



ECONOMIC RESEARCH
FEDERAL RESERVE BANK OF ST. LOUIS
WORKING PAPER SERIES

Expectations on Wealth Returns: Implications for Labor Supply During the Retirement Boom

Authors	Serdar Birinci, Miguel Faria e Castro, and Kurt See
Working Paper Number	2025-031B
Revision Date	April 2026
Citable Link	https://doi.org/10.20955/wp.2025.031
Suggested Citation	Birinci, S., Faria e Castro, M., See, K., 2026; Expectations on Wealth Returns: Implications for Labor Supply During the Retirement Boom, Federal Reserve Bank of St. Louis Working Paper 2025-031. URL https://doi.org/10.20955/wp.2025.031

Federal Reserve Bank of St. Louis, Research Division, P.O. Box 442, St. Louis, MO 63166

The views expressed in this paper are those of the author(s) and do not necessarily reflect the views of the Federal Reserve System, the Board of Governors, or the regional Federal Reserve Banks. Federal Reserve Bank of St. Louis Working Papers are preliminary materials circulated to stimulate discussion and critical comment.

Expectations on Wealth Returns: Implications for Labor Supply During the Retirement Boom*

Serdar Birinci[†]
FRB St. Louis

Miguel Faria-e-Castro
FRB St. Louis

Kurt See
National University of
Singapore

April 2026

Abstract

We use an overlapping-generations model with incomplete markets and a frictional labor market to study how assumptions about agents' expectations of changes in returns to wealth affect labor supply and retirement decisions. Focusing on 2020–23, when returns fluctuated sharply and retirements rose above trend, we find that when individuals internalize the dependence of returns on wealth and view changes in returns as persistent, the model generates counterfactual labor-market outcomes. Retirements fall because expectations of persistently high returns boost labor supply, outweighing wealth effects, and the model predicts retirements concentrated among the very wealthy, contrary to the microdata.

Keywords: Labor Supply, Retirement, Heterogeneous Returns, Incomplete Markets
JEL Classification: E24, G11, J21, J22, J26

*This research was supported through computational resources provided by the Big-Tex High Performance Computing Group at the Federal Reserve Bank of Dallas. The views expressed here are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of St. Louis or the Federal Reserve System. Birinci: serdar.birinci@stls.frb.org, Faria-e-Castro: miguel.fariaecastro@stls.frb.org, See: kurtgerrardsee@gmail.com.

[†]Corresponding author. Address: Broadway and Locust, 1 Federal Reserve Bank Plaza, St. Louis, MO 63102.

1 Introduction

A growing literature in macroeconomics and household finance documents that returns on wealth vary considerably across households (Bach et al., 2020; Fagereng et al., 2020; Smith et al., 2022; Ozkan et al., 2023), and that differences in returns to wealth are an important feature that can help heterogeneous-agent models match empirically observed levels of wealth inequality (Benhabib and Bisin, 2018; Benhabib et al., 2019; Xavier, 2021). However, the implications of heterogeneous asset returns for labor supply and retirement decisions, and how agents form expectations about changes in those returns, have not been widely studied. The post-pandemic period from 2020 to 2023, marked by volatile and heterogeneous returns to wealth alongside a surge in retirements, provides a unique setting to study these questions.

This paper makes a concise theoretical and quantitative contribution. First, we study how expectations about the future path of asset returns affect labor supply, with a focus on retirement decisions. Second, we quantitatively assess alternative assumptions about these expectations and determine which best account for the observed aggregate and distributional changes in retirement patterns during the period of interest. To do so, we employ a heterogeneous-agent overlapping-generations (OLG) model with incomplete markets, in which asset returns depend on wealth and age, and a frictional labor market. We show that expectations about the persistence of return fluctuations are central to labor supply responses. When agents view higher returns as *transitory*, wealth effects dominate and labor supply contracts. In contrast, if higher returns are seen as *permanent*, agents expand labor supply to accumulate more wealth. Thus, alternative expectations about changes in asset returns lead to opposite labor supply outcomes.

Our model incorporates realistic life-cycle income dynamics, unemployment insurance (UI), and social security (SS) benefits. We calibrate the model’s steady state to the U.S. in 2019 and study transitions for 2020–23—a period of great interest given the large observed changes in both wealth returns and aggregate labor supply, especially among older workers. Our analysis feeds in exogenous shocks that could help explain the observed changes in retirement. These shocks capture (i) the heterogeneous movements in returns to wealth, (ii) the heterogeneous rise in job-separation rates across the labor income distribution, (iii) economic impact payment programs, (iv) the expansion of UI, and (v) the increase in mortality risk, which was steeper for older people. Birinci et al. (2025) show that when agents expect changes in returns along the transition to be transitory, the model captures the observed changes in aggregate labor market moments during 2020–2023 and that its predictions are consistent with observed patterns of retirement across wealth and income levels in the micro-data. Birinci et al. (2025) also find that changes in asset returns explain one-fifth of excess

retirements in 2022–23.

In this paper, we contrast this result and show that when agents expect changes in asset returns to be permanent, the model predicts the opposite for retirement decisions. Instead of increasing the fraction of retirees in the economy, elevated returns encourage greater labor supply, as individuals work more to accumulate wealth faster. This “substitution effect” more than offsets the wealth effects of elevated returns on labor supply. We show that this prediction is counterfactual, at least during the period of analysis. Beyond the opposite aggregate response, the predicted distribution of new retirees—contrary to the observed patterns in the microdata—is composed almost exclusively of individuals in the top quintile of the asset distribution, as other agents find it worthwhile to keep working and accumulate wealth. Overall, our results serve as a cautionary note for models incorporating heterogeneous and time-varying returns on wealth, since the assumption about expectations regarding asset returns greatly matters for labor supply.

The paper proceeds as follows: Section 2 describes the model and its calibration, which are based in Birinci et al. (2025). Section 3 describes the main quantitative experiment, and Section 4 discusses the role of temporary vs. permanent changes in asset returns, the main result in this paper; Section 5 concludes.

2 Model

The model we use to answer the main research question of this paper is based on the framework developed and calibrated in Birinci et al. (2025). In this section, we present the model, and then discuss its calibration and quantification. The model combines a partial-equilibrium heterogeneous-agents incomplete markets OLG setting with a frictional labor market that captures the joint distribution of retirement, income, and wealth in 2019 and can be used to quantify contributions of various factors to the rise in the retired share during 2020–2023.

2.1 Environment

Time is discrete and infinite. There is a stationary mass of overlapping generations of agents, who are born at age 25 and die with certainty at age 90. These agents are characterized by the following state variables: physical age in months $j \in \{25 \times 12, 25 \times 12 + 1, \dots, 90 \times 12\}$, age of retirement in years $k \in \emptyset \cup \{62, \dots, 70\}$ (where \emptyset denotes that the agent has not yet retired), wealth $a \in [-\underline{a}, \infty)$, labor force status $\ell \in \{E, U, N\}$ (employed, unemployed, out of the labor force, respectively), and wage $w \in \mathbb{R}^+$ for those employed or last wage for those who are not currently employed. Agents face a conditional probability of death that depends on their age and employment status, $1 - \pi(j, \ell)$.

Flow utility is derived from the function $u(c, \ell, j) = \frac{c^{1-\sigma}}{1-\sigma} - \mathbb{I}[\ell = E]\phi^E(j) - \mathbb{I}[\ell = U]\phi^U(j)$,

where σ is the coefficient of relative risk aversion, $\phi^E(j)$ is the age-dependent disutility of working, and $\phi^U(j)$ is the disutility of unemployment. Agents can access a risk-free asset that yields return $r(a, j)$ on savings (i.e., $a \geq 0$) and a constant return r^b when borrowing (i.e., $a < 0$). Dependence of the return on savings on age and wealth is a tractable way of capturing portfolio heterogeneity with a single-asset model.

We model wage and labor income dynamics as in French (2005) and Blandin, Jones, and Yang (2023). Labor income has two components: first, there is an age-specific wage profile $\psi(j)$. Additionally, this profile is augmented by a stochastic wage component w that follows a persistent process $F(w'|w)$. Let the total labor income for a worked aged j be denoted by $\mathcal{W}_j = w_j \times \psi(j)$. Then:

$$\begin{aligned} \log \mathcal{W}_j &= \log \psi(j) + \log w_j \\ \log w_j &= \rho_w \log w_{j-1} + \varepsilon_j^w \\ \varepsilon_j^w &\sim \mathbb{N}(0, \sigma^\varepsilon), \text{ i.i.d.} \\ \log w_0 &\sim \mathbb{N}(0, \sigma^{w_0}), \end{aligned} \tag{1}$$

where $\log \psi(j) = \psi_0 + \psi_1 j + \psi_2 j^2$ is a quadratic function of age.

The timing within each period is as follows: at the beginning of each period, agents age linearly, the stochastic wage component w realizes, and agents choose their labor market status among the options that are currently available, given their states. This choice determines the relevant value function for the period, as well as the age of retirement state that we describe in more detail below. Agents then choose their consumption and savings. Finally, at the end of the period, mortality shocks and labor market shocks realize. Employed agents may lose their jobs with probability $\delta(w, j)$, and they can choose to become unemployed or non-participants at the start of the next period. If they keep their jobs, they learn their updated wage w in the next period and choose to stay in their jobs or exit to non-employment (either as unemployed or non-participant). Similarly, unemployed and non-participant agents may find a job with probability f and γf , respectively, at the end of the period. Then, at the start of the next period, they draw wage w from a distribution and then choose to become employed or non-employed (either as unemployed or non-participant). If they cannot find a job, they can switch between unemployment and non-participation freely.

We classify individuals 62 and older who are non-participants as retired. Age 62 is the minimum eligibility age for Social Security (SS) benefits in the U.S., making it the earliest point at which retirement meaningfully differs from non-participation.¹ SS benefits are

¹For tractability we are, in practice, conflating two decisions: that to stop participating after age 62, and that to start claiming SS benefits.

denoted by $\bar{y}^{ss}(w, j, k, \ell)$ and depend on the stochastic wage, age, age of retirement, and labor force status. We discuss this function in more detail in Section 2.4)

It is important to specify the law of motion for the age of retirement, k , which matters for the calculation of SS benefits: the age of retirement is $k = \emptyset$ until the first time an individual becomes a non-participant on or after the age of 62, at which point it becomes $k = \text{age}(j)$.² We set $k = 70$ for all individuals who did not retire before the age of 70.³ Formally, we write the law of motion for k as:

$$k' = \begin{cases} \emptyset, & \text{if } \text{age}(j') < 62 \vee (\ell' \neq N \wedge k = \emptyset) \\ \text{age}(j'), & \text{if } \text{age}(j') \in \{62, \dots, 69\} \wedge \ell' = N \wedge k = \emptyset \\ 70, & \text{if } \text{age}(j') \geq 70 \wedge k = \emptyset \\ k, & \text{if } k \neq \emptyset. \end{cases} \quad (2)$$

2.2 Value functions

As explained in the previous section, individuals' labor market status choices at the beginning of each period determine which value function they face in that period. We now describe each of the value functions in detail: employed, unemployed, and non-participant.

2.2.1 Employed individuals

We write the problem for an individual who is currently employed as:

$$\begin{aligned} V^E(j, k, a, w) = \max_{c, a'} & u(c, \ell = E, j) + \beta \pi(j, \ell) \left[\delta(w, j) \max\{V^U(j', k', a', w), V^N(j', k', a', w)\} \right. \\ & \left. + [1 - \delta(w, j)] \int_{w'} \max\{V^E(j', k', a', w'), V^U(j', k', a', w), V^N(j', k', a', w)\} dF(w'|w) \right] \\ \text{s.t. } & c + a' = y + a + T(y, j, a) \\ & a' \geq -\underline{a} \\ & y = w \times \psi(j) + \bar{y}^{ss}(w, j, k, \ell = E) + r(a, j) \times a, \end{aligned}$$

where k' evolves according to Equation (2), $j' = j + 1$, \underline{a} is the borrowing constraint, and $T(y, j, a)$ are government transfers, which depend on total income, age, and wealth. An employed individual has total income y , consisting of labor income $\mathcal{W}_j = w \times \psi(j)$,

² $\text{age}(j)$ is a function that converts physical age in months j to physical age in years, as we express k in years for computational reasons.

³As we explain in Section 2.4, premia for late retirement are maxed out at age 70 and so the age of retirement no longer matters for the calculation of benefits of those who have not yet retired at this point.

SS income $\bar{y}^{ss}(w, j, k, \ell = E)$, and capital income. At the end of the period, she may exogenously separate from her job with probability $\delta(w, j)$, which depends on the stochastic component w of the income process as well as her age j . If a separation occurs, she can choose to become unemployed or leave the labor force. If no exogenous separation takes place, she can choose to either stay in the current job or quit to non-employment ($\ell = U$ or $\ell = N$). Note that we still keep track of the last employment wage w for non-employed individuals, as it affects the amount of UI and SS income.

2.2.2 Unemployed individuals

Instead of labor income, unemployed individuals derive income from home production and UI. We allow the home production level $h(j)$ to depend on age and the UI replacement rate $b(w, j) \in [0, 1]$ to depend on last labor income \mathcal{W}_j , i.e., w and j . Thus, UI benefits for an unemployed with last wage w and age j are equal to $b(w, j) \times w \times \psi(j)$. The problem of this individual is:

$$\begin{aligned}
V^U(j, k, a, w) &= \max_{c, a'} u(c, \ell = U, j) + \beta \pi(j, \ell) \left[(1 - f) \max\{V^U(j', k', a', w), V^N(j', k', a', w)\} \right. \\
&\quad \left. + f \int_{w'} \max\{V^E(j', k', a', w'), V^U(j', k', a', w), V^N(j', k', a', w)\} dF(w'|w) \right] \\
\text{s.t. } c + a' &= y + a + T(y, j, a) \\
a' &\geq -\underline{a} \\
y &= b(w, j) \times w \times \psi(j) + h(j) + \bar{y}^{ss}(w, j, k, \ell = U) + r(a, j) \times a.
\end{aligned}$$

At the end of the period, an unemployed individual receives a job offer with probability f . If an offer is received, she draws a wage w' from F at the start of the next period and decides whether to become employed with labor income $w' \times \psi(j')$, remain unemployed, or leave the labor force. If no offer is received, she can still choose to leave the labor force.

2.2.3 Non-participant individuals

Individual who are out of the labor force derive income from home production $h(j)$, but are ineligible for UI benefits. To capture direct transitions from non-participation to employment in the data, we assume that a non-participant receives a job offer with probability $\gamma \times f$, with $\gamma < 1$. If an offer is received, they can choose to become employed, unemployed, or non-participant. If no offer is received, they can still choose to transition to unemployment.

The problem of a non-participant is given by:

$$\begin{aligned}
V^N(j, k, a, w) = \max_{c, a'} & u(c, \ell = N, j) + \beta\pi(j, \ell) \left[(1 - \gamma f) \max\{V^U(j', k', a', w), V^N(j', k', a', w)\} \right. \\
& \left. + \gamma f \int_{w'} \max\{V^E(j', k', a', w'), V^U(j', k', a', w), V^N(j', k', a', w)\} dF(w'|w) \right] \\
\text{s.t. } & c + a' = y + a + T(y, j, a) \\
& a' \geq -\underline{a}. \\
& y = h(j) + \bar{y}^{ss}(w, j, k, \ell = N) + r(a, j) \times a.
\end{aligned}$$

2.3 Death and birth

At age $j = 91$, all individuals die with probability 1 and obtain zero value, $V^\ell(j = 91, k, a, w) = 0, \forall(k, a, \ell, w)$. They are replaced with newborns, who enter the model at age $j = 25$, drawing their initial wealth from a distribution $Q(a)$ and initial wage w_0 from Equation (1). We assume that individuals enter the model as unemployed.

2.4 Calibration

We set some parameters externally, and internally calibrate most key parameters that relate to labor markets, demographics, and the income and wealth distributions. The goal of the model is to interpret labor market dynamics between 2020 and 2023, and so we take the model's steady state to be the U.S. economy in 2019. A period is a month and the numeraire is set to be 2019 dollars. We describe the external calibration, the internal calibration, and validation at the steady state.

2.4.1 Functional forms and external parameters

We assume that disutility functions for the employed and unemployed are linear in age, $\phi^\ell(j) = \phi_0^\ell + \phi_1^\ell \times j, \ell = E, U$. The functional form for the job-separation rate over labor income is:

$$\delta(w, j) = \bar{\delta} \times \exp \left[\eta_w^\delta \times \frac{w \times \psi(j) - \bar{W}}{\bar{W}} \right], \quad (3)$$

where \bar{W} is the average wage in the economy.⁴ The formula for the replacement rate is linear in labor income, $b(w, j) = b_0 + b_1 \times w \times \psi(j)$, and the value of home production is given by $h(j) = \bar{h}_0[1 + \bar{h}_1 \times \mathbb{I}[j \geq 62]]$. The fiscal transfer function $T(y, j, a)$ is set to zero at the stationary state, and described in detail in Section 3. Newly born agents draw their

⁴Shimer (2005) uses a similar functional form when defining how the aggregate job-separation rate changes with productivity over time.

initial wealth from the distribution $Q(a)$, which is log-normal with parameters (μ_a, σ_a) , and we choose the mean and standard deviation to match the wealth distribution of 25-year olds from the SCF. The resulting values are $\mu_a = \$8,685.32$ and $\sigma_a = \$39,597.24$. We set the coefficient of relative risk aversion σ to a standard value of 2.

We now discuss in detail the external calibration of the following key inputs: (i) the stochastic process and life-cycle profile for labor income \mathcal{W}_j ; (ii) the asset return function $r(a, j)$; (iii) the survival probabilities $\pi(j, \ell)$; (iv) the home production function $h(j)$; and (v) the SS income function $\bar{y}^{ss}(w, j, k, \ell)$.

Labor income process. Using monthly labor earnings data from the SIPP, we estimate the parameters of the life-cycle labor income process by closely following French (2005) and Blandin et al. (2023).⁵ The estimated persistence of the stochastic wage component is $\rho_w = 0.961$, with a standard deviation of $\sigma^\epsilon = 0.027$. The estimated dispersion of the initial wage distribution is $\sigma^{w_0} = 0.596$. For the life-cycle profile, we obtain $\psi_0 = 6.979$, $\psi_1 = 0.054$, $\psi_2 = -0.001$. Using these estimated parameters, we simulate the labor income process while incorporating life-cycle dynamics and unemployment risk, and derive an estimate for $\bar{\mathcal{W}}$, the average real labor income in the economy used as a parameter for $\delta(w, j)$. This approach yields $\bar{\mathcal{W}} = \$3,395$.

Asset returns. We parameterize the return function $r(a, j)$ using estimated returns on net worth. To do so, we follow an imputation procedure that combines the 2019 SCF with data on aggregate returns across different asset classes. This procedure assumes that the composition of asset portfolios in the 2019 SCF remains unchanged and that households are fully diversified within each asset class. We calculate returns only for changes in net worth that stem from asset classes for which data on realized returns are available.⁶

For calibration, we use the monthly return on net worth for each month of 2019. We restrict attention to households with a net worth-to-annual income ratio between 0 and 15 in 2019. This excludes households with negative net worth, since the model distinguishes between borrowing and saving rates, and also omits very wealthy households, as the model is not intended to capture extremely high levels of wealth. For this sample, we estimate:

$$r_{i,\tau}^{NW} = \beta_0 + \beta_1 \text{age}i + \beta_2 \text{age}i^2 + \beta_3 \text{age}i^3 + \beta_4 \left(\frac{NW_i}{12 \times \bar{\mathcal{W}}_{25y}} \right) + \varepsilon_i, \quad (4)$$

where $r_{i,\tau}^{NW}$ denotes the return on net worth in month τ of 2019, $\text{age}i$ is the individual's age in years, and $\left(\frac{NW_i}{12 \times \bar{\mathcal{W}}_{25y}} \right)$ represents the ratio of net worth to the average annual labor

⁵See Birinci et al. (2025) for details on the estimation.

⁶See Birinci et al. (2025) for details on the estimation.

income of a 25-year-old. We then take the average of all coefficients across the months of 2019. The borrowing rate is set equal to $\max_{a,j} r(a, j)$ plus a monthly spread of 0.005: the maximum return on savings to rule out arbitrage, along with an annualized borrowing spread of 6%.

Survival probabilities. To calibrate $\pi(j, \ell)$, we use the 2019 Actuarial Life Table from the Social Security Administration (SSA), which provides conditional death probabilities for males and females across age groups. We take an equally weighted average of men and women for each age group and convert these annual conditional death probabilities into monthly probabilities. In the steady state, there is no dependence on employment status ℓ .

Home production. We assume that income from home production is equal to a constant \bar{h}_0 for agents younger than 62, after which it rises to $1.15 \times \bar{h}_0$, i.e., $\bar{h}_1 = 0.15$. This value is drawn from [Dotsey et al. \(2014\)](#), who show that consumption of home goods among older workers begins to increase around age 60 and is approximately 25% higher by age 90. We adopt an average increase of 15% for individuals older than 62. The parameter \bar{h}_0 is calibrated internally in [Section 2.4.2](#).

Social Security income. To parameterize and calibrate the Social Security income function $\bar{y}^{ss}(w, j, k, \ell)$, we closely follow actual U.S. regulations, as in [French \(2005\)](#). This function consists of two components. The first is the Primary Insurance Amount (PIA), a piecewise concave function of a measure of past earnings, subject to a cap. To maintain tractability, we proxy past earnings using the product of the last realization of the stochastic wage component w before retirement and an average of the life-cycle component $\psi(j)$. The cap on this earnings measure, as well as the bend points that generate concavity, are set to their 2019 values.

The second component is a modifier that depends on the retirement age: individuals may start collecting Social Security benefits at age 62 but incur penalties if they retire before the full retirement age, which varies by birth cohort. We set the full retirement age (FRA) to 66, consistent with 2019. Individuals receive increased benefits if they retire after this age, up to age 70. We implement this modifier following the exact 2019 Social Security rules. We also incorporate the earnings test for individuals at or below the FRA, in line with current SSA regulations. Unemployed or non-participating agents do not face these penalties.⁷ A complete description of the Social Security income function, along with the calibration of its parameters, is provided in [Appendix A.1](#).

⁷We model this earnings test as a pure tax, consistent with findings from the empirical literature ([Gelber et al., 2020](#)). Additionally, we abstract from the personal tax treatment of Social Security benefits ([Jones and Li, 2018](#)).

Table 1: Internally calibrated parameters

Parameter	Value	Moment	Source	Data	Model
β	0.996	Fraction of population w/ NW ≤ 0 under 62	SCF	0.116	0.116
\underline{a}	-7894.46	Median credit limit/quarterly labor income	SCF	0.740	0.720
\bar{h}_0	1000.01	Retired share	CPS	0.213	0.230
b_0	0.774	Average UI replacement rate	SIPP	0.400	0.371
b_1	-1.25×10^{-4}	Q1/Q5 ratio of UI replacement rate	SIPP	2.015	1.789
ϕ_0^E	1.37×10^{-4}	Unemployment rate, all ages	CPS	0.030	0.069
ϕ_1^E	7.10×10^{-8}	Unemployment rate, over 55	CPS	0.027	0.018
ϕ_0^U	1.26×10^{-4}	LFPR, all ages	CPS	0.646	0.754
ϕ_1^U	5.03×10^{-7}	LFPR, over 55	CPS	0.389	0.464
γ	0.20	Ratio of monthly NE and RE rate to total monthly job-finding rate	CPS	0.202	0.256
f	0.361	Total monthly job-finding rate	CPS	0.439	0.457
$\bar{\delta}$	0.017	Total monthly job-separation rate	CPS	0.034	0.041
η_w^δ	-0.156	Q1/Q5 ratio of monthly E to U or N or R rate	CPS	2.889	2.322

Note: This table provides a list of internally calibrated parameters. The model frequency is monthly. SCF refers to the 2019 Survey of Consumer Finances. CPS refers to averages over the 12 months of 2019 for the Current Population Survey. All moments computed for a population over the age of 25, excluding armed forces, unless otherwise noted.

2.4.2 Internal calibration

We internally calibrate the remaining 13 parameters. The full set of parameters and respective targeted data moments are summarized in Table 1.

The discount factor β is chosen to match the share of individuals under age 62 with non-positive net worth in the SCF. The borrowing limit is set to target the median ratio of credit limits to quarterly labor income, following [Kaplan and Violante \(2014\)](#) and using SCF data. The level of home production income \bar{h}_0 is calibrated to match the fraction of retired individuals. This implies a home production share of GDP of approximately 13%, which is reasonably close to the BEA’s 2019 estimate of 21%. Finally, the slope of the UI replacement rate, b_1 , is chosen to match the Q1/Q5 ratio of replacement rates when individuals are ranked by their pre-unemployment labor income, as in [Birinci and See \(2023\)](#), while the level parameter b_0 is set to match the average replacement rate.

The level and slope of the employment disutility function are calibrated to match the overall unemployment rate and the unemployment rate for individuals aged 55 and above, respectively. The level and slope of the unemployment disutility function are determined similarly, but to match the labor force participation rate (LFPR) for the overall population and for those aged 55 and above.⁸ The parameter γ , which influences the job-finding probability of non-participants, is chosen to match the ratio of monthly non-participation (including retirement) to employment relative to the total monthly job-finding rate (from non-employment). The job-finding probability for the unemployed, f , is set to match the overall job-finding rate, defined as the sum of the average transition rates from unemploy-

⁸Because j denotes monthly age and consumption is measured in 2019 dollars, the estimated slope parameters of the disutility functions are small.

ment, non-participation, and retirement into employment. The level parameter of the job-separation rate, $\bar{\delta}$, is calibrated to match the monthly transition rate out of employment in a similar way. Finally, the slope parameter of the job-separation rate, η_w^δ , is chosen to match the Q1/Q5 ratio of the monthly job-separation rate in the data when employed individuals are ranked by their labor income, as in [Birinci and See \(2023\)](#).

3 Quantitative Experiment

Using the calibrated model, we now evaluate whether it can reproduce the observed changes in aggregate labor market outcomes between 2020 and 2023. We proceed in two steps. First, we describe how we measure the shocks in the data and translate them into model inputs. Second, we report the results from our baseline experiment, in which we feed all shocks into the model and assess whether it generates the empirical movements in the retired share, unemployment rate, and employment-to-population ratio. Because these aggregate dynamics are not targeted in the calibration, the model’s performance along these dimensions provides moments to validate the predictions of our model.

3.1 Shocks

Starting from the stationary state, we introduce five sequences of exogenous shocks: (i) a shock to the return on savings, varying by wealth and age; (ii) a shock to job-separation rates for employed workers, varying by labor income; (iii) a shock to lump-sum transfers, depending on age and total income; (iv) a shock to UI benefits received by the unemployed; and (v) a shock to mortality rates, varying by age and employment status. [Figure 1](#) shows the time paths of these shocks. Below, we describe in detail how each data-driven impulse is mapped into the model. At each period, agents treat certain shocks (e.g., job separation and mortality) as permanent changes to parameters, which requires fully solving the model and substantially increases computational costs.

Asset returns. Elevated asset returns during 2020–2023 may have generated wealth effects that both increased entries into retirement and reduced exits from retirement. To capture this channel, we re-estimate [Equation \(4\)](#) separately for each month from January 2020 through December 2023. Since estimated returns vary considerably at the monthly frequency, we smooth the coefficients using six-month moving averages and input the resulting series into the model as exogenous shocks. [Figure 1\(a\)](#) shows the mean and median paths of the implied monthly (annualized) return function. Both rise sharply early in the pandemic, with the mean and median exceeding 20% and 15% in 2021, respectively, before declining and turning negative in 2022 and early 2023, and then recovering to positive values later in 2023.

In the model, we incorporate this shock by replacing the stationary return function

Figure 1: Time series paths for exogenous shocks



Note: Panel (a) plots the mean and median paths of the estimated monthly return (annualized) function $r_t(a, j)$. We only plot the mean and median values at each month for expositional purposes. Panel (b) plots percent changes in the job-separation rate at each month $\delta_t(w, j)$ relative to the stationary state by quintiles of the labor income distribution. Panel (c) presents shocks to the economic impact payments $T_t(y, j, a)$ for eligible individuals. Panel (d) plots the shocks to UI benefit amount b_t . Panel (e) plots percent changes in mortality rates $\pi_t(j, l)$ at each month relative to the stationary state by age and employment status. Shocks in Panels (a) and (b) are smoothed by taking six-month moving averages.

$r(a, j)$ in the budget constraint (for agents with positive wealth) with a time-varying function $r_t(a, j)$. The key question concerns how this shock is implemented in the model. Specifically, we consider two alternative assumptions. First, we can assume that changes in the return function are unanticipated and transitory: agents expect the return on savings to revert to its stationary level in all future periods. Second, we can assume that agents perceive the changes in the return function as permanent throughout the transition.

As our baseline, we adopt the first assumption: these return shocks are unanticipated and treated as transitory, meaning agents expect the return on savings to revert to its stationary function in all future periods. This implementation is equivalent to a lump-sum windfall that does not alter savings decisions.⁹ It captures the unexpected nature of the large return fluctuations and prevents counterfactual changes in consumption and saving that might otherwise influence labor supply, such as inducing agents to work more and accumulate wealth in response to temporarily high returns.

⁹The implied lump-sum gain (or loss) is given by $a_t \times \frac{r_t(a_t, j_t) - r(a, j)}{1 + r(a, j)}$. This approach preserves decision distortions through wealth effects, as intended.

Importantly, in Section 4, we demonstrate that under the second assumption—where agents perceive the changes in the return function as permanent—the model generates counterfactual dynamics relative to the data, both at the aggregate level and across the cross-section.

Job-separation rates. The 2020–2023 period saw a sharp rise in aggregate job separations. Additionally, the COVID-19 shock disproportionately increased separation rates among low-income workers, while separations for higher-paying jobs rose less. Higher separations can reduce labor force participation, as unemployed workers are more likely to move into non-participation than employed workers (Hobijn and Şahin, 2021). We capture both the aggregate magnitude and cross-sectional heterogeneity by introducing into the model exogenous paths for job-separation rates that vary by labor-earnings quintile. Using the CPS, we compute the monthly job-separation rate as the fraction of employed individuals in one month who become non-employed in the next. We calculate this rate for each month from 2019 to 2023 separately by earnings quintile, with quintiles defined based on current labor income.¹⁰

For each month from 2020 to 2023, we then compute the percent change in separations relative to the 2019 average, again by earnings quintile. To reduce pronounced month-to-month volatility, we smooth these percent changes using six-month moving averages. Additionally, since the shocks become negligible after October 2021, we set them to zero beyond this date for tractability. Panel (b) of Figure 1 displays the resulting series, which we use as period-by-period shocks to the stationary job-separation rate $\delta(w, j)$.¹¹ These paths capture both the large spike in separations and the pronounced heterogeneity across the earnings distribution, with larger and more persistent effects for lower-quintile workers. Reflecting heightened uncertainty about the duration of the public health emergency and its labor market impacts, we assume that agents perceive these shocks as permanent at each point in time.

Economic impact payments. In the U.S., the COVID-19 pandemic triggered an unprecedented fiscal response, with substantial support targeted toward lower-income households (Faria-e-Castro, 2021). A key element of this response was the economic impact payments, consisting of three rounds of lump-sum transfers to eligible households. We model these payments as increases in government transfers $T(y, j, a)$ and directly map both their tim-

¹⁰At the start of the pandemic, temporary separations surged, but most affected workers were later recalled. Since the model does not distinguish temporary unemployment with recall from standard unemployment, we exclude temporary separations when computing monthly separation rates in the data.

¹¹For instance, in mid-2020 the job-separation rate for workers in the bottom two quintiles rises by over 60% relative to their stationary-state levels, while the top quintile rises by roughly 30%.

ing and dollar amounts into the model. For each round, households were ineligible if their adjusted gross income (AGI) exceeded \$80,000. IRS data from 2019 filed returns indicate that this threshold corresponds approximately to the 80th percentile of the AGI distribution. Accordingly, we set the eligibility cutoff at the 80th percentile of the stationary-state AGI distribution, defining AGI in the model as total income y .

The first round, under the CARES Act, occurred in March 2020 and provided \$1,200 per adult plus \$500 per child under 17. The second round, through the Tax Relief Act of 2020, was distributed in December 2020 and provided \$600 per adult plus \$600 per child under 17. The third round, initiated by the American Rescue Plan Act of 2021 in March 2021, provided \$1,400 per adult plus \$1,400 per dependent. Consequently, the presence of dependents could significantly increase the total transfer received by a household.


To incorporate these transfers into the model, we explicitly account for how household structure and the number of dependents vary with the age of the household head. Using the 2019 CPS ASEC, which reports the number of individuals under 18 by the head’s age, we construct an age-dependent transfer modifier. The procedure is detailed in Appendix A.2. Panel (c) shows the resulting effective transfers over time as a function of age. Because these payments were likely perceived as one-time events, we assume the associated shocks are unanticipated and occur for a single period only.

UI benefits. Another major component of pandemic-era household income support was the expansion of UI benefits. Additional payments of \$600 per week (on top of pre-pandemic benefits) were provided from March 2020 through June 2020, followed by \$300 per week from July 2020 to approximately June 2021.¹² We map these supplemental benefits into the model by assuming four weeks per month. The resulting UI benefit path is shown in Panel (d). Consistent with the data, these additional benefits are implemented as a lump-sum transfer to the unemployed: individuals receive their standard UI benefits plus the supplemental amount in months when it is provided. As with economic impact payments, we assume agents treat these shocks as temporary—unexpected and lasting only one period.

Mortality rates. The final shock involves changes in mortality rates, $\pi(j, \ell)$. The goal is not to exactly replicate realized mortality but to adjust agents’ perceived mortality risk during 2020. This channel is important because both perceived and actual increases in mortality effectively alter discounting and can influence labor force participation, especially among older agents. Moreover, unlike in the stationary state, we allow mortality to vary by labor force status, capturing the possibility that employment involving physical contact elevated COVID-19 transmission risk.

¹²States phased out benefits at different times; for simplicity, we end the program in June 2021.

Figure 2: Changes in aggregate labor market moments: Model vs data



Note: This figure plots the paths of the aggregate retired share (i.e., the fraction of retirees in the population) (Panel (a)), unemployment rate (Panel (b)), and employment-to-population ratio (Panel (c)) in the data and the model. We take six-month moving averages both in the data and in the model, and plot the percentage point deviation from the 2019 average in the data and stationary state of the model. Since the model is not designed to capture the sizable rise in temporary layoffs during COVID-19, our data benchmark for the unemployment rate is net of temporary unemployment, as classified in the CPS.

We model higher mortality by assuming that, at the start of 2020, agents perceive their mortality rates to rise to the levels reported in the SSA life tables. At the beginning of 2021, these rates are updated again, and they return to baseline in 2022. As with the labor market shocks, we assume agents perceive each change as permanent to reflect uncertainty about the duration of the public health emergency. We also incorporate an additional mortality increase for employed agents. To calibrate this differential, we combine estimates from [Eichenbaum et al. \(2021\)](#) with 2020 Census data: for employed workers over age 50, the probability of death rises by 2.2% more than for non-employed individuals, while for employed workers under 50 it rises by 0.08% more.¹³ Panel (e) shows the percent changes in mortality rates, by age and employment status, relative to the stationary state.

3.2 Aggregate labor market moments: model vs data

We now present the results from our main experiment. Starting from the stationary state, we introduce all five shocks into the model simultaneously and compare the resulting aggregate labor market dynamics over the transition with their empirical counterparts from 2020 to 2023. Figure 2 shows the data and model paths for the aggregate retired share (Panel (a)), unemployment rate (Panel (b)), and employment-to-population ratio (Panel (c)).¹⁴

For the retired share in the data, we examine deviations of the CPS fraction of retirees from its trend. We calculate six-month moving averages for both the data and the model and report percentage-point deviations from the 2019 average in the data and from the stationary state in the model. The model successfully captures both the magnitude and persistence of

¹³Details on obtaining these numbers are provided in Appendix A.3.

¹⁴In this exercise, agents who die are replaced by new 25-year-olds, keeping total population constant. We tested alternative assumptions (e.g., not replacing deceased agents) and found minimal quantitative impact.

the increase in the retired share: it rises in 2020, peaks at about 0.7 pp in 2021, and then gradually declines.

For the unemployment rate and employment-to-population ratio, we again compute six-month moving averages and report percentage-point deviations from their 2019 data averages or, in the model, from their stationary-state values. For unemployment, we note that the model is not designed to capture the large rise in temporary layoffs during COVID-19, so we compare it to a CPS measure that excludes temporary unemployment. The model accurately reproduces both the magnitude and timing of the increase in unemployment. For the employment-to-population ratio, the model slightly understates the decline by about 0.7 pp but closely captures the overall dynamics: the sharp initial drop followed by a gradual recovery. Both model and data indicate that the employment-to-population ratio remains roughly 0.5 pp below its 2019 level at the end of 2023. Overall, these results demonstrate that the model performs very well in replicating untargeted aggregate dynamics over 2020–2023.

4 Temporary vs. permanent changes in asset returns

In Section 3.2, we assume that agents expect future returns to follow the steady-state function $\bar{r}(a, j)$. Thus, changes to this function are viewed as temporary surprises, with elevated returns acting as windfalls that generate wealth effects without distorting labor supply or savings incentives.


As discussed in Section 3.1, an alternative assumption about the implementation of the return shock could be that agents perceive the changes in the return function to be permanent along the transition. That is, in each transition period t , agents observe the realized function $r_t(a, j)$ and expect the change in period t to be permanent. Panel (a) of Figure 3 shows the mean and median paths of the estimated return function $r_t(a, j)$.

In this section, we show that, under this alternative assumption, the model yields counterfactual labor market dynamics relative to the data, both at the aggregate level and across the cross-section.

4.1 Aggregate implications

Starting with the aggregate implications of this alternative assumption, Panel (b) of Figure 3 compares excess retirements in the data (orange line, as defined above) with changes in the retired share from our baseline exercise (blue line) and from the alternative exercise assuming permanent changes in returns (green line). While our baseline exercise matches the data very well, the alternative exercise generates counterfactual predictions for the retired share. At its peak in late 2021, the excess retired share (i.e., the rise in the retired share from its steady-state level) was close to 0.7 pp, whereas the alternative exercise predicts

Figure 3: Temporary vs. permanent changes in asset returns and retirement outcomes




Note: Panel (a) plots the mean and median paths of the estimated monthly return (annualized) function $r_t(a, j)$. We show only the mean and median values for expositional purposes. Panel (b) shows percentage-point deviations in the retired share (i.e., the fraction of retirees in the population) from the 2019 average in the data and the stationary state of the model. Panel (c) plots percent changes in average net worth in the model relative to the stationary state. Panels (b) and (c) present results from two model simulations: the baseline, where agents view changes in asset returns as temporary (blue lines), and the alternative, where changes in returns are perceived as permanent (green lines).

−1.2 pp. Furthermore, while the excess retired share starts falling in 2022, the alternative exercise predicts a counterfactual boom in retirements during this time. Importantly, the counterfactual patterns in retirement outcomes in the alternative exercise mirror the path of returns in Panel (a). As agents perceive increased returns in 2020–21 to be permanent, they choose to work more to exploit those elevated returns. These incentives reverse in 2022, when returns fall below the steady state, which explains the protracted retirement boom during this period. Panel (c) shows that the increase in average net worth is much larger in the alternative exercise, as agents decide to work and save more to exploit what they perceive as permanently elevated returns. In particular, average net worth rises by around 70% relative to the steady state, double that of the baseline exercise. This is purely the result of different behavior, as the same return shocks are fed to both experiments.

4.2 Micro implications

Turning to a distributional comparison, we examine the predictions of these two models in terms of retirement patterns across the wealth distribution using microdata from the Survey of Income and Program Participation (SIPP) 2020, 2021, and 2022 panels (covering monthly data between 2019 and 2021). In the data, for 2019, we define new retirees as individuals who report being labor force participants in one month and then indicate their retirement for the first time in the following month. We then categorize each new retiree into quintiles of the wealth distribution of employed individuals aged 62 to 72. This allows us to determine where each new retiree in 2019 falls within the wealth distribution of older employed workers eligible for retirement benefits, which is the key demographic for our analysis. We repeat this process for 2020 and 2021 to examine how retirement patterns based on wealth holdings

Figure 4: New retirees by wealth: Data vs model



Note: Panel (a) shows the fraction of new retirees across wealth quintiles using data from the SIPP, separately for those retiring in 2019 and 2020–2021. Panels (b) and (c) repeat the same calculations in the model, using results from the baseline and alternative exercises, respectively. The baseline exercise assumes that agents view changes in asset returns as temporary, whereas the alternative exercise assumes that changes in returns are perceived as permanent.

changed during the pandemic. Panel (a) of Figure 4 plots the fractions of new retirees during each period (2019 or 2020–21) who are in each wealth quintile. It shows that the distribution was relatively flat, slightly increasing with wealth in 2019. Importantly, this pattern remained similar in 2020–21.

Panel (b) plots the same moments in the model for the baseline exercise, showing that it does well in matching these two facts: (i) a relatively flat distribution that is slightly increasing with wealth in 2019, and (ii) no significant changes in this pattern along the transition. Finally, Panel (c) plots the same results for the alternative exercise. It matches the first fact by construction, as both models have the same steady state, but fails to match the second fact. In particular, it produces a steep relationship between wealth and retirement along the transition. It predicts no retirements in the bottom quintile of the wealth distribution, and nearly 70% of all new retirees originate from the top quintile. Low-wealth individuals find it worthwhile to remain in the labor force and accumulate wealth, leveraging the high returns. Only sufficiently wealthy agents find it worthwhile to retire during this period—at odds with the empirical evidence in Panel (a).

5 Conclusion

We use an OLG, incomplete-markets model with a frictional labor market and labor force participation decisions to analyze the effects of expectations over changes in asset returns on labor supply, with an emphasis on retirement. In the baseline exercise, where return fluctuations are seen as temporary, labor supply follows traditional wealth effects, decreasing when returns are high. Conversely, in an alternative exercise where return fluctuations are expected to be permanent, a substitution effect outweighs this wealth effect: individuals

increase labor supply when returns are high, prioritizing earnings and wealth accumulation to capitalize on higher returns.

We find that the alternative exercise generates predictions that are at odds with both the aggregate data and microdata for the 2020–23 period. It predicts a 1.2 pp decline in the retired share—while in reality it increased by 0.7 pp—and would have concentrated the distribution of new retirees among high-wealth individuals, which is at odds with the micro evidence. We conclude that expectations about returns are crucial in models with endogenous labor supply.

References

- BACH, L., L. E. CALVET, AND P. SODINI (2020): “Rich pickings? Risk, return, and skill in household wealth,” *American Economic Review*, 110, 2703–47.
- BENHABIB, J. AND A. BISIN (2018): “Skewed wealth distributions: Theory and empirics,” *Journal of Economic Literature*, 56, 1261–91.
- BENHABIB, J., A. BISIN, AND M. LUO (2019): “Wealth distribution and social mobility in the US: A quantitative approach,” *American Economic Review*, 109, 1623–1647.
- BIRINCI, S., M. FARIA-E-CASTRO, AND K. SEE (2025): “Dissecting the great retirement boom,” *Journal of Monetary Economics*, forthcoming.
- BIRINCI, S. AND K. SEE (2023): “Labor market responses to unemployment insurance: The role of heterogeneity,” *American Economic Journal: Macroeconomics*, 15, 388–430.
- BLANDIN, A., J. B. JONES, AND F. YANG (2023): “Marriage and work among prime-age men,” Working Papers 2313, Federal Reserve Bank of Dallas.
- DOTSEY, M., W. LI, AND F. YANG (2014): “CONSUMPTION AND TIME USE OVER THE LIFE CYCLE,” *International Economic Review*, 55, 665–692.
- EICHENBAUM, M. S., S. REBELO, AND M. TRABANDT (2021): “The macroeconomics of epidemics,” *The Review of Financial Studies*, 34, 5149–5187.
- FAGERENG, A., L. GUIISO, D. MALACRINO, AND L. PISTAFERRI (2020): “Heterogeneity and persistence in returns to wealth,” *Econometrica*, 88, 115–170.
- FARIA-E-CASTRO, M. (2021): “Fiscal policy during a pandemic,” *Journal of Economic Dynamics and Control*, 125.
- FRENCH, E. (2005): “The effects of health, wealth, and wages on labour supply and retirement behaviour,” *The Review of Economic Studies*, 72, 395–427.
- GELBER, A. M., D. JONES, AND D. W. SACKS (2020): “Estimating Adjustment Frictions Using Nonlinear Budget Sets: Method and Evidence from the Earnings Test,” *American Economic Journal: Applied Economics*, 12, 1–31.
- HOBIJN, B. AND A. ŞAHİN (2021): “Maximum employment and the participation cycle,” Working Paper 29222, National Bureau of Economic Research.
- JONES, J. B. AND Y. LI (2018): “The effects of collecting income taxes on Social Security benefits,”

- Journal of Public Economics*, 159, 128–145.
- KAPLAN, G. AND G. L. VIOLANTE (2014): “A model of the consumption response to fiscal stimulus payments,” *Econometrica*, 82, 1199–1239.
- OZKAN, S., J. HUBMER, S. SALGADO, AND E. HALVORSEN (2023): “Why are the wealthiest so wealthy? A longitudinal empirical investigation,” Working Papers 2023-004, Federal Reserve Bank of St. Louis.
- SHIMER, R. (2005): “The cyclical behavior of equilibrium unemployment and vacancies,” *American Economic Review*, 95, 25–49.
- SMITH, M., O. ZIDAR, AND E. ZWICK (2022): “Top wealth in America: New estimates under heterogeneous returns,” *The Quarterly Journal of Economics*, 138, 515–573.
- XAVIER, I. (2021): “Wealth inequality in the US: The role of heterogeneous returns,” Working paper, Federal Reserve Board of Governors.

A Model Appendix

A.1 SS income function

As in French (2005), we approximate the current SSA formula for SS benefits using a truncated linear function. SS benefits are computed as a product of two variables: the Primary Insurance Amount (PIA), which is a concave function of past earnings, and an adjustment factor that is based on the distance of one’s retirement age from the Full Retirement Age (FRA, also known as the Normal Retirement Age), i.e., the age at which a person can retire and claim full benefits. The PIA depends on the calendar year, while the FRA depends on a person’s birth year.

PIA. The main input to the computation of PIA is the average indexed monthly earnings (AIME). The AIME is calculated as the minimum between social security maximum taxable income \bar{y}^{\max} and an average of a worker’s 35-year highest indexed monthly labor earnings. We proxy for this average by taking the product between the last observation of the stochastic wage component w before retirement and the average of the lifecycle profile $\bar{\psi}$.¹⁵ Thus, the relevant measure of earnings for someone who decides to retire is the AIME, which is given by

$$AIME(w) = \min\{\bar{y}^{\max}, w \times \bar{\psi}\}.$$

Monthly social security maximum taxable income was $\bar{y}^{\max} = \$11,075$ in 2019. The PIA is equal to 90% of AIME up to a first bend point; plus 32% of AIME between the first point and a second bend point; plus 15% of AIME above the second bend point. Since the model steady state is calibrated to 2019, we use the 2019 bend points to calibrate the SS income function: \$960 and \$5,785, respectively. We use them to the model as parameters $\bar{y}_1 = \$960$ and $\bar{y}_2 = \$5,785$, respectively. Thus, the PIA formula in the model is:

$$PIA(w) = 0.9 \times \min\{\bar{y}_1, AIME(w)\} + 0.32 \times \max\{0, \min\{\bar{y}_2, AIME(w)\} - \bar{y}_1\} \\ + 0.15 \times \max\{0, AIME(w) - \bar{y}_2\}.$$

FRA modifier. The FRA depends on a person’s birth cohort. To keep the analysis tractable, we calibrate the FRA modifier to that of someone born between the years of 1943 and 1954, which is likely to represent the majority of normal-age retirees for the period we are focusing on. For someone born on these dates, the FRA is 66: this is the age at which

¹⁵If the worker has worked less than 35 years, the SS formula assigns zeros to the non-work years. We abstract from keeping track of the worker’s 35-year highest indexed monthly labor earnings for computational simplicity.

someone can retire and earn 100% of the benefits they are entitled to. This person can retire and start receiving benefits at any point after they turn 62, but the benefits will be scaled down by a penalty that is a function of the number of months between the retirement date and the date at which they reach 66. Similarly, this person can postpone retirement and increase their benefits by a factor that is a function of the same distance and capped at the age of 70. The SSA publishes formulas for these penalties and bonuses as a function of birth year and distance from the FRA. For early retirement, the penalty is given by

$$\text{penalty} = \begin{cases} \frac{5}{9} \times 0.01 \times 36 + \frac{5}{12} \times 0.01 \times (t - 36) & \text{if } t > 36 \\ \frac{5}{9} \times 0.01 \times t & \text{if } 0 \leq t \leq 36, \end{cases} \quad (5)$$

where t is the distance, in months, from the age of retirement to the FRA. The premium for delayed retirement is equal to 8%/12 per month past the FRA, and capped when the retiree reaches the age of 70.

In the model, we write the FRA modifier as:

$$\tau^{FRA}(k) = \begin{cases} 0 & \text{if } \text{age}(k) < 62 \\ -1.625929 + 0.005331 \times k & \text{if } \text{age}(k) \in [62, 66) \\ 1 & \text{if } \text{age}(k) = 66 \\ 1 + (0.08/12) \times (k - 66 \times 12) & \text{if } \text{age}(k) \in (66, 70) \\ 1 + (0.08/12) \times (70 \times 12 - 66 \times 12) & \text{if } \text{age}(k) \geq 70, \end{cases}$$

where age of retirement k is measured in months, and the formula for those aged between 62 and 66 is obtained by approximating the early retirement penalty in Equation (5) using a linear regression.

Benefit for non-employed. For agents who do not work, the SS benefit is then equal to the product of the PIA and the FRA modifier:

$$\bar{y}^{SS}(w, j, k, \ell) = PIA(w) \times \tau^{FRA}(k), \quad \ell = U, N.$$

Work penalty. As in the data, people may receive social security benefits while working, but these benefits may be reduced. In particular, benefits are reduced for people earning above a certain limit, before or on their FRA (which is 66 in our model). These annual income limits are known as the Earnings Test Annual Exempt Amount and were equal to \$17,640 and \$46,920 in 2019, respectively. For someone under the FRA, the SS benefit is reduced by \$1 for every \$2 earned above the limit. For individuals who will reach their FRA

in the same calendar year, the SS benefit is reduced by \$1 for every \$3 earned above the limit. While in reality this is defined at a monthly frequency, we assume that people at the NRA face the test, i.e., all those aged 66. We map these limits to the model as $\bar{y}_a = \$17,640/12$ and $\bar{y}_b = \$46,920/12$. For someone aged j , with the current wage w , the effective SS benefit is then computed as

$$\begin{aligned} \bar{y}^{SS}(w, j, k, E) = & \bar{y}^{SS}(w, j, k, N) - \mathbb{I}[j < 66] \times 0.5 \times \max\{w \times \psi(j) - \bar{y}_a, 0\} \\ & - \mathbb{I}[j = 66] \times 0.33 \times \max\{w \times \psi(j) - \bar{y}_b, 0\}. \end{aligned}$$

In reality, the earnings test is not a pure tax: it involves withholding benefits that are credited in the future in an actuarially fair manner (called “benefit enhancement”). We model it as a tax for two reasons. First, the empirical literature has found that people do react to the earnings test as if it were a tax, and that there is bunching at the earnings test kinks. [Gelber et al. \(2020\)](#) show this, and offer several potential explanations for why this may be the case: individuals may expect their life-span to be shorter than average, meaning that they will not fully enjoy the offsetting credits; individuals may be liquidity constrained or discount faster than average; finally, some individuals may not understand the benefit enhancement system. Second, the withholding and crediting system is cumbersome to model and would significantly complicate the model.

Note also that regulations do not count UI benefits as earnings. Finally, we abstract from taxation issues related to SS benefits ([Jones and Li, 2018](#)).

A.2 Economic impact payments

Here, we provide details on how we measure economic impact payments in the data and map them into our model as shocks.

There were three rounds of economic impact payments (EIP) after COVID-19. For all three rounds, transfer amounts include a supplement associated with the number of children under the age of 17 or number of dependents in the household. For simplicity, we treat all dependents as children under the age of 17. This supplement amount could be substantial, equating the size of the base transfer in the case of the second and third round of payments. This requires us to adjust transfer amounts based on the size of the household. To do this, we rely on data from the Census Bureau on the average number of people under and over age 18 per household, by the age of householder, for 2019.¹⁶ For each age group for the householder, we divide the average number of people under age 18 by the average number of people who are at least 18 years old. We use this ratio as a modifier for how much of the

¹⁶Please refer to America’s Families and Living Arrangements: 2019 from <https://www.census.gov/data/tables/2019/demo/families/cps-2019.html>.

Table 2: Effective transfers for each age group of householder

Age of householder	Modifier	1st round	2nd round	3rd round
25-29 years	0.34	1353.16	793.76	1769.89
30-34 years	0.61	1486.51	953.78	2126.70
35-39 years	0.78	1571.20	1055.41	2353.30
40-44 years	0.64	1502.36	972.80	2169.11
45-49 years	0.43	1399.79	849.72	1894.66
50-54 years	0.22	1296.05	725.23	1617.09
55-59 years	0.11	1241.44	659.69	1470.96
60-64 years	0.08	1222.83	637.36	1421.15
65-74 years	0.05	1209.94	621.89	1386.66
75 years and over	0.03	1198.20	607.81	1355.26

Note: This table provides a modifier (second column) for how much of the dependent supplement a householder of a certain age group (first column) should receive. Model counterparts of effective transfer amounts of economic impact payments from the first, second, and third rounds of payments are provided in the last three columns.

dependent supplement a householder of a certain age group receives. The 2019 dependent modifiers are provided in the second column of Table 2. The effective transfer per eligible individual is then the adult transfer plus dependent supplement times the modifier for that individual’s age.

First round. The first round of transfers was associated with the CARES Act and took place in March 2020. These transfers consisted of \$1,200 per person plus \$500 per child under 17. Using CPI deflators $P_{2020}^{2019} = 1.012$ and $P_{2021}^{2019} = 1.059$, we obtain the following amounts for adults and children:

$$T_{2020m3}^{adult} = 1200/1.012 = 1185.77$$

$$T_{2020m3}^{child} = 500/1.012 = 494.07.$$

The effective transfer is then computed as the adult transfer plus the relevant modifier times the dependent transfer. For example, for a household between 25-29 years of age, the effective transfer amounts from the first round is computed as $1185.77 + 494.07 \times 0.34 \simeq 1353.2$, which is shown in the third column of Table 2.

Second round. The second round of transfers was deployed in December 2020 as a part of the Tax Relief Act of 2020 and consisted of \$600 per person plus \$600 per child under the age of 17:

$$T_{2020m12}^{adult} = 600/1.012 = 592.89$$

$$T_{2020m12}^{child} = T_{2020m12}^{adult}.$$

Third round. The third round came in March 2021 with the American Rescue Plan and consisted of \$1,400 per person plus \$1,400 per dependent:

$$T_{2021m3}^{adult} = 1400/1.059 = 1322.00$$

$$T_{2021m3}^{child} = T_{2021m3}^{adult}.$$

A.3 Impact of employment on mortality rates

In this Appendix, we explain how we discipline the mortality rate shock such that it features higher death probability for employed relative to non-employed, capturing the potential increase in COVID-19 transmission rates from working in activities that involve physical contact.

First, we describe the key data inputs to our calculations. Eichenbaum et al. (2021) calibrate an increase in probability of infection from work-related activities of 17 percent. This is not sufficient for our purposes, as we need to convert this into a probability of dying from infection, which may be different across age groups. For simplicity, we divide the population into those 49 years old and younger and those 50 years old and older. In the 2020 U.S. Census, 64.4% of the U.S. population was 49 years old and younger. From the Centers of Disease Control and Prevention, 6.32% of all COVID-related deaths were for people 49 years old and younger.¹⁷ Finally, the World Health Organization calculated the cumulative case fatality rate (CFR) from COVID-19 in the U.S. in 2020 to be 4.92% (i.e., the percentage of people who died conditional on infection, note that this is higher than the cumulative CFR of around 1% through 2025).¹⁸

Our goal is to compute the object $\Pr(\text{COVID death}|\text{age} \geq 50)$. This is equal to $\Pr(\text{COVID death} \& \text{age} \geq 50) / \Pr(\text{age} \geq 50)$. The denominator is equal to 0.356, from the Census data. Using Bayes' Theorem, we can write

$$\Pr(\text{age} \geq 50|\text{COVID death}) = \Pr(\text{COVID death}|\text{age} \geq 50) \times \frac{\Pr(\text{age} \geq 50)}{\Pr(\text{COVID death})}.$$

We can then rearrange and solve for our object of interest:

$$\begin{aligned} \Pr(\text{COVID death}|\text{age} \geq 50) &= \Pr(\text{age} \geq 50|\text{COVID death}) \times \frac{\Pr(\text{COVID death})}{\Pr(\text{age} \geq 50)} \\ &= (1 - 0.0632) \times \frac{0.0492}{1 - 0.644} = 0.1295. \end{aligned}$$

¹⁷See https://www.cdc.gov/nchs/nvss/vsrr/covid_weekly/index.htm.

¹⁸See <https://ourworldindata.org/grapher/covid-cfr-exemplars>.

Finally, we can infer the probability of COVID death for those under the age of 50 by solving:

$$\begin{aligned}\Pr(\text{COVID death}|\text{age} < 50) &= \frac{\Pr(\text{COVID death}) - \Pr(\text{COVID death}|\text{age} \geq 50) \times \Pr(\text{age} \geq 50)}{\Pr(\text{age} < 50)} \\ &= 0.0048.\end{aligned}$$

Thus, the added probability of dying given employment is equal to 0.17 times 0.1295 for those over the age of 50 and 0.17 times 0.0048 for those under the age of 50. Notice that we assume equal infection rates for both age groups, which is a reasonable assumption as 32% of all COVID-19 cases in the US were for people over the age of 50 as of 2023—a similar fraction to their share of the population.