Dissecting the Great Retirement Boom[⋆]

Serdar Birinci[a,*], Miguel Faria-e-Castro[a], Kurt See[b]

[a]*Federal Reserve Bank of St. Louis, United States*
[b]*National University of Singapore, Singapore*

Abstract

Between 2020 and 2023, the fraction of retirees in the working-age population in the U.S. increased above its pre-pandemic trend. Several explanations have been proposed to rationalize this gap, including increases in net worth, the deterioration of the labor market with higher job separations, the expansion of fiscal transfer programs, and higher mortality risk. We develop an incomplete markets, overlapping generations model with a frictional labor market to quantitatively study the interaction of these factors and decompose their contributions to the rise in retirements. We find that new retirements were concentrated at the bottom of the income distribution, and the most important factors driving the rise in retirements were higher job separations and the expansion of fiscal transfers. We show that our model's predictions on aggregate labor market moments and cross-sectional moments on retirement patterns across income and wealth distributions are in line with the data.

[⋆]We thank Sam Jordan-Wood for outstanding research assistance and Jennifer Bernstein for excellent editorial assistance. We also thank participants at many seminars and conferences for comments and suggestions. 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.

[*]Corresponding author

Email addresses: `serdar.birinci@stls.frb.org` (Serdar Birinci), `miguel.fariaecastro@stls.frb.org` (Miguel Faria-e-Castro), `kurtgerrardsee@gmail.com` (Kurt See)

1. Introduction

The rise in the fraction of retirees in the working-age population in the U.S. since the beginning of the COVID-19 pandemic has garnered attention from both researchers and policy-makers (Hobijn and Şahin, 2021; Montes et al., 2022). In late 2021, the fraction of retired individuals in the working-age population rose 0.7 percentage points (pp) over what the pre-pandemic trend predicts—close to 2 million excess retirements. This phenomenon slowed the recovery of the U.S. labor force participation rate (LFPR), which remained 0.8 pp below its pre-pandemic level in May 2024. Several factors, some of which have been individually studied, are natural candidates to explain this phenomenon: (i) wealth effects due to elevated returns on assets (French, 2005), (ii) poor labor market conditions due to higher job separations (Hobijn and Şahin, 2022), (iii) expansion of fiscal transfers (French and Jones, 2001), (iv) expansion of the unemployment insurance (UI) program (Veracierto, 2008), and (v) increased mortality risk (Blundell et al., 2016). In this paper, we develop a unified approach to quantitatively analyze the interaction of these factors and decompose their contributions to the rise in retirements. Our main finding is that initially higher job separations and the subsequent provision of economic impact payments were the key drivers of increased retirements, which predominantly came from low-income workers.

Our paper makes three contributions. First, we present novel empirical results regarding the relationship between retirement decisions, wealth, and labor income before and after the COVID-19 episode. Using microdata from the Survey of Income and Program Participation (SIPP), we find that, in 2019, the fraction of new retirees is only slightly increasing in wealth quintiles but strongly decreasing in income quintiles. Importantly, we also find that these distributional patterns are remarkably stable between 2020 and 2021. Overall, these observations are informative for the predictions of our quantitative model, as they suggest that increased retirements were not driven by wealthier individuals but by income-poor individuals.

Second, we construct a heterogeneous agents model that allows us to account for potential factors behind the rise in retirements. Our framework incorporates a frictional labor market in an otherwise standard incomplete markets, overlapping generations (OLG) model. Besides making a consumption/savings decision, agents also choose their employment and labor force participation status, endogenizing flows in and out of retirement. The model also features realistic life-cycle profiles for labor income, social security payments, heterogeneous returns on savings, and heterogeneous unemployment risk. Overall, the novelty of our framework is that it combines a search and matching model of the labor market with a lifecycle incomplete markets model, a serious modeling of retirement decisions, and various types of fiscal transfers. We calibrate this model to the U.S. economy in 2019, matching a series of moments related to the distributions of wealth and labor income, as well as labor market flows. We validate the predictions of this model at the stationary equilibrium against untargeted moments, showing, in particular, that it captures the shares of new retirees by wealth and income quintiles.

We use the model to quantitatively study recent labor market dynamics. This is important since, to the best of our knowledge, there is no high-frequency dataset that allows us to track monthly labor market flows and, at the same time, contain information on wealth, returns on wealth, eligibility and receipt of various fiscal transfers during the pandemic, and mortality outcomes. This makes it necessary to use a model to understand recent retirement dynamics. Our main exercise consists of feeding sequences of exogenous shocks that represent the five channels we focus on to the stationary state of the model. These shocks are measured from the data and mapped into the model without targeting any endogenous aggregate labor market moments or cross-sectional moments from the microdata during 2020-2023. These shocks capture (i) the heterogeneous movements in returns to wealth, (ii) the heterogeneous rise of job-separation rates across the labor income distribution, (iii) economic impact payment programs, (iv) expansion of UI, and (v) the increase in mortality risk that was steeper for older people.

Third, we use the model to decompose the importance of each channel between 2020 and 2023. We first demonstrate that the model captures well the changes in untargeted aggregate labor market moments in the data, such as excess retirements, the unemployment rate net of temporary unemployment, and the employment-to-population ratio. Next, our decomposition exercises reveal that three of the five channels we consider played an important role in explaining excess retirements, with higher job separations being a more important driver in 2020 and 2021 (explaining around 75% of the rise) and economic impact payments playing a larger role in 2022 and 2023 (explaining 64% and 81%, respectively). The rise in mortality risk attenuate the effects of the other forces, and is crucial to get the magnitudes right.

We also compare the cross-sectional predictions of the model along the transition to changes in relevant moments from the microdata between 2020-2023 relative to 2019. We find that the model is able to broadly account for the rise in the average wealth, the compression of the wealth distribution, changes in fractions of new retirees by wealth and income quintiles, and changes in monthly flow rates in and out of retirement. Importantly, as in the data before and after the pandemic, our model predicts that new retirees are typically income poor, but not necessarily wealth poor. We argue that this result is consistent with the predictions of our decomposition exercise, in that the increase in retirements did not come from relatively wealthy individuals, but from low-income individuals who experienced larger increases in job separations and were relatively more sensitive to fiscal transfers.

Related literature. This paper contributes to the literature on retirement patterns and economic decisions of retirees in terms of consumption and savings (French and Jones, 2001; De Nardi et al., 2010, 2016; Jones and Li, 2018) as well as labor supply (Cheng and French, 2000; Coronado and Perozek, 2003; Benson and French, 2011). Relative to this work, we develop an incomplete markets, OLG model with a frictional labor market, and a detailed social security system. This model allows us to analyze how changes in labor market frictions and fiscal transfers—that impact the magnitude of the surplus from employment relative to non-employment—affect retirement decisions in a realistic manner.

Our paper also contributes to a recent empirical literature that focuses on changes in labor market participation and retirement patterns after the pandemic (Hobijn and Şahin, 2021; Hobijn and Şahin, 2022; Nie and Yang, 2021; Faria-e-Castro, 2021b; Montes et al., 2022). These studies were very useful in guiding researchers and policy makers to understand underlying sources behind these patterns. Relative to this literature, we develop a unified approach using a structural model that allows us to study interactions of these potential sources and decompose their relative contribution to aggregate labor market moments. Importantly, we also compare predictions of our model against relevant moments from macro and micro data.

2. Excess retirements in the data

In this section, we describe empirical trends in the aggregate fraction of the population that is retired in the U.S. with a special focus on the 2020-23 period, and use microdata to study retirement patterns across the wealth and income distributions during the same period.

2.1. Aggregate trends

The U.S. LFPR experienced its largest drop on record at the onset of the COVID-19 pandemic in early 2020, falling from 63.3% in January 2020 to 60.1% in April 2020. While there was a quick rebound from this 50-year minimum, it has not fully recovered to its pre-pandemic levels: 62.5% as of May 2024. Most of this gap can be attributed to a persistent drop in the LFPR for those aged 55 and over (38.2% in May 2024 vs. 40.2% in January 2020), as the LFPR or prime-age workers has actually exceeded its pre-pandemic level. This pattern motivates us to focus on older workers.

Several studies have documented a large increase in the share of the population that is retired over the same period (Nie and Yang, 2021; Faria-e-Castro, 2021b). Figure 2.1(a) plots the retired share, measured as the fraction of individuals who report to be retired among all individuals (excluding those in armed forces) aged 16 and over, in the U.S. from 1995 to the end of 2023 using data from the Current Population Survey (CPS).[1] The retired share was roughly constant until the late 2000s, when it started growing at a roughly linear trend (dashed line), estimated between June 2008 and January 2020, the last full month before the effects of the pandemic were felt in the economy. The rise in the retired share is plausibly related to demographic factors: 2008 was the first year in which Baby Boomers became eligible to retire and collect Social Security benefits. There is a significant gap between the linear trend and the actual retired share between 2020 and 2023, plotted in Panel (b): the retired share increased by 0.7 pp above the trend in late 2021. This gap corresponds to close to 2 million people who were retired beyond what the pre-pandemic trend implies.[2] We refer to this gap as "excess retirements," and analyzing its drivers is the main focus of this paper.

2.2. Aggregate shocks and their effects on retirement

The 2020–23 period was marked by a public health emergency that triggered very large fluctuations in the U.S. economy, driven by both individuals' optimization behavior and by an unprecedented macroeconomic policy response.

[1] Appendix A.1 has details on the construction of the data and shows that our measurement is robust to alternative definitions of retirement.

[2] Other filters, such as the one proposed by Hamilton (2018), the HP filter, and other deterministic trends, also generate above-trend increases of similar magnitudes in 2020.

Figure 2.1: Excess retirements between 2020 and 2023

(a) Retired share and linear trend (b) Deviations from trend

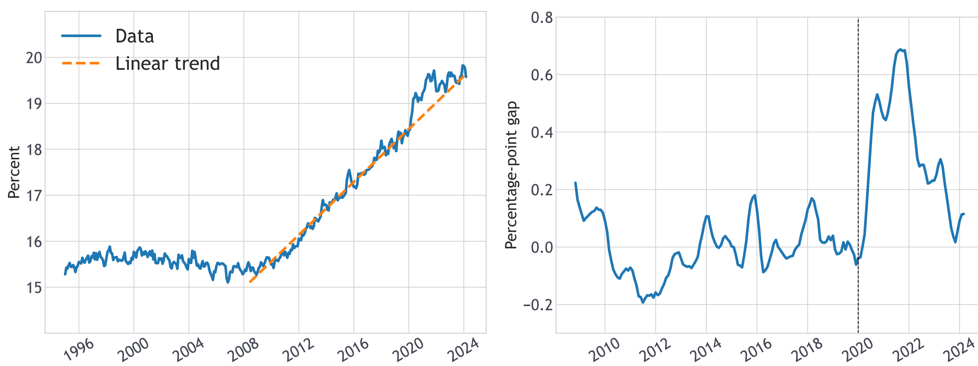

Note: Panel (a) plots the retired share in the U.S. which we calculate as the fraction of individuals who report to be retired in the CPS among all individuals aged 16 and over. The linear trend is estimated between June 2008 and January 2020. Panel (b) plots 6-month moving averages of deviations from trend.

This emergency consisted of an airborne respiratory disease that disproportionately put older people at risk of serious illness and death. According to the Centers for Disease Control and Prevention, 94% of all COVID-19-related deaths occurred among people aged 50 and older.[3] The heightened risk to older individuals plausibly influenced their labor force participation decisions, especially for those in occupations involving frequent physical contact. This motivates us to include mortality shocks in our model.

The early phase of the COVID-19 pandemic saw the highest unemployment rate in the postwar era—14.8% in April 2020. It is well documented that labor market conditions significantly affect labor force participation decisions (Hobijn and Şahin, 2021), which motivates us to include the state of the labor market as a potential driver of excess retirements in our model.

The swift economic recovery beginning in late 2020, combined with large-scale fiscal and monetary interventions, led to historically large returns across a variety of asset classes widely held by U.S. households.[4] Older households, being closer to retirement, tend to be wealthier and thus more likely to benefit from such returns. Moreover, the literature finds that households approaching retirement are more sensitive to wealth effects arising from changes in

[3]See `https://www.cdc.gov/nchs/nvss/vsrr/covid_weekly/index.htm`.

[4]See Appendix B.2 for time series of cumulative real returns for various asset classes during this period.

asset valuations (Cheng and French, 2000; French, 2005). We therefore consider fluctuations in asset returns as a potential source of wealth effects that could have contributed to the increase in retirements.

Finally, the fiscal policy response to the public health and economic emergency included household transfer programs that were large by historical standards. These took the form of three rounds of economic impact payments and a major expansion of unemployment insurance benefits. Such programs may have generated income and wealth effects that plausibly influenced individuals' labor force participation decisions (French and Jones, 2001; Veracierto, 2008).

2.3. Micro patterns

A key starting point to understanding the causes of this gap is identifying the worker groups where retirements were concentrated. In particular, given the potential effects of labor market disruptions, fiscal transfers, and asset returns on labor force participation decisions that we discussed above, we examine how retirement varied across both the wealth and income distributions. Later, we use these findings to validate model predictions.

As the CPS does not provide information on wealth holdings, we use data from the 2020, 2021, and 2022 panels (covering data from all months between 2019 and 2021) of the SIPP, which provide information on employment status, wealth, and labor income.[5] Our measure of wealth is total household net worth, while labor income is the total wages and salaries from all jobs.

Using this data, we identify new retirees in 2019 as those who report being in the labor force in a month in 2019 and report being retired for the first time in the following month. We then assign each new retiree to quintiles of the wealth distribution of employed individuals aged between 62 and 72. This allows us to calculate where each new retiree in 2019 sit within the wealth distribution of older employed workers eligible for retirement benefits—the

[5] We use CPS excess retirements as our baseline estimate for two reasons. First, monthly transition rates between employment statuses are underestimated in the SIPP relative to the CPS (Krusell et al., 2017; Birinci and See, 2023). Second, the most recent SIPP (2022) covers the reference period until December 2021, preventing us from studying aggregate retirement dynamics after 2021. Despite these limitations, the rise in the retired share in 2020-21 is also observed in the SIPP. This allows us to analyze the underlying retirement patterns across the wealth and income distributions during this period.

Figure 2.2: Retirement patterns in the micro data

(a) New retirees by wealth (b) New retirees by labor income

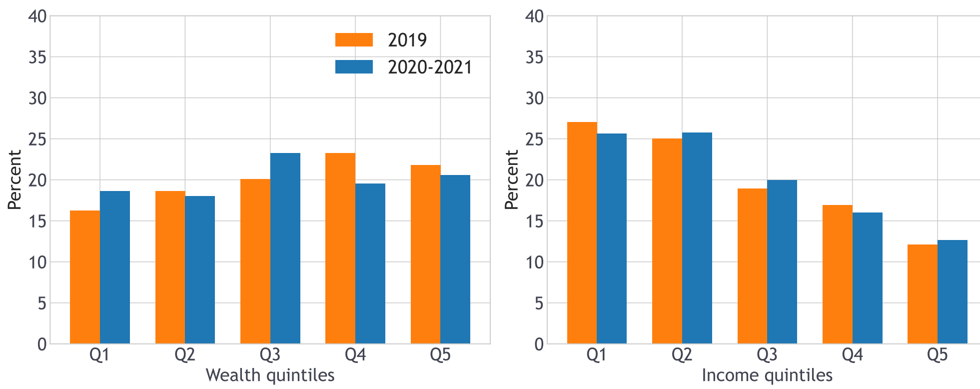

Note: Panel (a) shows the fraction of new retirees across wealth quintiles, separately for those retiring in 2019 and 2020–2021, using SIPP data. Panel (b) repeats this for labor income.

relevant demographic for our analysis. We then recompute the same moments between 2020 and 2021 to understand how retirement patterns by wealth holdings evolved during the pandemic.[6]

Figure 2.2(a) plots the fractions of new retirees during each period (2019 or 2020-21) who are in each wealth quintile. In 2019, the fraction of new retirees is slightly increasing in wealth quintiles, suggesting that new retirees are relatively wealthier, even though this relation is weak. Importantly, we find that this relationship remained mostly unchanged in 2020-21 relative to 2019. In other words, we find that the increase in retirements during the pandemic does not seem to be driven by wealthier people.

Figure 2.2(b) repeats the same exercise for labor income, using the distribution of labor income for those who are employed and aged between 62 and 72. For new retirees, labor income refers to earnings prior to retirement. In 2019, we find that new retirees typically have lower incomes. As with wealth, this pattern also changes little during 2020-21. Thus, most retirements in 2020-21 were still drawn from lower quintiles of the income distribution.

In sum, new retirees have lower income and are slightly wealthier relative to the employed workers at the age of retirement. This relationship did not change much during the pandemic.

[6] Appendix A.2 provides details on the data and construction of these moments.

3. Model

We now present a decision model of retirement that captures the joint distribution of retirement, income, and wealth in 2019. We combine a partial-equilibrium heterogeneous-agents incomplete markets OLG model with a frictional labor market to quantify contributions of various factors to the rise in the retired share between 2020 and 2023.

Environment. Time is discrete and infinite. The economy is populated by a stationary mass of overlapping generations of agents. Agents are born at age 25 and die with certainty at age 90. They are indexed by five 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 no retirement), wealth $a \in [-\underline{a}, \infty)$, employment status $\ell \in \{E, U, N\}$ (employed, unemployed, out of the labor force), and wage $w \in \mathbb{R}^+$ if employed or last wage if not employed. Additionally, they face an age- and employment-status-dependent probability of death, $1 - \pi(j, \ell)$.

Preferences are given by $u(c, \ell, j) = \frac{c^{1-\sigma}}{1-\sigma} - \mathbb{I}[\ell = E]\phi^E(j) - \mathbb{I}[\ell = U]\phi^U(j)$, where σ is the elasticity of intertemporal substitution, $\phi^E(j)$ is the disutility of working, and $\phi^U(j)$ is the disutility of looking for a job while unemployed. There is a risk-free asset that pays return $r(a, j)$ on savings ($a \geq 0$) and r^b on borrowings ($a < 0$). This is a single-asset model where the rate of return depends on the level of wealth and age: a tractable way of capturing portfolio heterogeneity across the wealth distribution.

Labor income depends on a stochastic wage w that evolves according to a persistent process $F(w'|w)$, and an age-specific profile $\psi(j)$. We follow French (2005) and Blandin, Jones and Yang (2023) in modeling income dynamics. Letting $\mathcal{W}_j = w_j \times \psi(j)$ denote the actual income of a worker aged j:

$$\begin{aligned} \log \mathcal{W}_j &= \log \psi(j) + \log w_j \\ \log w_j &= \rho_w \log w_{j-1} + \varepsilon_j^w \\ \varepsilon_j^w &\sim \mathbb{N}(0, \sigma^\epsilon), \text{i.i.d.} \\ \log w_0 &\sim \mathbb{N}(0, \sigma^{w_0}), \end{aligned} \tag{3.1}$$

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

At the end of each period, mortality shocks and labor market shocks, i.e., job separation and job-finding shocks, realize. At the beginning of the next period, agents age, the stochastic wage component w realizes, and then they choose their labor market status between available options. Their choice of labor market status determines the value function they experience, and, together with their age, also determines their age of retirement k as per Equation (3.2). Employed agents may lose their jobs with probability $\delta(w, j)$ at the end of the period, and they can choose to become unemployed or non-participant 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. Once labor force status decisions have been made, agents choose consumption and savings.

In the model, we classify individuals 62 and older who are out of the labor force as retired.[7] 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.[8] SS benefits are denoted by $\bar{y}^{ss}(w, j, k, \ell)$ and depend on the stochastic wage, age, age of retirement, and labor force status.

It is useful, at this point, 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 $k = \text{age}(j)$. Here, $\text{age}(j)$ is a function that converts physical age in months j to physical age in years, as we express k in years for computational reasons. We set $k = 70$ for all individuals who did not retire before the age of 70.[9] Formally, we can write:

[7] We have experimented with a stricter definition of retirement where we also require that agents never come back to the labor force to be considered as retired. Our quantitative results barely change under this alternative definition.

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

[9] As we explain in Section 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.

$$k' = \begin{cases} \emptyset, & \text{if } \text{age}(j') < 62 \vee (\ell' \neq N \wedge k = \emptyset) \\ \text{age}(j'), & \text{if } \text{age}(j') \in \{62, \ldots, 69\} \wedge \ell' = N \wedge k = \emptyset \\ 70, & \text{if } \text{age}(j') \geq 70 \wedge k = \emptyset \\ k, & \text{if } k \neq \emptyset. \end{cases} \tag{3.2}$$

Employed. The problem for an employed individual is given by:

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

where k' evolves according to Equation (3.2), $j' = j+1$, $\underline{a}$ is the borrowing constraint, and $T(y,j,a)$ are government transfers, which depend on total income, age, and wealth. An employed agent has total income y, consisting of labor income $\mathcal{W}_j = w \times \psi(j)$, SS income $\bar{y}^{ss}(w,j,k,\ell=E)$ (details of which are discussed in Section 4), 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 $\ell = N$). We note that, when individuals are non-employed, we still keep track of the last employment wage w as it affects the amount of UI and SS income.

Unemployed. Instead of labor income, unemployed agents 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 $b(w,j) \times w \times \psi(j)$. The problem of this agent is:

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

At the end of the period, an unemployed agent 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.

Non-participant. Agents who are out of the labor force receive 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 become unemployed. The problem of a non-participant is given by:

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

Death and birth. At age $j = 91$, all agents 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 (3.1). We assume that agents enter the model as unemployed.

4. Calibration

Our calibration strategy sets some parameters externally while internally calibrating most to match key moments related to labor market and demographic outcomes, as well as income and wealth distributions. Since we use our model to understand labor market dynamics between 2020–2023, we interpret the model's stationary state to be the U.S. economy at the end of 2019. A period is a month and the numeraire is set to be 2019 dollars.

4.1. Functional forms and external parameters

We assume that disutility functions for the employed and unemployed depend linearly with the individual's age, $\phi^\ell(j) = \phi_0^\ell + \phi_1^\ell \times j, \ell = E, U$. The job-separation rate varies with the labor income of the worker according to

$$\delta(w, j) = \bar{\delta} \times \exp\left[\eta_w^\delta \times \frac{w \times \psi(j) - \bar{\mathcal{W}}}{\bar{\mathcal{W}}}\right], \tag{4.1}$$

where $\bar{\mathcal{W}}$ is the average labor income in the economy. Shimer (2005) uses a similar functional form when defining how the aggregate job-separation rate changes with productivity over time. 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 5. The distribution of wealth for the newborn $Q(a)$ is log-normal with parameters (μ_a, σ_a); 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 also set the coefficient of relative risk aversion σ to 2, a standard value in this class of models.

Next, we describe in detail how we calibrate 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 data on labor earnings from the SIPP, we estimate the parameters of the life-cycle labor income process by closely following French (2005) and Blandin et al. (2023). Appendix B.1 provides details on the estimation. The estimated persistence for the stochastic wage component is $\rho_w = 0.961$, with a standard deviation of $\sigma^{\epsilon} = 0.027$. The estimated dispersion for the distribution of initial wage draws is $\sigma^{w_0} = 0.596$. For the life-cycle profile, we estimate $\psi_0 = 6.979, \psi_1 = 0.054, \psi_2 = -0.001$. With the estimated parameters, we simulate the labor income process taking into account life-cycle dynamics and unemployment risk, and obtain an estimate for $\bar{\mathcal{W}}$, the average real labor income in the economy that is used as a parameter for $\delta(w, j)$.[10] This procedure yields $\bar{\mathcal{W}} = \$3,395$.

Asset returns. We parametrize the return function $r(a, j)$ using estimated returns on net worth. To this end, we follow the imputation process that combines the 2019 SCF with data on aggregate returns for different asset classes. This imputation process assumes that the composition of asset portfolios in the 2019 SCF remains constant, and that households are perfectly diversified within each asset class. We compute returns only for changes in net worth that arise from asset classes for which we observe data on realized returns.[11]

For calibration purposes, we consider the monthly return on net worth for each month in 2019. We focus on households with a ratio of net worth to annual income between 0 and 15 in 2019. This excludes households with negative net worth, as our model differentiates between borrowing and saving rates. It also excludes the very wealthy, as the model is not designed to capture extremely high wealth levels. For this sample, we estimate:

$$r_{i,\tau}^{NW} = \beta_0 + \beta_1 \text{age}_i + \beta_2 \text{age}_i^2 + \beta_3 \text{age}_i^3 + \beta_4 \left(\frac{NW_i}{12 \times \bar{\mathcal{W}}_{25y}} \right) + \varepsilon_i, \qquad (4.2)$$

where $r_{i,\tau}^{NW}$ is the return on net worth during each month τ of 2019, age_i is the age of the individual in years, and $\left(\frac{NW_i}{12 \times \bar{\mathcal{W}}_{25y}} \right)$ is the ratio of net worth

[10]In particular, we simulate a simplified version of our model that incorporates the mortality parameters to capture life-cycle dynamics as well as the average job-finding and job-separation rates from the data. We do this to avoid having to calibrate the parameter $\bar{\mathcal{W}}$ internally, which would have required solving a fixed-point problem.

[11]Appendix A.3 provides more details about the data and Appendix B.2 presents the details of these calculations.

to the average annual labor income of a 25 year old. We then average all coefficients across months of 2019.[12] We set the borrowing rate to be equal to $\max_{a,j} r(a,j)$ plus a monthly spread of 0.005: the maximum returns on savings to prevent arbitrage, plus an annualized borrowing spread of 6%.[13]

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

Home production. We assume that income from home production is equal to a constant $\bar{h}_0$ for agents under 62, at which point it becomes equal to $1.15 \times \bar{h}_0$, i.e., $\bar{h}_1 = 0.15$. This value is taken from Dotsey et al. (2014), who show that home goods consumption for older workers starts increasing at around age 60, and is about 25% larger at age 90. We take an average of 15% for those older than 62. We internally calibrate $\bar{h}_0$ in Section 4.2.

Social Security income. To parametrize and calibrate the SS income function $\bar{y}^{ss}(w,j,k,\ell)$, we closely follow actual U.S. regulations, as in French (2005). This function is the product of two components. The first is the Primary Insurance Amount (PIA), a piece-wise concave function of a measure of past earnings, up to a limit. In order to keep the model tractable, we proxy past earnings by 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 measure of earnings as well as the bend points that generate concavity are all set to their 2019 values. The second component is a retirement-age-dependent modifier: individuals can begin collecting Social Security benefits at age 62 but face penalties if they retire before the full retirement age, which varies by birth cohort. We set the full retirement age (FRA) to 66, as in 2019. Additionally, they get a benefit if they retire past

[12] Among several other parametrizations, the specification in Equation (4.2) provided the best combination of simplicity and explanatory power.

[13] This falls in between the estimates of Lee et al. (2021) using Danish data (4%) and the implied borrowing spread used in Kaplan et al. (2018) (about 8%).

Table 4.1: 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.

this age, up to the age of 70. We follow the exact 2019 SS rules in setting up this modifier. We also follow current SSA regulations in imposing an earnings test for those on or under the FRA. Unemployed or non-participant agents receive no penalties.[14] A full description of the SS income function, as well as the calibration of its parameters can be found in Appendix B.3.

4.2. Internally calibrated parameters

We internally calibrate the remaining 13 parameters. The full set of parameters and respective targeted data moments are summarized in Table 4.1.

The discount factor β is chosen to match the fraction of individuals with non-positive net worth in the SCF under the age of 62. The borrowing limit is chosen to target the median value of the credit-limit-to-quarterly-labor-income ratio, as in Kaplan and Violante (2014) using the SCF. The level of home production income $\bar{h}_0$ is chosen to match the retired share.[15] This

[14]We model this earnings test as a pure tax, in line with the findings from the empirical literature (Gelber et al., 2020). Additionally, we abstract from the personal tax implications of SS benefits (Jones and Li, 2018).

[15]The retired share in Table 4.1 is for the population over the age of 25, which is different than the overall retired share that is shown in Figure 2.1. Our results in Section 2 remain unchanged if our earlier analysis was conducted for those over the age of 25.

results in a share of home production to GDP of around 13%, which is not far from the BEA's estimates for home production as a share of GDP in 2019 (21%). Finally, the slope of the UI replacement rate b_1 is set to match the Q1/Q5 ratio of replacement rates when individuals are ranked based on their labor income prior to unemployment, as in Birinci and See (2023), while the level b_0 is set to match the average replacement rate.

The level and slope of the employment disutility function are chosen to match the overall unemployment rate as well as the unemployment rate for those aged 55 and over, respectively. The level and slope of the unemployment disutility function are chosen in a similar way, but to match the LFPR of the population and those aged 55 and over.[16] The parameter γ that affects non-participants' job-finding probability is chosen to match the ratio of monthly non-participation (including retirement) to employment rate relative to the total monthly job-finding rate (out of non-employment). The probability of finding a job for the unemployed f is set to target the total job-finding rate, which is defined as the sum of the average flow rates from unemployment, non-participation, and retirement to employment. The level parameter of the job-separation rate $\bar{\delta}$ is chosen to match the monthly flow rate out of employment in an analogous manner. Finally, the slope parameter of the job-separation rate η_w^δ is chosen to target the Q1/Q5 ratio of the monthly job-separation rate in the data when employed individuals are ranked based on their labor income, as in Birinci and See (2023).

4.3. Model validation at the stationary state

The last two columns of Table 4.1 show that the model matches targeted data moments reasonably well. We now show that the model also captures untargeted data moments in 2019 that are relevant for the economic forces that we seek to analyze: the shape of the wealth distribution, and the wealth and income distributions of new retirees in the data reported in Section 2.

Unconditional wealth distribution. Figure 4.1(a) plots deciles relative to the median of the wealth distribution in the model's stationary state vs. the SCF and SIPP. To ensure comparability between the model and the data,

[16] Since j refers to monthly age and consumption is in units of 2019 dollars, the estimated slope parameters of disutility functions are small.

Figure 4.1: Validation of model predictions using microdata at the stationary state

(a) Wealth distribution (b) New retirees by wealth (c) New retirees by labor income

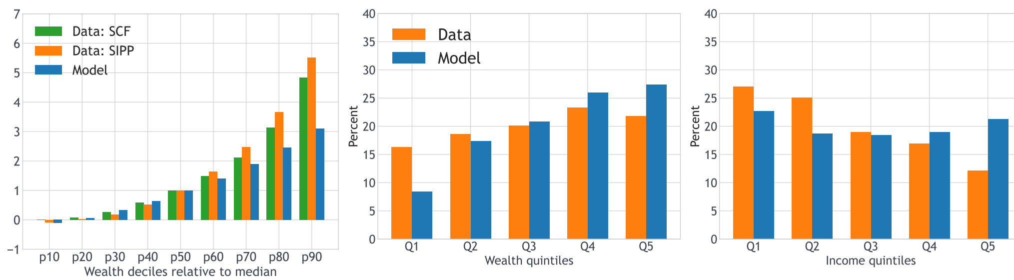

Note: Panel (a) presents deciles relative to median of the wealth distribution in the model's stationary state vs. 2019 SCF and 2019 SIPP. We exclude households with a ratio of net worth to average annual income greater than 15 from the data, as the model is not designed to capture very wealthy households. Panel (b) plots fractions of new retirees across wealth quintiles in the model's stationary state and in SIPP 2019. Panel (c) repeats the same calculations as in Panel (b) for labor income.

we report wealth deciles relative to median wealth. We also restrict the SCF and SIPP samples to households whose net worth to average annual income in 2019 is under 15, as the model is not designed to capture the very wealthy, in line with our sample used to estimate the asset returns in Equation (4.2). We find that the model does a good job of matching the overall shape of the wealth distribution, except that it does not fully capture the large wealth inequality driven by the very wealthy in the data, as expected.[17]

New retirees by wealth and labor income. Since our analysis is focused on the drivers of retirement patterns between 2020 and 2023, it is important that the model's stationary state generates the right patterns of retirement in 2019 in the data. Panels (b) and (c) of Figure 4.1 plot fractions of new retirees across quintiles of the wealth and income distributions in the model's stationary state vs. the 2019 SIPP data. We described how we computed these moments in the context of Figure 2.2 in the data, and implement the same calculations in the model. We find that the model broadly matches the patterns in the data. Specifically, the model matches the negative dependence of retirement decisions on income, as well as the positive relationship

[17]We abstract from several mechanisms commonly used to match the right tail of the wealth distribution: entrepreneurship (Cagetti and De Nardi, 2006), heterogeneous discount factors (Hendricks, 2007), and rare large idiosyncratic shocks to labor productivity (Castaneda et al., 2003), among others.

with wealth. These results indicate that the model is able to capture both the small wealth effects of labor supply, with those who retire being only slightly more likely to be wealthy, and the opportunity cost effects, with those who retire being more likely to have lower labor income.

There are two small discrepancies between the results in the model and the data. First, the model overestimates the rise in fractions of new retirees in wealth quintiles.[18] Second, while the model generates a negative relationship between retirement decisions and income, it generates a relatively larger fraction of new retirees at the top income quintile (20% in the model vs 12% in the data). This is driven by our simplifying assumption on the SS income function, which is based on the last realization of the stochastic wage component w before retirement. This simplification gives agents the incentive to wait until they obtain a high enough w before deciding to retire.

5. Aggregate dynamics during 2020-2023

Using the calibrated model, we now ask whether the model can generate the observed changes in aggregate labor market moments between 2020 and 2023. First, we describe how we measure and map the shocks to the model. Second, we present the results of our main experiment, where we feed in all these shocks and analyze whether the model generates the empirical changes in the retired share, unemployment rate, and employment-to-population ratio. As these movements are not targeted by our calibration, the model's fit in terms of these variables serves as yet another element of model validation.

5.1. Shocks

Starting from the stationary state, we introduce five shock sequences into the model: (i) a shock to the return on savings, which varies by wealth and age; (ii) a shock to job-separation rates for the employed, which varies by labor income; (iii) a shock to lump-sum transfers, which depends on age and total income; (iv) a shock to UI benefits for the unemployed; and (v) a shock

[18]Despite the presence of a stronger relationship between the level of wealth and incentives to retire in the model relative to data, our decomposition exercise in Section 6.1 shows that increases in asset valuations during 2020-2023 played a small role in explaining excess retirements.

Figure 5.1: Time series paths for exogenous shocks

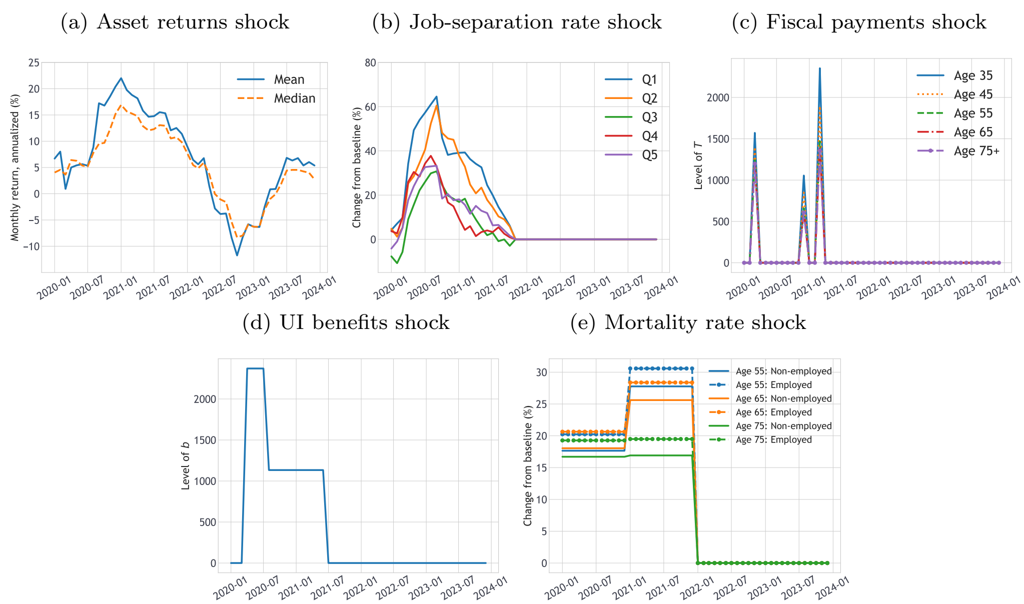

Note: Panel (a) plots the mean and median paths of the estimated monthly return (annualized) function $r_t(a, j)$. We only plot the mean and median values at each month for expositional purposes. Panel (b) plots percent changes in the job-separation rate at each month $\delta_t(w, j)$ relative to the 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.

to mortality rates, which varies by age and employment status. The time series of these shocks are presented in Figure 5.1. Below, we describe in detail how we map each of these impulses from the data to the model. At each date, agents treat certain shocks (e.g., job separation and mortality) as permanent changes to parameter values, which requires a full solution of the problem—substantially increasing the computational burden.

Asset returns. Elevated asset returns during 2020-2023 may have triggered wealth effects that led to above-average movements into retirements and also retained individuals already in retirement. To capture this channel, we estimate Equation (4.2) for each month from January 2020 to December 2023. Due to significant month-to-month variation in returns, we take six-month moving averages of the estimated coefficients and feed to the model as exogenous shocks. Figure 5.1(a) plots the mean and median paths of the estimated monthly return (annualized) function: both the mean and median increase in the early months of the pandemic, surpassing 20% and 15% in 2021, respectively. They then fall and become negative in 2022 and early 2023, but recover to positive levels later in 2023.[19]

For implementation, we replace the return function $r(a, j)$ in the budget constraint for each agent with positive wealth with $r_t(a, j)$. These return shocks are unexpected and assumed to be transitory. That is, individuals expect the return on savings to be the stationary function in all following periods. This is therefore equivalent to a lump-sum windfall that does not distort individual savings decisions.[20] This reflects the unexpected nature of these large movements, and prevents counterfactual changes in consumption and savings behavior that could affect labor supply by inducing agents to work more and accumulate wealth to take advantage of elevated returns.

Job-separation rates. The 2020-23 period was marked by a large increase in the aggregate job-separation rate. In addition, the COVID-19 episode induced a much larger increase in job-separation rates of low-income workers, while those who were employed at relatively higher-paying jobs experienced smaller increases in their job-separation rates. The rise in job separations may have negatively impacted labor force participation as unemployed workers are more likely to flow into non-participation than are employed workers (Hobijn and Şahin, 2021). We capture both the magnitude and heterogeneity in separations by feeding exogenous paths of job-separation rates that vary by quintiles of labor earnings. To this end, using the CPS, we first calculate the monthly job-separation rate as the fraction of employed individuals in one month who become non-employed in the next. We compute this rate separately for each month from 2019 to 2023 and by quintiles of the earnings distribution, where individuals are assigned to quintiles based on their current labor income.[21] We then calculate percent changes in job-separation rates for

[19]Appendix C.1 presents heterogeneity in these estimated asset returns by age, showing that younger individuals experienced wider return fluctuations during 2020-2023.

[20]The amount of lump-sum income (or loss) is equal to $a_t \times \frac{r_t(a_t,j_t)-r(a,j)}{1+r(a,j)}$. As such, this experiment preserves distortion of decisions through wealth effects (as it is intended).

[21]At the onset of the pandemic, the fraction of employed who were temporarily separated from their job increased substantially. However, most of these workers were later recalled to their jobs. Because our model does not feature elements to meaningfully differentiate between temporarily unemployed with a recall option from the regular unemployed,

each month in 2020-23 relative to the average job-separation rate in 2019, separately for each quintile of labor earnings. Due to sizable fluctuations in monthly rates, we compute six-month moving averages of these changes. Additionally, all these shocks becomes negligible after October 2021, and so we set them to zero after this date, for tractability.[22] Panel (b) of Figure 5.1 plots the series that we feed to the model as period-by-period shocks to the job-separation rate at the stationary state $\delta(w, j)$.[23] These series reflect both the sharp rise in separation rates and the substantial heterogeneity across labor income quintiles, with lower-quintile workers being more affected and experiencing a slower recovery to 2019 levels. We assume that these shocks are perceived by the agents to be permanent at each point in time, given the uncertainty surrounding the duration of the public health emergency and its effects on the labor market.

Economic impact payments. The COVID-19 episode in the U.S. triggered an unprecedented fiscal response that involved large scale support for households with relatively lower levels of income (Faria-e-Castro, 2021a). A large part of fiscal support programs to households was economic impact payments, which consisted of three rounds of lump-sum transfers to eligible households. We model these payments as increases in government transfers $T(y, j, a)$ in our model. We map the dollar value and timing of the transfers directly to the model. For each of the three rounds of transfers, households were ineligible if their adjusted gross income (AGI) exceeded $80,000. 2019 IRS data on the distribution of AGI for filed returns establishes that this value is close to the 80th percentile of the AGI distribution. Thus, we set the eligibility cutoff for transfers as the 80th percentile of the stationary state AGI distribution. We define AGI in the model as total income y.

The first round of transfers was associated with the Coronavirus Aid, Relief, and Economic Security (CARES) Act and took place in March 2020,

when calculating the monthly job-separation rates in the data, we do not include temporary job separations.

[22]Appendix C.2 presents the historical time series of these shocks and shows that these shocks in the pre-2020 period were typically stable and that they mostly return to their pre-2020 levels by the end of 2021.

[23]For example, the job-separation rate of those at the bottom two quintiles increased in mid 2020 by over 60% relative to their respective stationary state levels, while the separation rate of those at the top quintile increased at that time by around 30%.

consisting of $1,200 per person plus $500 per child under the age of 17. The second round of transfers was triggered by the Tax Relief Act of 2020 and took place in December 2020, consisting of $600 per person plus $600 per child under the age of 17. The American Rescue Plan Act of 2021 initiated a third round of transfers in March 2021, which consisted of $1,400 per person plus $1,400 per dependent. Thus, the presence of dependents could considerably increase the effective transfers earned by households.

To map the size of the effective transfers to the model, we explicitly account for the fact that household structure and the number of dependents may depend on the age of the household head. We use data from the 2019 Annual Social and Economic Supplement (ASEC) of the CPS, which provides the number of individuals under 18 by the head of household's age. This allows us to impute a transfer modifier that depends on the age of the head. The procedure is explained in detail in Appendix C.3. The effective transfer amounts over time, as a function of age, is plotted in Panel (c). Since these transfers were plausibly perceived to be one-time events, we assume that these shocks are unexpected and expected to last for a single period.

UI benefits. The other major component of household income support during the COVID-19 episode was the expansion of UI benefits. These extra benefits were $600 weekly (on top of pre-pandemic benefits) between March 2020 and June 2020, and then $300 weekly from July 2020 to about June 2021.[24] We map these extra benefits to the model by assuming four weeks per month. The path of UI benefits that we input in the model is plotted in Panel (d). Just as in the data, these benefits are modeled as a lump-sum transfer for the unemployed. That is, unemployed individuals receive their regular UI benefits, calculated with regular replacement rates, and these additional UI benefits in months when they are provided by the government. As with economic impact payments, we assume that agents perceive these shocks to be temporary: they are unexpected and expected to last only for a single period.

Mortality rates. The last shock we consider is a change in mortality rates $\pi(j, \ell)$. The goal is not to exactly match actual mortality patterns, but rather

[24] In practice, different states phased out benefits at different points around that time, and we choose to end them in June 2021 for simplicity.

to shock agents' perceived mortality risk during 2020. This is potentially an important channel given that perceived and realized increases in mortality operate as changes in the discount factor that may affect participation decisions especially for older agents. Additionally, different from the stationary state of the model, we now allow mortality rates to depend on labor force status, reflecting the potential increase in COVID-19 transmission rates from employment activities that involve physical contact.

To model the rise in mortality rates, we assume that at the beginning of 2020, agents perceive their mortality rate to have risen to the levels empirically observed in the SSA life tables. At the beginning of 2021, those rates change again, and they return to their baseline levels in 2022. Similar to the labor market shocks, and in order to capture the uncertainty about the duration of the public health emergency, we assume that agents perceive each of these changes to be permanent. We assume an additional increase in mortality for employed agents. To calibrate this increase, we combine estimates from Eichenbaum et al. (2021) with 2020 Census data: the probability of death for an employed worker over the age of 50 increased by 2.2% more relative to a non-employed, while the probability of death for an employed below 50 increased by 0.08% more relative to a non-employed. We describe how we obtain these numbers in Appendix C.4. The percent changes in mortality rates in each month relative to the stationary state by age and employment status are plotted in Panel (e).

5.2. Aggregate labor market moments: model vs data

Next, we present the results of our experiment, in which we introduce all shocks simultaneously starting from the model's stationary state and compare the resulting aggregate labor market dynamics along the transition to their empirical counterparts from 2020 to 2023. Figure 5.2 plots the data and the model paths for the aggregate retired share (Panel (a)), unemployment rate (Panel (b)), and employment-to-population ratio (Panel (c)).[25]

For the retired share in the data, we use the same definition as in Figure 2.1: the deviation of the actual fraction of retirees in the population in the

[25] In this exercise, agents who die are replaced by new 25-year olds and thus the total population is kept constant. We have experimented with alternative assumptions (i.e., not replacing agents who die) and found that this matters very little quantitatively.

Figure 5.2: Changes in aggregate labor market moments: Model vs data

(a) Retired share (b) Unemployment rate (c) Employment/population

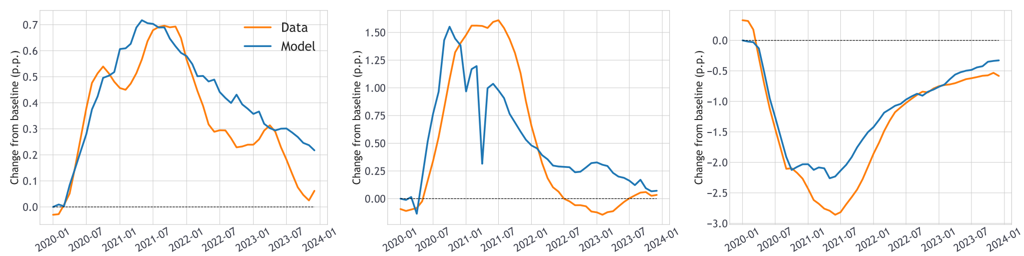

Note: This figure plots the paths of the aggregate retired share (i.e., the fraction of retirees in the population) (Panel (a)), unemployment rate (Panel (b)), and employment-to-population ratio (Panel (c)) in the data and the model. We take six-month moving averages both in the data and in the model, 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.

CPS relative to the trend. We take six-month moving averages both in the data and in the model, and plot the percentage-point (pp) deviation from the 2019 average in the data and stationary state of the model. The model matches well both the magnitude and persistence of the increase in the retired share: it rises in 2020, peaking at around 0.7 pp in 2021, and slowly declining thereafter.

Similarly, for both the unemployment rate and the employment-to-population ratio, we take six-month moving averages and plot the pp deviations from both the data average in 2019 or the model's stationary state. Starting with the unemployment rate, we note that since our model is not designed to capture the sizable increase in temporary layoffs during the COVID-19 episode, our data benchmark is the unemployment rate net of temporary unemployment, as classified in the CPS. We find that the model captures well both the magnitude and dynamics of the increase in the unemployment rate. Finally, the model slightly underestimates the decline in the employment-to-population ratio by about 0.7 pp, but matches its dynamics well: the initial drop and the subsequent slow recovery. In particular, both model and data are aligned with their prediction that the employment-to-population ratio is around 0.5 pp lower at the end of 2023 relative to the 2019 level. Taken together, these results suggest that the model does a very good job in capturing untargeted aggregate dynamics between 2020 and 2023.

6. Decomposing the retirement boom

Having shown that the model captures well the size and persistence of movements in key aggregate labor market moments, we now undertake a decomposition exercise where we quantify the importance of each of the five shocks in driving these movements during this episode.

6.1. Decomposing the increase in retired share

Panels (a) and (b) of Figure 6.1 offer two alternative decompositions that shed light in the importance of each exogenous force at each point in time on the increase in the retired share. Panel (a) plots the baseline (with all shocks included) and removes one shock at a time. Panel (b) adds only one shock at a time, starting from the stationary state (without any shock).

The results show that job-separation shocks, as shown by green circle lines, are the most important driver of the rise in the retired share in 2020. However, these shocks alone cannot explain the magnitude of the rise. Economic impact payments (purple square line) are important to get the magnitude of the rise right, and help explain the persistence of the retired share as shown in Panel (b). Asset returns and UI benefits contributed slightly to the increase in retirements. The expansion of UI creates an income effect on labor supply that leads unemployed workers to retire, but this effect is small as transitions between unemployment and retirement are infrequent. Albeit small in magnitude, UI expansion was relatively more important early on, 2020-21, while asset returns were more important at a later stage, 2022-23.

The mortality shock, represented by the light-gold line, counters the effects of these shocks in the aggregate and helps the model get the magnitudes right. The negative effect of the mortality shock on the retired share is mechanical: mortality risk rises by more for older people, who therefore die in greater numbers than younger people. Since a significant share of these agents are retired, this channel pushes the retired share down. Note that, as previously explained, we do explicitly account for greater risk of mortality from employment, which counteracts this mechanical effect of mortality shocks on retirement by inducing older people to retire. We find, however, that the inequality in mortality rates across ages is the dominating channel. Ultimately, the model requires all these shocks to adequately capture the retirement dynamics.

Panel A of Table 6.1 offers a formal decomposition to quantify the contri-

Figure 6.1: Decomposing movements in the retired share and unemployment rate

(a) Retired share: removing one shock (b) Retired share: adding one shock

(c) Unemployment rate: removing one shock (d) Unemployment rate: adding one shock

Note: Panels (a) and (c) plot the baseline (with all shocks included) and remove one shock at a time. Panels (b) and (d) add only one shock at a time, starting from the stationary state (without any shocks). $r(j,a)$, $\delta(w,j)$, b, $T(y,j,a)$, and $\pi(\ell,j)$ refer to shocks to returns, separations, UI, transfers, and mortality.

bution of all five shocks on the rise in the retired share for each year between 2020 and 2023, where we compute the average individual percent contribution of each shock for these years (that is, we compare the lines in Panel (b) to the blue line in Panel (a)). The table quantifies the previous discussions: 75% of the excess rise in the retired share in 2020 is accounted for by changes in job-separation rates. This share remains elevated throughout. Economic impact payments explain 40% in 2020-21, and become more important in 2022-23. Changes in asset returns play a small role in 2020-21, and account for a fifth of the share in 2022-23. Note that the sum of the contribution of the shocks adds up to more than 100%, which reflects not only interactions between the different mechanisms but also the importance of the offsetting effects of mortality shocks in order to adequately match the

Table 6.1: Decomposition of changes in the retired share and unemployment rate

	Asset returns	Job separations	UI benefits	Transfers	Mortality	Model (pp.)	Data (pp.)
				A. Retired share			
2020	1.0%	75.0%	11.7%	40.0%	-24.8%	0.25	0.28
2021	7.1%	74.4%	8.5%	42.4%	-30.3%	0.66	0.60
2022	17.4%	67.4%	9.5%	64.3%	-60.2%	0.46	0.34
2023	20.0%	78.8%	14.3%	80.7%	-93.2%	0.29	0.18
				B. Unemployment rate			
2020	-0.51%	109.00%	-9.74%	-7.18%	7.61%	0.68	0.44
2021	-8.35%	97.97%	15.69%	-5.69%	15.96%	0.85	1.43
2022	-1.57%	17.32%	33.36%	5.59%	44.28%	0.33	0.12
2023	22.82%	19.20%	17.86%	3.01%	44.99%	0.19	-0.03

Note: This table presents the average percentage change in the retired share (Panel A) and unemployment rate (Panel B) that is explained by feeding one shock (presented in columns) at a time, separately for each year. Due to interactions and averaging, values may not sum up to 100%. The last two columns present the average percentage-point (pp) changes in each variable during each year in the model and the data.

retirement dynamics along the transition.[26] In sum, job separations were a major factor throughout the period under analysis, while transfers became more significant in explaining the dynamics of excess retirements later on.

The importance of job separations and fiscal transfers in explaining excess retirements suggests that the rise in retirements may have been driven by income-poor workers, who faced relatively worse labor market prospects and were eligible and more sensitive to income effects from transfers. The positive effects of asset returns also warrant an investigation on the role of wealth. We study the composition of new retirees in more detail in Section 6.3.

6.2. Decomposing the increase in unemployment rate

Panels (c) and (d) of Figure 6.1 and Panel B of Table 6.1 repeat the same exercise for the unemployment rate. There are four key takeaways. First, the unemployment rate dynamics are almost completely explained by separation shocks. Second, mortality shocks play somewhat of a role in explaining the rise in unemployment, again due to larger mortality risk among older agents,

[26] A part of the increase in returns is driven by house price appreciation. One potential concern is that housing is a less liquid asset and thus capital gains should generate weaker wealth effects on labor supply. We analyze this point in Appendix C.5.

who tend to be employed or retired. Third, UI benefits are moderately important, especially in 2022. Fourth, asset returns and transfers play a relatively small role in driving the unemployment rate.

The last two columns of Table 6.1 present the average percentage-point (pp) changes in each variable during each year in the model and the data. We find that the model captures the increases in both the retired share and the unemployment rate well for each year, consistent with our results in Figure 5.2.

6.3. Model validation along the transition

We have shown that the model broadly matches the behavior of aggregate variables of interest along the transition. Does it also align well with microdata that are relevant for the mechanisms of interest? Comparing the outcomes from the model along the transition against the microdata also reinforces the credibility of our quantitative decomposition on the sources of changes in aggregate variables. In this section, we show that the model delivers three key predictions that are broadly in line with the microdata. In particular, the model matches changes in the wealth distribution and the distributions of new retirees by *both* wealth and income quintiles during 2020-2021 relative to 2019. Moreover, Appendix C.6 provides two additional results by comparing changes in monthly flow rates into and out of retirement, as well as average wealth over the transition. We show that the model's outcomes on these moments closely align with the empirical observations.

Changes in the distribution of net worth. The model captures the key movements in the wealth distribution. Table 6.2 presents the evolution of percentiles of the distribution relative to median in 2020-21 from the SIPP data (Panel A) and the model (Panel B). In the data, percentiles below the median increase relative to the median over time while percentiles above the median fall, suggesting a compression of the wealth distribution over time. The model captures the exact same pattern, with the bottom percentiles rising relative to the median and the top percentiles falling. Specifically, the magnitudes of the decline between 2021 and 2019 in percentiles above the median are almost identical in the model and the data, but the model slightly overestimates the magnitudes of the rise in percentiles below the median.

Overall, the model reproduces the overall dynamics of the wealth distribution between 2019 and 2021, which involved an increase in the average net

Table 6.2: Wealth distribution over time: Data vs model

Relative to median	p10	p20	p30	p40	p50	p60	p70	p80	p90
					A. Data				
2019	-0.08	0.03	0.18	0.52	1.00	1.64	2.47	3.65	5.51
2020	-0.04	0.04	0.21	0.53	1.00	1.58	2.34	3.37	4.91
2021	-0.02	0.05	0.23	0.55	1.00	1.54	2.23	3.12	4.51
					B. Model				
2019	-0.11	0.06	0.33	0.64	1.00	1.41	1.89	2.46	3.11
2020	0.01	0.23	0.48	0.73	1.00	1.32	1.70	2.15	2.64
2021	0.06	0.30	0.55	0.77	1.00	1.26	1.56	1.90	2.14

Note: This table presents the value of deciles relative to the median of the wealth distribution in the SIPP data (Panel A) and the model (Panel B), separately for 2019, 2020, and 2021.

worth (shown in Appendix C.6) and a reduction of inequality in net worth. The fact that the model matches these empirical patterns is important if we are to give wealth effects a chance to explain retirement dynamics during this period.

Changes in new retirees by wealth and labor income. Panels (a)-(b) and (c)-(d) of Figure 6.2 compare changes in fractions of new retirees in the data and model across the wealth and labor income distributions, respectively. Calculations of these moments follow the same steps as before. As discussed in Section 2, Panel (a) reveals that the post-COVID-19 episode is not characterized by a rise in the fraction of new retirees with high levels of wealth. If anything, retirements during 2020-2021 were slightly tilted toward people with low levels of wealth, and there is slightly less heterogeneity in fractions of new retirees across wealth quintiles in the 2020-2021 episode when compared with the same distribution in 2019. Panel (b) shows that the model reproduces the same patterns: retirements during 2020-2021 were not tilted toward wealthy individuals and changes in fractions of new retirees by wealth quintiles in 2020-2021 relative to 2019 were quite limited.

Panels (c) and (d) show that, in the data and the model, fractions of new retirees by labor income quintiles also change little over time, with the majority of new retirees continuing to come from the lower quintiles. This makes sense in light of our decomposition, which reveals that most new retirements were due to a deterioration of labor market conditions with increased job sep-

Figure 6.2: Validation of model predictions using microdata along the transition

Note: Panels (a) and (b) plot fractions of new retirees by wealth quintiles, separately for those who retire in 2019 and those who retire between 2020 and 2021 using data from the SIPP and from the model, respectively. Panels (c) and (d) repeat the same calculations for labor income.

arations especially for low-income workers and economic impact payments to which low-income individuals are more sensitive.

In summary, we show that the model not only matches the rise in the retired share during this episode but also generates fractions of new retirees by wealth and income groups as well as monthly flow rates into and out of retirement (shown in Appendix C.6) that are in line with the microdata.

7. Conclusion

In this paper, we develop an incomplete markets, OLG model combined with a frictional labor market to understand the rise in retirements experienced in the U.S. after 2019. We analyze the ability of five different channels to explain excess retirements during 2020-2023: elevated asset returns, increased job separations, provision of economic impact payments, expansion of UI benefits, and increased mortality risk. In a quantitative exercise that maps these shocks to the calibrated model, we show that the model is able to match the magnitude and persistence of excess retirements when all these forces are active. In a decomposition exercise, we show that the fluctuations in job separations and economic impact payments are the main drivers behind the excess retirements in 2020-23. On the other hand, increased mortality risk during COVID-19 mitigated the effects of the other forces. Fluctuations in asset returns and changes in UI benefits also contributed to the dynamics of excess retirements, but to a lesser extent.

The fact that increased job loss risk and economic impact payments conditional on income explain the bulk of excess retirements suggests that these were concentrated in lower-income individuals. We show that this prediction of the model is corroborated in the microdata: fractions of new retirees by wealth and income groups changed little during this period, and most new retirees came from lower income quintiles.

References

Benson, D., French, E., 2011. How do sudden large losses in wealth affect labor force participation? Chicago Fed Letter , 282.

Birinci, S., See, K., 2023. Labor market responses to unemployment insurance: The role of heterogeneity. American Economic Journal: Macroeconomics 15, 388–430.

Blandin, A., Jones, J.B., Yang, F., 2023. Marriage and work among prime-age men. Working Papers 2313. Federal Reserve Bank of Dallas.

Blundell, R., French, E., Tetlow, G., 2016. Retirement incentives and labor supply, in: Handbook of the Economics of Population Aging. 1 ed.. Elsevier. volume 1 of *Handbook of the Economics of Population Aging*. chapter 0, pp. 457–566.

Cagetti, M., De Nardi, M., 2006. Entrepreneurship, frictions, and wealth. Journal of Political Economy 114, 835–870.

Castaneda, A., Diaz-Gimenez, J., Rios-Rull, J.V., 2003. Accounting for the U.S. earnings and wealth inequality. Journal of Political Economy 111, 818–857.

Cheng, I.H., French, E., 2000. The effect of the run-up in the stock market on labor supply. Economic Perspectives 25, 48–65.

Coronado, J.L., Perozek, M.G., 2003. Wealth effects and the consumption of leisure: Retirement decisions during the stock market boom of the 1900s. FEDS 2003-20. Board of Governors of the Federal Reserve System.

De Nardi, M., French, E., Jones, J.B., 2010. Why do the elderly save? The role of medical expenses. Journal of Political Economy 118, 39–75.

De Nardi, M., French, E., Jones, J.B., 2016. Savings after retirement: A survey. Annual Review of Economics 8, 177–204.

Dotsey, M., Li, W., Yang, F., 2014. Consumption and time use over the life cycle. International Economic Review 55, 665–692.

Eichenbaum, M.S., Rebelo, S., Trabandt, M., 2021. The macroeconomics of epidemics. The Review of Financial Studies 34, 5149–5187.

Faria-e-Castro, M., 2021a. Fiscal policy during a pandemic. Journal of Economic Dynamics and Control 125.

Faria-e-Castro, M., 2021b. The COVID retirement boom. Economic Synopses , 1–2.

French, E., 2005. The effects of health, wealth, and wages on labour supply and retirement behaviour. The Review of Economic Studies 72, 395–427.

French, E., Jones, J.B., 2001. The effects of health insurance and self-insurance on retirement behavior. Working Paper Series WP-01-19. Federal Reserve Bank of Chicago.

Gelber, A.M., Jones, D., Sacks, D.W., 2020. Estimating adjustment frictions using nonlinear budget sets: Method and evidence from the earnings test. American Economic Journal: Applied Economics 12, 1–31.

Hamilton, J.D., 2018. Why you should never use the Hodrick-Prescott filter. The Review of Economics and Statistics 100, 831–843.

Hendricks, L., 2007. How important is discount rate heterogeneity for wealth inequality? Journal of Economic Dynamics and Control 31, 3042–3068.

Hobijn, B., Şahin, A., 2021. Maximum employment and the participation cycle. Working Paper 29222. National Bureau of Economic Research.

Hobijn, B., Şahin, A., 2022. Missing workers and missing jobs since the pandemic. Technical Report. National Bureau of Economic Research.

Jones, J.B., Li, Y., 2018. The effects of collecting income taxes on social security benefits. Journal of Public Economics 159, 128–145.

Kaplan, G., Moll, B., Violante, G.L., 2018. Monetary policy according to HANK. American Economic Review 108, 697–743.

Kaplan, G., Violante, G.L., 2014. A model of the consumption response to fiscal stimulus payments. Econometrica 82, 1199–1239.

Krusell, P., Mukoyama, T., Rogerson, R., Şahin, A., 2017. Gross worker flows over the business cycle. American Economic Review 107, 3447–3476.

Lee, S., Luetticke, R., Ravn, M.O., 2021. Financial frictions: micro vs macro volatility. Working Paper Series 2622. European Central Bank.

Montes, J., Smith, C., Dajon, J., 2022. "The great retirement boom": The pandemic-era surge in retirements and implications for future labor force participation. FEDS 2022-081. Board of Governors of the Federal Reserve System.

Nie, J., Yang, S.K.X., 2021. What has driven the recent increase in retirements? Federal Reserve Bank of Kansas City, Economic Bulletin, August 11.

Shimer, R., 2005. The cyclical behavior of equilibrium unemployment and vacancies. American Economic Review 95, 25–49.

Veracierto, M., 2008. On the cyclical behavior of employment, unemployment and labor force participation. Journal of Monetary Economics 55, 1143–1157.

Supplementary Material for "Dissecting the Great Retirement Boom"

Appendix A. Data

In this Appendix, we provide details on our empirical analysis to supplement the discussions in the main text and provide additional results from the data.

Appendix A.1. CPS

Our CPS sample consists of individuals aged 16 and over who are not in the armed forces. In our baseline analysis, we define retirees based on whether they identify themselves as retired, `EMPSTAT` equal to 36. We define the retired share as the weighted sum of all retirees divided by the weighted sum of all persons in our sample. We seasonally adjust the retired share by regressing it on month dummies.

We have also experimented with alternative definitions of retirement. Figures Appendix A.1 and Appendix A.2 replicate Figure 2.1 for two such alternative definitions. Figure Appendix A.1 considers a stricter definition where a person is considered retired if `EMPSTAT` is equal to 36 and age is at least 62. This is a strict subset of our baseline definition as it only considers people who identify themselves as retired and are old enough to be eligible for Social Security benefits. Figure Appendix A.2, on the other hand, considers a slightly broader definition of retirement: `EMPSTAT` is equal to or greater than 30 and age is at least 62. This means that we define retirees as non-participants who are at least 62 years old. Figures Appendix A.1 and Appendix A.2 show that our measure of the retired share (i.e., excess retirement share) is robust to alternative definitions of retirement.

Figure Appendix A.1: Alternative retirement definition: Retirees over 62

(a) Retired share and linear trend (b) Deviations from trend

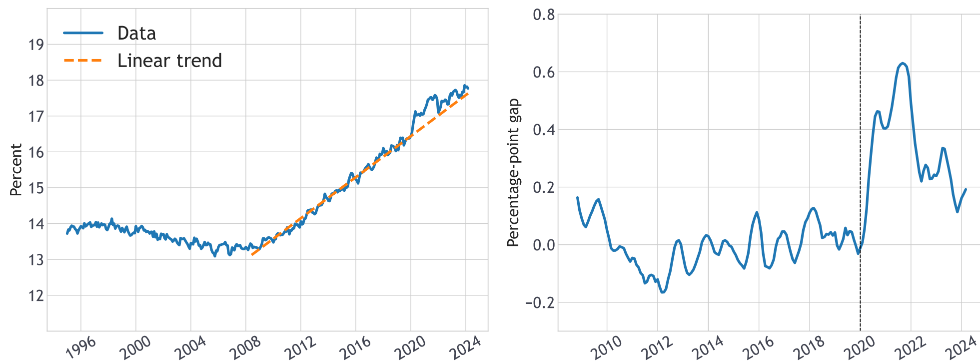

Note: Panel (a) plots the retired share in the U.S., calculated as the fraction of individuals who report to be retired in the Current Population Survey (CPS) and are at least 62 years old among all individuals (excluding those in armed forces) aged 16 and over. Linear trend is estimated between June 2008 and January 2020. Panel (b) plots deviations from trend by taking 6-month moving averages.

Appendix A.2. SIPP

We use the SIPP data for three purposes. First, we calculate the wealth distribution for each year between 2019 and 2021. These results are presented in Panel (a) of Figure 4.1 and in Table 6.2. Second, we calculate fractions of new retirees by wealth and labor income quintiles, separately for those who retire in 2019 and those who retire between 2020 and 2021. These results are presented in Figure 2.2. Finally, we estimate the parameters of the lifecycle labor income process using the SIPP data, as discussed in Section 4.1. In this Appendix, we provide details on calculations of the first two moments. Appendix B.1 provides details on the last one.

For these calculations, we use SIPP 2020, 2021, and 2022 panels covering data from the start of 2019 to the end of 2021.[27] Our sample consists of all individuals (excluding those in armed forces) aged 25 and over.

Wealth distribution. The SIPP provides values of assets across detailed asset categories at individual and household levels for each year. We obtain the value of total net worth for each household as follows.

We first calculate the gross liquid wealth for each household. This is

[27]Later panels of SIPP are not yet available as of this writing.

Figure Appendix A.2: Alternative retirement definition: Non-participants over 62

(a) Retired share and linear trend

(b) Deviations from trend

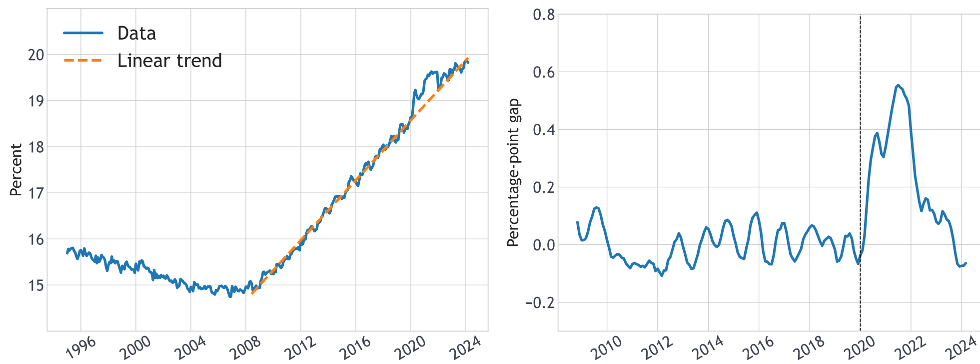

Note: Panel (a) plots the retired share in the U.S., calculated as the fraction of individuals who report to be out of the labor force in the Current Population Survey (CPS) and are at least 62 years old among all individuals (excluding those in armed forces) aged 16 and over. Linear trend is estimated between June 2008 and January 2020. Panel (b) plots deviations from trend by taking 6-month moving averages.

given by the household-level sum of (i) value of assets held at financial institutions `THVAL_BANK`, (ii) value of other interest-earning assets `THVAL_BOND`, (iii) value of stocks and mutual funds `THVAL_STMF`, and (iv) value of other assets `THVAL_OTH`. Next, we obtain the net liquid wealth as the gross liquid wealth minus the household-level sum of value of amount owed on all unsecured debt `THDEBT_USEC`. Our measure of household-level net worth is then given by the net liquid wealth plus the sum of household-level (i) value of retirement accounts `THVAL_RET`, (ii) equity in primary residence `THEQ_HOME`, (iii) equity in rental properties `THEQ_RENT`, (iv) equity in other real estate `THEQ_RE`, and (v) equity in vehicles `THEQ_VEH`.

We calculate household-level net worth for all households, separately using the SIPP 2019, 2020, and 2021 data. We further restrict our sample to households whose net worth to average annual income in 2019 is under 15, as the model is not designed to capture the very wealthy, in line with our sample used to estimate the asset returns in Equation (4.2). Next, we divide the household-level net worth by two for married households to obtain individual-level net worth. Then, for each year, we calculate the average and various percentiles of the net worth distribution using weights.

Fraction of new retirees by wealth quintiles. The SIPP also provides individual-level information on weekly employment status. For each of the five possible weeks in a month, this information is recorded in RWKESR1 to RWKESR5. We use this information to classify individuals into one of the three employment statuses each month as follows. If an individual reports having no job or business and that she is not looking for work and not on layoff in at least one week of a given month, we classify her as non-participant (i.e., out of labor force) in that month. That is, $\texttt{RWKESRj} = 5$ for at least one $j \in \{1, 2, 3, 4, 5\}$. If she reports having a job or business and either working or absent without pay (but not on layoff) in all weeks of that month, we classify her as employed in that month. That is, $\texttt{RWKESRj} \leq 2 \, \forall j \in \{1, 2, 3, 4, 5\}$. For all other cases with any other potential combination of employment statuses across weeks, we classify individuals as unemployed (i.e., those who report to have a job or business but on layoff or those who do not have a job or business and are looking for work).

Given this information on monthly employment status, we identify new retirees in 2019 as those who report as employed or unemployed (i.e., in the labor force) in a month in 2019 and report as retired for the first time in the next month in 2019.[28] Then, we assign each new retiree in 2019 into quintiles of the wealth distribution in 2019 (as calculated above) for those who are employed and aged between 62 and 72 using their own level of net worth. These steps allow us to calculate the fraction of new retirees in 2019 at each quintile among all new retirees in 2019. We repeat the same procedure to calculate the same moments for new retirees between 2020 and 2021.

Fraction of new retirees by labor income quintiles. We also obtain the fraction of new retirees by labor income quintiles following the same procedure as above except that we use total labor income (instead of net worth) to classify individuals into quintiles of the labor income distribution. We measure labor income as the sum of (i) total weekly wage or salary earnings across the weeks of the month from the first job and the second job and (ii) profits or losses a business made after correcting for any salary or wages that may have been paid to the owner.[29]

[28] The EEVERET variable in SIPP provides information on whether an individual is ever retired from a job or business. We use this variable to identify first time retirees.

[29] For the first job, weekly earnings are given by TJB1_WKSUM1 to TJB1_WKSUM5. For the second job, they are given by TJB2_WKSUM1 to TJB2_WKSUM5. Business profits or losses from the first and the second business are provided by TJB1_PRFTB and TJB2_PRFTB, respectively.

Appendix A.3. SCF

We use the 2019 wave of the SCF, downloaded from the website of the Federal Reserve Board, for two purposes. First, we compute the average net worth. Our definition of total assets covers the following variables: `equity` measures total direct and indirect holdings of stocks; housing is measured as `houses` + `oresre` + `nnresre`, which is the value of the primary residence plus other residential property and net equity in non-residential real estate; and government bond holdings are computed as `notxbnd` + `mortbnd` + `govtbnd` + `savbnd` + `tfbmutf` + `gbmutf`, which is tax exempt bonds plus mortgage-back bonds plus U.S. government and agency bonds plus savings bonds plus tax-free and government bond mutual funds. Corporate bond exposure is equal to `obnd` + `obmutf`, which is corporate and foreign bonds plus other bond mutual funds. Private business interests are measured as `bus`. The difference between `asset` and these assets is classified as other assets. Finally, debt is measured directly as `debt`. Net worth is measured as `asset` − `debt`. Second, we estimate how returns on savings change based on the level of net worth and age, where we use `age` as the age of the head of household. Similar to our sample in SIPP, we restrict our SCF sample to households whose net worth to average annual income in 2019 is under 15.

Appendix B. Calibration

This Appendix provides more details on some aspects of the calibration: the estimation of life-cycle labor income process, the calculation of asset returns in the data, the procedure to impute returns to the SCF net worth data, and a detailed explanation of the SS income function.

Appendix B.1. Labor income process

We estimate the parameters of the life-cycle labor income precess given in Equation (3.1) by closely following French (2005) and Blandin et al. (2023). To do so, we use the SIPP 2004 panel, covering a period of stable non-recessionary labor markets in the U.S. We focus on monthly labor earnings of a sample of individuals whose real wage is above 1/3 of the federal minimum wage at the time, whose usual weekly hours worked is at least 20, and who are at least 25 years old. Using this sample, we estimate a regression of the logarithm of monthly labor earnings (adjusted by the CPI) on age and age squared with individual-fixed effects and weights. This regression yields our estimates for ψ_0, ψ_1, and ψ_2. Then, using the predicted and the observed values of the logarithm of monthly labor earnings, we obtain a panel of residuals for labor earnings $\{\hat{w}_{i,j}\}_{i,j}$. Next, under the same stochastic process of labor earnings residuals as in Blandin et al. (2023), we obtain the autocorrelation of the stochastic wage component ρ_w as follows:

$$\rho_w = \frac{cov(\hat{w}_{i,j}, \hat{w}_{i,j+3})}{cov(\hat{w}_{i,j}, \hat{w}_{i,j+2})}.$$

Given ρ_w, we calculate the standard deviation of the innovations σ^ϵ as follows:

$$\sigma^\epsilon = \sqrt{\frac{cov(\hat{w}_{i,j}, \hat{w}_{i,j+2})(1-\rho_w^2)}{\rho_w^2}}.$$

Finally, the standard deviation of initial wage draws σ^{w_0} is simply the standard deviation of the residuals for those who are 25 years old.

Appendix B.2. Asset returns and SCF imputation

We use data on realized asset returns for various asset classes between 2020 and 2023 in order to impute returns on net worth for households in the 2019 SCF data. We explicitly consider returns on the following asset classes: stocks, private businesses, real estate, corporate bonds, and government bonds. All other asset classes are assumed to have zero real returns during this period.

All monthly series for asset returns are taken from FRED, from where we report the mnemonics. For stocks and private businesses, we use the S&P 500 (`SP500`); for housing, we use the S&P CoreLogic Case-Shiller U.S. National Home Price Index (`CSUSHPISA`); for corporate bonds, the ICE BofA US Corporate Index (`BAMLCC0A0CMTRIV`); and for government bonds, we construct a return index based on the 10-year Treasury rate (`DGS10`). Finally, we deflate all indices using the CPI (`CPIAUCSL`) and normalize them to one in December 2019. The cumulative return series are shown in Figure Appendix B.1.

We now provide details on how we impute returns in the SCF, which are used in Equation (4.2). The net worth for household i at the beginning of

Figure Appendix B.1: Cumulative real returns on selected asset classes

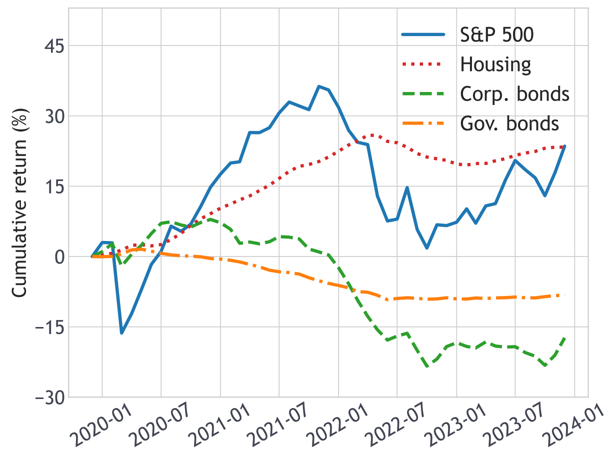

Note: This figure provides cumulative real returns on selected asset classes relative to 2019. We assume that the return on private businesses is the same as for stocks, proxied by the S&P 500.

2019 is given by

$$NW_{i,2019m1} = \sum_{k \in K} A_i^k - B_i,$$

where A_i^k is the dollar value of assets of type k and B_i is debt owed by the household in dollars. The asset classes k that we consider are the ones described above: stocks and private businesses, real estate, corporate bonds, government bonds, and other assets. We proxy for R_τ^k using the publicly available return data described above. Then, given data on realized returns for each of these returns over some period τ, we estimate the net worth over this period as follows:

$$NW_{i,\tau} = \sum_{k \in K} R_\tau^k A_i^k - B_i.$$

This procedure allows us to compute the net return on net worth over the same period as follows:

$$r_{i,\tau}^{NW} = \frac{NW_{i,\tau}}{NW_{i,2019m1}} - 1.$$

We note that this imputation procedure assumes that households are perfectly diversified within each asset class and the composition of asset portfolios in the 2019 SCF remains constant.

Appendix B.3. 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 $\bar{\psi}$.[30] 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:

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

[30] If 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.

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 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

$$\text{penalty} = \begin{cases} \frac{5}{9} \times 0.01 \times 36 + \frac{5}{12} \times 0.01 \times (t-36) & \text{if } t > 36 \\ \frac{5}{9} \times 0.01 \times t & \text{if } 0 \leq t \leq 36, \end{cases} \tag{Appendix B.1}$$

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:

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

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 (Appendix B.1) 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 full 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).

Appendix C. Quantitative results

In this Appendix, we provide details and present additional results related to the estimation of shocks and the main findings presented in Sections 5 and 6.

Figure Appendix C.1: Time series paths for median returns by age

(a) Young agents (b) Old agents

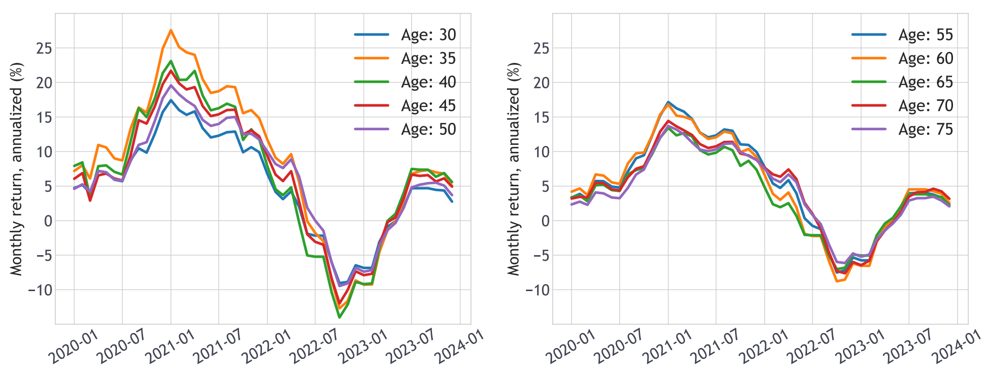

Note: This figure plots median imputed returns for agents of different ages, computed from the SCF.

Appendix C.1. Returns by age

In Section 5.1, we present estimated mean and median of asset return shocks. Here, in Figure Appendix C.1, we provide median returns for agents of different ages, with Panel (a) focusing on younger agents (30 to 50) and Panel (b) focusing on older agents (55 to 75). We show that there is large heterogeneity by age and that younger agents tend to experience higher returns along the transitions than older ones. This is primarily due to the fact that younger agents tend to own larger shares of their wealth portfolio in assets that appreciated substantially during this period, such as housing and stocks, and these agents tend to have more leveraged portfolios (i.e., more debt).

Appendix C.2. Shocks before 2020

In Section 5.1, we present all five shocks after 2019. Here, we provide these shocks prior to 2020. Figure Appendix C.2 plots the asset return and job-separation shocks pre-2020. The key insight from this figure is that both returns and separation shocks are relatively stable prior to the COVID-19 pandemic, which validates our decision to use the pre-pandemic period as a steady state for the model. In addition, job-separation shocks become mostly zero after 2021. By construction, the other shocks are not active during this period, since there were no economic impact payments, additional UI transfers, or additional mortality risk from any source.

Figure Appendix C.2: Time series paths for shocks pre-2020

(a) Asset returns shock

(b) Job-separation rate shock

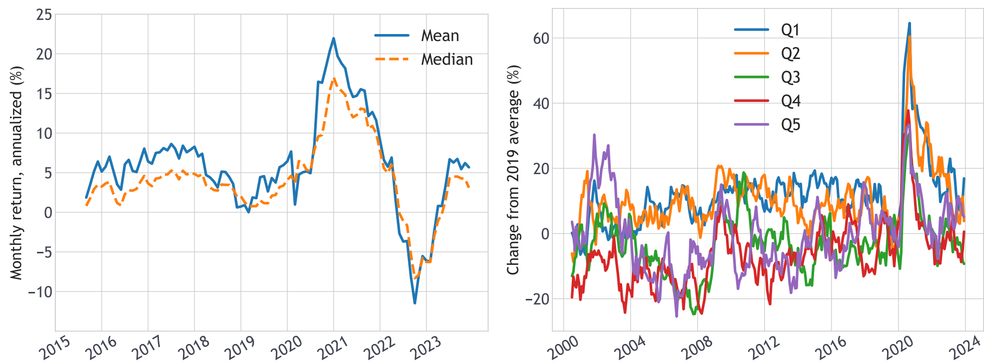

Note: This figure plots series for the estimated shocks prior to 2020, for the mean and median of the asset return shock and the job-separation shock.

Appendix C.3. 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.[31] 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 dependent supplement a householder of a certain age group receives. The 2019 dependent modifiers are provided in the second column of Table Appendix C.1. The effective transfer per eligible individual is then the adult transfer plus dependent supplement times the modifier for that individual's age.

[31] Please refer to America's Families and Living Arrangements: 2019 from `https://www.census.gov/data/tables/2019/demo/families/cps-2019.html`.

Table Appendix C.1: 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.

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 Appendix C.1.

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:

$$T^{adult}_{2021m3} = 1400/1.059 = 1322.00$$
$$T^{child}_{2021m3} = T^{adult}_{2021m3}.$$

Appendix C.4. 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.[32] 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).[33]

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

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

[32] See `https://www.cdc.gov/nchs/nvss/vsrr/covid_weekly/index.htm`.

[33] See `https://ourworldindata.org/grapher/covid-cfr-exemplars`.

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

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

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.

Appendix C.5. Results without housing returns

In our baseline exercise, we compute returns shock using the observed changes in returns for liquid assets such as bonds and stocks and for illiquid assets such as housing. While the appreciation of house prices should create some wealth effects on labor supply, it is insightful to analyze results in this exercise without taking into account house price appreciation during this period, as people may have not realized and/or internalized such capital gains.

In this section, we repeat our exercise but excluding housing returns from the estimated $r_t(a, j)$ function. We present the results for the aggregate labor market moments in Figure Appendix C.3. Clearly, excluding housing appreciation from the exercise slightly moderates the increase in the retired share and therefore the drop in employment-to-population ratio. There is very little effect on the unemployment rate, which is consistent with our baseline results that returns do not seem to play an important role in driving unemployment dynamics.

Figure Appendix C.3: Changes in aggregate labor market moments: No housing returns

(a) Retired share | (b) Unemployment rate | (c) Employment/population

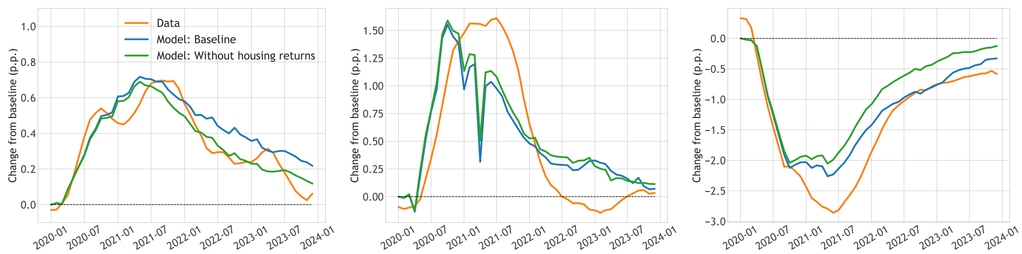

Note: This figure plots the paths of the aggregate retired share (i.e., the fraction of retirees in the population) (Panel (a)), unemployment rate (Panel (b)), and employment-to-population ratio (Panel (c)) in the data and the model. We provide results from two different exercises in the model: the baseline exercise (blue lines) and a version where we do not consider returns on housing (green lines). 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.

Appendix C.6. Model validation along the transition

Section 6.3 in the main text provides results to compare the predictions of our model along the transition with changes in outcomes in the data. In doing so, we discuss two additional results that are not presented in Section 6.3. Here, we now provide these two results. In particular, we compare changes in the average net worth and changes in monthly flow rates into and out of retirement in the model and the data during 2020-2023.

Change in the average net worth. Figure Appendix C.4 plots the average net worth in the SCF for the period in analysis, computed using the imputation procedure described in Section 4.1, and the equivalent wealth series in the model along the transition.[34] We plot percent changes relative to the baseline, which is the average net worth in the 2019 SCF for the data and the stationary state for the model. The model captures the broad movements in the average net worth, slightly overstating its rise after 2021. This result signals that our estimated return function does a good job of matching the evolution of net worth during this period.

[34]In the model, we follow a similar imputation procedure in order to make the model outcome comparable with the data, taking the initial joint distribution of age and net worth, and iterating forward using the estimated return function $r_t(a, j)$. In particular, for the purposes of this figure only, we do not use the model's decision rules as we cannot account for changes in consumption/savings behavior either in the data.

Figure Appendix C.4: Change in average wealth along the transition: Data vs model

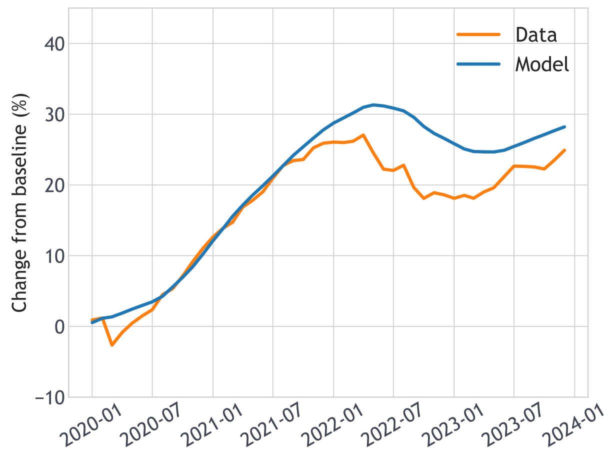

Note: This figure plots the change in the average net worth during 2020-2023 in the data and the model. The data series are computed using the imputation procedure described in Section 4.1. The model series is obtained under a similar imputation procedure to make the two series comparable.

Changes in monthly flow rates into and out of retirement. Figure Appendix C.5 compares changes in monthly flow rates into (Panel a) and out of (Panel b) retirement in the data and model. To compute the monthly flow rate into retirement in the data, we use CPS and measure the ratio of the number unemployed or employed individuals in a given month t who become retired in the next month $t+1$, to the total number of unemployed or employed individuals in t. Similarly, we compute the monthly flow rate out of retirement by calculating the ratio of the number of retired individuals in t who become unemployed or employed in $t+1$, to the total number of retired individuals in t. We then repeat these calculations for each month. We compute the same moments in the model using the same steps. We then compute pp changes from the average flow rates in 2019 in the data and from the average flow rates in the stationary state of the model.

Panel (a) plots changes in the monthly flow rate into retirement in the data and model.[35] The model replicates well the initial spike in 2020, matching both the level and the dynamics.

[35]Monthly flow rates into and out of retirement in the model are volatile during the transition period mostly because of observed fluctuations in job-separation rate shocks by quintiles of the labor income distribution, as shown in Panel (b) of Figure 5.1.

Figure Appendix C.5: Changes in flow rates into and out of retirement: Data vs model

(a) Flow rate into retirement (b) Flow rate out of retirement

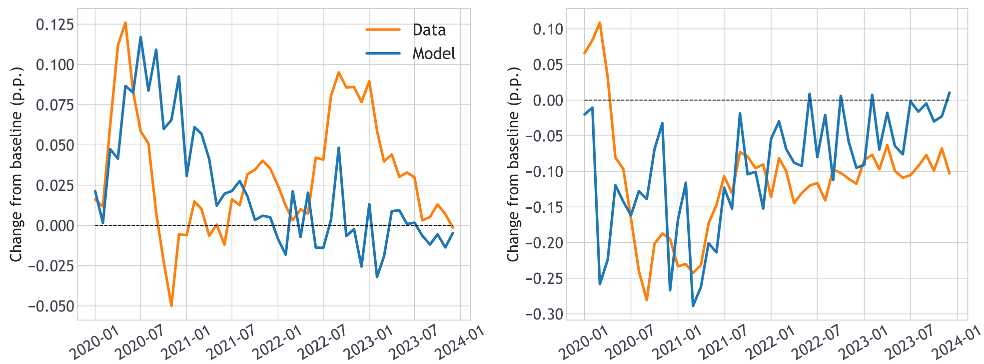

Note: This figure compares changes in monthly flow rates into (Panel a) and out of (Panel b) retirement in the data and model. To compute the monthly flow rate into retirement in the data, we use CPS and measure, for each month, we compute the ratio of the number unemployed or employed individuals in a given month who become retired in the next month, to the total number of unemployed or employed individuals in that month. We obtain the monthly flow rate out of retirement in a similar manner. The model calculations follow the same steps. These figures present pp changes from the average flow rates in 2019 in the data and from the average flow rates in the stationary state of the model.

The model fails to account for the observed rise in late 2022. Notice that this rise in the flow rate into retirement in the data is reflected in Panel (a) of Figure 5.2 where the retired share in the data starts to rise after 2022 until early 2023. The model is unable to capture this increase in the data because it is driven by people younger than 62 retiring, and our definition of retirement in the model includes only people older than 62. To see why, compare the evolution of excess retirements per our baseline definition in Panel (b) of Figure 2.1 (based on self-reported retirement in the CPS) to that of Panel (b) of Figure Appendix A.2, where we consider an alternative definition of retirement that includes non-participants aged 62 and older. Notice that while the baseline definition features an increase in the retired share in late 2022, the alternative definition does not. Given the definition of retirement in the model, we are therefore unable to capture this rise by construction.

Similarly, Panel (b) shows that the model does a good job in matching flows out of retirement: the initial decrease in 2020, and then slow recovery back to the baseline (steady state/pre-pandemic) level.