular information regarding the differences between occupations. These plots can capture various scenarios such as differences between early career and mid career saving rate strategies among various occupations; which provide more tailored strategies for short-term consumption security and healthy long-term pension finances. 22 parallel initially identical cohorts are simulated for 1000 weeks of agent-time in order to generate resulting tables and plots, which results in 40M agent time-step samples.

### 4.1 Labour, Income, Consumption and Wealth

Fig. 3 reflects a similar shape of average simulated income, consumption, and wealth accumulation and decrease over the life cycle compared to Coco et al. [10]. The simulated income trajectory is a reflection of the observed data, which is used for calibration of the environment, and the shape of decrease by retirement age is due to the retirement income being defined as a fraction of last income, which then gradually decreases. The consumption trajectory during the work-life reflects saving choices of the population. The agent saves during work-life for financing potential unemployment periods and for retirement finances. The pension income and consumption at retirement age of 65 converges to the determined retirement income percentage of 80% of latest salary. The data becomes noisy for older ages of 80, which might be due to significantly smaller sample size.

During retirement the pension income is supposed to come from pension savings that have been non-liquid during work-life, but if the pension savings are depleted any liquid savings can be used to finance the retirement income on Fig. 15. An interesting outcome of mandating pension income at retirement to be 80% of employment income is comparatively lower consumption during employment, which might not be desirable; but our optimiser were forced to high saving rates due to 80% mandate, which is stipulated by literature, detailed information can be found in previous sections focusing on literature. There might be various solutions to this problem that are out of our scope, such as easing pension level mandate, or government contributions, or higher returns of investment; these potential solutions can be investigated in future research.

### 4.2 Saving Profiles

The evolution of occupational income in a time frame of nearly 20 years on Fig. 5 reflect different characteristics for each occupational group, occupations such as "Sales and Related" and "Transportation and Material Moving" reflect significantly lower mean income with lower variance characteristics; on the contrary, occupations such as "Legal" and "Management", which

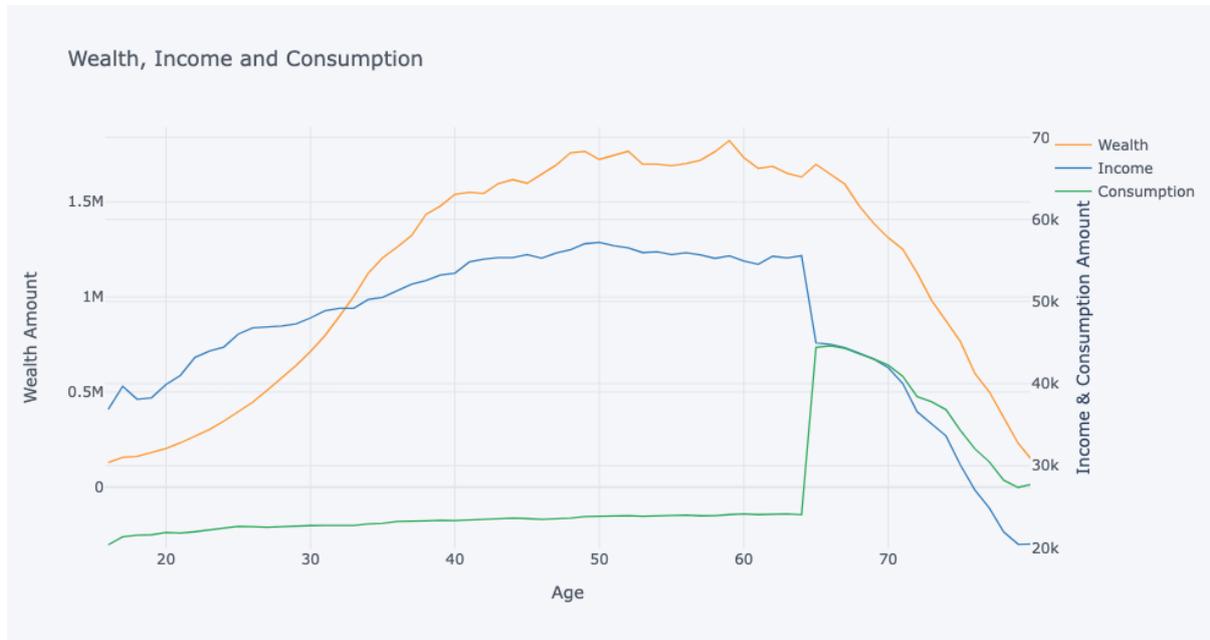
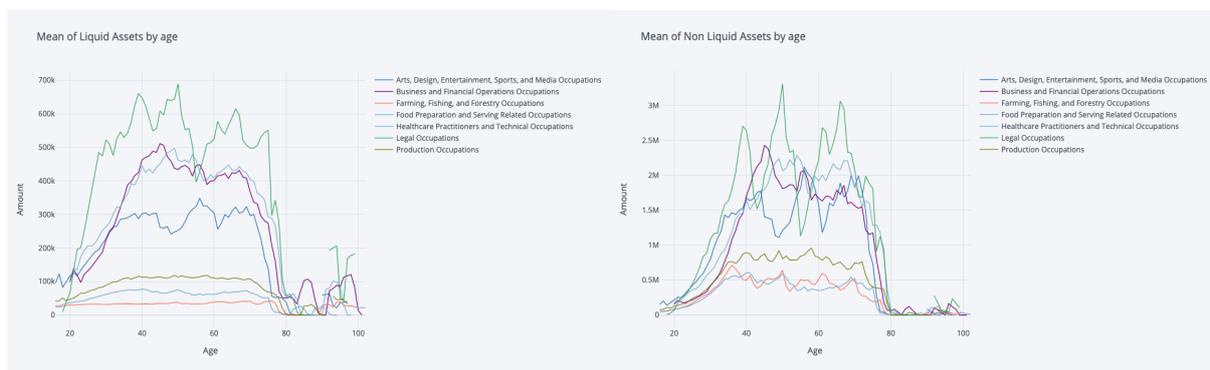


Figure 3: Wealth, consumption and labour income vs age plot



(a) Liquid Asset

(b) Non-Liquid Asset

Figure 4: Liquid and Nonliquid asset amounts by occupation at age, where only a selection of occupations are depicted on plot for clear visibility. The different characteristics of occupation groups are reflected by plots.

reflect highest mean income and high variance of income for each occupation group. This plot reflects even at the simplest level that the income characteristics of each occupation can differ greatly. The unemployment characteristics reflect great diversity, where occupations such as "Farming, Fishing and Forestry" possess greater and fluctuating risk profiles, which might be partially due to the characteristics of seasonality in these specific occupations. No obvious dependence of saving rate or non-liquid investment rate on age or income level can be found in the analysis, showing the complexity of the decision making happening in the system.

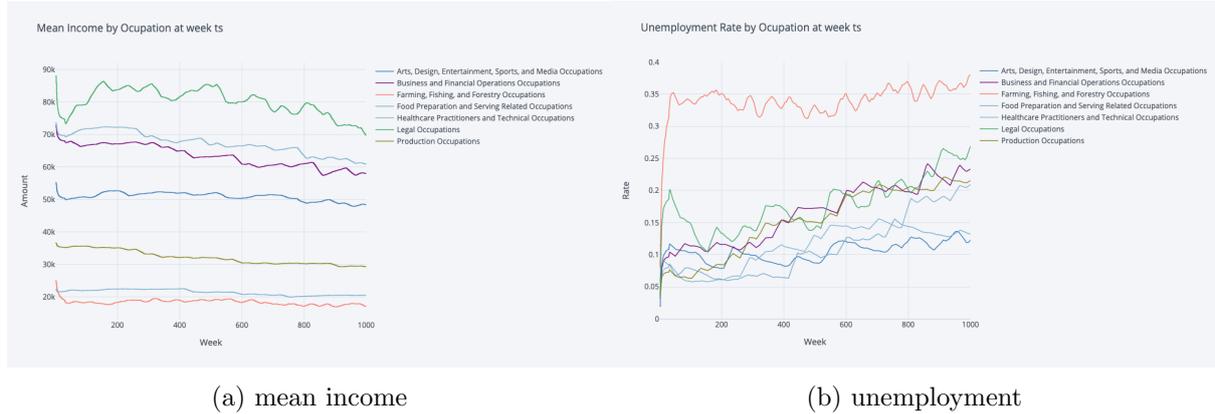


Figure 5: Mean of income and unemployment by occupation at week ts; the values are smoothed by 30-step moving average and only a selection of occupations are depicted on plot for clear visibility. The different characteristics of occupation groups are reflected by plots.

The savings profiles in Fig. 6 reflect heterogeneous characteristics, where at the same total wealth the saving rate differs greatly, which can be due to different income levels and unemployment risks of occupations. The saving rate plot shows increasing noise at higher wealth levels near 10M, and a much clearer trajectory at lower wealth. An interesting insight is that at the lowest wealth levels, all occupations display similar saving rates. Minimum consumption requirement has a direct consequence of lower saving rates by occupations with low income occupations such as "Farming, Fishing and Forestry", "Building and Grounds Cleaning and Maintenance", "Personal Care and Services", "Food Preparation and Serving Related" occupations have very low saving rate due to their difficulties to finance minimum consumption. Some general patterns can be identified, such as lower income occupations tend to have lower saving rates, but it does not imply that income itself can explain saving decisions; as we can observe varying saving rates among "Healthcare Practitioners", "Legal Professionals", and "Business and Financial Operations".

### 4.3 Portfolio Allocation

Compared to the literature on the share of non-liquid assets and age distribution, Fig. 7 reflects similar shape and rates, where Campanale et al. [11] plots a very similar life-cycle stock share model profiles when the illiquid and liquid assets are differentiated with transaction costs for switching between them, our model reflect similar characteristics when the transaction costs are high. The share of non-liquid asset according to total current wealth also reflects a similar shape of initial increase followed by plateau.

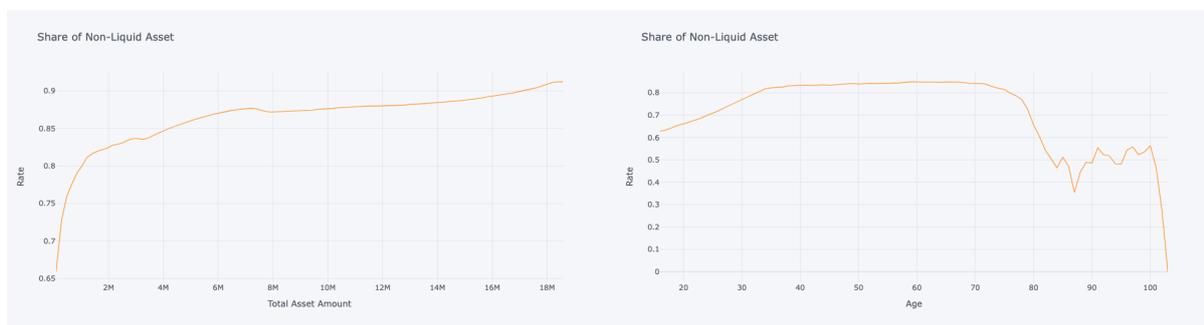
The relationship between share of non-liquid asset and age inferred in our model is in par with the literature, furthermore its shape is similar to empirical data presented by Campanale et al. [11]. Furthermore, Fig. 8 reflects a more heterogeneous relationship with higher granularity in comparison to the literature. Our model provides similar income, consumption, and wealth curves with respect to the literature, and portfolio allocation strategies with higher granularity suitable for heterogeneity and various income processes. A downside of the model is the



(a) Week

(b) Total Asset Amount

Figure 6: Saving rate by occupation at week ts and saving rate by occupation at amount capped at 10M, the values are smoothed by 30-step moving average and only a selection of occupations are depicted on plot for clear visibility.



(a) Total Asset Amount

(b) Age

Figure 7: Nonliquid asset share vs total asset amount and age

high level of noise for the portfolio allocation 3d surface 8, which might arise from the higher complexity of the model.

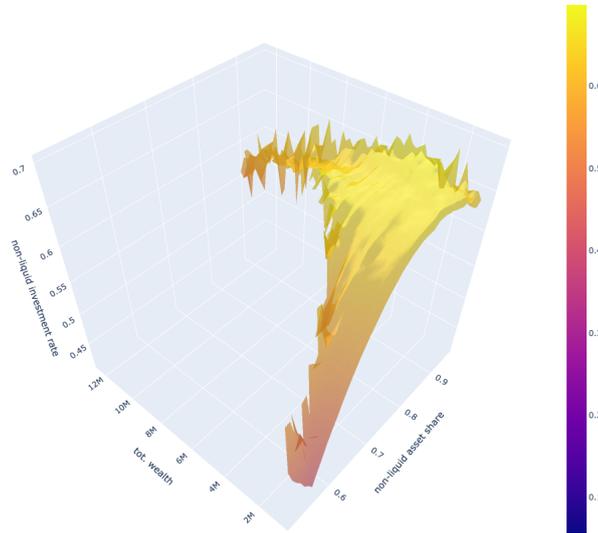


Figure 8: 3d Surface Plot of share of non-liquid assets in x-axis, with respect to total asset wealth in y-axis, and corresponding decision of non-liquid asset investment rate in z-axis, the values are smoothed with 9 step moving average for clearer visibility



Figure 9: Non liquid investment rate by occupation at amount capped at 10M, the values are smoothed by 30-step moving average and only a selection of occupations are depicted on plot for clear visibility.

Contrasting general saving rate and the non-liquid investment rate characteristics of occupations with respect to total assets result in interesting findings. The non-liquid investment rate by total asset among occupations diverges less than the saving rate by total asset, but still the characteristically differentiating investment strategies are evident in Fig. 9 Saving rate by total asset generally increases for all occupations with more asset, with exponential like increase, then it plateaus and slightly varies with noise. Saving rates by highest total asset amounts fluctuate greatly, which might be due to different dynamics governing their decisions such as capital income or behavioural parameters weighing more themselves rather than income being the determinant of the decisions.

Non-liquid investment rate per total asset is much higher for low income occupations, and

only exceptions is for very high amounts of total asset, where higher income occupations can surpass the low income occupations by non-liquid investment rate. This difference is due to a necessity for low-income occupations, which results in lower buffer savings in liquid assets that can be used to finance short-term consumption during periods of unemployment. This finding can be a guidance for policy makers to come up with instruments to mitigate the risk for income occupations to face higher risk of consumption crisis, such as unemployment benefits, or early career pension contributions from the government to support low-income early career workers. Non-liquid investment rate seems to have an increasing trend for low-income occupations with respect to total asset. On the other hand, for some of the mid and high occupations, the trend is first a decrease followed by an increase.

Comparing our results with those of Campanale et al. in Table 1 with high transaction cost makes it clear that our model refers to a higher non-liquid investment share with respect to total asset portfolio, but there are exceptions where Campanale et al. determine a high non-liquid asset share. Furthermore, profile specific differences between occupations is evident from varying non liquid asset shares with respect to income quartile and age group.

#### **4.4 Effects of Behavioural Parameters**

The behavioural parameters of the agents are the consumption utility factor, the shock sensitivity factor, and the individuality factor. These factors capture the behaviour of agents, and they impact how agents perceive, understand, and act in their environment. The consumption utility factor is necessary for quantifying how agents value immediate consumption, which can be interpreted as level of consumerism, or temporal preference and eagerness. The shock sensitivity factor is a parameter helpful for capturing the agent's perception of the consumption change, which can amplify the effects of the changes and force agents to avoid abrupt changes, and an alternative interpretation can be as risk aversion modifier that augments the utility. The individuality factor models the level of influence an agent's social network exerts on the agent. This is achieved by factoring in the information transmitted from the neighbourhood. The increase in the liquid assets reflect a linear increase; on the contrary, the increase of non-liquid assets is exponential due to interest income of the assets. The distribution of outcomes reflect heterogeneous characteristics according to behavioural parameters and the relationship between parameters and outcomes are non-linear. The details can be found in the Appendix.

Occupation	Quart 1.	Quart 2.	Quart 3.	Quart 4.
<hr/>				
Arts-Design-Entertainment-Sports-and Media				
20-30	0.681	0.701	0.705	0.754
30-40	0.702	0.769	0.803	0.849
40-50	0.696	0.745	0.805	0.863
50-60	0.721	0.771	0.828	0.871
60-70	0.539	0.752	0.833	0.863
<hr/>				
Business and Financial Operations				
20-30	0.626	0.648	0.615	0.632
30-40	0.653	0.671	0.712	0.774
40-50	0.659	0.700	0.793	0.848
50-60	0.658	0.692	0.775	0.848
60-70	0.547	0.695	0.759	0.841
<hr/>				
Farming-Fishing-and Forestry				
20-30	0.774	0.810	0.856	0.901
30-40	0.772	0.893	0.946	0.965
40-50	0.758	0.854	0.938	0.965
50-60	0.751	0.843	0.922	0.963
60-70	0.547	0.799	0.907	0.959
<hr/>				
Healthcare Practitioners and Technical				
20-30	0.644	0.655	0.653	0.680
30-40	0.674	0.674	0.714	0.793
40-50	0.674	0.716	0.774	0.832
50-60	0.692	0.738	0.793	0.851
60-70	0.663	0.746	0.797	0.857
<hr/>				
Legal				
20-30	0.644	0.614	0.607	0.632
30-40	0.659	0.649	0.717	0.802
40-50	0.664	0.709	0.760	0.827
50-60	0.670	0.675	0.730	0.828
60-70	0.630	0.699	0.791	0.859
<hr/>				
Production				
20-30	0.667	0.704	0.751	0.826
30-40	0.682	0.770	0.843	0.900
40-50	0.694	0.779	0.849	0.909
50-60	0.684	0.771	0.858	0.912
60-70	0.589	0.757	0.839	0.904
<hr/>				
All Occupations				
20-30	0.680	0.723	0.743	0.735
30-40	0.707	0.769	0.806	0.830
40-50	0.708	0.774	0.819	0.847
50-60	0.703	0.774	0.817	0.856
60-70	0.610	0.755	0.815	0.860
<hr/>				
Campanale et al. TC high				
20-30	0.077	0.471	0.467	0.577
30-40	0.575	0.591	0.547	0.739
40-50	0.539	0.621	0.757	0.704
50-60	0.70	0.765	0.791	0.698
60-70	0.735	0.767	0.751	0.706
70-80	0.562	0.701	0.756	0.667

Table 1: Occupation and Age vs Share of Non-Liquid Investments for Wealth Quartiles, the results from Campanale et al. with high Transaction Costs(TC) are used for comparison, in our model there are no transfers between non-liquid and liquid assets before retirement so high transaction cost results are relatively compatible with our model

## 5 Conclusion

We modelled a pension ecosystem, where heterogeneous contributors make consumption and investment decisions with Deep RL, which advances available models by providing better granularity and accounting for profile heterogeneity.

We provide a novel methodology to optimise agent behaviour for consumption and investment between pension savings and liquid cash buffer, which is flexible and can be calibrated to work in various scenarios and capture agent heterogeneity. Our model does not need an explicit formulation of the income process and can work with empirical data.

Our research represents a first step of end-to-end modelling of pension ecosystem, to provide a highly capable model to optimise contributor and other actor behaviour, in the existence of high variability and heterogeneity, as well as dynamic environment. We successfully devised optimal contributor portfolio allocation strategies between nonliquid pension savings and liquid cash buffer, as well as optimal consumption decision, which can be calibrated with behavioural parameters of agents. We accomplish this by minimising the consumption crisis periods of agents and maximising the retirement savings.

## 6 Future Work

Future work can be planned in four dimensions. One dimension is the representativeness and information richness, which can be enhanced by more in depth shock embedding for capturing how a shock affects a specific agent, considering shocks such as financial crashes, epidemics, and natural catastrophes such as earthquakes. A second dimension is a more in depth representation of each profession using vector representation of professions, which will capture the job similarity and to some extent the relationship between professions. Utilising a more comprehensive graph encompassing real geographical neighbourhoods, peer network, or employment network, as well as synthetic networks with well-known structures such as small world will provide more extensive capabilities to model the world. The third dimension of future work is introducing more actors in the multi-agent model including companies as employers and business-to-business relationships, which will provide a model less leaning on statistical calibration via real-world data, and capture second-order effects and endogenous properties. Another pivotal agent that can increase the impact and self-consistency of the simulated environment is the government acting as a tax regulator. The fourth dimension is technical improvement to the simulation system to use vectorised parallel architectures such as GPUs, which will enable us to test more extensive market conditions more efficiently and in shorter time.

## References

- [1] OECD. *Pensions at a Glance 2019: OECD and G20 Indicators*. OECD Pensions at a Glance. OECD, November 2019.
- [2] OECD. Pension Markets in Focus 2019, 2019.
- [3] OECD. Pension Markets in Focus 2020, 2020.
- [4] Office for National Statistics. Occupational pension schemes in the UK. 2019.
- [5] Lauren Wilkinson and John Adams. What impact has the COVID-19 pandemic had on underpensioned groups? *Pensions Policy Institute*, (ISBN 978-1-914468-05-6), December 2021.
- [6] Fatih Ozhamaratli, Oleg Kitov, and Paolo Barucca. A generative model for age and income distribution. *EPJ Data Science*, 11(1):4, January 2022.
- [7] Paul A. Samuelson. Lifetime Portfolio Selection By Dynamic Stochastic Programming. *The Review of Economics and Statistics*, 51(3):239–246, 1969.
- [8] Robert C. Merton. Lifetime Portfolio Selection under Uncertainty: The Continuous-Time Case. *The Review of Economics and Statistics*, 51(3):247–257, 1969.
- [9] Robert C Merton. Optimum consumption and portfolio rules in a continuous-time model. *Journal of Economic Theory*, 3(4):373–413, December 1971.
- [10] João F. Cocco, Francisco J. Gomes, and Pascal J. Maenhout. Consumption and Portfolio Choice over the Life Cycle. *Review of Financial Studies*, 18(2):491–533, 2005.
- [11] Claudio Campanale, Carolina Fugazza, and Francisco Gomes. Life-cycle portfolio choice with liquid and illiquid financial assets. *Journal of Monetary Economics*, 71:67–83, April 2015.
- [12] Larry G. Epstein and Stanley E. Zin. Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework. *Econometrica*, 57(4):937–969, 1989.
- [13] Magnus Dahlquist, Ofer Setty, and Roine Vestman. On the Asset Allocation of a Default Pension Fund. *SSRN Electronic Journal*, 2016.
- [14] Stephan Zheng, Alexander Trott, Sunil Srinivasa, Nikhil Naik, Melvin Gruesbeck, David C. Parkes, and Richard Socher. The AI Economist: Improving Equality and Productivity with AI-Driven Tax Policies. *arXiv:2004.13332 [cs, econ, q-fin, stat]*, April 2020. arXiv: 2004.13332.
- [15] J. K Terry, Benjamin Black, Nathaniel Grammel, Mario Jayakumar, Ananth Hari, Ryan Sullivan, Luis Santos, Rodrigo Perez, Caroline Horsch, Clemens Dieffendahl, Niall L Williams, Yashas Lokesh, Ryan Sullivan, and Praveen Ravi. Pettingzoo: Gym for multi-agent reinforcement learning. *arXiv preprint arXiv:2009.14471*, 2020.
- [16] John W. Pratt. Risk Aversion in the Small and in the Large. *Econometrica*, 32(1/2):122–136, 1964.
- [17] James Cane, Denise O’Connor, and Susan Michie. Validation of the theoretical domains framework for use in behaviour change and implementation research. *Implementation Science*, 7(1):37, April 2012.

- [18] 2019 annual averages - household data - tables from employment and earnings, 2019. publisher: U.S. Bureau of Labor Statistics.
- [19] ASPE. 2019 poverty guidelines, 2019.
- [20] Actuarial life table - SSA, 2017. publisher: Social Security Administration U.S.A.

## 7 Appendix

### 7.1 Model Card

Simulation Parameters	
Parallel Environment Count	32
Income Calibration Data	USA CPS 2019 Median weekly earnings
Unemployment Duration Data	USA CPS 2019 Unemployment duration table
use_min_max_scaler	1
time steps	1000
consumption_crisis_penalty	100000
invalid_action_penalty_modifier	1000
retirement_age	65
retirement_salary_multiplier	0.8
death_rate	USA SSA Actuarial Life Table
Agent States	"OCC_CODE" , "age", "income", "consumption_utility_factor", "shock_sensitivity_factor", "individuality_factor", "non_liquid_asset", "liquid_asset"]
Market Parameters	
monthly_market_interest_rate	0
CPI	0
monthly_non_liquid_asset_return_rate	0.0125
monthly_liquid_asset_return_rate	0.0025
monthly_minimum_consumption	1073 (2021 USA Poverty Guidelines)
monthly_minimum_wage	1160
ML Parameters	
batch_size	14656
n_lstm	128
n_steps	1

Table 2: Parameters

## 7.2 Nomenclature

### List of Symbols

The next list describes several symbols that will be later used within the body of the document

$\delta(x, y)$  Kronecker Delta

$\eta$  Degree of non-linearity of crra function

$\kappa_{i,t}$  Shock Sensitivity Factor

$C_{i,t}$  Instant Consumption

$E_{i,t}$  Current Earnings

$\psi$  Consumption Crisis Penalty Modifier

$\zeta$  Invalid Action Penalty Modifier

$A_{a,b}$  1 if there is an edge  $a \rightarrow b$

$c_{i,t}$  Accumulated consumption value

$m_{i,t}$  Minimum consumption

$P_{i,t}$  Instant Penalty for Consumption Crisis

$p_{i,t}$  Accumulated penalty

$q_{i,t}$  Consumption Utility Factor

## 7.3 Neural Architecture

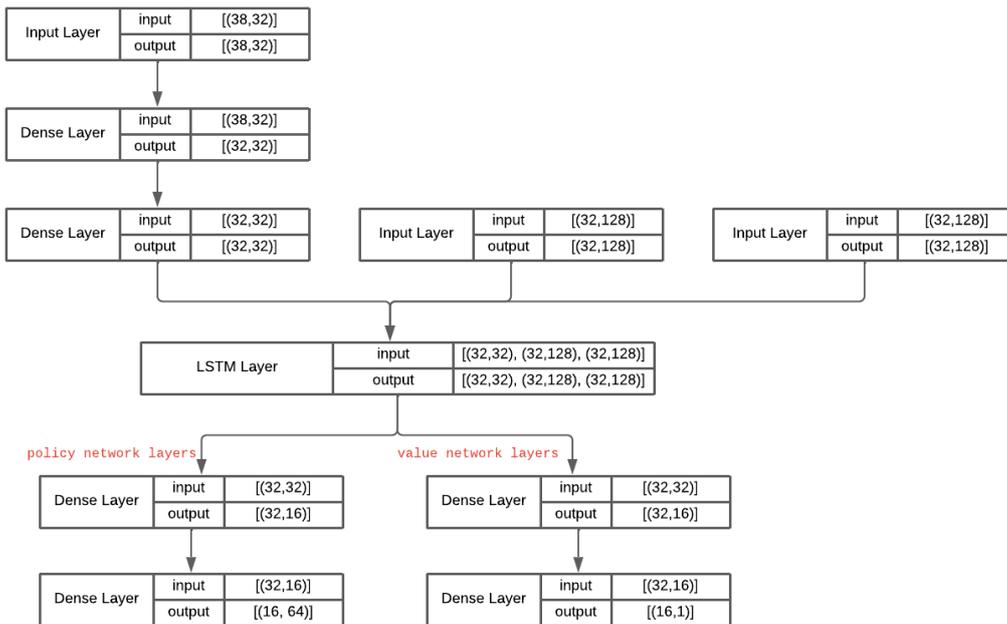


Figure 10: Neural Architecture

## 7.4 Graph Plot

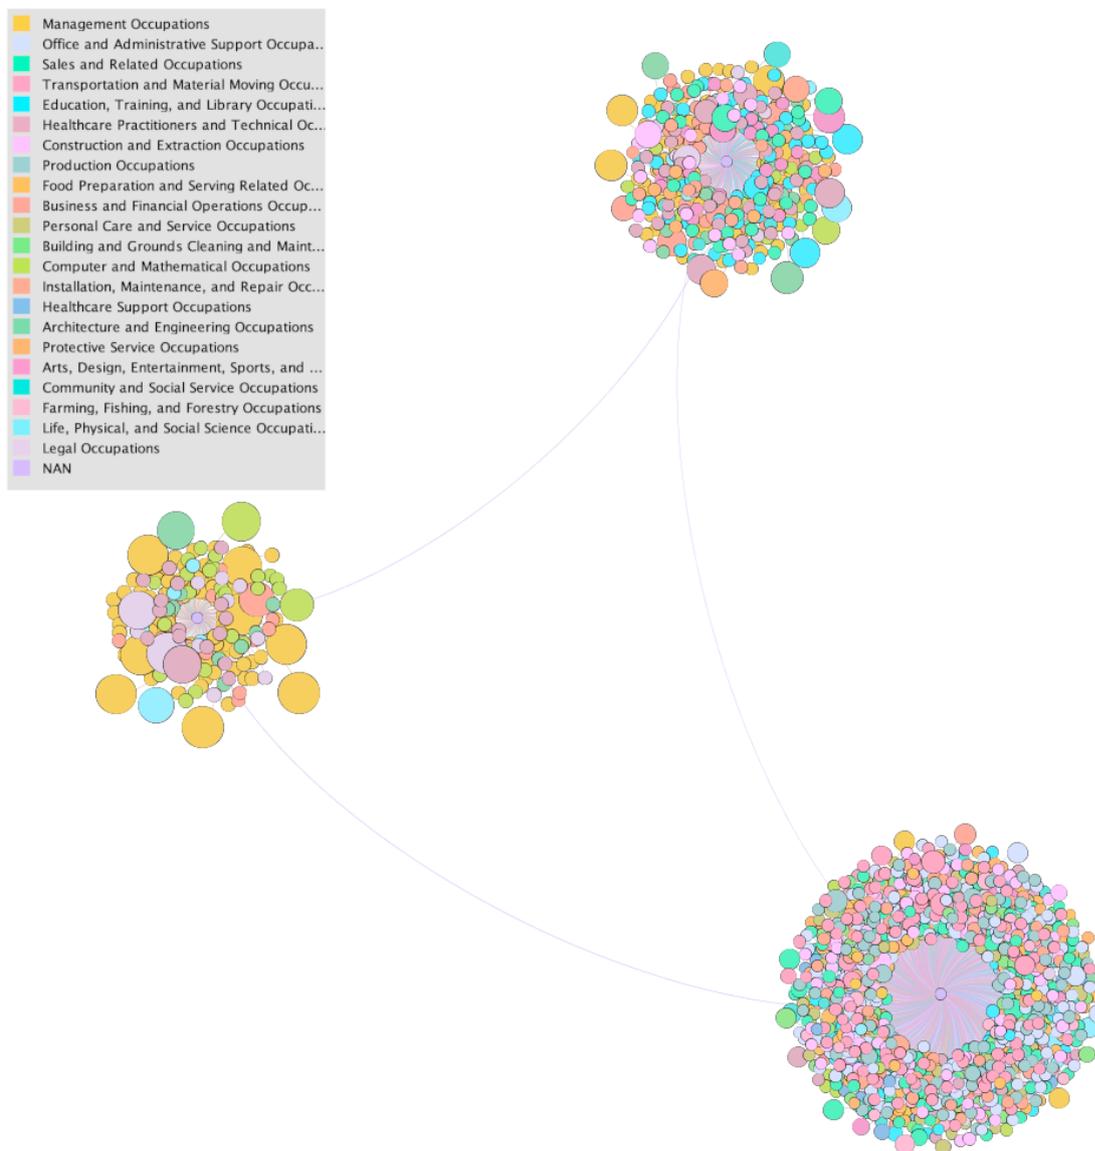
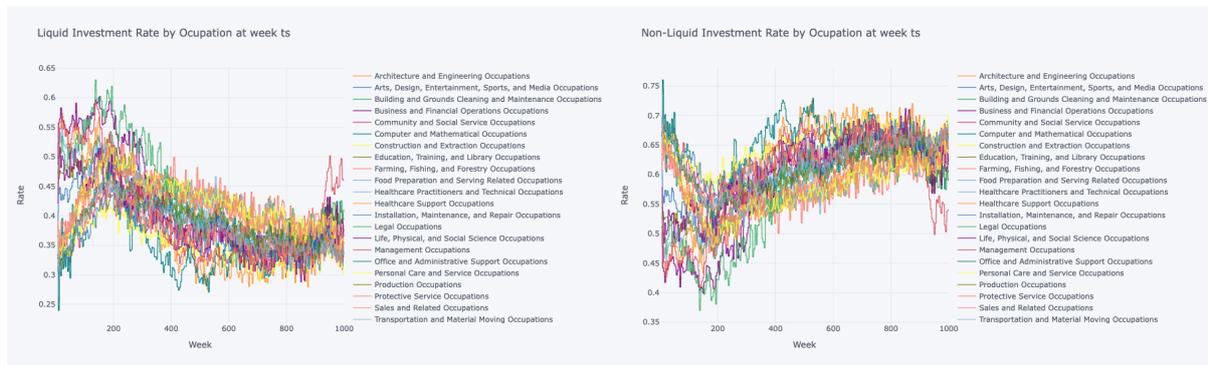


Figure 11: Graph Plot where occupations are reflected with colours and income is reflected with the size of nodes, the graph consists of three sub-groups representing three neighbourhood with three differing income levels as high mid and low income, central nodes of each group are representing the neighbourhoods.

## 7.5 Cross-sectional analysis



(a) Liquid Investment Rate

(b) Non-Liquid Investment Rate

Figure 12: Liquid and Nonliquid Investment rate by Occupation at week ts

The Fig. 13 reflects the relationship between liquid investment rate and wealth for amounts less than 5M USD. 3 distinctive behaviours are observable, one is Computer and Mathematical Occupations, which start at lowest liquid investment rate, the other group represents the majority of the occupations representing most of low- and mid-income occupations, which start at nearly 35% but then lower their liquid investment rates when the total wealth increases, the third group consist mostly of the high income occupations such as Management Occupations which increase their liquid investment rate with total asset increase until nearly 500K USD, at that point they start to decrease their liquid investment rate.



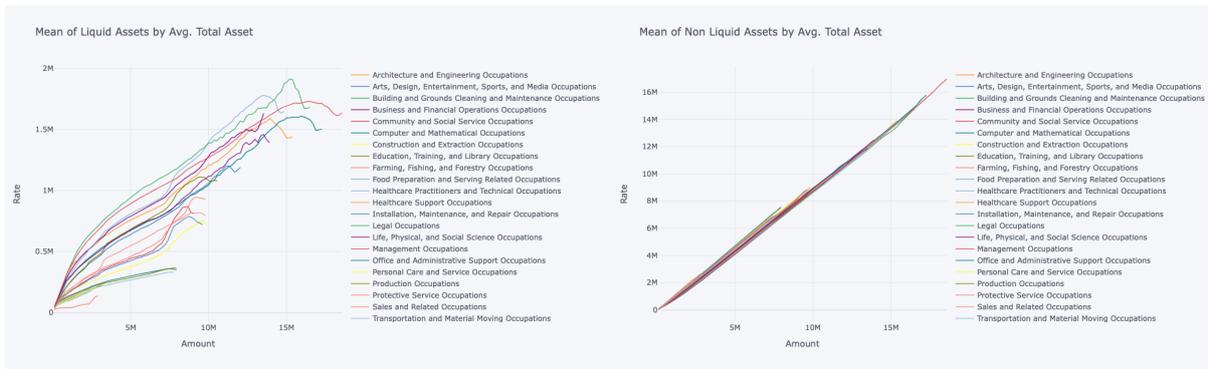
(a) Liquid Investment Rate

(b) Non-Liquid Investment Rate

Figure 13: liquid and nonliquid investment rate by occ at amount

Fig. 14 shows that the increase in liquid assets slows with increasing total wealth, which reflects the fact that the need for security buffer savings decrease and reward of illiquid asset is higher. On the contrary, the increase of no liquid assets with respect to the total wealth increase speeds up at higher amounts and converges to a stable linear trajectory.

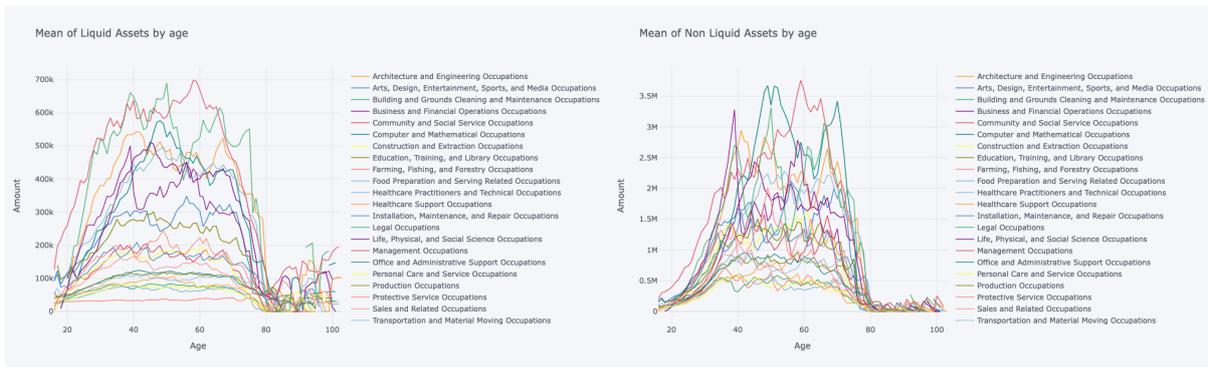
The distribution of assets with respect to age on Fig. 15 highly differentiates according to the occupation, where Management and Legal Occupations possess highest value of assets and Farming, Fisheries and Food Preparation Occupations possess lowest level of assets and also reflect lowest change of assets with age, which points out that asset differentiation with respect to age depends heavily on the occupation type, where some occupations show greatly changing income asset values and some occupations can provide minimal savings opportunity, due to income merely sufficing to finance current consumption during workforce participation.



(a) Liquid Asset

(b) Non-Liquid Asset

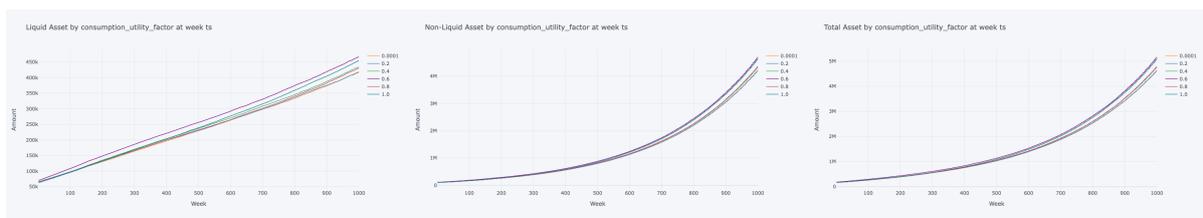
Figure 14: Liquid and Nonliquid Assets by occupation at Total Amount



(a) Liquid Asset

(b) Non-Liquid Asset

Figure 15: Liquid and Nonliquid assets by occupation at age



(a) Liquid Asset

(b) Non-Liquid Asset

(c) Total Asset

Figure 16: Total Asset, Nonliquid Asset, Liquid Asset at week ts by Consumption Utility factor

## 7.6 Tables

occupation title	saving rate	non liquid investment rate	share of non liquid investments
Architecture and Engineering	0.578	0.607	0.730
Arts, Design, Entertainment, Sports, and Media	0.422	0.587	0.756
Building and Grounds Cleaning and Maintenance	0.207	0.576	0.777
Business and Financial Operations	0.512	0.557	0.714
Community and Social Service	0.363	0.623	0.773
Computer and Mathematical	0.596	0.643	0.769
Construction and Extraction	0.374	0.649	0.787
Education, Training, and Library	0.404	0.574	0.743
Farming, Fishing, and Forestry	0.160	0.569	0.845
Food Preparation and Serving Related	0.218	0.574	0.790
Healthcare Practitioners and Technical	0.533	0.576	0.737
Healthcare Support	0.210	0.576	0.777
Installation, Maintenance, and Repair	0.381	0.635	0.786
Legal	0.519	0.550	0.703
Life, Physical, and Social Science	0.471	0.603	0.732
Management	0.590	0.584	0.722
Office and Administrative Support	0.274	0.608	0.778
Personal Care and Service	0.208	0.573	0.778
Production	0.267	0.608	0.782
Protective Service	0.396	0.622	0.777
Sales and Related	0.282	0.578	0.767
Transportation and Material Moving	0.269	0.613	0.789

Table 3: Occupation vs Rates, the saving rate denotes to the average monthly saving rate of the members of each occupation, the non liquid investment rate denotes the average of the decided rate of allocating monthly savings to non liquid investments for each occupation, the share of non liquid investments denotes the share of non liquid assets with respect to all of the investments averaged for each occupation

age	non liquid investment rate	share of non liquid investments	consumption rate	Coco et al. consumption rate
21	0.551	0.677	0.648	0.884
22	0.546	0.684	0.634	0.915
23	0.553	0.690	0.629	0.948
24	0.554	0.701	0.628	0.976
25	0.556	0.710	0.623	0.996
26	0.558	0.717	0.611	0.998
27	0.562	0.728	0.608	0.999
28	0.567	0.736	0.607	0.999
29	0.571	0.743	0.609	0.998
30	0.572	0.750	0.606	0.996
31	0.579	0.756	0.598	0.987
32	0.584	0.764	0.595	0.979
33	0.587	0.771	0.596	0.972
34	0.590	0.777	0.591	0.966
35	0.592	0.781	0.591	0.962
36	0.594	0.782	0.588	0.960
37	0.596	0.783	0.583	0.959
38	0.597	0.786	0.578	0.959
39	0.597	0.786	0.575	0.960
40	0.600	0.789	0.571	0.962
41	0.597	0.787	0.563	0.963
42	0.601	0.785	0.562	0.965
43	0.603	0.785	0.561	0.966
44	0.603	0.787	0.562	0.966
45	0.604	0.786	0.559	0.966
46	0.605	0.788	0.561	0.966
47	0.604	0.787	0.559	0.965
48	0.605	0.789	0.557	0.964
49	0.603	0.788	0.555	0.963
50	0.605	0.788	0.555	0.963
51	0.609	0.789	0.557	0.964
52	0.607	0.787	0.558	0.966
53	0.607	0.785	0.564	0.970
54	0.605	0.785	0.563	0.976
55	0.606	0.786	0.568	0.985
56	0.604	0.787	0.569	0.996
57	0.604	0.787	0.569	1.011
58	0.604	0.790	0.573	1.029
59	0.606	0.790	0.572	1.051
60	0.606	0.789	0.578	1.077
61	0.609	0.788	0.578	1.107
62	0.606	0.788	0.572	1.142
63	0.607	0.784	0.572	1.180
64	0.606	0.785	0.570	1.223

Table 4: Age vs Rates: Consumption rates are defined as consumption amount divided to income. Consumption rates are compared to literature by extracting values from plots of Coco et al. [10], their research differs by our work such that the income values exclude contributions toward pension income, and savings are used as a mean to finance consumption deficit especially during retirement, so during retirement there are positive consumption rates, which mean that the pension deficit is financed by spending savings. This definition difference causes consumption rates to be much higher.

Occupation	20-30	30-40	40-50	50-60	60-70
Architecture and Engineering	0.669	0.68	0.658	0.668	0.325
Arts, Design, Entertainment, Sports, and Media	0.465	0.448	0.445	0.484	0.236
Building and Grounds Cleaning and Maintenance	0.237	0.237	0.236	0.235	0.118
Business and Financial Operations	0.556	0.602	0.594	0.599	0.281
Community and Social Service	0.447	0.422	0.433	0.41	0.2
Computer and Mathematical	0.639	0.672	0.671	0.674	0.319
Construction and Extraction	0.419	0.435	0.432	0.42	0.188
Education, Training, and Library	0.456	0.471	0.472	0.436	0.211
Farming, Fishing, and Forestry	0.181	0.18	0.176	0.184	0.09
Food Preparation and Serving Related	0.236	0.235	0.235	0.236	0.114
Healthcare Practitioners and Technical	0.598	0.615	0.591	0.578	0.308
Healthcare Support	0.236	0.235	0.234	0.234	0.12
Installation, Maintenance, and Repair	0.446	0.427	0.434	0.412	0.212
Legal	0.603	0.584	0.592	0.573	0.322
Life, Physical, and Social Science	0.56	0.549	0.556	0.535	0.286
Management	0.683	0.692	0.692	0.68	0.346
Office and Administrative Support	0.31	0.315	0.317	0.311	0.152
Personal Care and Service	0.238	0.236	0.237	0.236	0.116
Production	0.309	0.305	0.31	0.303	0.147
Protective Service	0.429	0.437	0.445	0.426	0.212
Sales and Related	0.334	0.313	0.323	0.312	0.157
Transportation and Material Moving	0.307	0.302	0.315	0.298	0.152

Table 5: Saving Rate by Occupation and Age

7.7 Raw Plots and Behavioural Parameter Effect Plot

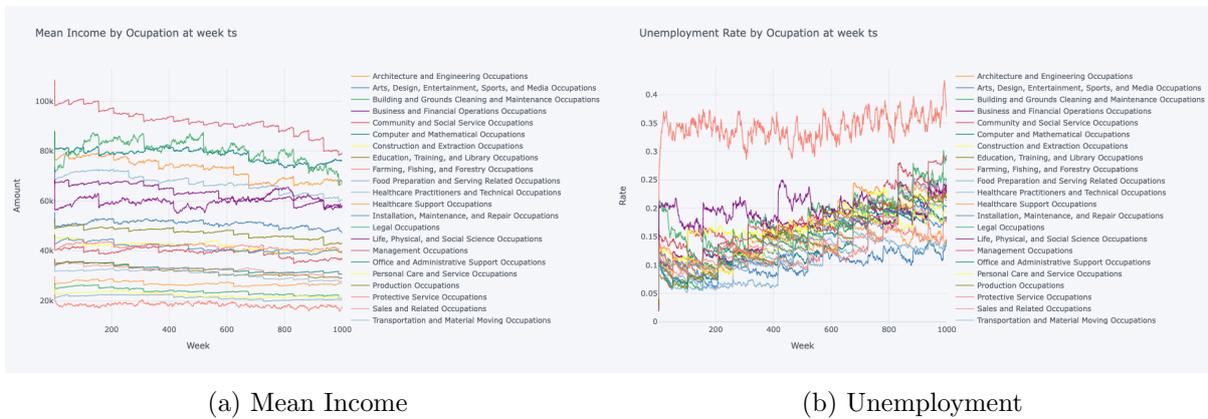


Figure 17: Mean Income and Unemployment by Occupation at week ts



Figure 18: Saving Rate by Occupation at week ts and Saving Rate by Occupation at amount capped at 10M

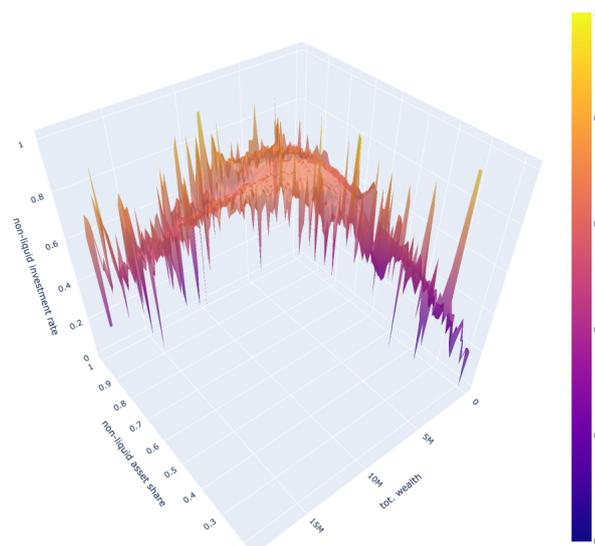


Figure 19: 3d Surface Plot of share of non-liquid assets in x-axis, with respect to total asset wealth in y-axis, and corresponding decision of non-liquid asset investment rate in z-axis

## Scatter Plots

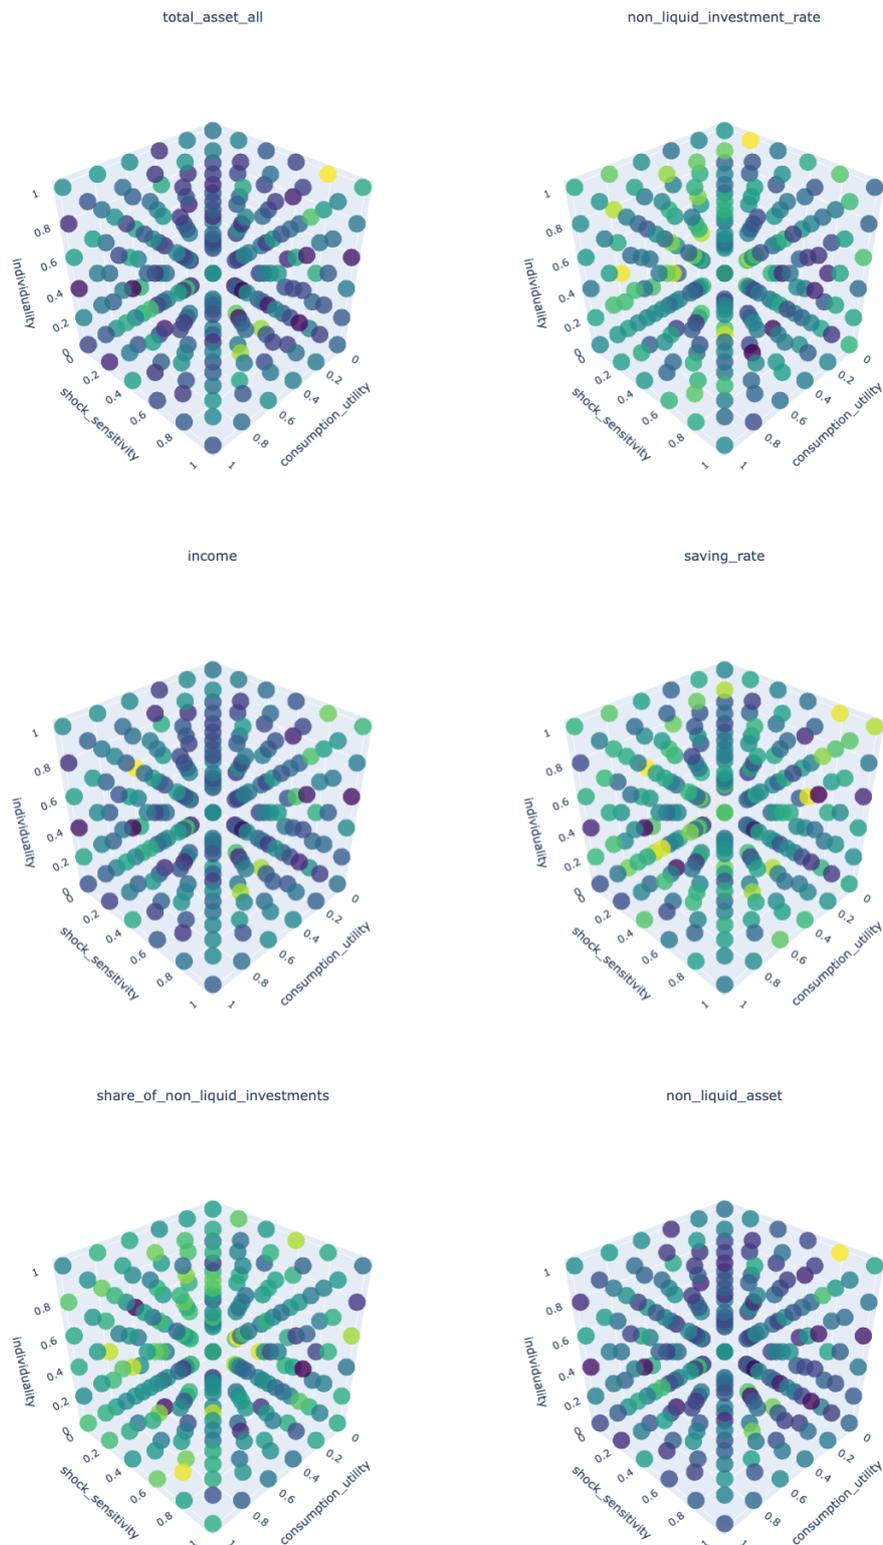


Figure 20: 3D Scatter Plot of each indicator relative to the behavioural parameters of the agents, where dark blue indicate lower values and light yellow indicate higher values, which reflect how the parameters are effecting the values such as accumulated assets, investment rates or share of non liquid assets. The income vs parameters plot is provided for convenience, the income itself is not effected by the behavioural parameters