History Dependence of Pension Systems

Sahber Ahmadi-Renani*
Tehran Institute for Advanced Studies,
Khatam University

First Version: January 23, 2023
This Version: November 18, 2025

Please click here for the most recent version.

Abstract

In defined-benefit pension systems such as U.S. Social Security, retirement benefits depend on a history-dependent transformation of past earnings: the system condenses a worker's entire earnings trajectory into the Average Indexed Monthly Earnings (AIME), typically by averaging the top 35 years of earnings. This earnings-summarization rule embeds both redistribution and behavioral incentives, yet its quantitative effects are difficult to characterize because AIME is a highly nonlinear, non-smooth function of lifetime earnings histories. To address this challenge, I develop a structural life-cycle model of labor supply, retirement, and earnings risk, solved using a novel deep neural network architecture that directly learns the mapping from full earnings histories to retirement benefits. Using the estimated model, I evaluate alternative benefit rules that vary the number of counted earnings years—top 20, top 5, and lifetime-average earnings. Reducing the number of counted years modestly raises consumption and redistributes gains toward workers with volatile or intermittent earnings, while lifetime averaging reverses these patterns, sharply increasing labor supply and retirement ages and amplifying inequality by exposing benefits to low-earning years. These results show that the structure of the earnings-summarization formula plays a central role in shaping redistribution, insurance, and work incentives in public pension systems.

*This paper is a substantial extension of the second chapter of my Ph.D. dissertation at the University of Wisconsin–Madison. I am deeply grateful to my advisors—Matthew Wiswall, Naoki Aizawa, and Corina Mommaerts—for their guidance and support. I thank Zahra Delbari for outstanding research assistance. I also thank Jeffrey Smith, Christopher Taber, Jesse Gregory, and Yadollah Yaghoobzadeh for many valuable comments and suggestions. All the errors are my own.

1 Introduction

The modest numbers on the bottom right of social security checks are the main sources of income for many American retirees and the last barrier for keeping them out of poverty. In addition, the incentive system that they build shape the leading public policy in influencing the retirement age of workers, as well as their labor supply in old age. A retiree's public pension benefit is based on the history of her earnings before retirement. In the US, like most countries, two functions working in the composite determine the amount of pension benefits. One function summarizes all of the history of earnings into one outcome, and the other finds the amount of benefit based on that outcome. In the US, with some additional details, the history of earnings will be summarized by taking the average of the top 35 years of earnings of a worker. Then, this result will yield the retirement benefit after passing through a progressive benefit function. Most countries, however, summarize lifetime earnings very differently. Many systems base benefits on the last years of earnings (e.g., France, Greece, Portugal, Spain), while an increasing number of OECD and developing countries now rely on full lifetime averaging, in which every year of earnings enters the pension calculation.

The history-dependent part of pension systems is vastly understudied. Although many papers have researched design of the benefit function, design of the summarizing function has been mostly ignored[1]. Studying the ramifications of the design of the history-dependent part of the pension system and exploring the paths to improve it is the main focus of this paper.

It is crucial that this design be guided by how it affects the redistribution, insurance, and incentives system that it creates for workers. Different countries institute different sets of rules to calculate the pension benefits from the history of earnings. This creates a different balance of these often conflicting criteria in each country:

Redistribution. The earnings–summarization rule of a pension system shapes how resources are redistributed across workers with different life-cycle earnings profiles. When benefits are based on a shorter set of top earnings years, the system implicitly favors individuals whose histories contain many low-earning or non-working years but at least a few reasonably strong peak years. This includes workers with lower attachment to the labor force and many low- or zero-earning spells, as well as some low-volatility workers whose earnings are generally modest but feature a handful of particularly good years. In contrast, workers with long, stable, and consistently high earnings—often college-educated and strongly attached to the labor force—gain relatively little from further shortening the window At the opposite extreme, lifetime averaging reverses this pattern. By incorporating every year of earnings into the benefit formula, it heavily penalizes individuals with prolonged low-earning periods or intermittent participation—such as low-attachment and many low-volatility workers—while relatively favoring those with smoother or more strongly upward-sloping careers. Nevertheless, this tighter lifetime-based linkage may also be normatively appealing to some, since benefits more closely reflect individuals' total earnings contributions and are therefore perceived as more closely aligned with what workers are "entitled'' to receive.

Insurance. The pension formula also embeds an implicit insurance mechanism against idiosyncratic earnings risk. Rules that down-weight or exclude low-earning years insure workers against transitory downturns,

[1] It is noteworthy that this concern was raised by Peter Diamond in his Presidential Address on social security delivered during American Economic Association meeting in 2004 ((Diamond (2004)): "Depending on the nature of the underlying stochastic process of wage rates, both under weighting early years (relative to the use of interest rates) and not counting some low years may or may not help with insuring lifetime earnings. This is not an area that has received much research attention."

early-career learning periods, unemployment spells, and family-related interruptions: a sequence of bad years leaves retirement benefits largely determined by the better parts of the earnings history. By contrast, lifetime averaging eliminates much of this insurance: all fluctuations in earnings, whether arising from shocks, unemployment, or human-capital investment phases, feed directly into the pension base. While lifetime averaging provides some smoothing by distributing shocks over a larger number of years, it still transmits every low-earning spell into the pension base and therefore offers weaker insurance than any top-years scheme.

Incentives. In old age, decisions about how many hours to work and how long to remain in the labor force depend critically on how additional working years affect pension benefits. When the benefit formula places substantial weight on a limited set of top earnings years, the incentive to work late in life depends on whether those late-career earnings are themselves among the individual's highest. If a worker has already "filled'' the relevant top-years window with earlier peak earnings, then additional years—especially if earnings decline with age—contribute little to the AIME, weakening incentives to keep working. By contrast, if late-career earnings are near or above past peak levels, the top-years system rewards extended work because those years can replace lower past observations in the benefit calculation. The smaller the window, the more pronounced this pattern becomes. Lifetime averaging generates a very different incentive structure. Because every year of earnings enters symmetrically into the pension base, any additional year with earnings above one's current lifetime average increases future benefits. Workers cannot rely on a few strong peak years to determine their pensions; instead, improving the average requires sustained labor-force participation.

To study these trade-offs quantitatively, I develop a structural life-cycle model in which the full history dependence of the Social Security formula directly shapes workers' savings, labor supply, and retirement decisions. Individuals face stochastic earnings and endogenously choose hours, consumption, savings, and the timing of retirement. Pension benefits are computed through the nonlinear mapping of the entire earnings history into the AIME-based top-years average, allowing the model to capture the precise way in which alternative summarization rules alter incentives over the life cycle. The model is calibrated to match key features of U.S. data, including retirement hazards and the evolution of hours worked over the life cycle.

I then introduce a general deep learning framework for solving dynamic economic models characterized by history dependence, discrete choices, and life-cycle transitions. Unlike traditional numerical dynamic programming methods that approximate policy functions on fixed grids, our approach parameterizes all policy functions with deep neural networks trained end-to-end using automatic differentiation. The architecture introduces three methodological innovations. First, it handles *history dependence* through attention-based sequence summarization: a transformer-style encoder compresses entire earnings histories into compact differentiable state representations, allowing the model to learn which past events matter for future decisions. Second, it models *life-cycle dynamics* through a multitask highway network that shares parameters across ages but allows for residual age-specific flexibility, combined with an adaptive early-retirement (ER) gate that captures sharp behavioral regime shifts near retirement. Third, it resolves the *non-differentiability of discrete choices* (retirement, labor participation) using a dynamic Gumbel–Softmax relaxation, ensuring smooth gradients during training while preserving exact discreteness in final simulations. Together, these components produce a neural dynamic programming framework capable of learning complex policy rules from high-dimensional, non-stationary environments with realistic institutional discontinuities.

While developed to study the history dependence of pension systems, the framework provides a general blueprint for using deep learning to solve high-dimensional dynamic economic models. By integrating attention mechanisms, multitask highway architectures, and differentiable discrete-choice approximations within

a unified end-to-end training pipeline, the method overcomes the three major bottlenecks of structural modeling—curse of dimensionality, non-smooth decision boundaries, and time heterogeneity—without relying on handcrafted functional forms or coarse grid approximations. Beyond retirement and social insurance applications, the same modeling strategy can be directly adapted to heterogeneous-agent macroeconomic models, search and matching models, human capital accumulation problems, or dynamic models of firm entry and exit where past histories, regime changes, and discrete actions jointly determine behavior. The framework thus bridges modern deep learning techniques and structural economic modeling, offering a scalable and flexible computational tool for dynamic analysis.

The counterfactual analysis reveals that the way a pension system summarizes lifetime earnings plays a central role in shaping redistribution, insurance, and labor-supply incentives. Shortening the averaging window—from the current top-35 rule to top-20 or top-5 years—raises benefits and compresses inequality by discarding many low-earning years, with the largest gains accruing to workers with intermittent careers, low wage volatility, or weaker labor-force attachment. Those with stable, higher, or more volatile earnings histories gain little, since they have fewer low years to exclude and limited dispersion among their peak years. At the opposite extreme, shifting to lifetime averaging reverses these patterns entirely: by incorporating every low-earning spell, it reduces benefits for workers with prolonged or repeated periods of low wages while favoring individuals with smooth, continuously high earnings. These mechanical changes lead to starkly different behavioral responses: lifetime averaging generates strong incentives to work longer and delay retirement, prompting large reallocations of labor supply and consumption over the life cycle.

While the analysis provide a benchmark for life-cycle behavior under history-dependent pension systems, several important extensions remain outside the scope of this paper. First, the analysis does not explicitly model women's labor-force attachment, despite its central importance for understanding the aggregate and distributional consequences of pension reforms. A significant share of women exit the labor market during child-rearing years, and the interaction between such employment interruptions and history-dependent benefit formulas may generate gender-specific distortions. Second, the model abstracts from another dimension of history dependence in real-world Social Security systems: eligibility rules that require a minimum number of years with sufficient earnings to qualify for benefits. These rules link access to Social Security directly to individuals' work histories and may have meaningful labor-supply and distributional implications, especially for workers with intermittent or low-earning careers. Incorporating these eligibility margins—and the behavioral responses they induce—is an important avenue for future research.

1.1 Related Literature

This paper mainly contributes to the literature on the design of pension systems (Diamond and Mirrlees (1978), Diamond and Mirrlees (1986) Huggett and Parra (2010), Golosov, Shourideh, Troshkin, and Tsyvinski (2013), Shourideh and Troshkin (2017), Grochulski and Kocherlakota (2010) and Ndiaye (2020)). None of these papers study the design of the history-dependence part of the pension systems. This paper also relates to the literature on labor supply and retirement literature decisions of workers (Fan, Seshadri, and Taber (2022), and Borella, Nardi, and Yang (2023), O'Dea (2017), French (2005)). None of these papers consider the history-dependence of the pension systems or the optimal design of pension systems. I contribute to this literature by incorporating the full history-dependent structure of the Social Security benefit formula into a structural life-cycle model—an element entirely absent from existing work. Current quantitative retirement

and pension models typically approximate AIME as a simple average of lifetime earnings or as a smooth function of permanent income, effectively abstracting from the discrete, nonlinear way in which real systems count and discard earnings years. As a result, they miss the key redistribution and incentive forces created by the actual earnings-summarization rule—forces that depend on the timing, volatility, and intermittency of work. My framework embeds the true AIME mechanism and quantifies how alternative benefit formulas differentially reward or penalize volatile, low-attachment, or steadily employed workers, generating new insights into the behavioral and welfare consequences of pension-system design.

History-Dependent public transfer systems have been studied in the optimal taxation literature. Batzer (2021) studies optimal tax schedule when the tax function depends on all the history of income in a general equilibrium OLG model similar to Heathcote, Storesletten, and Violante (2017). Kapička (2022) studies the history dependence of taxation in a life cycle model and finds that a simple tax function that depends on the few past years can increase welfare. In the literature on income averaging and taxation on lifetime earnings, work of Vickrey (1947) is notwithstanding. Vickrey (1947) was concerned with the impact of progressive annual taxes on those with fluctuating incomes relative to those with constant incomes. However, using a longer period for determining taxes is likely to reduce the built-in stabilization from the income tax and lessen the easing of borrowing constraints. In Diamond (2009) absence of age and history-related tax rules counts as one of the reasons for a mandatory retirement pension System to redistribute based on lifetime between fluctuating and constant income individuals. I contribute to this literature by providing the first quantitative analysis of the history-dependent component of public pension systems. By embedding this non-smooth, highly nonlinear rule into a structural life-cycle model, I show how the earnings-summarization mechanism redistributes resources across workers with different earnings volatility and attachment, and how it shapes labor-supply and retirement incentives through the marginal value of additional working years.

Recent advances in computational economics have increasingly leveraged deep learning and machine learning techniques to solve high-dimensional structural models that were previously intractable using traditional numerical methods. The approach gained prominence with the unified deep learning framework of Maliar, Maliar, and Winant (2021), which demonstrated that neural approximations can efficiently solve dynamic stochastic general equilibrium (DSGE) and heterogeneous-agent models. Building on this, Fernández-Villaverde, Nuño, and Perla (2024) and Ebrahimi Kahou, Fernández-Villaverde, Perla, and Sood (2021) show that deep neural networks can tame the curse of dimensionality in dynamic programming and equilibrium computation. Further methodological innovations include bias-corrected Monte Carlo operators (Pascal, 2024) and symmetry-based deep architectures (Ebrahimi Kahou et al., 2021). Applications to macro-finance and policy problems have been developed by Valaitis and Villa (2024), Lepetyuk, Maliar, and Maliar (2020), and Beck, Garcia Sanchez, Moura, Pascal, and Pierrard (2024), while Kase, Melosi, and Rottner (2025) introduce neural estimation for nonlinear HANK models. In parallel, reinforcement learning approaches—ranging from deep multi-agent systems (Curry, Trott, Phade, Bai, and Zheng, 2022) to adaptive bounded-rational agents in DSGE environments (Chen, Cont, Joseph, Kumhof, Pan, Xiong, and Zhou, 2025)—have emerged as powerful tools for global solution and estimation of equilibrium behavior (Atashbar and Shi, 2022; Chen, 2025; Evans and Ganesh, 2024). Extensions to search and matching frameworks (Payne, Rebei, and Yang, 2024) and life-cycle or overlapping-generation models (Azinović and Zemlička, 2023; Kwon, 2025; Sun, 2023) further demonstrate the versatility of deep learning for structural economic analysis. Together, these contributions suggest that machine learning, particularly deep reinforcement learning and neural approximation, is reshaping the computational frontier of quantitative economics.

I contribute to this emerging literature by introducing a new deep learning framework for solving dynamic economic models with rich state dependence (i.e. history dependence) and discrete choices. Unlike existing approaches that typically approximate value or policy functions directly through static feedforward networks, my method integrates three key innovations: (1) Attention mechanisms to summarize the full earnings history; (2) Smooth treatment of discrete retirement choices using a dynamically annealed Gumbel–Softmax relaxation =. Together, these elements provide a unified and computationally efficient solution method for dynamic structural models featuring long memory, discrete state transitions, and heterogeneous life-cycle behavior; (3) Life-cycle network architecture that combines highway-style residual connections with regime change adapter gates. Together, these elements provide a unified and computationally efficient solution method for dynamic structural models featuring long memory, discrete state transitions, and heterogeneous life-cycle behavior. Beyond pension systems, this framework can be readily applied to a wide class of dynamic economic environments—such as education, macro, and finance—where agents' past histories and combination of discrete and continuous choices shape current behavior.

The remainder of the paper is organized as follows. Section 2 reviews the history dependence of pension systems in the United States and across other countries, motivating the relevance of incorporating lifetime earnings histories into dynamic models. Section 3 presents the structural life-cycle model of consumption, labor supply, and retirement decisions under history-dependent pensions. Section 4 introduces the deep learning framework developed to solve the model efficiently, including its treatment of history dependence, discrete choices, and life-cycle transitions. Section 5 describes the empirical strategy and calibration procedure. Section 6 reports the main quantitative results under the status quo pension policy. Section 7 evaluates the effects of alternative pension policy reforms. Finally, Section 8 concludes.

2 History Dependence of Pension Systems Around the World

In this section, I explain in more detail the history-dependence part of the US pension system and compare it to defined-benefit pension plans in other countries around the world.

2.1 United Satates

In the US, the two functions that determine an individual's Social Security benefits are Average Indexed Monthly Earnings (AIME) and the Primary Insurance Amount (PIA). The top 35 years of earnings are used to calculate the AIME. At first, each year's earnings are indexed to reflect the growth in the economy and wage levels during workers' employment years. After the top 35 years of indexed earnings are determined, the indexed earnings will be summed and divided by the total number of months. Then, the average amount will be rounded down to the next lower dollar amount. The result is called AIME.

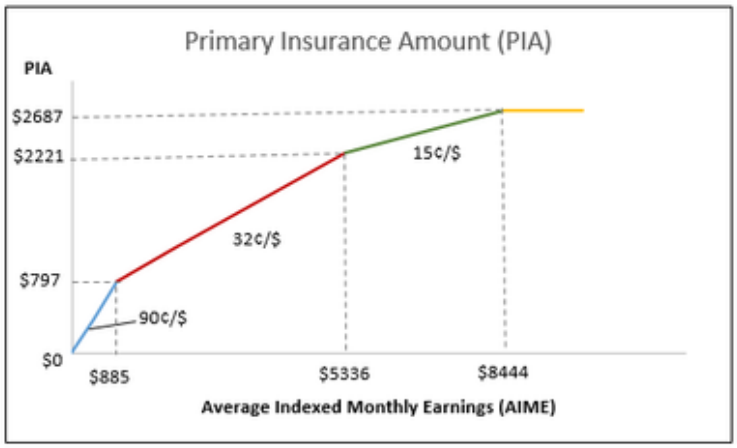
$$\underbrace{\text{AIME}}_{\substack{\text{Average Indexed}\\ \text{Monthly Earnings}}} = f(y_1, \cdots, y_T) = \frac{1}{12} \times \frac{1}{35} \times \sum_{\substack{\text{Top 35}\\ y_i \in \{\text{years of}\}\\ \text{earnings}}} (y_i). \tag{1}$$

The second function takes AIME as given and determine the amount of monthly benefit (PIA).

$$\underbrace{\text{PIA}}_{\substack{\text{Primary Insurance}\\ \text{Amount}}} = g(AIME). \tag{2}$$

The current function g in the US is presented in 1.

Figure 1: Function g in the US: Determining PIA from AIME



2.2 Other Countries

There are three types of $f(.)$ that are most common across countries.

1. **Average Earnings of the Last L Years:** Although various OECD countries have moved from this method, this method is still employed among countries like France, Greece, Portugal, Spain, Norway, and Sweden. There are also numerous countries among developing countries that utilize this method. Thailand, Brazil, and Iran are some examples of developing counties.

2. **Average Life Time Earning, Excluding Some of the Lowest Earnings:** This method is used in the US in Canada. In the US, the highest 35 years determine the pension and in Canada, the lowest 15 % of earnings are excluded from the history.

3. **Average Life Time Earning:** In this method, the average of all of the earnings over the lifetime is used to calculate the pension benefit. Currently, this method is the most common method in OECD countries. Over the past decades, most OECD countries have moved from the second method (counting the last 10-25 years) to this method. Among developing countries, China, Indonesia, and Vietnam utilize all of their lifetime earnings.

Although the focus of this paper is public pension systems, it is also noteworthy that the first method is the most common method among private defined-benefit pension systems (like teacher unions).

3 Model

3.1 Setup

The economy consists of a continuum of males who enter the model at age 22 ($t = 1$) and live deterministically until age 82 ($T_D = 82$). Time is discrete and one period corresponds to one year. Individuals are indexed by i and characterized by education type $edu_i \in \{0, 1\}$, where $edu_i = 1$ denotes college graduates. Each individual faces idiosyncratic productivity shocks that evolve stochastically throughout the life cycle and influence labor earnings.

At each age t, the individual chooses consumption c_t, next-period assets a_{t+1}, and, while working, hours of labor supply h_t. The working phase begins at age 22 and continues until the individual either retires or reaches the late-retirement cutoff. From the early-retirement age ($T_{ER} = 62$) to the late-retirement age ($T_{LR} = 70$), individuals endogenously choose whether to retire. Upon retirement, they permanently exit the labor force and start receiving a Social Security–style pension benefit that depends on their lifetime earnings history and the chosen retirement age. If they continue working beyond the full-retirement benchmark ($T_{FR} = 67$), benefits increase according to delayed-retirement adjustments in the pension formula.

After retirement, individuals consume out of accumulated assets and pension benefits, earning a fixed return r on their savings. At death ($T_D = 82$), any remaining wealth is left as a bequest, which provides utility to the individual's estate according to a warm-glow bequest motive following De Nardi (2004). This feature allows the model to capture late-life saving behavior and the distribution of assets among the elderly.

3.2 Preferences

At each age t, instantaneous utility is separable in consumption and labor, with a fixed cost of participation:

$$u_t(c_t, h_t) = \frac{c_t^{1-\gamma}}{1-\gamma} - \psi \frac{(h_t/\bar{h})^{1+\frac{1}{\eta}}}{1+\frac{1}{\eta}} - \phi\, \mathbb{I}\{h_t > 0\}, \tag{3}$$

where γ is the coefficient of relative risk aversion, η the Frisch elasticity of labor supply, ψ the disutility weight on the intensive margin, and ϕ the fixed cost of work participation. The lifetime objective discounts future utility by β^t and includes a warm-glow bequest motive following standard formulations:

$$U = \sum_{t=1}^{T_L} \beta^{t-1} u_t(c_t, h_t) + \theta_b \frac{(a_{T_L} + q_b)^{1-\gamma}}{1-\gamma}, \tag{4}$$

where $\theta_b > 0$ scales the strength of the bequest motive and $q_b > 0$ captures the non-homotheticity of the utility from bequests.

3.3 Budget Constraints

During the working phase ($t < t_R$), consumption equals disposable income plus assets carried from the previous period:

$$c_t = y_t - T^{inc}(y_t) - T^{SS}(y_t) + a_t - \frac{1}{1+r} a_{t+1}, \qquad y_t = w_t h_t, \quad a_t \geq 0. \tag{5}$$

Labor earnings w_t evolve according to both deterministic age-education profiles and persistent idiosyncratic shocks:

$$\log w_t = \mu(edu, t) + \theta_t, \tag{6}$$

$$\theta_{t+1} = \theta_t + \epsilon_{t+1}, \qquad \epsilon_t \sim N(0, \sigma_\epsilon^2), \tag{7}$$

where $\mu(edu, t)$ is a quadratic function of age and education capturing deterministic wage growth, and θ_t represents an individual's stochastic productivity component.

After retirement ($t \geq t_R$), individuals no longer work and receive a pension benefit b_t determined by the Social Security rules:

$$c_t = b(t_R, \{y_s\}_{s=1}^{t_R}) + a_t - \frac{1}{1+r} a_{t+1}. \tag{8}$$

Benefits are computed as

$$b(t_R, \{y_s\}) = \delta(t_R)\, \mathcal{B}(S_t)\,, \qquad S_t = \frac{1}{T_S} \sum_{s \in \mathrm{Top}(T_S)} y_s, \tag{9}$$

where $\mathcal{B}(\cdot)$ is the piecewise-linear benefit function over earnings bend points, $\delta(t_R)$ the actuarial adjustment for retirement age, and S_t the average of the top T_S earnings years (default $T_S = 35$). This structure captures the strong dependence of pensions on past income histories.

3.4 Individual Problem

Prior to early retirement ($t < T_{ER}$), the individual remains in the working state and solves

$$V_t^W(\{y_s\}_{s=1}^{t-1}, a_t, \theta_t, edu) = \max_{c_t, h_t} \Big\{ u_t(c_t, h_t) + \beta\, \mathbb{E}_t[V_{t+1}^W(\{y_s\}_{s=1}^{t}, a_{t+1}, \theta_{t+1}, edu)] \Big\}, \tag{10}$$

$$a_{t+1} = (1+r)\big[a_t + w_t h_t - T^{inc}(y_t) - T^{SS}(y_t) - c_t\big]\,. \tag{11}$$

Between the early and full retirement ages ($T_{ER} \leq t < T_{FR}$), individuals choose whether to retire or continue working:

$$V_t^E(\cdot) = \max\{\, V_t^W(\cdot),\, V_t^R(\cdot)\,\}, \tag{12}$$

where the retirement hazard $\Pr(R_t = 1)$ is endogenously determined within the model.

After retirement ($t \geq t_R$), the problem simplifies to consumption–saving under fixed pension benefits:

$$V_t^R(AIME_{t_R}, a_t, t_R) = \max_{c_t} \Big\{ u_t(c_t, 0) + \beta\, V_{t+1}^R(AIME_{t_R}, a_{t+1}, t_R) \Big\}, \tag{13}$$

$$a_{t+1} = (1+r)\big[a_t + b_t - c_t\big], \tag{14}$$

with terminal utility $V_{T_L+1}^R(a_{T_L}) = \theta_b((a_{T_L} + q_b)^{1-\gamma})/(1-\gamma)$.

4 Deep Learning Framework

A central challenge in modeling the history-dependent effects of pension systems is that the entire earnings history enters the state space of the life-cycle model. This creates both a high-dimensional and expanding

state space, making the problem practically unsolvable using traditional numerical methods in economics—such as grid-based approximations of policy functions. To overcome this challenge, I approximate the policy functions using deep neural networks (DNNs). Deep neural networks have recently been applied to solve a wide range of structural and quantitative dynamic economic models ((Azinovic, Gaegauf, and Scheidegger, 2022; Fernández-Villaverde et al., 2024; Maliar et al., 2021)). However, a particular difficulty in this project arises from the expanding nature of the state space due to history dependence in pension benefits, which renders most existing approaches unsuitable.

To address this, I exploit the similarity between history-dependent state spaces and sequence models in machine learning. Specifically, I employ *attention* mechanisms—originally introduced in Transformer architectures (Vaswani, Shazeer, Parmar, Uszkoreit, Jones, Gomez, Kaiser, and Polosukhin, 2017)—to summarize the entire earnings history into a fixed-length latent representation. The use of attention layers has also proven effective in time-series econometrics and forecasting (Lim and Zohren, 2021).

Another key challenge is the coexistence of both continuous (e.g., consumption and labor supply) and discrete (e.g., retirement) decisions. To maintain differentiability and exploit the efficiency of gradient-based optimizers such as Adam (Kingma and Ba, 2015a), I model all components within a continuous, differentiable framework. This is achieved through the *Gumbel–Softmax* reparameterization trick (Jang, Gu, and Poole, 2017; Maddison, Mnih, and Teh, 2017), widely used in deep reinforcement learning to approximate discrete choices while preserving gradient flow.

Finally, the life-cycle structure of the problem introduces both gradual behavioral evolution and sharp regime shifts—most notably at the retirement transition—which require architectural flexibility. To address this, I adopt a *highway-network design* (Srivastava, Greff, and Schmidhuber, 2015) combined with an *adopter gate* that activates near retirement ages. The highway structure enables the model to share parameters across adjacent life stages, reducing the number of parameters relative to a fully separate network for each age, while the residual connections capture smooth year-to-year variations in behavior. The adopter gate, in turn, allows for abrupt adjustments in policy functions when new decision regimes (such as retirement) become active, ensuring both parsimony and flexibility in representing life-cycle heterogeneity.

This section introduces the deep learning framework used to approximate policy functions in the life-cycle model. We first present the general DNN setup and notation, then discuss how the model handles history dependence, discrete retirement decisions, and life-cycle transitions. We then describe the overall neural architecture and training algorithm.

4.1 Model Setup

To solve the dynamic programming problem, I approximate the household policy functions using deep neural networks (DNNs) and express the model in its extensive form, where analytical expectations are replaced by Monte Carlo simulations over stochastic shocks. For each draw of exogenous shock paths, I recursively simulate endogenous state variables based on the exogenous processes and the approximated policy functions. The resulting simulated sample is used to evaluate the ex-ante welfare implied by the current network parameters. The loss function is defined as the negative of this expected welfare. In this formulation, the household's maximization problem is transformed into a loss-minimization problem: the objective becomes to estimate the neural network parameters that best approximate the optimal policy functions.

Extensive-form problem. The baseline life-cycle problem can be written as:

$$\max_{\{c_t\}_{t=1}^{T},\,\{R_t\}_{t=t_{ER}}^{t_R},\,\{h_t\}_{t=1}^{t_R}} \quad \mathbb{E}_{\epsilon_1,\ldots,\epsilon_T}\left[\sum_{t=1}^{T}\beta^{t-1}\left(\frac{c_t^{1-\gamma}}{1-\gamma}-\psi\frac{(h_t/\bar{h})^{1+\frac{1}{\eta}}}{1+\frac{1}{\eta}}-\phi\,\mathbb{1}\{h_t>0\}\right)\right] \tag{15}$$

$$\text{s.t.} \quad (t<t_R)\; c_t = y_t - \tau^{inc}(y_t) - \tau^{SS}(y_t) + a_t - \frac{1}{1+r}a_{t+1}, \quad y_t = w_t h_t,\; a_t \geq 0,\; a_1 = 0, \tag{16}$$

$$(t \geq t_R)\; c_t = b(t_R, \{y_s\}_{s=1}^{t_R}) + a_t - \frac{1}{1+r}a_{t+1}, \tag{17}$$

$$(t \leq t_R)\; \log w_t = \mu(edu, t) + \theta_t, \quad \theta_{t+1} = \theta_t + \epsilon_{t+1}, \quad \epsilon_t \sim \mathcal{N}(0, \sigma_\epsilon^2),\; \theta_0 = 0, \tag{18}$$

$$c_t > 0, \quad h_t \in (0, 2860), \tag{19}$$

$$R_t \in \{0,1\}, \quad t_R = \min\{\arg\min_t [R_t = 1],\, T_{LR}\}. \tag{20}$$

DNN-based approximation. Let M_t denote the state vector at time t and $\pi_t(M_t)$ the true policy function mapping states to decisions. I approximate $\pi_t(M_t)$ by a neural network $\pi_t(M_t; \Omega_t^{\lambda_t})$ parameterized by weights $\Omega_t^{\lambda_t}$ and architecture λ_t.

The optimal weights solve:

$$\Omega^* = \arg\max_{\Omega} V(\Omega) = \mathbb{E}_{\epsilon_{1:T},\, edu}\left[\sum_{t=1}^{T}\beta^{t-1}u\big(\pi_t(M_t;\Omega_t)\big)\right], \tag{21}$$

$$u\big(\pi_t(M_t;\Omega_t)\big) = \frac{[c_t(M_t;\Omega_t)]^{1-\gamma}}{1-\gamma} - \psi\frac{(h_t(M_t;\Omega_t)/\bar{h})^{1+\frac{1}{\eta}}}{1+\frac{1}{\eta}} - \phi\,\mathbb{1}\{h_t(M_t;\Omega_t)>0\}. \tag{22}$$

For $t < t_R$,

$$c_t(M_t;\Omega_t) = y_t - \tau^{inc}(y_t) - \tau^{SS}(y_t) + a_t - \frac{1}{1+r}a_{t+1}(M_t;\Omega_t), \tag{23}$$

$$y_t = w_t h_t(M_t;\Omega_t), \quad a_1 = 0, \tag{24}$$

$$M_t = \{\{y_l\}_{l=1}^{t-1}, a_t, \theta_t, edu\}, \tag{25}$$

$$\pi_t(M_t;\Omega_t) = \begin{cases} \{c_t(M_t;\Omega_t), h_t(M_t;\Omega_t)\}, & t < T_{ER}, \\ \{c_t(M_t;\Omega_t), h_t(M_t;\Omega_t), R_t \in \{0,1\}\}, & T_{ER} \leq t < t_R. \end{cases} \tag{26}$$

For $t \geq t_R$,

$$M_t = \{b_t, a_t\}, \qquad c_t = B(t_R, \{y_s\}_{s=1}^{t_R}) + a_t - \frac{1}{1+r}a_{t+1}, \tag{27}$$

$$B_{t_R} = \mathbb{b}(t_R, S), \quad S = \mathbb{s}(\{y_l\}_{l=1}^{t_R-1}), \quad B_t = B_{t-1} \;\text{ for } t > t_R. \tag{28}$$

Finally, wages evolve as in (18), and all feasible allocations satisfy $c_t > 0$, $h_t \in (0, 2860)$, $R_t \in \{0,1\}$ and $t_R = \min\{\arg\min_t[R_t = 1],\, T_{LR}\}$.

4.2 Modeling Life Cycle DNN

A distinctive feature of life-cycle models is that individual behavior evolves smoothly over time yet may undergo abrupt shifts at specific transitions, such as entry into retirement. Modeling this pattern poses a

trade-off between flexibility and parsimony. A separate neural network for each age or period would capture heterogeneity accurately but would entail an explosion in parameters and a loss of statistical efficiency. Conversely, a single network with fully shared parameters would impose excessive smoothness, failing to reflect age-specific differences in preferences and constraints. The proposed architecture resolves this tension through a combination of *highway-network design*, *task-specific multitask learning*, and a *gate adopter* mechanism. Together, these elements make the deep learning model well suited to dynamic settings characterized by both continuous aging processes and discrete regime transitions.

At each age t, the model maps the current-period state vector X_t into a shared representation via a compact *general layers*

$$s_t \;=\; G(X_t;\Theta_G), \tag{29}$$

implemented as a small stack that may combine Multi-Layer Perceptron (MLP) blocks and attention layers. From s_t, the network branches into two *task-specific* heads for labor hours and consumption share. We use standard hard-parameter-sharing multitask learning (Caruana, 1997), so $G(\cdot)$ is shared while each task has its own head; this improves data efficiency and inductive bias across related decisions (hours and consumption), consistent with recent multitask learning (MTL) theory and practice (Kendall, Gal, and Cipolla, 2018; Sener and Koltun, 2018; Standley, Zamir, Chen, Guibas, Malik, and Savarese, 2020).

Each task head uses a *highway* decomposition with simple addition: a shared *core* mapping and a year-specific *residual* mapping are summed and squashed,

$$\widehat{y}_t^{(q)} \;=\; \sigma\Big(f_{\text{core}}^{(q)}\big(s_t^{\text{ER},(q)};\Theta_{\text{core}}^{(q)}\big) \;+\; f_{\text{res},t}^{(q)}\big(s_t;\Theta_{\text{res},t}^{(q)}\big)\Big), \qquad q\in\{\text{hours worked},\text{consumption share}\}. \tag{30}$$

The *highway design*, inspired by Srivastava et al. (2015), decomposes each policy head into a shared *core component* and a year-specific *residual adjustment*. The core component captures persistent behavioral patterns common across the life cycle, while the residual layer models incremental deviations associated with aging or evolving state variables. This structure allows the model to represent year-to-year differences without fitting an entirely separate DNN for each period. By explicitly combining shared and residual pathways, the architecture achieves parameter efficiency and smooth behavioral transitions, reflecting the economic intuition that preferences and responses evolve gradually over the life cycle.

While the highway design captures smooth within-stage variation, it remains necessary to accommodate sharp regime changes—most notably the transition to retirement, where constraints and objectives shift discontinuously. To model this, a small task-specific *gate adopter* MLP is introduced, which activates only in the age window of retirement decision. The adopter modulates the latent representation before the core head:

$$s_t^{\text{R},(q)} \;=\; \begin{cases} A^{(q)}(s_t;\Theta_A^{(q)}), & t\in[t_{\text{ER}},\,t_{\text{LR}}],\\ s_t, & \text{otherwise}, \end{cases} \tag{31}$$

where $A^{(q)}(\cdot)$ is a lightweight feed-forward network. This mechanism enables the model to preserve continuity before the retirement decision window and to adapt its functional form thereafter, thus capturing both gradual and abrupt life-cycle changes within a unified framework. A separate retirement head reads directly from s_t during the retirement window.

Figure 2 provides a schematic overview of the entire life-cycle DNN. The network begins with a general representation layer $G(X_t)$ that encodes the contemporaneous state vector. From this layer, information

branches into two task-specific heads—one for labor hours and another for consumption share—representing a hard-parameter-sharing multitask architecture. Within each task, an *R adopter* (a small MLP) acts only in the retirement window to modify the input representation passed to the *core* path. The *residual* path always receives the unmodified state representation, allowing it to learn smooth age-specific adjustments. The outputs of the core and residual sub-networks are then summed and passed through a sigmoid activation to yield bounded predictions for each decision variable. A separate retirement head can use the shared summary to generate a retirement head during the retirement window (not shown in the graph). The design thus mirrors the economic environment: smooth within-life-cycle adaptation combined with targeted flexibility at key transition points.

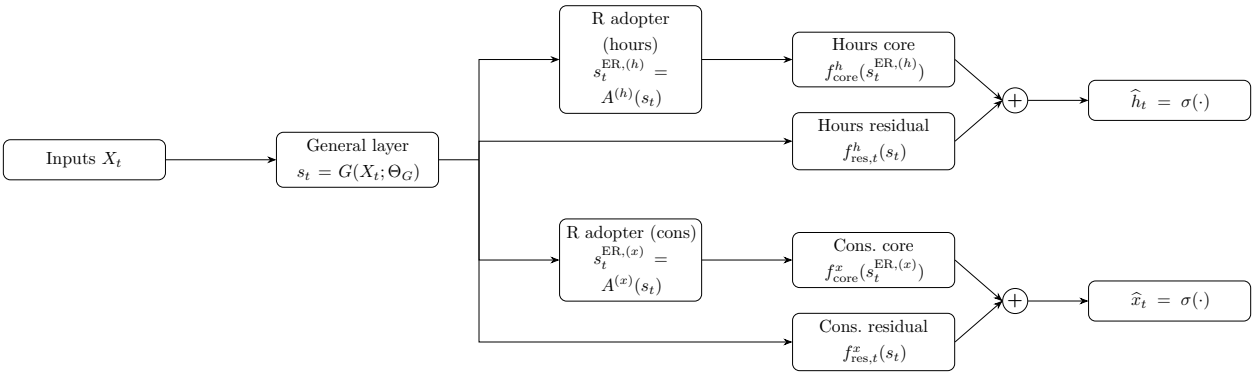


Figure 2: Core architecture at age t: Inputs and General layer are vertically aligned; a short connector then fans out to two task heads (hours and consumption). In each head, an ER adopter $A^{(q)}$ applies only to the core path within the ER window, while the residual path always receives s_t directly. Core and residual outputs merge at a centered "+" node and pass through a sigmoid.

4.3 Modeling History Dependence

One of the main challenges in dynamic life-cycle modeling of the rich features of the social security system is the treatment of history dependence. In addition to retirement, in many other economic settings individuals' current decisions depend not only on their current states but also on a long sequence of past realizations, such as lifetime earnings or employment histories. However, explicitly including the full history $\{y_1, \ldots, y_{t-1}\}$ as part of the state vector leads to an exponentially expanding state space, rendering standard numerical methods infeasible. To overcome this "curse of history", we approximate the information content of the earnings path using modern *attention-based sequence modeling* techniques drawn from deep learning.

In recent years, the *attention mechanism* (Vaswani et al., 2017) has become the dominant framework for modeling sequential data in natural language processing, time-series analysis, and reinforcement learning. The key idea is that, rather than summarizing a sequence through fixed-length recurrent states (as in recurrent neural networks), attention allows the model to assign flexible, data-driven weights to all past observations when forming a summary representation. Given an input sequence of vectors $\{v_1, \ldots, v_T\}$, the attention layer computes a weighted average:

$$\mathrm{Attn}(Q, K, V) = \mathrm{softmax}\left(\frac{QK^\top}{\sqrt{d_k}}\right) V, \tag{32}$$

where $Q, K, V \in \mathbb{R}^{T \times d_k}$ are "query," "key," and "value" projections of the sequence into a common latent space. The softmax weighting determines which past elements the model should "attend to" when summarizing information relevant to the current decision. The architecture can easily be extended to multiple parallel "heads"—known as *multi-head attention*—which allows the model to capture heterogeneous patterns of dependence (for example, earnings persistence versus recent shocks) at different scales.

In our implementation, each element of the earnings history is treated as a *token*. Each token encodes both the value of past earnings and several additional contextual features through a process known as *tokenization* and *embedding*. Embeddings are low-dimensional vector representations learned jointly with the model (Devlin, Chang, Lee, and Toutanova, 2018; Mikolov, Sutskever, Chen, Corrado, and Dean, 2013); they transform discrete or scalar inputs into a continuous representation that preserves relevant structure. For example, other state variables (i.e. education, productivity type, and asset) can be embedded into latent vectors of the same dimensionality as the income tokens, allowing all features to interact in the same representational space.

Here we explain our method in practice.

1. **Tokenization and embeddings.**

   Each element of the past earnings sequence is treated as a *token* that encodes not only the realized income level but also its temporal position and relative magnitude within the individual's earnings history. Formally, for each period $\tau \leq t-1$ we define the feature vector

   $$\phi_\tau = \big(y_\tau,\, p_\tau,\, r_\tau\big), \tag{33}$$

   where y_τ is earnings, p_τ is a positional index indicating the point in the life cycle, and r_τ is a *soft rank* of y_τ within $\{y_1, \ldots, y_{t-1}\}$.

   Including p_τ and r_τ alongside y_τ serves an economic as well as computational purpose. From an economic perspective, the life-cycle decision rules are highly nonlinear in earnings because institutional features such as social-security benefits depend on the ordering and position of past incomes (for example, whether a year belongs to the highest-earning subset used in benefit calculations). Providing the network with both the position and the rank allows it to learn these nonlinearities directly, rather than requiring the model to infer them implicitly. From a computational standpoint, the *soft rank* is used instead of a discrete ranking so that gradients can be propagated through it; this ensures that the entire history-encoding process remains differentiable.

   Each token is mapped to a d-dimensional embedding,

   $$v_\tau = \psi_y(\phi_\tau) \in \mathbb{R}^d, \tag{34}$$

   where ψ_y denotes a shared projection function (e.g., linear or shallow nonlinear mapping) with output dimension d.

   Static characteristics such as education, assets, and idiosyncratic productivity type are collected in a vector $s \in \mathbb{R}^m$. These variables are projected to the same representational space to allow joint interactions with income tokens:

   $$u = \psi_s(s) + E_{\text{type}}(\text{type}) \in \mathbb{R}^d, \tag{35}$$

where $\psi_s : \mathbb{R}^m \to \mathbb{R}^d$ is a learnable projection and E_{type} is a discrete *type embedding*. The type embedding introduces these structural differences as learnable offsets in the shared representation, enabling the model to condition its policy mapping on individual-level heterogeneity.

2. **Soft top-years weighting.**

   Pension systems often introduce sharp discontinuities: only a fixed number of highest-earning years (e.g., the top 35 in the U.S. Social Security system) contribute to benefit determination. To help the neural network learn these discontinuous yet economically structured nonlinearities, we introduce a smooth "soft top-years" weighting that biases attention toward higher income ranks:

$$\pi_\tau = \frac{\exp(r_\tau/\tau_T)}{\sum_{j=1}^{t-1}\exp(r_j/\tau_T)}, \qquad \tilde{v}_\tau = \pi_\tau\, v_\tau. \tag{36}$$

   When the temperature parameter τ_T is small, attention concentrates sharply on the highest-ranked incomes (approximating a discrete top-k rule). Larger τ_T values distribute attention more evenly across years. During training, τ_T follows an annealing schedule that gradually decreases from a high to a low value, allowing the network to first learn broad patterns and then focus progressively on the most relevant (top-income) observations. This dynamic temperature annealing follows the general principle introduced by Jang et al. (2017); Maddison et al. (2017), ensuring a differentiable approximation of discrete top-year selection.

3. **Concatenation and multi-head attention.**

   Let $V = [\tilde{v}_1, \ldots, \tilde{v}_{t-1}]^\top \in \mathbb{R}^{(t-1)\times d}$ denote the reweighted income tokens and $U \in \mathbb{R}^{n_s \times d}$ the static token matrix. We concatenate them into a single sequence,

$$Z = \begin{bmatrix} V \\ U \end{bmatrix} \in \mathbb{R}^{T' \times d}, \qquad T' = (t-1) + n_s. \tag{37}$$

   A multi-head self-attention (MHA) layer (Vaswani et al., 2017) computes

$$\mathrm{MHA}(Z) = \mathrm{Concat}\big(\mathrm{head}_1, \ldots, \mathrm{head}_H\big) W^O, \qquad \mathrm{head}_h = \mathrm{softmax}\Big(\frac{Q_h K_h^\top}{\sqrt{d_k}}\Big) V_h, \tag{38}$$

   where $Q_h = ZW_h^Q$, $K_h = ZW_h^K$, and $V_h = ZW_h^V$ are the query, key, and value projections for head h. Each attention head learns a distinct form of temporal dependence over the income path—such as long-term averages, volatility, or recent trends—while jointly conditioning on static heterogeneity through U. In economic terms, this mechanism learns an endogenous sufficient statistic summarizing the history relevant for current-period decisions.

4. **History summary**

   The MHA output is reduced to a fixed-width vector through a small feed-forward and pooling operator,

$$s_t^{\text{hist}} = \rho\big(\mathrm{MHA}(Z)\big) \in \mathbb{R}^{d_s}. \tag{39}$$

   This representation acts as a differentiable summary of the entire earnings history. It is then combined with contemporaneous states to form the input to the *general layer*:

$$X_t = \big[s_t^{\text{hist}},\, a_t,\, \theta_t,\, \mathit{edu},\, \mathit{type}\big], \qquad s_t = G(X_t; \Theta_G). \tag{40}$$

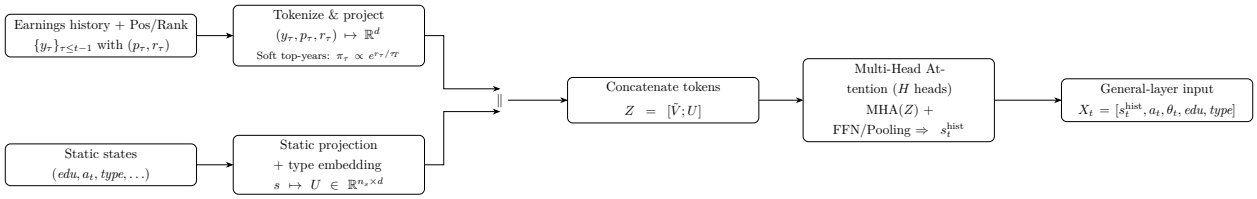


Figure 3: Compact history encoder. The earnings path with positional and rank features is tokenized and projected with a soft top-years weighting $\pi_\tau \propto \exp(r_\tau/\pi_T)$. Static states are mapped to the same space via a projection and type embedding. The sequences are concatenated into Z, processed by multi-head attention and a feed-forward pooling step to form the history summary s_t^{hist}, which feeds the general-layer input X_t.

> The general layer produces a compact shared representation s_t that captures both static and historical information and serves as a common input to all task-specific policy heads. The output s_t from the general layer forms the entry point for the multitask life-cycle DNN described in Section 4.2

Figure 3 expands the *General layer* box from Figure 2. It details how the model constructs the shared representation used by both task heads. Past earnings $\{y_\tau\}_{\tau \le t-1}$ are turned into tokens by stacking the raw value with a position index p_τ and a soft rank r_τ, then projecting into a d–dimensional embedding. A temperature–controlled soft "top-years" gate reweights these income tokens toward higher earnings years. Time-invariant states (education, assets, type, etc.) are projected and type-embedded to live in the same space as income tokens. All tokens are concatenated and passed through a multi-head attention block, whose heads learn complementary dependence patterns. A small feed-forward "mixing" layer collapses the attended sequence into a compact history summary s_t^{hist}, which then enters the general layer together with contemporaneous states. This design bypasses the expanding state space by learning which aspects of history are most informative for decision-making, while preserving economic structure through rank-aware (top-years) weighting, static heterogeneity embeddings, and differentiable sequence summarization.

4.4 Modeling Discrete Choices

Deep models are trained by backpropagation, which requires differentiable computation graphs. Discrete choices (e.g., retire vs. work) break differentiability: the policy and value functions become non-smooth, and gradients vanish almost everywhere. A classic fix is to inject smoothing via uncertainty and then rely on numerical differentiation; however, modern deep learning frameworks exploit *analytical* differentiation through the chain rule, making differentiable relaxations more attractive.

I adopt a continuous, temperature-controlled relaxation of the binary retirement choice based on the *Gumbel–Softmax* distribution (Jang et al., 2017; Maddison et al., 2017). This delivers unbiased forward passes that approximate the discrete decision and provides stable surrogate gradients for learning. The relaxation is only used during training; the final evaluation and counterfactual simulations use the *hard* discrete policy.

(1) Annealed Gumbel–Softmax for the retirement decision

Let $\ell_t \in \mathbb{R}$ be the *logit* produced by a small retirement head from the shared representation s_t (the general layer output): $\ell_t = f_{\text{ret}}(s_t)$. Define class probabilities $p_t = \sigma(\ell_t)$ for retire and $1 - p_t$ for work. The Concrete relaxation draws Gumbel noise $g_1, g_0 \overset{\text{i.i.d.}}{\sim} \text{Gumbel}(0,1)$ and forms the *soft* one-hot for retire:

$$\tilde{R}_t(\tau) \;=\; \frac{\exp\{(\log p_t + g_1)/\tau\}}{\exp\{(\log p_t + g_1)/\tau\} + \exp\{(\log(1-p_t) + g_0)/\tau\}} \;\in\; (0,1), \tag{41}$$

where $\tau > 0$ is a *temperature*. As $\tau \downarrow 0$, $\tilde{R}_t(\tau) \to \mathbb{1}\{g_1 - g_0 > \log\frac{1-p_t}{p_t}\}$ almost surely, i.e. it becomes a hard Bernoulli draw; as $\tau \uparrow \infty$, it approaches $1/2$.

To stabilize learning early and sharpen the decision later, we use the Dynamic temperature (annealing) method and *anneal* the temperature:

$$\tau_e \;=\; \max\Big\{\tau_{\min},\; \tau_0 \cdot \exp(-\kappa\, e)\Big\}, \qquad e = 0, 1, 2, \ldots \tag{42}$$

with initial temperature $\tau_0 > 0$, floor $\tau_{\min} > 0$, and decay rate $\kappa > 0$. The schedule follows the recommendations in Jang et al. (2017); Maddison et al. (2017) (see also discussion of temperature annealing therein): early epochs keep τ high to provide smooth gradients; later epochs reduce τ so the relaxed policy concentrates near $\{0, 1\}$ ("soft top-1"). We use $(\tau_0, \tau_{\min}, \kappa)$ in (42). Typical values are $\tau_0 \in [1, 5]$, $\tau_{\min} \in [0.1, 0.5]$, κ chosen so that τ_e reaches $\tau_{\min}$ over the first third to half of training.

If desired, at training time one may use the *straight-through* estimator: use the hard decision $R_t^{\text{hard}} = \mathbb{1}\{\tilde{R}_t(\tau) > 1/2\}$ in the forward pass but backpropagate gradients through $\tilde{R}_t(\tau)$ (the soft variable). This preserves discrete semantics with stable gradients; in our baseline we keep the fully soft forward pass and switch to hard choices only at evaluation. For the detailed formulation of the corresponding value functions, see Appendix A.

(2) Soft state transitions and value mixture

The relaxation lets us write retirement-period objects as convex combinations weighted by $\tilde{R}_t(\tau)$:

$$\begin{aligned}
\text{flow utility at } t: \quad u_t \;&=\; \tilde{R}_t(\tau)\, u_t^R(a_t^R, B_t) \;+\; \big(1 - \tilde{R}_t(\tau)\big)\, u_t^W(a_t^W, h_t^W), \\
\text{next states}: \quad a_{t+1}^R \;&=\; \tilde{R}_t(\tau)\, a_{t+1}^{R|R} \;+\; \big(1 - \tilde{R}_t(\tau)\big)\, a_{t+1}^{R|W}, \\
a_{t+1}^W \;&=\; a_{t+1}^{W|W}, \qquad B_{t+1} \;=\; \tilde{R}_t(\tau)\, B_t \;+\; \big(1 - \tilde{R}_t(\tau)\big)\big(t, \sim(y_{1:t-1})\big),
\end{aligned} \tag{43}$$

and for the cumulative probability of having retired by t:

$$P_{t+1}^R \;=\; P_t^R \;+\; \big(1 - P_t^R\big)\, \tilde{R}_t(\tau), \qquad P_{T_{ER}}^R = 0. \tag{44}$$

Eqs. (43)–(44) mirror your recursive formulation while keeping everything differentiable with respect to network parameters. As $\tau \to 0$ and at evaluation time (when we take hard decisions), they converge to the discrete model.

(3) Architectural adjustments

To accommodate the relaxation while preserving economic structure, the DNN uses the following heads from the shared representation s_t:

1. *Retirement head* $f_{\text{ret}}(s_t)$ outputs ℓ_t (logit). The Concrete variable $\tilde{R}_t(\tau)$ is computed via (41) and is used only for ages $t \in [T_{ER}, T_{LR}]$; outside this window we set $\tilde{R}_t(\tau) \equiv 0$.

2. *Working branch* (hours & assets): a highway-style *core+residual* head (with the R *adopter* on the core path, as described earlier) outputs $(h_t^W, a_{t+1}^{W|W})$.

3. *Retirement branch* (assets & benefit): a light head outputs $(a_{t+1}^{R|W}, a_{t+1}^{R|R})$ and uses the deterministic benefit mapping B_t or $(t, \sim(y_{1:t-1}))$.[2]

The overall forward pass forms the *soft* mixture in (43), then computes value and continuation terms. This matches your code's practice of (i) keeping the general layer shared, (ii) branching into task-specific heads, and (iii) conditioning the retirement branch activation on the retirement window through a simple MLP adopter.

(4) Objective and Convergence

Let Ω be all network parameters. The training objective maximizes expected discounted utility under the relaxed dynamics:

$$\max_{\Omega} \mathbb{E}\left[\sum_{t=1}^{T} \beta^{t-1} u_t\big(\tilde{R}_t(\tau), h_t^W, a_{t+1}^{W|W}, a_{t+1}^{R|W}, a_{t+1}^{R|R}, B_t\big)\right], \tag{45}$$

subject to the (soft) transitions in (43)–(44) and the exogenous shock process for θ_t. With temperature annealing (42), the learned policy approaches a near-discrete boundary as training proceeds. At test time we set $\tau \to 0$ and take the hard action $R_t \in \{0, 1\}$, then simulate the *original* discrete model. In practice, because benefits and constraints induce strong nonconvexities around the retirement window, annealing improves optimization: high τ smooths the landscape for coarse learning; low τ sharpens the policy near discrete optima.

4.5 Overall DNN Design

Having discussed the main components and innovations of the deep learning framework in the previous sections—namely, history dependence, discrete choices, and life-cycle modeling—this section summarizes the complete architecture, training procedure, and hyperparameters used to solve the model.

Summary of the architecture. The overall neural network architecture is designed to approximate policy functions $\pi_t(M_t; \Omega_t)$ for each age t while allowing efficient parameter sharing and flexible adaptation at key transitions.

- **Input.** At each age t, the state vector is

$$X_t = [s_t^{\text{hist}}, a_t, \theta_t, \textit{edu}, \textit{type}],$$

  where s_t^{hist} is the attention-based embedding of past earnings, and $(a_t, \theta_t, \textit{edu}, \textit{type})$ are current and static states.

[2] Benefits are a known institutional function of the earnings history summary; no additional parameters are needed.

- **General layer $G(\cdot)$.** Two shared MLP blocks (width 64) with ReLU activations, layer normalization, and optional dropout. The output is a shared representation $s_t = G(X_t)$.

- **Task heads.** For each task $q \in \{\text{hours}, \text{consumption}\}$:

$$\widehat{y}_t^{(q)} = \sigma\Big(f_{\text{core}}^{(q)}(s_t^{\text{ER},(q)}; \Theta_{\text{core}}^{(q)}) + f_{\text{res},t}^{(q)}(s_t; \Theta_{\text{res},t}^{(q)})\Big),$$

  where the *R-adopter* modifies the core path only for $t \in [T_{ER}, T_{LR}]$, with $s_t^{\text{ER},(q)} = A^{(q)}(s_t)$ and $s_t^{\text{ER},(q)} = s_t$ otherwise.

- **Output constraints.** Outputs are squashed to feasible domains:

$$h_t \in \{0\} \cup [\underline{h}, \overline{h}], \quad x_t \in (0,1), \quad a_{t+1} \geq 0.$$

  The architecture enforces these via sigmoid (for shares), softplus (for nonnegativity), or by the budget identity.

- **Parameter sharing.** The general layer $G(\cdot)$ and task-core parameters are shared across ages t, while the residual heads are age-specific. This design captures smooth life-cycle variation with parsimonious heterogeneity.

/paragraphRetirement-only DNN. For ages $t > T_{FR}$, the decision problem simplifies to a single continuous choice over assets a_{t+1}, given pension benefit B_t and current assets a_t. This stage is approximated by a lightweight, standalone feed-forward network distinct from the working/retirement joint DNN. The retirement-only DNN takes

$$X_t^R = [a_t, B_t]$$

as input and outputs $\widehat{a}_{t+1} = f_{\text{ret-only}}(X_t^R; \Theta_R)$, where $f_{\text{ret-only}}$ is a two-layer MLP (width 16) with ReLU activations and softplus on the final output to ensure $a_{t+1} \geq 0$. This network is trained jointly with the rest of the model using the same Adam optimizer (Kingma and Ba, 2015b) but with independent parameters Θ_R and no dropout.

4.6 Training Algorithm

The DNN is trained to maximize expected lifetime utility under the model's stochastic environment. The optimization relies on fully differentiable simulation of the life-cycle problem via the annealed Concrete relaxation for retirement and Monte Carlo approximation of expectations.

Algorithm 1 End-to-end training procedure for the life-cycle DNN

Initialize network parameters Ω_0 and temperature schedule $\tau_0, \tau_{\min}, \kappa$.

Simulate exogenous states:

- Draw J Monte Carlo paths $\{\epsilon_{1:T}^{(j)}\}_{j=1}^{J}$, simulate $\theta_{t+1}^{(j)} = \theta_t^{(j)} + \epsilon_{t+1}^{(j)}$.
- Draw $edu^{(j)}$ from its Bernoulli distribution.

for each training epoch $e = 1, \ldots, E$ **do**
  Update temperature $\tau_e = \max\{\tau_{\min}, \tau_0 e^{-\kappa e}\}$.
  for each path $j = 1, \ldots, J$ **do**
    Initialize $a_1^{(j)} = 0$, $P_{T_{ER}}^{R,(j)} = 0$.
    for $t = 1$ to T **do**
      Construct state $X_t^{(j)} = [s_t^{\text{hist},(j)}, a_t^{(j)}, \theta_t^{(j)}, edu^{(j)}, type^{(j)}]$.
      Compute $s_t^{(j)} = G(X_t^{(j)}; \Theta_G)$.
      Evaluate task heads to obtain $h_t^{(j)}$, $x_t^{(j)}$, and $a_{t+1}^{(j)}$.
      If $t \in [T_{ER}, T_{LR}]$, compute soft retirement $\tilde{R}_t(\tau_e)$ using Gumbel–Softmax.
      Update endogenous states using budget constraints and (if in ER window) the soft mixture rule.
    end for
  end for
  Compute the loss $\mathcal{L}(\Omega_e)$ as negative discounted utility plus regularizers.
  Update parameters $\Omega_{e+1} \leftarrow \text{Adam}(\Omega_e, \nabla_\Omega \mathcal{L}, \eta_e)$ with gradient clipping and LR schedule.
end for
Evaluation: Fix $\tau = 0$, replace $\tilde{R}_t$ with hard $R_t \in \{0, 1\}$, and simulate discrete model for welfare and policy analysis.

Regularization and dropout. We apply moderate dropout for both the residual and attention components to prevent overfitting: No additional jump or residual-scaling penalties are used. This dropout-based regularization improves generalization while preserving numerical stability in gradient propagation.

Seaming mechanism. To ensure a smooth connection between working and retirement regimes during the early-retirement transition, we employ a retirement *seaming mechanism*. It softly blends the predicted outcomes from the working and retirement heads during a warm-up phase. Let α_e denote the blending weight at epoch e, During the first retire blend warmup epochs:

$$\alpha_e = \alpha_0 \exp(-\lambda_\alpha e),$$

The final blended output is

$$\tilde{y}_t^{\text{blend}} = \alpha_e \, y_t^R + (1 - \alpha_e) \, y_t^W,$$

where y_t^R and y_t^W are the outputs from the retirement and working heads, respectively. This mechanism stabilizes the training around the discontinuity of retirement entry and aligns both heads before they fully decouple.

Optimization. All parameters Ω are updated via the Adam optimizer (Kingma and Ba, 2015b), which adaptively rescales learning rates for each parameter using running estimates of first and second moments of gradients:

$$m_t = \beta_1 m_{t-1} + (1-\beta_1) g_t, \qquad v_t = \beta_2 v_{t-1} + (1-\beta_2) g_t^2,$$
$$\hat{m}_t = \frac{m_t}{1-\beta_1^t}, \qquad \hat{v}_t = \frac{v_t}{1-\beta_2^t}, \qquad \Omega_t = \Omega_{t-1} - \eta \frac{\hat{m}_t}{\sqrt{\hat{v}_t} + \epsilon}.$$

Adam combines the advantages of momentum and RMSProp, allowing stable convergence even in high-dimensional stochastic gradients. We use the default coefficients $\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 10^{-8}$, and learning rate $\eta = 5 \times 10^{-5}$.

4.7 Hyperparameters

Table 4.7 reports the full set of architectural, optimization, and relaxation hyperparameters used in the deep learning framework. These values govern the configuration of all model components, including the multitask life-cycle DNN, the attention-based history encoder, the Gumbel–Softmax relaxations (for both top-year summarization and retirement hazard), and the post-retirement continuation network. Layer normalization is applied within each MLP block, which stabilizes gradients across simulated paths and prevents batch-size sensitivity common in BatchNorm. Optimization relies on the Adam algorithm with a fixed learning rate and gradient clipping for stability, while regularization is achieved primarily through moderate dropout in the residual and attention layers. Parameters are initialized using Kaiming initialization for ReLU layers and Xavier initialization for linear heads, ensuring stable gradient magnitudes at the start of training. In addition, losses for hours and consumption are normalized by their sample standard deviations to balance gradients across tasks. Table 4.7 ensures reproducibility and provides a transparent reference for future implementations of similar hybrid structural–deep learning models.

Table 1: Final hyperparameters of the DNN life-cycle model

Component	**Parameter**	**Value / Setting**
General & Task Architecture	General layer depth	2 layers (MLP)
	General layer width	32 neurons per layer
	Task-core layer depth	2 layers
	Task-residual layer depth	1 layer
	R adapter scale cap	0.25
	Activation function	ReLU
	Normalization	LayerNorm
Attention Encoder (History)	Number of heads H	4
	Embedding dimension d	32
	Feed-forward width	64
	Dropout (attention)	0.10
	Soft-rank temperature	0.10
Retirement-only DNN ($t > T_{FR}$)	Input features	$[a_t, B_t]$
	Depth	2 layers (MLP)
	Width	8 neurons per layer
	Activation	ReLU (hidden), Softplus (output)
	Dropout	None
Retirement Head and Seaming	Blend warm-up epochs	30
	Blend initial α_0	0.8
	Blend decay λ_α	0.05
Gumbel–Softmax Relaxation: Top Years	min temperature	0.50
	decay rate	1×10^{-5}
	update every	500 steps
Gumbel–Softmax Relaxation: Retirement	min temperature	0.35
	decay rate	5×10^{-5}
	update every	25 steps
Optimization	Optimizer	Adam ($\beta_1 = 0.9$, $\beta_2 = 0.999$)
	Learning rate	5×10^{-5}
	Epochs	200
	Batch size	256
	Gradient clipping	$\|\nabla\| \leq 1$
Regularization	Dropout (residual)	0.15
	Weight decay	1×10^{-5}
Simulation	Monte Carlo paths J	100,000

5 Calibration

The model is estimated annually, with each period corresponding to one year. A model period begins at age 22 ($t = 1$) and ends at death at age 82 ($T_D = 82$). The early, full, and late retirement ages are set to 62, 67, and 70, respectively. All individuals face deterministic death at T_D and enter the model with zero assets and initial productivity $\theta_0 = 0$. The estimation combines parameters estimated from the data with values taken from the literature, summarized in Table 2.

5.1 Data

To estimate the empirical model, I use data from the Panel Study of Income Dynamics (PSID). The PSID is the longest running longitudinal household survey in the world. It began as a nationally representative sample of 5,000 families in the United States. I use all the observations between 1996-2017 and create the sample of families from the PSID in the following way. I only consider families who will have at most two children throughout their life. I use the couples who are married after 1970. The start of each family is from when the father is at least 22 years old until he is 75 years old. For the couples who get divorced, I drop the observations after the divorce. I also drop the households with multiple children born on the same year.

5.2 Preferences

The baseline preference parameters follow standard values in the life-cycle literature. The annual discount factor is set to $\beta = 0.988$, and the coefficient of relative risk aversion to $\gamma = 1.66$, consistent with Huggett and Parra (2010). The Frisch elasticity of labor supply is fixed at $\eta = 0.5$, a common benchmark in structural labor models (e.g. Hosseini, 2017). Two parameters governing the disutility of labor are estimated within the model to match labor-supply moments for men above age 50: the intensive-margin parameter $\psi = 0.0055$ and the fixed cost of work $\phi = 0.0013$. Together, they allow the model to reproduce both the mean annual working hours and the participation rate near retirement.

The bequest motive follows the warm-glow specification in De Nardi (2004), calibrated to capture realistic late-life wealth accumulation. Specifically, bequest utility is parameterized as

$$u^B(a_{T_D}) = \theta_b \frac{(a_{T_D} + q_b)^{1-\gamma}}{1-\gamma},$$

with $\theta_b = 2419$ and $q_b = 215$. These values match the empirical evidence on bequest behavior among older households, ensuring that savings remain positive late in life.

5.3 Budget Constraint and Returns

Total available time each year is normalized to 8,760 hours. The discrete hours grid is defined by $\{0, 1300, 2080, 2860\}$ corresponding to non-work, part-time, full-time, and overtime schedules. The risk-free interest rate is fixed at $r = 0.042$ per year, in line with Huggett and Parra (2010). The minimum admissible consumption level is set to $\underline{c} = 2000$.

5.4 Earnings Process

The deterministic component of log wages depends quadratically on age and education:

$$\mu(age, edu) = \beta_0^w + \beta_1^w age + \beta_2^w age^2 + \beta_3^w edu + \beta_4^w edu \times age + \beta_5^w edu \times age^2, \tag{46}$$

where $(\beta_0^w, \beta_1^w, \beta_2^w, \beta_3^w, \beta_4^w, \beta_5^w) = (1.6698,\ 0.0605,\ -0.0006,\ -0.3780,\ 0.03214,\ -0.0002)$. The idiosyncratic productivity component follows a random-walk process

$$\theta_{t+1} = \theta_t + \epsilon_{t+1}, \qquad \epsilon_t \sim N(0, \sigma_\epsilon^2),$$

with $\sigma_\epsilon^2 = 0.0260$. The probability of being college-educated is set to $P(edu = 1) = 0.19$, matching the sample share of college men in the calibration data.

5.5 Taxes and Social Security

Taxes and pension benefits follow the U.S. institutional rules around the year 2000. The income tax function takes the non-linear form

$$T^{inc}(y_t) = y_t - \kappa\, y_t^{(1-\tau_T)},$$

with parameters $\kappa = 2.716$ and $\tau_T = 0.103$, estimated from TAXSIM data following Heathcote et al. (2017). The Social Security payroll tax is proportional up to the taxable maximum:

$$T^{SS}(y_t) = \tau_{SS} y_t, \quad \tau_{SS} = 0.106, \quad \text{for } y_t \le \bar{y} = 76{,}200.$$

Pension benefits are determined by the Primary Insurance Amount (PIA) formula with three bend points $a = \{6{,}372,\ 38{,}422,\ 76{,}200\}$ and marginal replacement rates $b = \{0.9,\ 0.32,\ 0.15\}$. The benefit is further adjusted by the actuarial factor $\delta(t_R)$ capturing early or delayed claiming incentives, with

$$\delta(t_R) \in \{0.70, 0.75, 0.80, 0.867, 0.933, 1.00, 1.08, 1.16, 1.24\} \quad \text{for } t_R \in \{61, \ldots, 70\}.$$

Average Indexed Monthly Earnings (AIME) are computed as one-twelfth of the mean of the top 35 annual earnings years, consistent with U.S. Social Security rules.

Table 2: Calibrated Parameters

Category	Parameter	Value	Source / Target
Preferences	β	0.988	Huggett and Parra (2010)
	γ	1.66	Huggett and Parra (2010)
	η	0.50	Hosseini (2017)
	θ_b	2419	De Nardi (2004)
	q_b	215	De Nardi (2004)
Budget / Returns	r	0.042	Huggett and Parra (2010)
Wage Process	$(\beta_0^w, \ldots, \beta_5^w)$	see text	GMM estimates
	σ_ϵ^2	0.0260	GMM estimates
	$P(edu = 1)$	0.19	Data share
Taxes / Social Security	κ	2.716	TAXSIM, Heathcote et al. (2017)
	τ_T	0.103	TAXSIM, Heathcote et al. (2017)
	τ_{SS}	0.106	SSA 2000 rules
	Wage base	76,200	SSA 2000 rules
	Bend points (a)	$\{6{,}372, 38{,}422, 76{,}200\}$	SSA 2000 rules
	Slopes (b)	$\{0.9, 0.32, 0.15\}$	SSA 2000 rules
	Actuarial adj. $\delta(t_R)$	see text	SSA 2000 rules
	Top earnings years T_S	35	SSA 2000 rules

5.6 Estimation via the Simulated Method of Moments (SMM)

We estimate the structural parameters of the model using a Population-Based Training (PBT) algorithm ((Jaderberg, Dalibard, Osindero, Czarnecki, Donahue, Razavi, Vinyals, Green, Dunning, Simonyan, Fernando, and Kavukcuoglu, 2017)) integrated with the Simulated Method of Moments (SMM). The estimation targets two key structural parameters that govern labor supply and participation behavior:

$$\phi \quad \text{(fixed cost of work)}, \qquad \psi \quad \text{(disutility of work hours)}.$$

These parameters jointly determine the model's ability to match the empirical distribution of working hours and retirement ages observed in the data.

Let $M_{\text{model}}(\theta, \phi, \psi)$ denote the vector of simulated moments generated by the deep life-cycle model—given network weights θ and structural parameters (ϕ, ψ)—and let M_{data} denote the corresponding empirical moments from the data. Defining W as the positive semi-definite weighting matrix, the SMM objective is to minimize the quadratic form:

$$\mathcal{L}_{\text{SMM}}(\theta, \phi, \psi) \;=\; c(\theta, \phi, \psi)^\top c(\theta, \phi, \psi) \;=\; \big(M_{\text{model}} - M_{\text{data}}\big)^\top W \big(M_{\text{model}} - M_{\text{data}}\big). \tag{47}$$

PBT evolves a population of candidate models trained in parallel, jointly optimizing network weights θ through stochastic gradient descent and exploring structural parameters (ϕ, ψ) via mutation and selection. Each PBT round alternates between two complementary phases:

- **Exploit:** copy the trained neural network parameters θ from top-performing models (lowest moment residuals) to weaker models, ensuring rapid knowledge transfer.
- **Explore:** perturb the structural parameters (ϕ, ψ) of the copied models to encourage diversity:

$$(\phi, \psi) \leftarrow (\phi, \psi)_{\text{donor}} + u \cdot \text{scale} \cdot \big[(\phi_{\max}, \psi_{\max}) - (\phi_{\min}, \psi_{\min})\big], \qquad u \sim U[-1, 1], \tag{48}$$

  where the mutation scale decays geometrically across rounds—large initially for exploration, smaller later for fine-tuning.

The following algorithm summarizes the estimation procedure combining Population-Based Training (PBT) with the Simulated Method of Moments (SMM).

Algorithm 2 Population-Based Training with Simulated Method of Moments (PBT–SMM)

Initialize: Generate K candidate models with distinct parameter pairs (ϕ_i, ψ_i) sampled from a Sobol sequence over $[\phi_{\min}, \phi_{\max}] \times [\psi_{\min}, \psi_{\max}]$.
for each training round $s = 1, \ldots, S$ **do**
  for each candidate $i = 1, \ldots, K$ **do**
    Inner training: Train DNN parameters θ_i for E epochs using the primary loss $\mathcal{L}_{\text{primary}}$ (welfare maximization), keeping (ϕ_i, ψ_i) fixed.
    Moment evaluation: Compute simulated moments $M_{\text{model}}(\theta_i, \phi_i, \psi_i)$ using Common Random Numbers (CRN) for all i.
    Compute SMM loss:

$$\mathcal{L}_{\text{SMM}}(\theta_i, \phi_i, \psi_i) = \big(M_{\text{model}}(\theta_i, \phi_i, \psi_i) - M_{\text{data}}\big)^{\top} W \big(M_{\text{model}}(\theta_i, \phi_i, \psi_i) - M_{\text{data}}\big).$$

  end for
  Selection: Rank candidates by $\mathcal{L}_{\text{SMM}}$.
  Exploitation: Copy network weights θ from top-performing models to the weakest ones.
  Exploration: Mutate (ϕ, ψ) of replaced candidates with annealed jitter:

$$(\phi, \psi) \leftarrow (\phi_{\text{donor}}, \psi_{\text{donor}}) + u \cdot \text{scale} \cdot [(\phi_{\max} - \phi_{\min}), (\psi_{\max} - \psi_{\min})], \quad u \sim U[-1, 1].$$

end for
Output: Select the top-n candidates minimizing $\mathcal{L}_{\text{SMM}}$ as final estimates.

5.7 Estimated Structural Parameters

The empirical moments and parameter estimates reported in Table 3 are obtained using the simulated method of moments, matching the data moments for mean working hours and average retirement age.

Table 3: Estimated Labor Disutility Parameters

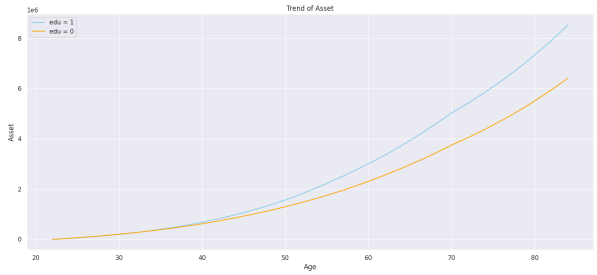
Moment	**Model Target**	**Value**
Mean hours worked (men)	$\mathbb{E}[h_t \mid h_t > 0]$	2,034.4
Mean retirement age	$\mathbb{E}[t_R]$	66.83
Estimated Parameters		
Intensive-margin disutility (ϕ)		0.005534
Extensive-margin fixed cost (ψ)		0.001266

6 Results

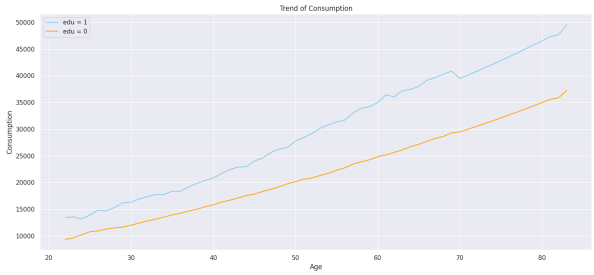
6.1 Status Quo Policy Simulation

This section presents the results of the benchmark simulation under the status quo Social Security policy, where benefits are based on the top 35 years of earnings. The model is solved using the numerical method described in Section 4 and simulated for 25,000 individuals under the estimated parameters (ϕ, ψ) and other calibrated values reported in Section 5.

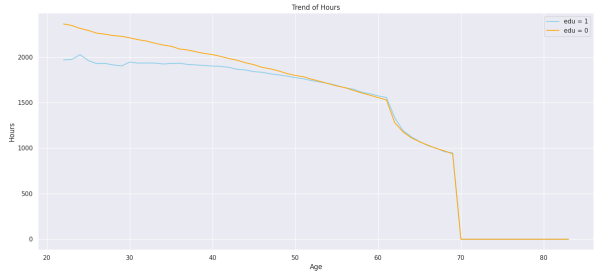
Figure 4 shows the simulated life-cycle profiles of assets, income, and labor supply, separately for college-educated ($edu = 1$) and non-college ($edu = 0$) individuals. The results reproduce the main empirical regularities observed in U.S. micro data.



(a) Asset accumulation by age



(b) Mean Consumption profile



(c) Median income profile



(d) Median hours of work

Figure 4: Simulated life-cycle profiles under the status quo (top-35-year) pension policy.

Panel (a) shows that asset accumulation rises steadily over the working years and accelerates toward the end of the career as individuals prepare for retirement. College-educated individuals accumulate substantially more wealth, reflecting both their higher lifetime earnings and their lower marginal propensity to consume. After retirement, asset decumulation is gradual rather than abrupt, consistent with the bequest motive embedded in the model, which dampens dissaving at older ages. Panel (b) displays the evolution of consumption over the life cycle. Consumption increases steadily with age, mirroring the rise in income and asset holdings.[3]

Panel (c) illustrates labor supply behavior. Hours worked remain close to full-time early in the life cycle and decline gradually as individuals age. A pronounced drop occurs during the early-retirement window (ages 62–67), consistent with strong incentives to claim benefits once eligible. This decline is smoother for college-educated individuals, who earn higher wages and tend to retire later.

Panel (d) plots mean income profiles. Earnings exhibit the well-documented hump-shaped life-cycle pattern: rising rapidly in early adulthood, peaking in middle age, and falling sharply after the early 60s. The post-62 decline reflects the shift from labor earnings to pension benefits and the reduction in mean labor income as many workers reduce labor supply or exit the labor force. College-educated workers earn considerably more at all ages, consistent with the returns to education embedded in the estimated wage process.

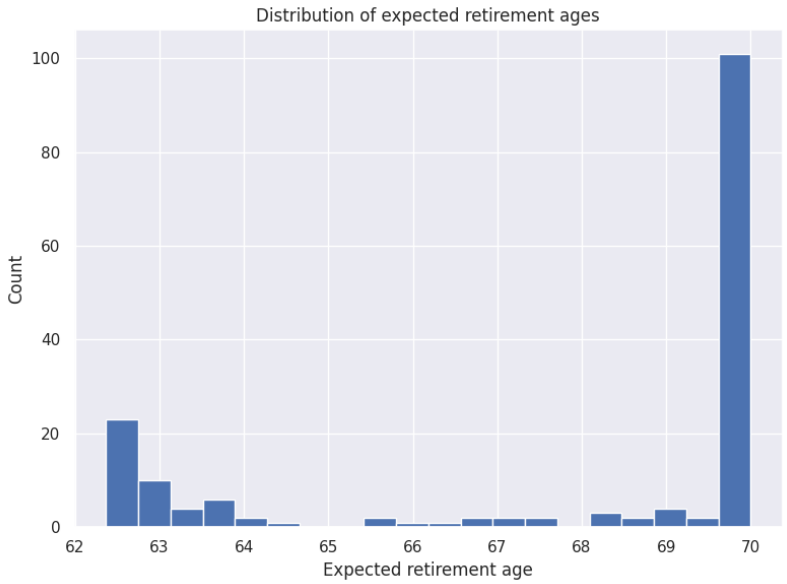


Figure 5: Distribution of expected retirement ages.

Figure 5 reports the distribution of expected retirement ages. The simulated distribution is distinctly bi-modal, with a large mass retiring shortly after the early retirement age (around 62) and another pronounced spike at the full-benefit age (around 70). Only a relatively small share of individuals choose retirement ages in between, reflecting heterogeneity in preferences, earnings histories, and productivity shocks that lead some workers to delay retirement in exchange for higher future benefits or continued earnings. This pattern shows that the model concentrates retirement choices at salient policy ages—capturing the discontinuous

[3]Because the model abstracts from family-size dynamics, it does not generate the hump-shaped consumption profile typically observed in household data.

nature of retirement behavior—while still allowing for smooth optimization during training through the Gumbel–Softmax relaxation mechanism. The average expected retirement age in the simulation is 66.8.

Table 4 summarizes economic outcomes under the status-quo pension rule, which computes AIME using the top 35 earnings years. The table highlights substantial heterogeneity across population subgroups. College-educated workers exhibit markedly higher consumption, earnings, and retirement benefits than non-college workers, reflecting their steeper lifetime wage profiles. Wage-volatility groups differ sharply as well: high-volatility individuals earn more on average but also face greater dispersion in both consumption and benefits, while low-volatility workers display much lower levels and variation. Work-attachment groups show the strongest contrasts. High-attachment individuals have higher consumption, higher benefits, and work substantially more hours at all ages, while low-attachment workers have much lower lifetime earnings, retire earlier, and experience lower welfare levels.

Table 4: Status quo pension rule: counting top 35 earnings years (levels)

	All	Education		Wage volatility		Work attachment	
		College	No college	High	Low	High	Low
Status quo: Counting top 35 years (levels)							
Mean(c_{it})	28684.86	42711.32	25404.10	40528.95	16840.76	39931.53	17409.36
Mean(c_{it}|working)	24574.83	37057.43	21652.52	33890.24	14582.60	33360.70	15191.90
Mean(c_{it}|retired)	39654.22	57842.13	35410.45	60761.17	22158.81	59684.59	22694.40
sd(c_{it})	25053.64	36626.64	20098.31	30351.77	7319.47	30466.06	8512.96
sd(c_{it}|working)	20827.53	31098.03	16245.50	25043.06	5782.86	25151.67	7139.86
sd(c_{it}|retired)	31301.37	45044.89	25293.41	35690.31	7790.93	36010.43	9167.59
Mean(b_i)	17515.11	19287.54	17101.56	22825.70	13113.21	22540.88	13259.77
sd(b_i)	7455.68	7001.94	7497.59	5755.46	5608.22	6122.49	5605.37
Mean($a_i t$|$Age = Age_R$)	4586916.39	6806664.01	4067722.67	6850798.19	2323034.60	6755664.32	2412609.35
sd($a_i t$|$Age = Age_R$)	3762713.48	5468489.48	3008809.42	4147205.79	930799.63	4197018.18	1117616.95
Mean(h_{it}|Age< 70)	2005.59	1984.59	2010.50	2131.33	1879.86	2155.34	1855.46
Mean(h_{it}|62 ≤Age< 70)	924.85	952.56	918.36	1179.92	669.77	1187.89	661.13
Mean(h_{it}|Age< 62)	2221.74	2191.00	2228.93	2321.61	2121.87	2348.83	2094.33
Mean(Age_i^R)	67.10	67.13	67.09	68.68	65.52	68.52	65.68
Mean(Wel_i)	-0.32	-0.27	-0.33	-0.27	-0.37	-0.28	-0.37
Mean(CV)	0.00	0.00	0.00	0.00	0.00	0.00	0.00

Notes: The table reports means and standard deviations of consumption c_{it}, retirement benefits b_i, assets at retirement a_{it} evaluated at Age_i^R, hours worked h_{it} over different age ranges, retirement age Age_i^R, retirement utility Wel_i, and mean consumption equivalence CV. All statistics are in levels under the status-quo pension rule counting the top 35 earnings years. Education groups: edu=1 vs. edu=0. Wage-volatility and work-attachment groups split at medians.

7 Counterfactual

This section discusses the quantitative effects of alternative pension-benefit rules—replacing the status-quo "top 35 years" with counting the top 20 years, top 5 years, or lifetime-average earnings. I show the counterfactual results for consumption, hours worked, retirement behavior, and welfare. We examine how each counterfactual affects aggregate consumption, labor supply, and retirement timing. To understand the heterogeneous effects of these policy changes—and the redistribution, insurance, and incentive consequences they generate—we analyze how the reforms impact workers with different characteristics: high vs. low edu-

cation, high vs. low wage volatility, and high vs. low attachment to the labor force. We define *high-volatility workers* as those whose standard deviation of pre-retirement wages lies above the median of the population wage-volatility distribution. *Low-attachment workers* are defined as individuals whose number of full-time working years before retirement lies below the median of full-time years among all workers.

Table 5: Counterfactual pension rules (retirement benefits only, no behavioral response)

	All	Education		Wage volatility		Work attachment	
		College	No college	High	Low	High	Low
Status quo: Counting top 35 years (levels, no behavioral response)							
Mean(b_t)	18239.98	19974.14	17833.84	23730.33	13691.45	23320.93	13934.08
sd(b_t)	7106.77	6531.73	7174.52	4498.32	5470.17	4830.51	5754.56
Counterfactual: Counting top 20 years (percent change w.r.t. status quo, no behavioral response)							
Mean(b_t)	5.12	2.38	5.83	1.29	10.60	1.46	10.30
sd(b_t)	-6.35	-6.02	-6.11	-4.05	-0.33	-3.66	-1.78
Counterfactual: Counting top 5 years (percent change w.r.t. status quo, no behavioral response)							
Mean(b_t)	10.18	4.63	11.64	1.81	22.20	2.28	21.39
sd(b_t)	-13.99	-12.56	-13.75	-4.98	-2.93	-6.60	-5.29
Counterfactual: Lifetime-average earnings (percent change w.r.t. status quo, no behavioral response)							
Mean(b_t)	-13.50	-8.22	-14.89	-6.99	-22.86	-7.31	-22.29
sd(b_t)	6.39	14.79	3.84	13.11	-12.23	12.18	-9.13

Notes: The table reports the mean and standard deviation of retirement benefits b_t for the status-quo Top-35 rule (levels) and for alternative pension rules (percentage deviations from Top-35). In all cases, labor supply, savings, and retirement decisions are fixed at the Top-35 equilibrium (no behavioral responses), and only the pension benefit rule changes. Education groups: edu=1 vs. edu=0. Wage-volatility and work-attachment groups split at medians.

To isolate the effects of behavioral responses in saving, labor supply, and retirement decisions from the mechanical changes in pension benefits, Table 5 reports how retirement benefits change when the AIME formula is altered while holding all behavioral decisions fixed. This serves as a preliminary evaluation of the redistribution and insurance properties of each counterfactual rule. In the next subsections, we examine each reform separately—Top-20, Top-5, and lifetime-average earnings—incorporating the full behavioral responses of the model to assess their total equilibrium effects.

Reducing the number of top earnings years—from the current Top-35 to Top-20 and then to Top-5—raises average benefits and compresses their dispersion. The Top-20 rule increases mean benefits by about 5%, and the Top-5 rule by about 10%, while substantially reducing benefit inequality as reflected in 6–14% declines in their standard deviation. These aggregate changes mask sizable heterogeneity across groups. Non-college, low-wage-volatility, and low-attachment workers consistently gain the most: the Top-5 rule raises their benefits by 12%, 22%, and 21%, respectively. For non-college and low-attachment workers, this reflects the removal of many low-earning years from the AIME calculation. For low-volatility workers, the mechanism is different: they have few extremely high-earning years and few very low-earning years, so shrinking the averaging window amplifies the influence of their limited peak years while diminishing the impact of their limited bad years, mechanically boosting their AIME. By contrast, college-educated, high-volatility, and high-attachment workers benefit far less when the averaging window is shortened. For high-volatility individuals, the small gains arise because their earnings include both many very low and many very high years; although removing the low years increases AIME, the dispersion in their top years is not especially large, so the Top-20 or Top-5 formulas do not raise their assessed lifetime earnings by much. High-attachment and

college-educated workers also gain little because their earnings histories contain fewer low-earning spells to discard.

Switching to a lifetime-average formula reverses all these patterns. Because lifetime averaging includes every low-earning year, it reduces benefits on average (by 13.5%) and heavily penalizes groups with consistently modest or intermittent earnings, such as low-volatility and low-attachment workers, whose benefits fall by more than 22%. Meanwhile, benefit dispersion rises sharply for college-educated and high-volatility workers, reflecting the wider spread of lifetime earnings in these groups.

7.1 Changing the Number of Years in AIME

The results of the two counterfactuals that change the AIME rule—from counting the top 35 years to counting the top 20 and top 5 years—are presented in Table 6[4].

7.1.1 From Top 35 to Top 20 Years

Lowering the number of earnings years used in AIME—from the current Top-35 rule to a Top-20 rule—has modest but systematic effects on lifetime outcomes. By shortening the averaging window, the reform increases the weight placed on higher-earning years and reduces the influence of low-earning spells in the benefit formula. On average, this leads to higher assessed lifetime earnings and retirement benefits, with the largest gains accruing to workers with low wage volatility and weaker labor-force attachment, while individuals with highly volatile earnings or strong attachment experience only small changes. These shifts in retirement benefits are the primary channel through which the reform affects consumption, labor supply, and retirement timing.

Across the full population, consumption rises by 1.8% under the Top-20 reform, reflecting parallel increases during both working years (+1.9%) and retirement (+1.5%). The pattern closely mirrors the change in retirement benefits, which rise by 4.4% on average and constitute the key driver of the consumption response. Workers with low wage volatility experience the strongest gains in both consumption (+2.4%) and benefits (+10.3%), since removing many low-earning years compresses their AIME calculation upward. High-volatility workers, by contrast, see smaller increases in consumption (+1.6%) and almost no change in benefits (+0.3%). The same mechanism appears when comparing work-attachment groups: those with low attachment—who experience a 10.0% rise in benefits—see the largest consumption gains, while high-attachment workers, whose benefits barely change (−0.6%), register minimal consumption gains. Differences by education are similarly modest: college workers gain +1.7% in consumption and non-college workers +1.9%, reflecting two offsetting forces. College-educated individuals benefit from placing more weight on their higher peak earnings, while non-college individuals benefit from eliminating early low-income spells; the net effect is a small and balanced increase for both groups.

The Top-20 reform generates a small overall increase in labor supply. The strongest responses come from low-volatility and low-attachment workers, who benefit from removing many low-earning years from the

[4]Note that changes in the pension rule affect both the level of benefits and their distribution across individuals. To abstract from changes in the overall benefit level, Appendix B presents budget-neutral counterfactuals that isolate the distributional and behavioral effects of alternative pension formulas.

AIME calculation. Their retirement benefits rise substantially under the reform which strengthens their incentives to work more at younger ages in order to reach the higher-benefit range. However, once they approach retirement eligibility, the same increase in benefits weakens their incentives to remain in the labor force. Consistent with this mechanism, these two groups are the only ones who reduce their hours between ages 62 and 70 (by 1.75% and 1.09%, respectively), and they also exhibit slightly earlier retirement (0.20 and 0.11 years earlier).

Table 6: Counterfactual pension rules: counting top earnings years (percent change vs status quo)

	All	Education		Wage volatility		Work attachment	
		College	No college	High	Low	High	Low
Counterfactual: Counting top 20 years (percent change w.r.t. status quo)							
Mean(c_{it})	1.83	1.72	1.88	1.60	2.39	1.06	1.23
Mean(c_{it}|working)	1.86	1.68	1.93	1.71	2.16	1.31	1.06
Mean(c_{it}|retired)	1.51	1.58	1.50	1.13	2.47	0.04	1.36
sd(c_{it})	1.00	0.95	0.97	0.91	2.37	0.33	-0.86
sd(c_{it}|working)	1.20	1.15	1.20	1.04	2.82	0.53	-0.65
sd(c_{it}|retired)	0.55	0.45	0.49	0.68	1.01	0.27	-2.66
Mean(b_i)	4.37	1.73	5.07	0.30	10.27	-0.50	10.12
sd(b_i)	-5.84	-5.49	-5.64	-0.58	-1.11	1.25	-2.08
Mean($a_i t$|$Age = Age_R$)	0.81	0.99	0.74	0.65	1.27	-0.23	0.47
sd($a_i t$|$Age = Age_R$)	0.49	0.43	0.39	0.56	0.71	0.44	-2.59
Mean(h_{it}|Age< 70)	1.72	1.73	1.72	1.38	2.11	0.97	2.33
Mean(h_{it}|62 ≤Age< 70)	2.12	1.72	2.22	4.31	-1.73	2.98	-1.00
Mean(h_{it}|Age< 62)	1.69	1.74	1.68	1.09	2.35	0.77	2.54
Mean(Age_i^R)	-0.18	-0.14	-0.18	-0.14	-0.21	-0.31	-0.11
Mean(Wel_i)	0.14	0.40	0.10	0.63	-0.21	0.69	0.17
Mean(CV)	-0.01	-0.03	-0.01	-0.03	0.01	-0.03	0.01
Counterfactual: Counting top 5 years (percent change w.r.t. status quo)							
Mean(c_{it})	0.33	-1.05	0.87	0.18	0.68	-0.57	-0.83
Mean(c_{it}|working)	0.11	-1.83	0.88	-0.05	0.36	-0.44	-1.43
Mean(c_{it}|retired)	0.22	-0.07	0.36	0.14	0.70	-1.72	-0.14
sd(c_{it})	0.32	0.98	0.39	0.42	1.15	-0.56	-0.44
sd(c_{it}|working)	0.19	0.90	0.52	0.24	1.32	-0.70	-1.64
sd(c_{it}|retired)	0.06	0.44	-0.03	0.32	0.23	-0.10	-0.36
Mean(b_i)	8.16	2.98	9.53	0.14	19.82	-0.90	19.76
sd(b_i)	-11.61	-9.80	-11.52	-0.22	-4.05	0.62	-4.62
Mean($a_i t$|$Age = Age_R$)	-0.93	-0.80	-0.98	-0.68	-1.66	-2.19	-1.78
sd($a_i t$|$Age = Age_R$)	-0.02	0.29	-0.17	0.08	-0.00	0.08	0.37
Mean(h_{it}|Age< 70)	-0.76	-2.35	-0.39	-0.42	-1.14	-0.75	-1.17
Mean(h_{it}|62 ≤Age< 70)	1.07	0.93	1.11	4.42	-4.82	1.84	-2.46
Mean(h_{it}|Age< 62)	-0.91	-2.63	-0.51	-0.91	-0.91	-1.02	-1.09
Mean(Age_i^R)	-0.30	-0.24	-0.32	-0.22	-0.39	-0.49	-0.21
Mean(Wel_i)	-0.50	-0.35	-0.52	-0.11	-0.78	0.18	-0.41
Mean(CV)	0.02	0.02	0.02	0.01	0.03	0.01	0.03

Notes: The table reports percentage deviations from the status-quo pension rule (Top-35) in the means and standard deviations of consumption c_{it}, retirement benefits b_i, assets at retirement a_{it} evaluated at Age_i^R, hours worked h_{it}, retirement age Age_i^R, retirement utility Wel_i, and mean consumption equivalence CV. Positive values denote increases relative to the status quo. Education groups: edu=1 vs. edu=0. Wage-volatility and work-attachment groups split at medians.

7.1.2 From Top 35 to Top 5 Years

Further reducing the number of earnings years used in AIME—from the Top-20 rule to an extreme Top-5 rule—amplifies the mechanisms already observed when shortening the window from Top-35 to Top-20. By placing almost all weight on the very highest-earning years, the Top-5 formula substantially increases assessed lifetime earnings for individuals who have at least a few strong peak earnings years, while sharply reducing the importance of the remaining years in the earnings history. As shown mechanically in Table 5, this reform increases average pension benefits by about 10% and compresses their dispersion, with the largest mechanical gains accruing to workers with stable earnings or intermittent labor-force attachment. These mechanical changes remain the dominant force shaping behavioral and welfare outcomes under the Top-5 system.

Average consumption rises (+0.33%), with small and heterogeneous adjustments during both working years (+0.13%) and retirement (+0.19%). The mechanism mirrors the one observed under Top-20 but becomes stronger: because they have few very high-earning years and few very low ones, shrinking the window to only five years amplifies the influence of their limited peak years while eliminating the effects of their few low-earning spells. Low-attachment workers similarly benefit from removing long sequences of low-earning years accrued during periods of weak labor-force participation. College-educated workers see a modest decline in consumption (–1.05%), reflecting that their highest earnings are already concentrated among a relatively larger number of good years; reducing the window to only five does not raise their assessed earnings much and slightly shifts the timing of resources. High-volatility workers experience small changes in consumption (+0.19%) despite having many high and low years; the key reason is that the dispersion in their top five years is not especially large, limiting the effect of concentrating the window further. High-attachment workers experience small declines in consumption (–0.68%), as their earnings histories contain fewer low-earning years to discard, causing the Top-5 rule to generate only minimal changes in their AIME.

The labor-supply effects of the Top-5 rule are also very small. Hours worked before age 70 decline slightly (–0.76%), with minimal differences across subgroups. There is no clear systematic behavioral response analogous to the Top-20 case. For most groups, the Top-5 gain in benefits is too small to induce significant changes in work incentives. The only exception is a modest increase in hours among high-volatility and low-volatility workers at ages 62–69 (+4.44% and +1.07%, respectively), reflecting the fact that the Top-5 formula increases their peak-year valuation slightly. Retirement timing also shifts minimally: the average retirement age decreases by only 0.30 years, and subgroup differences are negligible.

7.2 Changing the AIME to Lifetime-Average Earnings

8 Conclusion

This paper has examined how the history-dependent component of public pension systems—the way in which lifetime earnings are summarized into a benefit base—shapes redistribution, insurance, and work incentives over the life cycle. I developed and estimated a structural life-cycle model in which Social Security benefits depend on a non-linear transformation of the entire earnings history through a top-years (AIME-style) formula, and workers make endogenous decisions about hours, savings, and retirement timing in response to

this rule. The model replicates key features of U.S. data and provides a laboratory for comparing alternative earnings-summarization regimes that are widely used or debated in policy discussions.

The counterfactual analysis shows that modest adjustments to the U.S. Top-35 rule have relatively small but systematic effects on individual behavior and welfare. Shortening the averaging window to the top 20 years raises benefits and consumption slightly and compresses their dispersion, with the largest gains accruing to workers with low wage volatility and weaker attachment to the labor force, for whom many low-earning years are removed from the calculation. Further shrinking the window to an extreme Top-5 rule mostly amplifies these redistributive patterns without generating strong additional behavioral responses: benefits rise more for non-college, low-volatility, and low-attachment workers, but aggregate consumption, hours, and retirement timing move only marginally. In contrast, replacing the Top-35 rule with a lifetime-average formula produces large and highly uneven effects on both behavior and welfare. Because lifetime averaging brings every low-earning spell back into the benefit base, it substantially reduces pension entitlements for individuals with prolonged or repeated periods of low wages or intermittent participation and tightens the link between realized earnings and benefits. Workers respond by increasing labor supply significantly and postponing retirement, especially at older ages and particularly in the groups hit hardest by the mechanical benefit cuts, such as low-volatility and low-attachment workers.

References

Tohid Atashbar and Aruhan Rui Shi. Deep reinforcement learning: Emerging trends in macroeconomics and future prospects. 2022. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4313684.

Marlon Azinović and Jan Zemlička. Economics-inspired neural networks with stabilizing homotopies. 2023. URL https://arxiv.org/abs/2303.14802.

Marlon Azinovic, Luca Gaegauf, and Simon Scheidegger. Deep equilibrium nets. *International Economic Review*, 63(4):1471–1525, 2022. doi: https://doi.org/10.1111/iere.12575. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/iere.12575.

Ross Batzer. Letting your past define your taxes: Optimal history-dependent income taxation in general equilibrium. *Working Paper*, pages 1–62, 2021.

Pierre Beck, Pablo Garcia Sanchez, Alban Moura, Julien Pascal, and Olivier Pierrard. Deep learning solutions of dsge models: A technical report. 2024. URL https://www.bcl.lu/en/publications/Working-papers/184/BCLWP184.pdf.

Margherita Borella, Mariacristina De Nardi, and Fang Yang. Are marriage-related taxes and social security benefits holding back female labor supply. *The Review of Economic Studies*, 90:102–131, 2023.

Rich Caruana. Multitask learning. *Machine Learning*, 28:41–75, 1997.

Mingli Chen, Rama Cont, Andreas Joseph, Michael Kumhof, Xinlei Pan, Wei Xiong, and Xuan Zhou. Deep reinforcement learning in a monetary model. 2025. URL https://www.bankofengland.co.uk/working-paper/2025/deep-reinforcement-learning-in-a-monetary-model.

Ruxin Chen. Deep reinforcement learning in a search-matching model of labor market fluctuations. *Economies*, 13(10):302, 2025. doi: 10.3390/economies13100302. URL https://doi.org/10.3390/economies13100302.

Michael Curry, Alexander Trott, Soham Phade, Yu Bai, and Stephan Zheng. Analyzing micro-founded general equilibrium models with many agents using deep reinforcement learning. 2022. URL https://arxiv.org/abs/2201.01163.

Mariacristina De Nardi. Wealth inequality and intergenerational links. *Review of Economic Studies*, 71(3):743–768, 2004. doi: 10.1111/0034-6527.00302.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*, 2018.

Peter Diamond. Social security. *American Economic Review*, 94(1):1–24, March 2004. doi: 10.1257/000282804322970670. URL https://www.aeaweb.org/articles?id=10.1257/000282804322970670.

Peter Diamond. Taxes and pensions. *Southern Economic Journal*, 76:2–15, 2009. ISSN 0038-4038. doi: 10.4284/sej.2009.76.1.2.

Peter A . Diamond and James A . Mirrlees. Payroll-tax financed social insurance with variable retirement. *The Scandinavian Journal of Economics*, 88:25–50, 1986.

Peter A Diamond and James A. Mirrlees. A model of social insurance with variable retirement. *Journal of Public Economics*, 10:295–336, 1978.

Mahdi Ebrahimi Kahou, Jesús Fernández-Villaverde, Jesse Perla, and Arnav Sood. Exploiting symmetry in high-dimensional dynamic programming. 2021. URL https://doi.org/10.3386/w28981.

Benjamin Patrick Evans and Sumitra Ganesh. Learning and calibrating heterogeneous bounded rational market behaviour with multi-agent reinforcement learning. 2024. URL https://arxiv.org/abs/2402.00787.

Xiaodong Fan, Ananth Seshadri, and Christopher Taber. Estimation of a life-cycle model with human capital, labor supply and retirement. *NBER Working Papers 29905*, 2022.

Jesús Fernández-Villaverde, Galo Nuño, and Jesse Perla. Taming the curse of dimensionality: Quantitative economics with deep learning. 2024. URL https://doi.org/10.3386/w33117.

Eric French. The effects of health, wealth, and wages on labour supply and retirement behaviour. *The Review of Economic Studies*, 72:395–427, 2005. URL https://academic.oup.com/restud/article/72/2/395/1558553.

Mikhail Golosov, Ali Shourideh, Maxim Troshkin, and Aleh Tsyvinski. Optimal pension systems with simple instruments. *American Economic Review: Papers Proceedings*, 103:502–507, 2013.

Borys Grochulski and Narayana Kocherlakota. Nonseparable preferences and optimal social security. *Journal of Economic Theory*, 145:2055–2077, 2010. ISSN 0022-0531. doi: 10.1016/j.jet.2010.01.012. URL http://dx.doi.org/10.1016/j.jet.2010.01.012.

Jonathan Heathcote, Kjetil Storesletten, and Giovanni L. Violante. Optimal tax progressivity: An analytical framework. *Quarterly Journal of Economics*, 132:1693–1754, 2017. ISSN 15314650. doi: 10.1093/qje/qjx018.

Roozbeh Hosseini. Adverse selection in the annuity market and the role for social security. *Journal of Political Economy*, 125(1):122–160, 2017. doi: 10.1086/689607.

Mark Huggett and Juan Parra. How well does the u . s . social insurance system provide social insurance ? *Journal of Political Economy*, 118:76–112, 2010.

Max Jaderberg, Valentin Dalibard, Simon Osindero, Wojciech M. Czarnecki, Jeff Donahue, Ali Razavi, Oriol Vinyals, Tim Green, Iain Dunning, Karen Simonyan, Chrisantha Fernando, and Koray Kavukcuoglu. Population based training of neural networks. *arXiv preprint arXiv:1711.09846*, 2017. URL https://arxiv.org/abs/1711.09846.

Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. In *International Conference on Learning Representations (ICLR)*, 2017.

Marek Kapička. Quantifying the welfare gains from history dependent income taxation. *Working Paper*, 2022.

Hanno Kase, Leonardo Melosi, and Matthias Rottner. Estimating nonlinear heterogeneous agent models with neural networks. 2025. URL https://www.bis.org/publ/work1241.pdf.

Alex Kendall, Yarin Gal, and Roberto Cipolla. Multi-task learning using uncertainty to weigh losses. In *CVPR*, 2018.

Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, 2015a.

Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *International Conference on Learning Representations (ICLR)*, 2015b.

Alexander Kwon. Using deep neural networks to solve discrete-continuous dynamic choice models in an overlapping generations setting. 2025. URL https://ssrn.com/abstract=5195102.

Vadym Lepetyuk, Lilia Maliar, and Serguei Maliar. When the u.s. catches a cold, canada sneezes: A lower-bound tale told by deep learning. *Journal of Economic Dynamics & Control*, 117:103926, 2020. doi: 10.1016/j.jedc.2020.103926. URL https://doi.org/10.1016/j.jedc.2020.103926.

Bryan Lim and Stefan Zohren. Temporal fusion transformers for interpretable multi-horizon time series forecasting. In *International Journal of Forecasting*, 2021.

Chris J. Maddison, Andriy Mnih, and Yee Whye Teh. The concrete distribution: A continuous relaxation of discrete random variables. In *International Conference on Learning Representations (ICLR)*, 2017.

Lilia Maliar, Serguei Maliar, and Pablo Winant. Deep learning for solving dynamic economic models. *Journal of Monetary Economics*, 122:76–101, 2021. doi: 10.1016/j.jmoneco.2021.07.004. URL https://doi.org/10.1016/j.jmoneco.2021.07.004.

Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. *Advances in Neural Information Processing Systems (NeurIPS)*, 2013.

Abdoulaye Ndiaye. Flexible retirement and optimal taxation. *Working Paper*, 2020.

Cormac O'Dea. Insurance , efficiency and design of public pensions. *Working Paper*, 2017.

Julien Pascal. Artificial neural networks to solve dynamic programming problems: A bias-corrected monte carlo operator. 2024. URL https://ssrn.com/abstract=4476122.

Jonathan Payne, Adam Rebei, and Yucheng Yang. Deep learning for search and matching models. 2024. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4768566.

Ozan Sener and Vladlen Koltun. Multi-task learning as multi-objective optimization. In *NeurIPS*, 2018.

Ali Shourideh and Maxim Troshkin. Incentives and efficiency of pension systems. *Working Paper*, 2017.

Rupesh Kumar Srivastava, Klaus Greff, and Jürgen Schmidhuber. Highway networks. In *arXiv preprint arXiv:1505.00387*, 2015.

Trevor Standley, Amir R. Zamir, Dawn Chen, Leonidas Guibas, Jitendra Malik, and Silvio Savarese. Which tasks should be learned together in multi-task learning? In *ICML*, 2020.

Jeffrey E. Sun. Continuation value is all you need: "drop-in" deep learning for heterogeneous-agent models with aggregate shocks. 2023. URL https://www.jeffreyesun.com/Continuation_Value_Is_All_You_Need.pdf.

Vytautas Valaitis and Alessandro T. Villa. A machine learning projection method for macro-finance models. *Quantitative Economics*, 15(1):145–173, 2024. doi: 10.3982/QE1403. URL https://doi.org/10.3982/QE1403.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.

William Vickrey. *Agenda for Progressive Taxation*. 1947.

A Modeling Discrete Choice in DNN

For $t = 1, \cdots, T_{ER} - 1$, the state variables are $M_t = \{\{y_l\}_{l=1}^{t-1}, a_t, \theta_t, edu\}$ and the control variables are $\pi_t(M_t, \Omega_t) = \{a_{t+1}(M_t, \Omega_t), h_t(M_t, \Omega_t)\}$ and the individual solves

$$
\begin{aligned}
\Omega_t^* = argmax_{\Omega_t} \mathbb{E} \mid_{\theta_{t+1}} V_t\big(\pi_t(M_t, \Omega_t) \mid M_t\big) = \\
u^t(a_{t+1}(M_t, \Omega_t), h_t(M_t, \Omega_t)) + \beta \mathbb{E} \mid_{\theta_{t+1}} V_{t+1}\big(\pi_{t+1}(M_{t+1}, \Omega_{t+1}), \Omega_{t+1}) \mid M_{t+1}\big) \\
s.t. \\
M_{t+1} = \{\{y_l\}_{l=1}^{t}, a_{t+1}, \theta_{t+1}, edu\} \\
a_{t+1}(M_t, \Omega_t) = (1+r)\Big(a_t + y_t - \tau^{inc}(y_t) - \tau^{SS}(y_t) - c_t\Big), y_t = \exp(\mu(edu, t) + \theta_t) \cdot h_t(M_t, \Omega_t) \\
\{y_l\}_{l=1}^{t} = \{\{y_l\}_{l=1}^{t-1}, \exp(\mu(edu, t) + \theta_t) \cdot h_t(M_t, \Omega_t)\} \\
\theta_{t+1} = \theta_t + \epsilon_{t+1},\ \epsilon_{t+1} \sim N(0, {\sigma_\epsilon}^2)
\end{aligned}
\tag{49}
$$

For $t = T_{ER}$, the state variables are $M_{T_{ER}} = \{\{y_l\}_{l=1}^{T_{ER}-1}, a_{T_{ER}}, \theta_{T_{ER}}, edu\}$ and the control variables are $\pi_{T_{ER}}(M_{T_{ER}}, \Omega_{T_{ER}}) = \{\mathbb{P}^R_{T_{ER}}(M_{T_{ER}}, \Omega_{T_{ER}}), a^W_{T_{ER}+1}(M_{T_{ER}}, \Omega_{T_{ER}}), a^R_{T_{ER}+1}(M_{T_{ER}}, \Omega_{T_{ER}}), h^W_{T_{ER}}(M_{T_{ER}}, \Omega_{T_{ER}})\}$, where $a^W_{T_{ER}+1}$ and $a^R_{T_{ER}+1}$ are the assets left for period $T_{ER}+1$ in case of working or retiring at period T_{ER}, respectively and $\mathbb{P}^R_{T_{ER}}$ is the probability of choosing to be retired at period T_{ER}. The individual solves

$$
\begin{aligned}
\Omega^*_{T_{ER}} = argmax_{\Omega_{T_{ER}}} \mathbb{E} \mid_{\theta_{T_{ER}+1}} V_{T_{ER}}\big(\pi_{T_{ER}}(M_{T_{ER}}, \Omega_{T_{ER}}) \mid M_{T_{ER}}\big) = \\
\mathbb{P}^R_{T_{ER}}(M_{T_{ER}}, \Omega_{T_{ER}}) \cdot u^{T_{ER}}\big(a^R_{T_{ER}+1}(M_{T_{ER}}, \Omega_{T_{ER}}), \mathbb{b}(T_{ER}, \mathbb{s}(\{y_l\}_{t=1}^{T_{ER}-1}))\big) \\
+ \Big(1 - \mathbb{P}^R_{T_{ER}}(M_{T_{ER}}, \Omega_{T_{ER}})\Big) \cdot u^{T_{ER}}\big(a^W_{T_{ER}+1}(M_{T_{ER}}, \Omega_{T_{ER}}), h^W_{T_{ER}+1}(M_{T_{ER}}, \Omega_{T_{ER}})\big) \\
+ \beta \mathbb{E} \mid_{\theta_{T_{ER}+1}} V_{T_{ER}+1}\big(\pi_{T_{ER}+1}(M_{T_{ER}+1}, \Omega_{T_{ER}+1}), \Omega_{T_{ER}+1}) \mid M_{T_{ER}+1}\big) \\
s.t. \\
M_{T_{ER}+1} = \{\mathbb{P}^R_{l<T_{ER}+1}, \{y_l\}_{l=1\ |_W}^{T_{ER}}, a_{T_{ER}+1|_W}, a_{T_{ER}+1|_R}, B_{T_{ER}+1|_R}, \theta_{T_{ER}+1}, edu\} \\
\mathbb{P}^R_{l<T_{ER}+1} = \mathbb{P}^R_{T_{ER}}(M_{T_{ER}}, \Omega_{T_{ER}}) \\
\{y_l\}_{l=1\ |_W}^{T_{ER}} = \{\{y_l\}_{l=1}^{T_{ER}-1}, \exp(\mu(edu, T_{ER}) + \theta_{T_{ER}}) \cdot h^W_{T_{ER}}(M_{T_{ER}}, \Omega_{T_{ER}})\} \\
a_{T_{ER}+1|_W} = a^W_{T_{ER}+1}(M_{T_{ER}}, \Omega_{T_{ER}}) = (1+r)\Big(a_{T_{ER}} + y^W_{T_{ER}} - \tau^{inc}(y^W_{T_{ER}}) - \tau^{SS}(y^2_{T_{ER}}) - c^W_{T_{ER}}\Big), y^W_{T_{ER}} = \\
\exp(\mu(edu, T_{ER}) + \theta_{T_{ER}}) \cdot h^W_{T_{ER}}(M_{T_{ER}}, \Omega_{T_{ER}}) \\
a_{T_{ER}+1|_R} = a^R_{T_{ER}+1}(M_{T_{ER}}, \Omega_{T_{ER}}) = (1+r)\Big(a_{T_{ER}} + \mathbb{b}(T_{ER}, \mathbb{s}(\{y_l\}_{t=1}^{T_{ER}-1})) - c^R_{T_{ER}}\Big) \\
B_{T_{ER}+1|_R} = \mathbb{b}(T_{ER}, \mathbb{s}(\{y_l\}_{t=1}^{T_{ER}-1})) \\
\theta_{T_{ER}+1} = \theta_{T_{ER}} + \epsilon_{T_{ER}+1},\ \epsilon_{T_{ER}+1} \sim N(0, {\sigma_\epsilon}^2)
\end{aligned}
\tag{50}
$$

For $t = T_{ER}+1, \cdots, T_{FR}-1$, the state variables are $M_t = \{\mathbb{P}^R_{l<t}, \{y_l\}_{l=1\ |_W}^{t-1}, a_{t|_W}, a_{t|_R}, B_{t|_R}, \theta_t, edu\}$, where $\mathbb{P}^R_{l<t}$ is the probability of entering period t while being retired (becoming retired at one time in the past). All

the variables that are conditional on being in the working stage ($\{y_l\}_{l=1|_W}^{t-1}$ for history of earning at t and $a_{t|_W}$ for asset at t) show state variables in t conditional on entering period t in working stage (not being retired). All the variables that are conditioned on being on retirement stage ($B_{t|_R}$ for the yearly pension benefit at t and $a_{t|_R}$ for asset at t) show state variables in t conditional on entering period t while in retirement stage (getting retired at $s < t$ in the past).

The control variables are

$$\pi_t(M_t;\Omega_t) = \{\mathbb{P}_t^R(M_t,\Omega_t), a_{t+1|_W}^W(M_t,\Omega_t), a_{t+1|_W}^R(M_t,\Omega_t), a_{t+1|_R}^R(M_t,\Omega_t), h_t^{W}{}_{|_W}(M_t,\Omega_t)\}, \tag{51}$$

where $a_{t+1|_W}^W(M_t$ and $\Omega_t), a_{t+1|_W}^R(M_t$ are the assets left for period $t+1$ conditional on being in working stage in t in case of working or retiring at period t, respectively. $h_t^{W}{}_{|_W}(M_t,\Omega_t)$ is the amount of work conditional on being still in working stage. $a_{t+1|_R}^R(M_t,\Omega_t)$ is the asset left for period $t+1$ conditional on being retired in age t. $\mathbb{P}_t^R(M_t,\Omega_t)$ is the probability of choosing to be retired at period t. The individual solves

$$\begin{aligned}
\Omega_t^* &= argmax_{\Omega_t}\mathbb{E}\,|_{\theta_{t+1}}\, V_t\big(\pi_t(M_t,\Omega_t)\mid M_t\big) = \\
&\quad \mathbb{P}_{l<t}^R\cdot u^t\big(a_{t+1|_R}^R(M_t,\Omega_t);a_{t|_R},B_{t|_R}\big) \\
&\quad +\big(1-\mathbb{P}_{l<t}^R\big)\cdot\Big[\mathbb{P}_t^R(M_t,\Omega_t)\cdot u^t\big(a_{t+1|_W}^R(M_t,\Omega_t),\mathbb{b}(t,\mathbb{s}(\{y_l\}_{l=1|_W}^{t-1}));a_{t|_W}\big) \\
&\quad +(1-\mathbb{P}_t^R(M_t,\Omega_t))\cdot u^t\big(a_{t+1|_W}^W(M_t,\Omega_t),h_t^{W}{}_{|_W}(M_t,\Omega_t);a_{t|_W}\big)\Big] \\
&\quad +\beta\mathbb{E}\,|_{\theta_{t+1}}\, V_{t+1}\big(\pi_{t+1}(M_{t+1},\Omega_{t+1}),\Omega_{t+1}\big)\mid M_{t+1}\big) \\
&s.t. \\
M_{t+1} &= \{\mathbb{P}_{l<t+1}^R,\{y_l\}_{l=1|_W}^t,a_{t+1|_W},a_{t+1|_R},B_{t+1|_R},\theta_{t+1},edu\} \\
\mathbb{P}_{l<t+1}^R &= \mathbb{P}_{l<t}^R+(1-\mathbb{P}_{l<t}^R)\cdot\mathbb{P}_t^R(M_t,\Omega_t) \\
\{y_l\}_{l=1|_W}^t &= \{\{y_l\}_{l=1}^{t-1},\exp(\mu(edu,t)+\theta_t)\cdot h_t^{W}{}_{|_W}(M_t,\Omega_t)\} \\
a_{t+1|_W} &= a_{t+1|_W}^W(M_t,\Omega_t) = (1+r)\Big(a_{t|_W}+y_t^W-\tau^{inc}(y_t^W)-\tau^{SS}(y_t^W)-c_t^{W}{}_{|_W}\Big), y_t^W = \\
&\quad \exp(\mu(edu,t)+\theta_t)\cdot h_t^{W}{}_{|_W}(M_t,\Omega_t) \\
a_{t+1|_R} &= \mathbb{P}_{l<t}^R\cdot a_{t+1|_R}^R(M_t,\Omega_t)+(1-\mathbb{P}_{l<t}^R)\cdot a_{t+1|_W}^R(M_t,\Omega_t), a_{t+1|_R}^R(M_t,\Omega_t) = \\
&\quad (1+r)\Big(a_{t|_R}+B_{t|_R}-c_t^{R}{}_{|_R}\Big), a_{t+1|_W}^R(M_t,\Omega_t) = (1+r)\Big(a_{t|_W}+\mathbb{b}(t,\mathbb{s}(\{y_l\}_{l=1|_W}^{t-1}))-c_t^{R}{}_{|_W}\Big) \\
B_{t+1|_R} &= \mathbb{P}_{l<t}^R\cdot B_{t|_R}+(1-\mathbb{P}_{l<t}^R)\cdot\mathbb{b}(t,\mathbb{s}(\{y_l\}_{l=1|_W}^{t-1})) \\
\theta_{t+1} &= \theta_t+\epsilon_{t+1},\ \epsilon_{t+1}\sim N(0,{\sigma_\epsilon}^2)
\end{aligned} \tag{52}$$

For $t = T_{FR}$, the state variables are $M_{T_{FR}} = \{\mathbb{P}_{l<T_{FR}}^R,\{y_l\}_{l=1}^{T_{FR}-1}{}_{|_W},a_{T_{FR}|_W},a_{T_{FR}|_R},B_{T_{FR}|_R}\}$, where $\mathbb{P}_{l<T_{FR}}^R$ is the probability of entering period T_{FR} while being retired (becoming retired at one time in the past). All the variables that are conditional on being in the working stage $(\{y_l\}_{l=1}^{T_{FR}-1}{}_{|_W}$ for history of earning at T_{FR} and $a_{T_{FR}|_W}$ for asset at T_{FR}) show state variables in T_{FR} conditional on entering period T_{FR} in working stage (not being retired). All the variables that are conditioned on being on retirement stage ($B_{T_{FR}|_R}$ for the yearly pension benefit at T_{FR} and $a_{T_{FR}|_R}$ for asset at T_{FR}) show state variables in T_{FR} conditional on entering period T_{FR} while in retirement stage (getting retired at $s < T_{FR}$ in the past).

The control variables are

$$\pi_{T_{FR}}(M_{T_{FR}};\Omega_{T_{FR}}) = \{a_{T_{FR}+1|_W}(M_{T_{FR}},\Omega_{T_{FR}}), a_{T_{FR}+1|_R}(M_{T_{FR}},\Omega_{T_{FR}})\}, \tag{53}$$

where $a_{T_{FR}+1|_W}(M_{T_{FR}}$is the asset left for period $T_{FR}+1$ conditional on being in working stage in $t = T_{FR}$ in case of working or retiring at period T_{FR}. $a_{T_{FR}+1|_R}(M_{T_{FR}},\Omega_{T_{FR}})$ is the asset left for period $T_{FR}+1$ conditional on being retired in age T_{FR}. The individual solves

$$\begin{aligned}
\Omega^*_{T_{FR}} &= argmax_{\Omega_{T_{FR}}} V_{T_{FR}}\big(\pi_{T_{FR}}(M_{T_{FR}},\Omega_{T_{FR}}) \mid M_{T_{FR}}\big) = \\
&\quad \mathbb{P}^R_{l<T_{FR}} \cdot u^{T_{FR}}\big(a_{T_{FR}+1|_R}(M_{T_{FR}},\Omega_{T_{FR}}); a_{T_{FR}|_R}, B_{T_{FR}|_R}\big) \\
&\quad + \big(1-\mathbb{P}^R_{l<T_{FR}}\big) \cdot u^{T_{FR}}\big(a_{T_{FR}+1|_W}(M_{T_{FR}},\Omega_{T_{FR}}), \mathbb{b}(T_{FR}, \mathbb{s}(\{y_l\}^{T_{FR}-1}_{l=1}{}_{|_W})); a_{T_{FR}|_W}\big) \\
&\quad + \beta V_{T_{FR}+1}\big(\pi_{T_{FR}+1}(M_{T_{FR}+1},\Omega_{T_{FR}+1}),\Omega_{T_{FR}+1}) \mid M_{T_{FR}+1}\big) \\
&\quad s.t. \\
M_{T_{FR}+1} &= \{a_{T_{FR}+1}, B_{T_{FR}+1}\} \\
a_{T_{FR}+1} &= \mathbb{P}^R_{l<T_{FR}} \cdot a_{T_{FR}+1|_R}(M_{T_{FR}},\Omega_{T_{FR}}) + (1-\mathbb{P}^R_{l<T_{FR}}) \cdot a_{T_{FR}+1|_W}(M_{T_{FR}},\Omega_{T_{FR}}), a_{T_{FR}+1|_R}(M_{T_{FR}},\Omega_{T_{FR}}) = \\
&\quad (1+r)\Big(a_{T_{FR}|_R} + B_{T_{FR}|_R} - c^R_{T_{FR}|_R}\Big), a_{T_{FR}+1|_W}(M_{T_{FR}},\Omega_{T_{FR}}) = \\
&\quad (1+r)\Big(a_{T_{FR}|_W} + \mathbb{b}(T_{FR}, \mathbb{s}(\{y_l\}^{T_{FR}-1}_{l=1}{}_{|_W})) - c^R_{T_{FR}|_W}\Big) \\
B_{T_{FR}+1|_R} &= \mathbb{P}^R_{l<T_{FR}} \cdot B_{T_{FR}|_R} + (1-\mathbb{P}^R_{l<T_{FR}}) \cdot \mathbb{b}(T_{FR}, \mathbb{s}(\{y_l\}^{T_{FR}-1}_{l=1}{}_{|_W}))
\end{aligned} \tag{54}$$

For $t > T_{FR}$, the state variables are $M_t = \{a_t, B_t\}$ and the control variables is $\{a_{t+1}\}$. The individual solves

$$\begin{aligned}
\Omega^*_t &= argmax_{\Omega_t} V_t\big(\pi_t(M_t,\Omega_t) \mid M_t\big) = \\
&\quad u^t\big(a_{t+1}(M_t,\Omega_t); a_t, B_t\big) \\
&\quad + \beta V_{t+1}\big(\pi_{t+1}(M_{t+1},\Omega_{t+1}),\Omega_{t+1}) \mid M_{t+1}\big) \\
&\quad s.t. \\
M_{t+1} &= \{a_{t+1}, B_{t+1}\} \\
a_{t+1} &= (1+r)\Big(a_t + B_t - c_t\Big) \\
B_{t+1|_R} &= B_t
\end{aligned} \tag{55}$$

B Counterfactual with Budget Neutrality

A key feature of the counterfactual exercises is that changing the history-dependent component of the AIME function alters the input into the PIA formula, and therefore the level of benefits that the system delivers. To isolate the pure effect of changing history-dependence—separating it from mechanical changes in the

generosity of benefits—we implement each counterfactual in a **budget-neutral** manner. Specifically, for each policy reform we scale benefits as

$$\tilde{b} = \lambda \times b(t_R, \{y_s\}_{s=1}^{t_R}),$$

where the scalar λ is chosen such that the present value of Social Security finances (payroll-tax revenues net of total benefit payments) remains equal to that of the status quo. Because behavioral responses to each reform affect earnings, retirement behavior, and therefore the system's finances, this policy-neutrality condition must be satisfied in equilibrium. Computationally, we determine the value of λ using a fixed-point algorithm that iterates on the benefit-scaling factor until the government's intertemporal budget constraint is restored.

Table 7: Budget-neutral counterfactual pension rules: counting top earnings years

	All	Education		Wage volatility		Work attachment	
		College	No college	High	Low	High	Low
Status quo: Counting top 35 years (levels)							
Mean(c_it)	28684.86	42711.32	25404.10	40509.49	16860.22	39925.01	17385.20
Mean(c_it|$working$)	24580.66	37054.29	21664.51	33880.53	14603.66	33367.80	15164.73
Mean(c_it|$retired$)	39646.40	57798.61	35395.19	60733.45	22176.61	59673.64	22674.10
sd(c_it)	25053.64	36626.64	20098.31	30362.74	7336.87	30455.65	8458.34
sd(c_it|$working$)	20860.33	31136.23	16289.89	25102.11	5802.59	25191.58	7072.42
sd(c_it|$retired$)	31253.11	44953.10	25238.99	35623.04	7814.37	35920.39	9115.71
Mean(b_t)	17527.64	19262.14	17121.42	22843.90	13123.33	22560.09	13262.83
sd(b_t)	7443.83	7016.87	7482.44	5716.04	5606.73	6084.49	5598.57
Mean(h_it|$Age < 70$)	2005.59	1984.59	2010.50	2131.33	1879.85	2155.10	1855.29
Mean(h_it|$62 <= Age < 70$)	924.85	952.56	918.36	1180.44	669.25	1187.01	661.29
Mean(h_it|$Age < 62$)	2221.74	2191.00	2228.93	2321.51	2121.97	2348.72	2094.09
Mean(Age_R)	67.11	67.09	67.11	68.69	65.53	68.55	65.67
Mean(Wel)	-0.32	-0.27	-0.33	-0.27	-0.37	-0.28	-0.37
Budget-neutral CF: Counting top 20 years (percent change w.r.t. status quo)							
Mean(c_it)	21.04	21.77	20.75	21.31	20.41	22.51	16.95
Mean(c_it|$working$)	-2.37	-7.38	-0.38	-3.25	3.24	-1.40	-1.34
Mean(c_it|$retired$)	51.37	62.23	47.25	46.17	45.65	45.36	45.45
sd(c_it)	42.56	49.41	39.61	47.26	59.06	45.77	49.33
sd(c_it|$working$)	-5.50	-3.09	-5.75	-4.00	2.72	-5.19	-8.45
sd(c_it|$retired$)	58.60	63.94	49.99	59.20	62.24	58.54	58.44
Mean(b_t)	-10.56	-12.99	-9.91	-21.96	3.68	-22.13	3.79
sd(b_t)	-36.54	-35.61	-36.72	-22.93	-27.83	-25.54	-26.67
Mean(h_it|$Age < 70$)	-5.54	-5.34	-5.59	-5.16	-5.98	-5.05	-6.19
Mean(h_it|$62 <= Age < 70$)	-18.74	-28.85	-16.29	-34.07	8.29	-33.79	8.36
Mean(h_it|$Age < 62$)	-4.44	-3.29	-4.71	-2.22	-6.88	-2.15	-7.11
Mean(Age_R)	-2.34	-2.30	-2.35	-4.09	-0.50	-4.35	-0.23
Mean(Wel)	1.21	3.46	0.79	1.37	1.10	0.49	1.87
Budget-neutral CF: Counting top 5 years (percent change w.r.t. status quo)							
Mean(c_it)	20.93	21.42	20.74	21.32	20.01	22.88	16.66
Mean(c_it|$working$)	-2.27	-7.61	-0.15	-3.11	3.25	-0.96	-1.36
Mean(c_it|$retired$)	49.92	60.58	45.90	44.93	43.80	44.62	43.78
sd(c_it)	41.99	48.76	39.14	46.54	57.38	45.27	48.23
sd(c_it|$working$)	-5.35	-3.12	-5.29	-3.88	2.95	-5.00	-8.20
sd(c_it|$retired$)	57.42	62.75	49.02	57.89	60.89	57.35	57.49
Mean(b_t)	-10.20	-14.72	-9.00	-25.25	10.14	-25.10	9.96
sd(b_t)	-44.90	-43.21	-45.04	-28.98	-33.79	-31.73	-33.84
Mean(h_it|$Age < 70$)	-6.42	-6.24	-6.47	-5.59	-7.37	-5.37	-7.64
Mean(h_it|$62 <= Age < 70$)	-23.75	-32.94	-21.52	-37.65	0.77	-37.43	0.96
Mean(h_it|$Age < 62$)	-4.98	-3.92	-5.23	-2.33	-7.89	-2.13	-8.18
Mean(Age_R)	-2.66	-2.61	-2.67	-4.40	-0.83	-4.65	-0.57
Mean(Wel)	1.10	3.37	0.66	1.37	0.90	0.40	1.59

Notes: The table reports means and standard deviations of consumption c_{it}, retirement benefits b_t, hours worked h_{it}, retirement age Age_R, and retirement utility Wel. Status-quo results (Top-35) are levels; counterfactuals show percentage deviations from Top-35. Education groups: edu=1 vs. edu=0. Wage-volatility and work-attachment groups split at medians.

Table 8: Budget-neutral counterfactual pension rule: lifetime-average earnings

	All	Education		Wage volatility		Work attachment	
		College	No college	High	Low	High	Low
Status quo: Counting top 35 years (levels)							
Mean(c$_i t$)	28684.86	42711.32	25404.10	40509.49	16860.22	39925.01	17385.20
Mean(c$_i t$|*working*)	24580.66	37054.29	21664.51	33880.53	14603.66	33367.80	15164.73
Mean(c$_i t$|*retired*)	39646.40	57798.61	35395.19	60733.45	22176.61	59673.64	22674.10
sd(c$_i t$)	25053.64	36626.64	20098.31	30362.74	7336.87	30455.65	8458.34
sd(c$_i t$|*working*)	20860.33	31136.23	16289.89	25102.11	5802.59	25191.58	7072.42
sd(c$_i t$|*retired*)	31253.11	44953.10	25238.99	35623.04	7814.37	35920.39	9115.71
Mean(b$_t$)	17527.64	19262.14	17121.42	22843.90	13123.33	22560.09	13262.83
sd(b$_t$)	7443.83	7016.87	7482.44	5716.04	5606.73	6084.49	5598.57
Mean(h$_i t$|$Age < 70$)	2005.59	1984.59	2010.50	2131.33	1879.85	2155.10	1855.29
Mean(h$_i t$|$62 <= Age < 70$)	924.85	952.56	918.36	1180.44	669.25	1187.01	661.29
Mean(h$_i t$|$Age < 62$)	2221.74	2191.00	2228.93	2321.51	2121.97	2348.72	2094.09
Mean(Age$_R$)	67.11	67.09	67.11	68.69	65.53	68.55	65.67
Mean(Wel)	-0.32	-0.27	-0.33	-0.27	-0.37	-0.28	-0.37
Budget-neutral CF: Lifetime-average earnings (percent change w.r.t. status quo)							
Mean(c$_i t$)	23.78	11.70	28.53	22.58	26.66	10.38	54.71
Mean(c$_i t$|*working*)	50.68	23.24	61.71	49.67	48.44	37.13	83.16
Mean(c$_i t$|*retired*)	-19.54	-9.30	-23.67	-23.23	-6.54	-33.46	10.27
sd(c$_i t$)	116.46	109.25	134.01	132.37	205.87	98.09	440.57
sd(c$_i t$|*working*)	200.21	184.04	235.97	218.88	350.60	176.29	644.93
sd(c$_i t$|*retired*)	-20.84	-18.36	-29.89	-18.77	57.62	-37.86	171.23
Mean(b$_t$)	1.45	0.09	1.78	7.03	-3.06	-1.73	4.84
sd(b$_t$)	5.51	-6.27	8.09	10.77	-22.85	16.22	12.90
Mean(h$_i t$|$Age < 70$)	9.92	6.36	10.74	15.57	3.51	17.12	1.51
Mean(h$_i t$|$62 <= Age < 70$)	-14.56	-71.30	-0.79	16.78	-69.84	19.69	-76.39
Mean(h$_i t$|$Age < 62$)	11.96	13.11	11.70	15.45	8.14	16.86	6.44
Mean(Age$_R$)	-2.36	-2.73	-2.28	-0.93	-3.87	-2.68	-2.04
Mean(Wel)	45.83	43.20	46.33	51.87	41.40	54.27	39.45

Notes: The table reports means and standard deviations of consumption c_{it}, retirement benefits b_t, hours worked h_{it}, retirement age Age_R, and retirement utility Wel. Status-quo results (Top-35) are levels; counterfactuals show percentage deviations from Top-35. Education groups: edu=1 vs. edu=0. Wage-volatility and work-attachment groups split at medians.