

The Price of Uniformity: Welfare Costs of One-Size-Fits-All Investment Plans under Income Risk

Paulo RODRIGUES

Finance Department, Maastricht University, Maastricht, Netherlands, 6211 LM,
(p.rodrigues@maastrichtuniversity.nl)

Peter SCHOTMAN

Finance Department, Maastricht University, Maastricht, Netherlands, 6211 LM,
(p.schotman@maastrichtuniversity.nl)

Abstract We study the optimal consumption and investment decisions of investors who are heterogeneous with respect to their income risk. We solve and compare the solutions of two market setups. In the first setup, the investor finds the individual's best investment and consumption strategy. In the second one, a central planner, e.g., a pension fund, sets up a one-size-fits-all investment plan for all investors, and the investors select their optimal consumption path. Instead of relying on a parametric model, we bootstrap economic scenarios and solve the life cycle problem using neural networks to find the optimal policy functions. We find that the one-size-fits-all solution can have costs of up to three percent of certainty equivalent consumption when compared to the individual optimal solution.

Key Words Optimal Portfolio, Dynamic Optimization, Neural Networks, Consumption Choice, Lifecycle Model, Collective Investment Strategy

JEL Classifications: G11, C45, C61

The Price of Uniformity: Welfare Costs of One-Size-Fits-All Investment Plans under Income Risk

Abstract We study the optimal consumption and investment decisions of investors who are heterogeneous with respect to their income risk. We solve and compare the solutions of two market setups. In the first setup, the investor finds the individual's best investment and consumption strategy. In the second one, a central planner, e.g., a pension fund, sets up a one-size-fits-all investment plan for all investors, and the investors select their optimal consumption path. Instead of relying on a parametric model, we bootstrap economic scenarios and solve the life cycle problem using neural networks to find the optimal policy functions. We find that the one-size-fits-all solution can have costs of up to three percent of certainty equivalent consumption when compared to the individual optimal solution.

Key Words Optimal Portfolio, Dynamic Optimization, Neural Networks, Consumption Choice, Lifecycle Model, Collective Investment Strategy

JEL Classifications: G11, C45, C61

1. INTRODUCTION

We study consumption and portfolio decisions of investors who are heterogeneous with respect to their labor income. We compare and evaluate two setups for the investment decisions. In the first case, all investors make their own optimal consumption and investment decisions. Second, we assume that the investment decisions are outsourced to a central planner who optimizes the portfolio for a collective of heterogeneous investors. This setup mimics, e.g., the institutional arrangements of most asset-backed pension systems, where a pension fund invests on behalf of its participants. Comparing both outcomes lets us quantify the welfare costs of a collective investment decision that arises from heterogeneity in income.

To derive our conclusions, we make two choices that are non-standard in the literature. First, instead of assuming a parametric model for the dynamics of the economy we use bootstrap methods to generate sample paths of the economy. Using a block bootstrap algorithm allows us to circumvent the necessity to estimate or calibrate the covariance matrix between income returns and capital market returns. This covariance drives in a large part the optimal portfolio decisions, as one of the main drivers is the possible hedging benefit for labor income risk that the optimal portfolio can offer. The second choice concerns the optimization algorithm. Here we train a neural network to find the optimal policy function that translates states of the world into economic decisions.

Our analysis connects to several strands of the literature. First, we connect to papers that study the relationship between income risk and financial market outcomes. An early example of this is Cocco, Gomes, and Maenhout (2005), who solved a life cycle model including a parametric income model and showed how income risks impact optimal portfolio and consumption choices. Catherine (2022) investigates how the relationship between idiosyncratic income skewness and market returns impacts portfolio choices of investors. Catherine, Sodini, and Zhang (2024) shows how tail risks in household income predict differences in portfolio holdings across households in Sweden. Eiling (2013) investigates the impact of heterogeneous human capital impacts the cross-section of stock returns. Schmidt (2025) shows that stock returns provide information about labor market risks and labor market uncertainty impact stock returns.

The second strand of literature this paper relates to is collective investments. More advanced studies here investigate optimal sharing rules using continuous time models that optimize utility of terminal wealth. Branger, Chen, Mahayni, and Nguyen (2023) investigate the problem under the assumption that investors are heterogeneous with respect to initial wealth as well as to the risk aversion coefficient. Jensen and Nielsen (2016) investigate the sub-optimality of a linear sharing rule when investors have differing risk attitudes. Both papers assume that the

planner optimizes a weighted average of the investor utilities. In contrast to that, Balter and Schweizer (2024) investigate the optimal decision rule for a planner that wants to maximize the distribution of certainty equivalents of the investors.

Finally, we also contribute to the literature that develops numerical solution algorithms for dynamic optimization problems. Using a simulation based approach to solve dynamic optimization problems was introduced in Brandt, Goyal, Santa-Clara, and Stroud (2005). This algorithm was used and improved by Koijen, Nijman, and Werker (2009) and Diris, Palm, and Schotman (2015). Papers that combine a simulation based approach with neural networks as a solution method to dynamic optimization problems are, e.g., Duarte, Fonseca, Goodman, and Parker (2021) or Druedahl and Røpke (2025).

Concretely, our contribution can be summarized as follows. First, we solve the life cycle problem using bootstrapped sample paths. Prior literature that used a bootstrap method in this context is, e.g., Anarkulova, Cederburg, and O’Doherty (2023). We differ to this study in that we also consider collective investment problems, analyze heterogeneous income processes, and optimize the consumption and portfolio decisions. This procedure distinguishes our study also from papers that use a parametric model for the economic dynamics. The main distinction to the collective investment literature is that our main source of heterogeneity is the income process of the investors instead of risk aversion, as is standard in the literature. A second main distinction is that we optimize the consumption and the investment decisions instead of optimizing utility of terminal wealth. Finally, our solution method can be used also for dynamics that allow for changing investment opportunities. Our methodological contribution consists in applying neural network methods to solve dynamic optimization problems for collective investment portfolios. This method is flexible enough to also cover heterogeneous risk aversion across investors and allows for more realistic economic dynamics than the models usually used.

We measure the heterogeneity in income by using data on wages and salaries provided by the Bureau of Economic Analysis through their National Income and Products Accounts. We separate the investors by income generated by working in the finance, manufacturing, construction, service, retail, government, transportation, mining, and wholesale sectors. We find that using a tailor made investment policy can generate welfare gains for up to three percent, when measured in terms of certainty equivalent consumption. This gain is measured for the retail sector. The minimum difference is 0.06% for the service sector. We also find in general more conservative investment strategies for the planner portfolio policy compared to the full individual optimal. This difference is particularly striking after retirement.

2. PROBLEM AND SOLUTION METHOD

The goal is to solve the consumption and portfolio problem for two setups. In the first setup, all investors can have tailor made (Tailor) decisions for consumption and portfolio plans. This provides the benchmark to which to evaluate the welfare costs of the collective investment strategy against. Second, we have a central planner (Planner) that decides on the investment decisions for a group of heterogeneous investors. This Planner, e.g., a pension fund, does not restrict the consumption and savings decision of the investors to be equal. The investors can have different consumption-wealth ratios over time and states of the world.

2.1 Problem

We assume that each investor ranks the outcomes of the lifetime investment and consumption decisions according to their expected utility values given by

$$U_i = E_{0i} \left[\sum_{t=0}^T \delta^t q_t \frac{C_{i,t}^{1-\gamma}}{1-\gamma} \right], \quad (1)$$

with $\delta = 0.96$ denoting the subjective discount factor, where q_t is the probability that the individual is alive at time t assuming she is alive at the beginning of the planning period ($q_0 = 1$), C_{it} denotes consumption at t . This setup is standard and is the base case in Cocco et al. (2005). In our calibration, we use the survival probabilities provided by the Royal Dutch Actuarial Association.¹

In case the investor optimizes the investment strategy independently, the budget constraint is given by

$$W_{it} = (W_{it-1} + Y_{it} - C_{it}) R_{it}^{(P)} \quad (2)$$

$$R_{it}^{(P)} = \sum_{j=1}^K \omega_{jit-1} R_{jt} \quad (3)$$

with W_{it} denoting wealth, Y_{it} denoting income, $R_{it}^{(P)}$ denoting the total return of the portfolio from $t-1$ to t , R_{jt} denoting the total return of asset j from $t-1$ to t , and ω_{jit-1} denoting the weight of asset j in the portfolio of investor i determined at $t-1$. The investor decides on C_{it} as well as $\omega_{j,i,t-1}$ to maximize Equation (1).

¹For the computational work we use the 2022 Dutch survival probabilities from the [Actuarieel Genootschap](#).

In case the investment decisions are outsourced to a central planner, e.g., a pension fund, we assume that the central planner follows a utilitarian approach and aims to maximize the functional given by

$$V(U_i) = \sum_i^N e_i U_i \quad (4)$$

with N denoting the number of distinct investors and e_i denoting the weight that the planner puts on each investor. Each weight is larger than zero, and the weights sum to one. If the investment decision is made by the planner, it is uniform across all investors. In this case the budget constraint for each investor is given by

$$W_{it} = (W_{it-1} + Y_{it} - C_{it})R_t^{(P)} \quad (5)$$

$$R_t^{(P)} = \sum_{j=1}^K \omega_{jt-1} R_{jt} \quad (6)$$

with $R_t^{(P)}$ the return of the portfolio of all investors and $\omega_{j,t-1}$ denoting the weight of asset j in the portfolio in percentage of investable wealth of all investors. The setup models the consumption decision of each investor and the portfolio decision of the Planner as a joint optimization problem. The portfolio decisions impact all U_i in Equation (4) and the consumption decisions only impact one U_i .

The investors are heterogeneous with respect to their income² Y_{it} , which is paid at the beginning of the period. This income is added to the wealth from the prior period, and all investors decide on the level of consumption independently of each other. In addition, we impose the following constraints on the problem for all investors i at each time point t :

$$\sum_{j=1}^K \omega_{ijt} = 1 \quad (7)$$

$$\omega_{ijt} \geq 0 \quad \forall j \quad (8)$$

$$C_{it} \leq W_{it-1} + Y_{it}. \quad (9)$$

With (7) ensuring that all wealth that is not consumed is invested, (8) impose no short sale constraints on the portfolio weights. Since all asset weights are non-negative, the constraints also rule out leverage. Finally, (9) restricts wealth to be non-negative at all times for all investors.

²In principal, we can also model investor heterogeneity across more dimensions than income. For example, the investors could differ with respect to their risk aversion or could also have different survival probabilities. However, we disregard these other sources for this analysis.

2.2 Solution Method

The solution to this problem consists of finding policy functions that determine the decisions at time t given the information set at the same time. Formally, we want to find

$$\omega_{it} = g_i^\omega(I_t) \tag{10}$$

$$C_{it} = g_i^c(I_t) \tag{11}$$

with I_t denoting the information set at time t , which includes the time to maturity $T - t$, and g^x and g^c denoting the policy functions that transform the information into the decisions. We include time to maturity $T - t$ as one argument in the function, since we have a finite-horizon problem and these decisions are, therefore, time dependent. Using this setup allows us to train one policy function across all time periods, which significantly simplifies the problem. Note that the policy function for the consumption decision is indexed by i , since each investor can decide on the consumption stream. In case the investment decision is made by the planner; the portfolio policy will be identical across all investors and we can drop the index i in Equation (10).

Determining these functions in closed form is only possible in a few models with restrictive assumptions. Therefore, the standard approach in the literature is to use numerical methods and find solutions using backward induction. Although computing power has been increasing, this method still poses severe computational hurdles for model setups of interest. For example, if the investor has a utility that is not homothetic in wealth, the optimal decisions depend on the wealth of the investor at time t , which is unknown when using backward induction. This problem can be solved by using a grid for wealth, however, this increases the computational costs considerably. In addition, each endogenous state variable increases the dimensionality of the grid, making the problem intractable.

We train a neural network to find the policy functions given in Equations (10) and (11). A neural network, in essence, consists of a chain of function transformations of input data that result in the desired output at the end of the chain. This flexible setup has been shown in, e.g., Hornik, Stinchcombe, and White (1989) or Leshno, Lin, Pinkus, and Schocken (1993), by proving a universal approximation theorem that shows it can approximate any Borel measurable function. This is the key difference to other proposals in the literature that used polynomials to approximate the policy function, as, e.g., Chacko, Desai, Golts, and Novakovsky (2005). These approximations are not flexible enough, a problem that is circumvented by using neural networks. The universal approximation theorem ensures that the method converges to the optimal solution if the number of transformations grows large. Figure 1 shows a graphical

representation of the framework. The figure shows the evolution of the network at each point t . At each $t = 0, \dots, T - 1$ the algorithm proceeds as follows:

1. Income is added to the wealth of each investor

$$W_{it} = W_{it-} + Y_{it}$$

with $W_{i,t-}$ denoting financial wealth at the beginning of the period.

2. The information set, which consists of the inputs to the policy function is determined, e.g. in our application, $I_t = (dp_t, R_t^Y, T - t)$, with dp_t denoting vector containing the dividend-price ratios and R_t^Y
3. Two hidden-layer transformations are applied in the next step

$$o_t^{(1)} = \psi_1(a_1 + A_1 I_t) \quad (12)$$

$$o_t^{(2)} = \psi_2(a_2 + A_2 o_t^{(1)}) \quad (13)$$

where a_1 and a_2 are $L_1 \times 1$ and $L_2 \times 1$ vectors of free parameters, respectively. A_1 is a $L_1 \times L_I$ matrix, and A_2 is a $L_2 \times L_1$ matrix of free parameters, respectively. We have that L_1 and L_2 denote the respective numbers of neurons in each layer, with L_I denoting the number of variables that are in the information set. The functions $\psi_1(x)$ and $\psi_2(x)$ indicate a non-linear transformation of the argument x , here the values resulting from the affine transformations. One popular option for this transformation is the Rectified Linear Unit (ReLU) function which would set $\psi(x) \equiv \max(x, 0)$.

4. The last output of the hidden layers is transformed to arrive at the decisions

$$\omega_t = \psi_x(a_x + A_x o_t) \quad (14)$$

$$c_t = \psi_c(a_c + A_c o_t) \quad (15)$$

with a_x and A_x denoting a $M \times 1$ vector and $M \times L_2$ matrix, respectively. Restricting the number of neurons to M ensures we can interpret the output x_t as a portfolio vector. To restrict the weights to be between zero and one and to sum to one, we use the Softmax function at this stage, i.e., $\psi(x_i) \equiv \frac{e^{x_i}}{\sum_i e^{x_i}} \forall i = 1, \dots, M$. The output of this layer is used to compute the portfolio return of all investors, since the investment decision is made by the planner.

For the consumption decisions, set a_c to a $N \times 1$ vector and A_c to a $N \times L_2$ matrix, i.e., as many neurons as investors for this function. As a non-linear transformation, we use

the Sigmoid function, i.e., $\psi_c(x) \equiv \frac{1}{1+e^{-x}}$, which ensures that the output is between zero and one and we interpret the result as the consumption-wealth ratios for each investor. To arrive at the consumption level we multiply the output by current wealth, i.e.,

$$C_t = c_t \odot W_t \quad (16)$$

with $\odot$ denoting the Hadamard product.

5. Wealth is carried forward using

$$W_{(t+1)-} = (W_t - C_t)\omega'_t R_{t+1}$$

Once we arrive at $T - 1$ we can compute the consumption outcomes for all investors. Repeating the procedure for S sequences, we can compute an estimator for the expected utility of each investor as follows

$$\hat{U}_i = \frac{1}{S} \sum_{s=1}^S \sum_{t=0}^T \delta^{t-1} q_t \frac{C_{it,s}^{1-\gamma}}{1-\gamma} \quad (17)$$

If the investor optimizes the portfolio positions, the loss is simply


$$\text{LOSS} = -\hat{U}_i \quad (18)$$

In case we are solving the problem where the planner does the investment decisions, the loss function to be minimized is then given by

$$\text{LOSS} = -\sum_{i=1}^N e_i \hat{U}_i \quad (19)$$

with e_i denoting the percentage of total employment in sector i .

Several points are worth mentioning. First, we have not specified the data-generating process of the economy nor the functional form of the investor utility. To use the algorithm, we need to be able to simulate the data-generating process and evaluate the realized utility. Second, instead of training the neural network to find one policy function at each point in time, as in equations (10) and (11), we include the time to maturity as a state variable and train the neural network to find only one function. Since the time to maturity is one input variable, the decisions become time-dependent, as is necessary for finite time consumption and portfolio plans. The main advantages of this approach are that we can significantly reduce the number of parameters to fit and we can find the policy for any decision interval in one go. Since time



The graph shows the structure of the recursive neural network containing inputs (W_t, Z_t) , a hidden layer with neurons in the vector o_t , and decision layers for consumption and portfolio weights. Consumption also depends on income, which is assumed exogenous, *i.e.* the innovation in the state variables Z_{t+1} affect income, but not income does not affect the macro states Z_{t+1} .

Figure 1: recursive NN with consumption

to maturity changes in different increments for different decision frequencies, we can train our network on annual data and apply the policy function for monthly increments. Third, we do not condition on the starting values in our sequences of the economy. This ensures that our solution can be applied to any starting value for the state of the economy. Fourth, in principle, the setup of the network can be made more complex. For example, we can connect more than one transformation in the hidden layer. Although, increasing the number of nodes should generate a fully flexible functional form, in practice the chaining of transformations achieves a better approximation than just increasing the number of layers. We have used only one transformation step in the description above but the generalization to more steps is straightforward. In fact, as shown in Hornik et al. (1989) or Leshno et al. (1993), neural networks can approximate any Borel measurable function if the number of hidden units is large enough. Fifth, this method easily takes care of endogenous state variables like, e.g., wealth in case the utility function uses a reference level. Since this algorithm evaluates the loss function by going forward in time, we do not need to solve these problems on a grid, as is usually the case when using backward induction. Finally, we can relax the no leverage or short sale constraint used in the description. This can be done by modifying the non-linear transformation for the portfolio decision step. However, it is not easy to guarantee that the resulting policy function does not violate the budget constraint, as in any simulation exercise of a discrete-time model.

3. ECONOMIC SCENARIOS

3.1 Data Construction

To use the neural network solution method, we need to feed the algorithm time series of asset returns, state variables, as well as income. To determine income, we use the data on wages and salaries provided by the Bureau of Economic Analysis (BEA). The BEA provides data on wages and salaries for sectors separated by SIC code from 1958-1998 and by NAICS code from 1998 on. Table 1 contains information on the BEA table names and line codes used when downloading the data. We use the 10 sectors identified by the first two digits of the SIC codes for the data before 1998 and allocate the NAICS codes for the data after 1998 as described in Table 1. We sum the wages and salaries for the quarters within each year to arrive at annual data. To construct the time series of income returns we follow Jagannathan and Wang (1996) and Eiling (2013) and use

$$R_{k,t}^{inc} = \frac{L_{k,t} + L_{k,t-1}}{L_{k,t-1} + L_{k,t-2}} \quad (20)$$

where $L_{k,t}$ denotes the total wages and salaries for sector k in year t .

Stock market data comes from the Center for Research in Security Prices (CRSP). We use the monthly stock returns from the new file format released in January 2025. We keep all common equity shares traded on the AMEX, NASDAQ, and NYSE. The risky part of the asset menu for the planner consists of the value-weighted CRSP stock market index. For the case in which the investor also optimizes the investment decisions individually, we construct two risky asset portfolios that the investor can invest in. The first is a value-weighted portfolio of all stocks that are not in the sector in which the investor works. For example, if the investor works in the finance industry, the return of this portfolio is computed by using the value-weighted average return of all stocks that have an SIC code that does not belong to the finance industry. We construct such a portfolio for each sector k under consideration as follows

$$R_{k,t}^{mkt} = \sum_{i \notin k} v_{i,k} R_{i,t} \quad (21)$$

with $v_{i,k}$ denoting the relative market value of stock i in a portfolio of stocks that does not contain the stocks of companies with SIC codes of sector k . The second risky asset in the menu of the investor is the value-weighted portfolio of the companies that have an SIC code of the sector in which the investor works. This sector portfolio is constructed following the methodology described in Fama and French (1992), i.e., we construct the sector portfolio at

Table 1: BEA LineCodes by Sector

Sector	SIC: SQINC7S	NAICS: SQINC7N
Forestry	100	100
Mining	200	200
Construction	300	400
Manufacturing	400	500
Transp., Comm., and Util.	500	300
		800
		900
Wholesale Trade	610	600
Retail Trade	620	700
Finance, Insurance, Real Estate	700	1000
		1100
Services	800	1200
		1400
		1500
		1600
		1700
		1800
		1900
Government	900	2000

This table shows the line codes we use to allocate the information from the BEA to the sectors. The column “SIC: SQINC7S” shows the information for the data period 1959-1998, where the BEA uses the SIC codes to identify sectors. The column “NAICS: SQINC7N” shows the information for the data from 1999 on, where the BEA uses the NAICS codes to identify sectors.

the end of June of each year and hold the portfolio for the following year.

The reason to set up the risky asset menu like this is to create the possibility to have an implicit short position in the own sector even if we impose short sale and leverage constraints. The short position is implicit, since the market index is computed without the companies operating in their own sector. The investor could now include their own sector return in the portfolio and re-create the total market portfolio, which is the usual option in the risky asset menu in the life-cycle literature. In addition, we include a risk-free asset into the asset menu of both problems. The returns of this asset are downloaded from the Fama-French factors via

WRDS.

Moreover, we use the consumer price index (CPI) from the Federal Reserve via the Federal Reserve Economic Data (FRED) database. This time series is used to convert all returns to real returns as follows

$$\tilde{R}_t = \frac{R_t}{I_t} \quad (22)$$

with $I_t = \frac{CPI_t}{CPI_{t-1}}$ denoting inflation.

Finally, we use dividend-price ratios as state variables. We denote the dividends received between $t - 1$ and t by D_t and compute the ratios as

$$dp_t = \left(\frac{P_t + D_t}{P_{t-1}} - \frac{P_t}{P_{t-1}} \right) \frac{P_{t-1}}{P_t} = \frac{R_t - R_t^{ex}}{R_t^{ex}} = \frac{D_t}{P_t} \quad (23)$$

with R_t denoting the total return and R_t^{ex} denoting the return excluding dividends. As this ratio is already in real terms, no further inflation adjustment is necessary.

3.2 Bootstrap Method

Instead of relying on a parametric model, we use bootstrap to generate sample paths of the economy. In particular, we use a moving block bootstrap algorithm with replacement with a block length of 10 years. This algorithm samples a 10 year block of all time series under consideration and connects them until the desired time length for the sample is reached. In our case, we generate sequences of 76 years of length. To generate these bootstrap samples, we use the Python package `tsbootstrap` (Gilda, Heidrich, and Kiraly (2024)).³

We collect the income and stock returns as well as the dividend-price ratios for the case under consideration and bootstrap 100 thousand sequences for the training and the validation of the neural network and another 100 thousand sequences for the testing of the policy functions. For example, if we want to find the optimal policy function for the investor working in the finance industry, we collect the data on the risk-free rate, the return on the market ex finance, the finance sector return, the dividend-price ratio of the market ex finance, the dividend-price ratio of the finance sector, as well as the income return of the finance sector. Using this, we construct the sample paths using the moving block bootstrap. For all other problems the data sample is constructed in a similar fashion.


Using the income returns, we can construct the time series of income by simply compounding

³The package documentation can be found at <https://tsbootstrap.readthedocs.io>.


Table 2: Descriptives of capital market returns

	Mean	Std. Dev.	Min	25%	50%	75%	Max
R_f	0.006	0.023	-0.061	-0.009	0.007	0.022	0.055
$R_{exFinan}$	0.082	0.172	-0.362	-0.015	0.119	0.219	0.315
R_{Finan}	0.089	0.207	-0.511	-0.046	0.101	0.217	0.477
$dp_{exFinan}$	0.027	0.012	0.009	0.017	0.026	0.034	0.055
dp_{Finan}	0.028	0.011	0.011	0.019	0.026	0.034	0.054
$R_{exManuf}$	0.079	0.175	-0.422	-0.053	0.126	0.198	0.344
R_{Manuf}	0.086	0.170	-0.352	-0.012	0.120	0.220	0.324
$dp_{exManuf}$	0.029	0.014	0.011	0.017	0.027	0.037	0.062
dp_{Manuf}	0.026	0.010	0.010	0.019	0.025	0.033	0.048
$R_{exConst}$	0.082	0.171	-0.390	-0.024	0.124	0.214	0.331
R_{Const}	0.092	0.288	-0.481	-0.110	0.098	0.251	1.060
$dp_{exConst}$	0.027	0.012	0.010	0.018	0.026	0.034	0.055
dp_{Const}	0.017	0.011	0.004	0.007	0.015	0.026	0.041
$R_{exServi}$	0.080	0.165	-0.385	-0.025	0.127	0.190	0.331
R_{Servi}	0.110	0.260	-0.439	-0.010	0.129	0.248	0.967
$dp_{exServi}$	0.028	0.011	0.012	0.020	0.027	0.035	0.056
dp_{Servi}	0.014	0.009	0.001	0.007	0.012	0.023	0.033
$R_{exRetai}$	0.081	0.170	-0.401	-0.023	0.125	0.207	0.346
R_{Retai}	0.106	0.226	-0.411	-0.028	0.105	0.269	0.605
$dp_{exRetai}$	0.028	0.012	0.011	0.018	0.027	0.035	0.056
dp_{Retai}	0.020	0.010	0.005	0.012	0.017	0.026	0.047
$R_{exTrans}$	0.086	0.177	-0.398	-0.026	0.128	0.215	0.337
R_{Trans}	0.063	0.161	-0.329	-0.033	0.069	0.166	0.367
$dp_{exTrans}$	0.025	0.010	0.009	0.017	0.022	0.032	0.048
dp_{Trans}	0.042	0.020	0.017	0.027	0.035	0.053	0.090
$R_{exMinin}$	0.082	0.173	-0.386	-0.027	0.123	0.215	0.337
R_{Minin}	0.079	0.239	-0.472	-0.083	0.068	0.253	0.563
$dp_{exMinin}$	0.027	0.012	0.010	0.018	0.026	0.034	0.055
dp_{Minin}	0.028	0.015	0.009	0.013	0.025	0.040	0.065
$R_{exWhole}$	0.082	0.171	-0.390	-0.025	0.124	0.216	0.332
R_{Whole}	0.075	0.205	-0.380	-0.046	0.102	0.205	0.644
$dp_{exWhole}$	0.027	0.012	0.010	0.018	0.026	0.034	0.055
dp_{Whole}	0.020	0.009	0.008	0.013	0.017	0.025	0.040


This table shows the descriptive statistics of capital market returns. R_f denotes the risk free rate, $R_{exSector}$ denotes the return of the market portfolio when we exclude the “Sector”, R_{Sector} denotes the return of the “Sector” portfolio, $dp_{exSector}$ denotes the dividend-price ratio of the market portfolio that excludes the “Sector”, and dp_{Sector} denotes the dividend-price ratio of the “Sector” portfolio. The column “Mean” shows the arithmetic average over time, the column “Std. Dev.” the standard deviation, “Min.” the minimum, and “Max.” the maximum. The columns “25%”, “50%”, and “75%” show the respective quantiles. All information is in annual terms and the data covers the period 1960-2024.




A: Construction



B: Manufacturing




C: Transportation




D: Retail

Figure 2: Income paths across sectors


This figure shows the income paths implied the income returns of the test dataset. The solid line shows the arithmetic average across all 100,000 paths in the test dataset. The other three lines show randomly selected paths from the same dataset. Panel “A” shows the paths for the construction sector, Panel “B” for the manufacturing sector, Panel “C” for the transportation sector, and Panel “D” for the retail sector.




A: Mining



B: Wholesale



C: Finance



D: Service

Figure 3: Income paths across sectors

This figure shows the income paths implied the income returns of the test dataset. The solid line shows the arithmetic average across all 100,000 paths in the test dataset. The other three lines show randomly selected paths from the same dataset. Panel “A” shows the paths for the mining sector, Panel “B” for the wholesale sector, Panel “C” for the finance sector, and Panel “D” for the service sector.

Table 3: Descriptives of income returns

	Mean	Std. Dev.	Min	25%	50%	75%	Max	Corr R_{ex}	Corr R_{Sector}
Finan	0.035	0.031	-0.054	0.015	0.040	0.054	0.110	0.055	0.135
Manuf	0.003	0.030	-0.063	-0.017	0.005	0.021	0.062	-0.012	-0.016
Const	0.023	0.046	-0.132	-0.011	0.037	0.052	0.098	0.015	0.066
Servi	0.044	0.022	-0.008	0.027	0.043	0.059	0.109	0.108	0.151
Retai	0.013	0.030	-0.085	-0.004	0.018	0.032	0.059	-0.020	-0.122
Gover	0.021	0.026	-0.041	0.008	0.020	0.037	0.076	0.111	nan
Trans	0.024	0.040	-0.051	0.007	0.020	0.036	0.207	0.128	0.084
Minin	0.011	0.067	-0.163	-0.026	0.014	0.051	0.140	-0.130	-0.227
Whole	0.020	0.023	-0.046	0.004	0.024	0.039	0.057	-0.032	0.065

This table shows the descriptive statistics of income returns for all sectors. The column “Mean” shows the arithmetic average over time, the column “Std. Dev.” the standard deviation, “Min.” the minimum, and “Max.” the maximum. The columns “25%”, “50%”, and “75%” show the respective quantiles. The columns “Corr R_{ex} ” and “Corr R_{Sector} ” shows the contemporaneous correlations between the income returns and the market returns ex sector and the sector portfolio, respectively. All information is in annual terms and the data covers the period 1960-2024.

the starting value over time. For the starting value, we use the estimation results in Cocco et al. (2005) for a single household. Figures 2 and ?? show the resulting income processes for the sectors under consideration. The figures show the average income at all ages over all sample paths in the test dataset as the solid line. In addition, we show three other paths that were randomly selected. As can be seen, there is considerable risk in the income path. Particularly an investor working in the mining industry is exposed to the possibility of decreasing income over time.

4. RESULTS

We calculate the certainty equivalent value of the outcomes from the optimal life cycle strategies by computing the constant consumption stream that would result in an equal expected utility as using the optimal decisions. The results can be found in Table 4. These results also provide an estimation for the welfare costs that each investor incurs when the investment strategy follows a one-size-fits-all approach.

We see that from the eight sectors under consideration, we have four sectors that can improve the outcomes from the one-size-fits-all strategy by over 1%, namely transportation, manufacturing, construction, and retail. These sectors can improve their outcomes by 1.2%, 2.1%, 2.6%, and 3.1%, respectively. These are all sizable effects. The sectors that are affected by losses below

Table 4: Certainty Equivalent Results

	Planner	Tailor	Increase	Weight
Minin	14.048	14.070	0.156	0.443
Const	18.990	19.484	2.603	6.281
Manuf	16.151	16.484	2.063	9.669
Trans	20.052	20.301	1.243	7.775
Whole	20.441	20.542	0.493	4.682
Retail	17.833	18.386	3.104	11.773
Finan	21.320	21.437	0.548	6.976
Servi	22.679	22.693	0.059	34.549

This table shows the certainty equivalent consumption results across sectors. The columns “Planner” shows the certainty equivalent value for the collective investment plan, the column “Tailor” shows the certainty equivalent value for the individually optimal investment policy, the column “Increase” shows the increase in percentage of the individual optimal outcome when compared to the collective investment policy, and the column “Weight” shows the weight in percentage that each sector has in the criterion function for the central planner given in Equation (4).

1% are mining, wholesale, finance, and services. They lose 0.16%, 0.49%, 0.55%, and 0.06%, respectively.

4.1 Consumption-Wealth Ratio

The resulting consumption-wealth ratios are presented in Figure 4, where Panel A displays the results when the planner assumes control of the investment strategy, and Panel B shows the results for the tailor-made optimisation. Three results are immediately visible from the figure. First, the optimal choice for many sectors is not to save at all in their early years. Second, all consumption-wealth ratios converge to roughly the same values after retirement. Third, if the investment strategy is the same across sectors, the investor starts to save earlier than when the investment strategy is tailor-made.

This third result is confirmed for each sector in Figures 5 and 6. We see that the consumption-wealth ratio for the tailor-made optimization in blue is never below the one-size-fits-all results in red. We see in both figures that the largest divergence between the ratios occur at young ages, with smaller differences in retirement.

Finally, two further takeaways from these results are worth mentioning. First, the decision on consumption-wealth depends largely on the age of the investor. We see almost no impact of the states in the economy, as the distribution in of the violin plots is tight around the expected value. At a very young age as well as in retirement the violin shape is very narrow. We see





Figure 4: Consumption-Wealth ratio

This figure shows the consumption-wealth ratios across investor age for all sectors under consideration. Panel “A” shows the results for the collective investment plan and Panel “B” shows the result for the individually optimal investment strategy. The legend within the figure shows which sector is shown in which color. To improve clarity of the figure, we only show the violin for every three years of the life cycle.


larger bodies of the violin plots for the ages at which the investor saves for retirement. Second, for the individually optimal life cycle decisions, the time at which the investor starts saving for retirement basically follows the rank of the expected wage growth, i.e., the larger the expected wage growth the later the investor saves for retirement. One exception to this rule is the investor working in the mining industry and the investor of the manufacturing industry. The investor in the mining industry has a very low wage growth rate and therefore needs to start saving at a very young age. Investors in the manufacturing industry have a lower expected wage growth but also a much lower standard deviation. In addition, the max wage loss for the mining industry is 16.3% and the max loss for manufacturing is 6.3%. The wage process in the mining industry is much riskier; therefore, the higher savings rate.

4.2 Investment Strategy


The optimal investment strategies, naturally, depend on the return characteristics of the assets, the risk aversion of the investor, and the income risk the investor is exposed to. The main source of heterogeneity across investors is the income risk, which can have different correlations to equity returns. If the investment decision is outsourced to a central planner, the investment decisions have to balance the income risk exposures of all investors under consideration. The results of the portfolio decisions for this case can be seen in Panel A of Figure 7. We see that the




A: Construction



B: Manufacturing




C: Transportation




D: Retail

Figure 5: Consumption-Wealth ratio across sectors


This figure shows the consumption-wealth ratios across investor age for the individual optimal solution in blue and the collective investment strategy in yellow. Panel “A” shows the paths for the construction sector, Panel “B” for the manufacturing sector, Panel “C” for the transportation sector, and Panel “D” for the retail sector. To improve clarity of the figure, we only show the violin for every three years of the life cycle.




A: Mining



B: Wholesale



C: Finance



D: Service

Figure 6: Consumption-Wealth ratio across sectors

This figure shows the consumption-wealth ratios across investor age for the individual optimal solution in blue and the collective investment strategy in yellow. Panel “A” shows the paths for the mining sector, Panel “B” for the wholesale sector, Panel “C” for the finance sector, and Panel “D” for the service sector. To improve clarity of the figure, we only show the violin for every three years of the life cycle.





Figure 7: Investment policy of the planner

This figure shows the violin plots for the portfolio decisions of the collective investment decision in Panel “A”. The violin plot in blue shows the distribution of the investments into the stock market index and the violin plot in green shows the distribution in the risk-free investment over the age of the investor. Panel “B” of the figure shows the violin plots of the cumulative returns of the investment portfolio of the central planner in red, the individual optimal portfolio for the service sector in blue, and the individual optimal portfolio for the retail sector in yellow. To improve clarity of the figure, we only show the violin for every three years of the life cycle.


violin plots at a young age are concentrated at 100% for equities and 0% in the risk-free asset. After retirement the equity exposure decreases and the risk-free component of the portfolio goes up.

In case the investor can optimally chose the equity exposure of the investment portfolio, we expect the investor to make investment decisions in reaction to the personal income risk. Recall from Table 3 that some sectors have negative correlations between income and sector returns. We would expect investors working in these sectors to have significant exposure to their own sector. In contrast, for sectors with significant higher income and sector portfolio correlation we expect to have a higher exposure to the market ex sector portfolio.


The investment decisions for the eight tailor made strategies are depicted in Figures 8 and 9. We see that the portfolio decisions are less dependent on the retirement event when compared to the outcomes for the planner. We observe a tendency to de-risk over time but the effect is more moderate than we observed for the planner. To see the impact of these different investment decisions on the realized portfolio returns, we show in Panel B of Figure 7 the violin plots of the cumulative returns of the planner portfolio in red, the service sector in blue and the retail sector in yellow. We see that returns are very similar at the beginning of the life cycle but the planner returns diverge lower after retirement, since the de-risking is more pronounced than for the tailor-made solution.




A: Construction



B: Manufacturing




C: Transportation




D: Retail

Figure 8: Investment policy across sectors


This figure shows the violin plots for the portfolio decisions for the individually optimal investment decisions. The violin plot in blue shows the distribution of the investments into the market index excluding the own sector, the violin plot in red shows the investments into the own sector portfolio, and the violin plot in green shows the distribution in the risk-free investment over the age of the investor. Panel “A” shows the paths for the construction sector, Panel “B” for the manufacturing sector, Panel “C” for the transportation sector, and Panel “D” for the retail sector. To improve clarity of the figure, we only show the violin for every three years of the life cycle.




A: Mining



B: Wholesale



C: Finance



D: Service

Figure 9: Investment policy across sectors

This figure shows the violin plots for the portfolio decisions for the individually optimal investment decisions. The violin plot in blue shows the distribution of the investments into the market index excluding the own sector, the violin plot in red shows the investments into the own sector portfolio, and the violin plot in green shows the distribution in the risk-free investment over the age of the investor. Panel “A” shows the paths for the mining sector, Panel “B” for the wholesale sector, Panel “C” for the finance sector, and Panel “D” for the service sector. To improve clarity of the figure, we only show the violin for every three years of the life cycle.

4.3 Consumption

The main variable of interest is the consumption path. The violin plots for the eight sectors can be found in Figures 10 and 11. The consumption patterns at a young age largely overlap. At this age the investor has very low savings, if at all, and all consumption is financed through wages, which are not affected by the investment strategy. The largest differences in consumption can be found for the retirement age. At this age, the tailor-made investment strategy can generate a higher investment level on average, as well as deliver higher minimum consumption levels.


We show the four sectors that benefit the most from the tailor-made investment strategy in Figure 10. From all four sectors shown in the figure, the largest divergence in average consumption level across states can be found for the retail sector, which is also the sector with the largest improvement in terms of certainty equivalent value. What is striking, however, in all four panels of the figure is that the main difference in the consumption distribution over time comes from the upper tail. That is, the tailor-made investment strategy generates a consumption distribution that is shifted upwards and has higher positive skewness.

The consumption plots for the sectors that lose the least from the planner investment strategy are shown in Figure 11. We see a similar pattern with respect to the average consumption pattern as in Figure 10. However, there is a larger overlap between the distributions for these four sectors. For the service sector, the consumption distribution is basically identical towards the end of the life cycle.


4.4 Wealth

The resulting wealth patterns of the consumption and investment decisions can be seen in Figures 12 and 13. We see that the wealth accumulation starts at a younger age for the solution using the investment strategy determined by the planner. This is a result of the lower consumption-wealth ratios and, therefore, higher savings rates at the beginning of the life cycle. We see that the tailor-made solution largely catches up with this wealth accumulation towards the middle of the life cycle and, in most cases, overtakes it.


We show that violin plots for the wealth patterns of the four sectors that benefit the most from a tailor-made investment strategy in Figure 12. We see that the wealth distribution for the planner and tailor-made strategies largely overlaps. The largest difference can be found for the retail sector, where the wealth level at the later stage of the life cycle is larger for the tailor-made strategy.




A: Construction



B: Manufacturing




C: Transportation




D: Retail

Figure 10: Consumption paths across sectors


This figure shows the violin plots for the consumption paths if the investment decision is individually optimal in blue and if the investment decision is made collectively in yellow. Panel “A” shows the paths for the construction sector, Panel “B” for the manufacturing sector, Panel “C” for the transportation sector, and Panel “D” for the retail sector. To improve clarity of the figure, we only show the violin for every three years of the life cycle.




A: Mining



B: Wholesale




C: Finance




D: Service

Figure 11: Consumption paths across sectors


This figure shows the violin plots for the consumption paths if the investment decision is individually optimal in blue and if the investment decision is made collectively in yellow. Panel “A” shows the paths for the mining sector, Panel “B” for the wholesale sector, Panel “C” for the finance sector, and Panel “D” for the service sector. To improve clarity of the figure, we only show the violin for every three years of the life cycle.




A: Construction



B: Manufacturing




C: Transportation




D: Retail

Figure 12: Wealth paths across sectors


This figure shows the violin plots for the wealth paths if the investment decision is individually optimal in blue and if the investment decision is made collectively in yellow. Panel “A” shows the paths for the construction sector, Panel “B” for the manufacturing sector, Panel “C” for the transportation sector, and Panel “D” for the retail sector. To improve clarity of the figure, we only show the violin for every three years of the life cycle.




A: Mining



B: Wholesale



C: Finance



D: Service

Figure 13: Wealth paths across sectors

This figure shows the violin plots for the wealth paths if the investment decision is individually optimal in blue and if the investment decision is made collectively in yellow. Panel “A” shows the paths for the mining sector, Panel “B” for the wholesale sector, Panel “C” for the finance sector, and Panel “D” for the service sector. To improve clarity of the figure, we only show the violin for every three years of the life cycle.

5. CONCLUSION

This article studies the optimal consumption and investment decisions for investors that are heterogeneous with respect to their income risk. We analyze two sets of results. First, we solve the life cycle problem for a situation in which the investors outsource the investment decision to a central planner, e.g., a pension fund. The asset menu in this case consists of a risk-free investment and a diversified stock market portfolio. This portfolio choice is equal across all investors and the investor decides on the optimal individual consumption path. Second, we let all investors choose their individual optimal investment and consumption plan. In this case, we let the investor choose to invest into the risk-free asset, a diversified stock market portfolio that excludes the companies from the sector that they work in, as well as the sector portfolio comprising of companies out of this sector.

We study these problems under the following assumptions. First, we do not assume a specific parametric model for the dynamics of the economy. Instead, we bootstrap economic paths from economic data that comprises of financial market returns, income returns, and state variables like dividend-price ratios. Second, the investors are homogeneous with respect to their utility function. In particular, they have CRRA utility with a risk aversion parameter of 10 and a subjective discount factor of 0.96. Third, we assume that the planner follows a utilitarian objective function and maximizes the weighted sum of investor utilities with the weights being their respective employment share. Our solution method is to train a neural network to find the optimal policy function.

We find that the improvements from the planner solution to the tailor-made strategies can be as large as 3% for the retail sector and as low as 0.5% for the service industry. The investment policies can differ significantly across sectors, depending on the correlations between income returns and capital market returns. We also find that the planner solution takes less risk than the individual optimal choice, especially after retirement. In terms of consumption streams, we find largely overlapping distributions over time, especially at the beginning of the life cycle. This is due to the low savings rate at early age and we see larger differences for the time in retirement.

References

- Anarkulova, A., Cederburg, S., and O’Doherty, M. S. (2023), “Beyond the status quo: a critical assessment of lifecycle investment advice,” .
- Balter, A. G. and Schweizer, N. (2024), “Robust decisions for heterogeneous agents via certainty equivalents,” *European Journal of Operational Research*, 317, 171–184.
- Brandt, M. W., Goyal, A., Santa-Clara, P., and Stroud, J. R. (2005), “A simulation approach to dynamic portfolio choice with an application to learning about return predictability,” *The Review of Financial Studies*, 18, 831–873.
- Branger, N., Chen, A., Mahayni, A., and Nguyen, T. (2023), “Optimal collective investment: an analysis of individual welfare,” *Mathematics and Financial Economics*, 17, 101–125.
- Catherine, S. (2022), “Countercyclical labor income risk and portfolio choices over the life cycle,” *The Review of Financial Studies*, 35, 4016–4054.
- Catherine, S., Sodini, P., and Zhang, Y. (2024), “Countercyclical income risk and portfolio choices: Evidence from sweden,” *The Journal of Finance*, 79, 1755–1788.
- Chacko, G., Desai, M., Golts, M., and Novakovsky, V. (2005), “A Forward-Solving Numerical Technique for Dynamic Consumption and Portfolio Allocation Problems,” SSRN Working Paper 687622.
- Cocco, J., Gomes, F., and Maenhout, P. (2005), “Consumption and Portfolio Choice over the Life Cycle,” *Review of Financial Studies*, 18, 491–533.
- Diris, B., Palm, F., and Schotman, P. (2015), “Long-Term Strategic Asset Allocation: An Out-of-Sample Evaluation,” *Management Science*, 61, 2185–2202.
- Druehdahl, J. and Røpke, J. (2025), “Deep Learning Algorithms for Solving Finite-Horizon Models,” Tech. rep., University of Copenhagen.
- Duarte, V., Fonseca, J., Goodman, A. S., and Parker, J. A. (2021), “Simple allocation rules and optimal portfolio choice over the lifecycle,” Tech. rep., National Bureau of Economic Research.
- Eiling, E. (2013), “Industry-Specific Human Capital, Idiosyncratic Risk, and the Cross-Section of Expected Stock Returns,” *The Journal of Finance*, 68, 43–84.
- Fama, E. F. and French, K. R. (1992), “The Cross-Section of Expected Stock Returns,” *Journal of Finance*, 47, 427–465.
- Gilda, S., Heidrich, B., and Kiraly, F. (2024), “tsbootstrap: Enhancing Time Series Analysis with Advanced Bootstrapping Techniques,” .
- Hornik, K., Stinchcombe, M., and White, H. (1989), “Multilayer feedforward networks are universal approximators,” *Neural networks*, 2, 359–366.

- Jagannathan, R. and Wang, Z. (1996), “The Conditional CAPM and the Cross Section of Expected Returns,” *Journal of Finance*, 51, 3–53.
- Jensen, B. A. and Nielsen, J. A. (2016), “How suboptimal are linear sharing rules?” *Annals of Finance*, 12, 221–243.
- Koijen, R. S. J., Nijman, T. E., and Werker, B. J. M. (2009), “When Can Life Cycle Investors Benefit from Time-Varying Bond Risk Premia?” *The Review of Financial Studies*, 23, 741–780.
- Leshno, M., Lin, V. Y., Pinkus, A., and Schocken, S. (1993), “Multilayer feedforward networks with a nonpolynomial activation function can approximate any function,” *Neural networks*, 6, 861–867.
- Schmidt, L. D. (2025), “Climbing and falling off the ladder: Asset pricing implications of labor market event risk,” *Journal of Financial Economics*, 172, 104131.