Dissecting the Great Retirement Boom*

Serdar Birinci
FRB St. Louis

Miguel Faria-e-Castro
FRB St. Louis

Kurt See
Bank of Canada

July 2024

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, such as the rise in net worth due to higher asset returns, the labor market's deterioration due to higher unemployment risk, the expansion of fiscal support programs, and increased mortality risk. We quantitatively study the interaction of these factors and decompose their relative contribution to the recent rise in retirements using an incomplete markets, overlapping generations model with a frictional labor market. We find that all of these channels contribute to excess retirements, with labor market conditions being a more important driver in 2020-2021 and fiscal programs playing a larger role in 2022-2023. We show that our model's predictions on aggregate labor market moments and cross-sectional moments on retirement patterns across the wealth distribution are in line with the data.

Keywords: Retirement, Labor Supply, Labor Flows, Financial Markets
JEL Classification: E24, G11, J21, J22, J26

*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, the Federal Reserve System, or the Bank of Canada. Birinci: serdar.birinci@stls.frb.org, Faria-e-Castro: miguel.fariaecastro@stls.frb.org, See: seek@bankofcanada.ca.

1 Introduction

The increase 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 policymakers (Hobijn and Şahin, 2021; Powell, 2022; Montes et al., 2022). In late 2021, the fraction of retired individuals in the working-age population increased 0.7 percentage points (pp) over what the pre-pandemic trend predicts—a gap corresponding close to 2 million more retirements. This phenomenon has hampered the recovery of the U.S. labor force participation rate (LFPR), which was still 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) changes in labor market conditions, (ii) wealth effects arising from elevated returns on assets, (iii) economic impact payment programs, (iv) expansion of unemployment insurance (UI) programs, and (v) increased mortality risk especially for older people. In this paper, we develop a unified approach to quantitatively study the interaction of these factors and decompose their relative contribution to the rise in retirements. In doing so, this paper offers three main contributions.

First, we present novel facts regarding the relationship between retirement decisions and wealth levels before and after the COVID-19 episode. In particular, using micro data, we calculate fractions of new retirees across wealth quintiles, separately for those who retired in 2019 and those who retired between 2020 and 2021. While the fraction of new retirees increases in wealth quintiles, we find that, compared with the distribution of new retirees across wealth quintiles in 2019, the post-COVID-19 episode is not characterized by an increase in the fraction of new retirees with higher levels of wealth. In fact, 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. This outcome occurred during a period marked by a large increase in average wealth and by a reduction in inequality in the wealth distribution. Overall, this empirical result is informative for our decomposition analysis using our quantitative model, as it suggests that the increase in retirements were not driven by wealthier individuals.

Second, we construct a heterogeneous agents model that allows us to account for potential factors behind the rise in retirements: an incomplete markets, overlapping generations model with a frictional labor market that incorporates direct flows in and out of retirement. We calibrate this model to the U.S. economy in 2019 and use it to implement quantitative experiments to study recent labor market dynamics. This is important since, to the extent of our knowledge, there is no relatively high frequency dataset that allow 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 COVID-19 episode, and mortality outcomes. This makes it necessary to use a model to understand recent retirement dynamics. However, we make sure to calibrate this model and validate its predictions before and after the COVID-19 episode using

available dimensions of micro data, including labor market flows, distributions of wealth and labor earnings, and the distribution of new retirees across wealth quintiles.

Third, we use the model to decompose the importance of each channel between 2020 and 2023. We first show that the model captures both the magnitude and the persistence of movements of untargeted aggregate labor market moments in the data, such as the retired share (i.e., the excess retirements relative to trend), the unemployment rate purged temporarily unemployed, and the employment-to-population ratio.[1] Next, our decomposition exercises reveal that four channels (all five channels we consider, except the expansion of the UI program) played a role in driving excess retirements, with the deterioration of labor markets due to the rise in separations being a more important driver in 2020-2021 and economic impact payments playing a larger role in 2022-2023. On the other hand, the expansion of the UI program was as important as the rise in separations in terms of generating the initial increase in the unemployment rate, while the rise in separations was key to match its persistence. Crucially, our model's cross-sectional predictions are also aligned with changes in untargeted moments from micro data during 2020-2023 relative to 2019. Specifically, the model broadly matches the rise in average net worth, the reduction of inequality in net worth driven primarily by a faster growth of net worth at the bottom of the distribution, and changes in the distribution of new retirees across the net worth distribution.

Underlying these results is a partial-equilibrium, incomplete markets, overlapping generations model where agents decide on their employment and participation status and face a consumption-savings problem. The model features a frictional labor market where agents (i) suffer a utility cost to elicit job offers that may occur with a certain probability and be associated with an uncertain wage and (ii) face exogenous job separation risk. Employed agents earn a constant wage that is drawn when a match is formed. They face an exogenous job separation risk, according to a certain rate that depends on their wage level, and also endogenously choose to stay employed or quit to any other employment status (i.e., unemployment, non-participation, or retirement). When agents choose not to search and drop out of the labor force, they pay no utility cost but face a lower probability of obtaining a job offer. Finally, when agents are old enough, they may choose to retire, in which case they do not participate in the labor market but become eligible for social security benefits that depend on their age, and may also receive job offers at a lower rate. As such, the model incorporates direct flows in and out of retirement. We focus on the decision problems of households and on the macroeconomic variables that are implied by the aggregation of these individual decisions under an endogenous distribution.

We calibrate this model to the U.S. economy in 2019. In particular, we choose parameters to match a series of moments that are related to distributions of wealth and labor earnings as well as the aggregate unemployment rate, LFPR, retired share, and labor market flow rates. Next,

[1]For brevity, from this point onward, when refer to the increase in the retired share, we mean the increase in the retired share relative to trend (i.e., the excess retirements).

we validate the predictions of the model's stationary state against untargeted moments from the micro data in 2019 that are relevant for the economic forces we consider. In particular, we show that the model captures the shape of the wealth distribution and distributions of new retirees across wealth and income distributions. This last point is especially important as it means that the model can adequately capture the financial characteristics of the agents who decide to retire.

Our main quantitative 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 micro data during 2020-2023. First, we capture the rise and heterogeneity in returns to wealth by replacing the baseline level of returns on savings with a function that determines returns on wealth depending on both the level of wealth and the age of the agent. We argue that this is a parsimonious way of capturing, in a single asset model, the heterogeneity in returns that results from the fact that households with different wealth and of different ages hold different assets in their portfolios, and whose returns differed greatly during 2020-2023. Overall, we show that average returns rose substantially during 2020-2021, fell during 2022, and then recovered in 2023. Second, we capture deterioration of labor market conditions early in the pandemic by feeding shocks to job separation rates. Importantly, we measure changes in separation rates by each wage quintile in the data, allowing us to capture the fact that labor market conditions deteriorated more for people with more physical contact-intensive jobs who tend to be at the bottom of the wage distribution. Indeed, we show that job separation rates increased by much more at the bottom two wage quintiles during 2020 when compared with those at higher quintiles, and they all declined sharply at the end of the year, only to then slowly stabilize over the next few years. Third, we map the two major fiscal transfer programs that were deployed during 2020-2023 into our model: the three rounds of economic impact payments and the expansion of UI benefits. Finally, we measure changes in mortality risk by age group in the data and feed them into the model.

Next, we analyze the predictions of the model for the movements of aggregate labor market moments that are driven by these five exogenous shocks. We show that the model matches the magnitude and dynamics of the movements in the data for the retired share, the unemployment rate purged temporarily unemployed, and the employment-to-population ratio. We see this as the second element of model validation, as none of these aggregates are targeted in our exercise. Given that the model broadly reproduces the changes in aggregate labor market moments, we proceed to decompose the increases in the retired share and the unemployment rate. Starting with the retired share, we find that labor market disruptions through increased job separation rates were the primary factor driving the increase in the retired share during 2020, while increased mortality rates were only essential to explain the initial increase during the first months of 2020. These forces are not sufficient, however, to explain why excess retirements remained high in

following years, as improving labor markets should have brought it down faster than what was observed in the data. We find that the persistence of the rise in the retired share is explained by the provision of economic impact payments and, to a lesser extent, elevated returns on assets, in spite of improving labor market conditions post-2021. Overall, we show that all these four forces are necessary to match the empirical dynamics in the retired share. In particular, changes in mortality and labor market dynamics alone cannot explain the persistence of the rise, while economic impact payments and changes in asset returns alone predict a much smaller increase during 2020-2021. Moving to the decomposition of the rise in the unemployment rate, we find that while the expansion of UI benefits plays a negligible role in explaining excess retirements, it was important to account for the initial increase in the unemployment rate. In this case, elevated job separation rates help explain the subsequent dynamics in the unemployment rate.

As a third element of model validation, we compare the cross-sectional predictions of the model along the transition to changes in relevant moments from the micro data between 2020-2023 relative to 2019. We find that the model is able to broadly account for the rise in the average net worth, the compression of the wealth distribution (driven mostly by the rise in net worth at the bottom of the distribution), and changes in the distribution of new retirees across the net worth distribution. The model predicts that the share of new retirees at the bottom quintile slightly increases, while the share of new retirees at the top quintile falls slightly, leading to less wealth inequality in the new cohorts of retirees relative to 2019, as in the data. 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 those facing worse labor market prospects and who were relatively more sensitive to fiscal transfers.

We conclude our analysis with a brief discussion on the implications of our results for incomplete markets models, where agents internalize the impact of their wealth in the rate of return and endogenously supply labor. We find that, for our exercise to generate reasonable predictions, it is important to account for return heterogeneity as a function of wealth, but in a way that agents do not internalize such dependence. We achieve this by assuming that agents are surprised by fluctuating ex-post returns but expect returns on wealth to be back to their constant, stationary state level in the future. This allows us to capture the wealth effects that result from heterogeneous returns, without significant distortions in labor supply and savings decisions. We show that if we repeat our baseline exercise while allowing agents to internalize that returns are increasing in wealth, we obtain macro and micro predictions that are at odds with the data: excess retirements fall substantially in the early stages of the transition and then rise by about 14 pp, versus an increase of about 0.7 pp in the data and the baseline model. This also generates counterfactual dynamics for average net worth, which rises by about 60% relative to the pre-pandemic baseline, as well as for the distribution of new retirees across wealth quintiles. This model also predicts that shares of new retirees across wealth quintiles in 2020-2021

are mostly equal, which is also at odds with the data and predictions of the baseline model.

Related literature. This paper contributes to an extensive literature on retirement patterns and economic decisions of retired individuals in terms of both consumption and savings (De Nardi et al., 2010, 2016) as well as labor supply. Our paper is particularly related to the strand of this literature that focuses on wealth effects of labor supply for those who are retired or close to retirement (Cheng and French, 2000; Coronado and Perozek, 2003; Benson and French, 2011). Relative to this literature, we develop an incomplete markets, overlapping generations model combined with a labor search framework. This allows us to analyze new channels that have not been the focus of previous studies on retirement decisions: how labor market frictions and changes in various fiscal transfers—that impact the magnitude of the surplus from employment relative to non-employment—affect retirement decisions.

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.

Finally, our paper contributes to an active and growing literature on heterogeneous returns on wealth, either documenting that they tend to be increasing in wealth in a variety of contexts (Bach et al., 2020; Fagereng et al., 2020; Smith et al., 2022; Ozkan et al., 2023) or studying their implications in the context of structural models without analyzing how heterogeneous returns affect labor supply decisions (Benhabib and Bisin, 2018; Benhabib et al., 2019; Xavier, 2021). We study a version of our model where agents understand that returns on their savings are a function of their level of wealth. We show that heterogeneity in returns substantially affects individual labor supply decisions and leads to counterfactual predictions when compared with our moments of interest in the data. As such, we view our results as a cautionary note for structural models with heterogeneous returns on wealth.

The rest of the paper is organized as follows. Section 2 analyzes recent changes in aggregate retirement dynamics and retirement patterns across the wealth distribution in the data. Section 3 presents the model, and Section 4 describes the calibration exercise and discusses validation of model predictions at the stationary state. Section 5 explains our main experiment and presents the results regarding the aggregate labor market dynamics. Section 6 presents the results from decomposition and validation exercises and validation of model predictions along the transition. Section 7 discusses implications of heterogeneous returns on wealth for models with endogenous labor supply, and Section 8 concludes.

2 Excess Retirements in 2020-23: Aggregate Dynamics and Micro Heterogeneity

In this section, we first present empirical trends in the aggregate fraction of the population that is retired in the U.S. since 2000 and document the rise in this fraction between 2020 and 2023. Next, to understand the underlying sources behind this rise, we study retirement patterns across the wealth distribution during this episode using micro data.

2.1 Increased retirements post-2020 in the aggregate

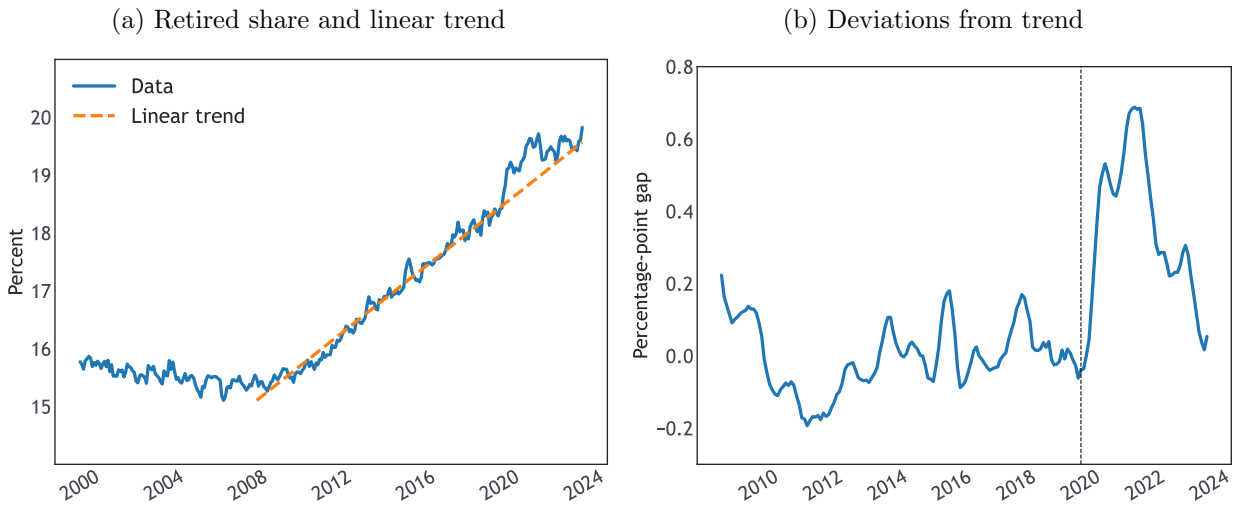
The U.S. LFPR experienced its largest drop on record at the onset of the COVID-19 pandemic in early 2020, falling from 63.3% in February 2020 to 60.1% in April 2020. While there was a quick rebound from this 50-year minimum, it has not fully recovered to its pre-pandemic levels, standing at 62.5% as of May 2024. The LFPR for prime-age workers (25-54 years) has actually exceeded its pre-pandemic level (83.6% in May 2024 vs. 83.1% in January 2020), and most of the 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. Thus, it is necessary to analyze participation dynamics for relatively older workers to understand the aggregate LFPR dynamics during this episode.

Several authors have documented a significant increase in the share of the population that is retired over the same period (Nie and Yang, 2021; Faria-e-Castro, 2021b), an apparent departure from what previous trends would suggest. Panel (a) of Figure 1 plots the retired share in the U.S. economy, from 2000 to the end of 2023, which we calculate as the fraction of individuals who report as retired in the Current Population Survey (CPS) among all individuals (excluding those in armed forces) aged 16 and over.[2] The figure shows that the percentage of this population that is retired was roughly constant (or slightly declining) up until the late 2000s, when it started growing at a seemingly linear trend (dashed line); this is estimated between June 2008, when the oldest baby boomers start retiring, and January 2020, the last full month before the effects of the COVID-19 pandemic were felt in the US economy. The rise in the retired share is plausibly related to demographic factors: 2008 was the first year in which members of the Baby Boomer generation became eligible to retire and collect Social Security benefits. Since U.S. population growth has slowed down in recent decades, this year marks the beginning of a period in which a significant share of the active population begins to retire.

There is a significant gap between the linear trend and the actual retired share between 2020 and 2023. Panel (b) of Figure 1 plots the gap between the two lines: while the linear trend is a good approximation until 2020, the retired share increased by 0.7 percentage points (pp)

[2]Appendix B.1 presents details on the construction of the data and also shows that our measure of retired share is robust to alternative definitions of retirement: (i) focusing only on retirees who are old enough to be eligible for social security benefits, and (ii) focusing on all social security-eligible people who are out of the labor force. There results are provided in Figures B.1 and B.2 respectively.

Figure 1: Excess retirements between 2020 and 2023



Note: Panel (a) plots the retired share in the U.S. economy which we calculate as the fraction of individuals who report to be retired in the Current Population Survey (CPS) 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.

over the trend in late 2021. This gap corresponds to close to 2 million people who were retired beyond what pre-pandemic trend predicts. We call this gap between the two lines as "excess retirements." Analyzing the drivers of this gap is the main focus of this paper.

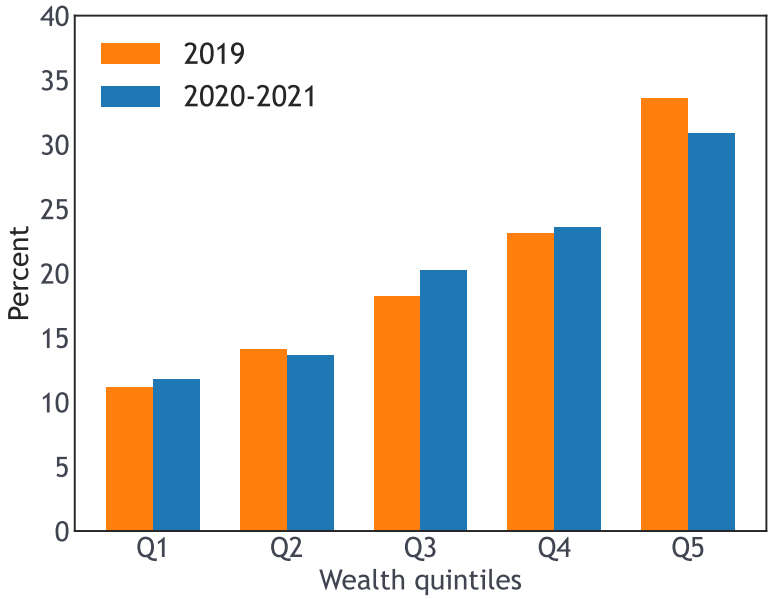
2.2 Post-2020 retirements across the wealth distribution

To quantify the role of potential drivers on this gap, we need to analyze which worker groups experienced a larger increase in their retired shares. For instance, one of the key potential drivers of excess retirements in the post-2020 period is changes in net worth, which may have triggered wealth effects and led people to retire or remain in retirement (Faria-e-Castro and Jordan-Wood, 2024). We now present empirical results on the relationship between wealth holdings and retirement patterns during this period. Later, we will use these empirical findings to validate the predictions of our quantitative model.

As the CPS does not provide information on wealth holdings, we use data from the 2020, 2021, and 2022 panels (covering data from the start of 2019 to the end of 2021) of the Survey of Income and Program Participation (SIPP) collected by the US Census Bureau, which provide information on employment status and wealth holdings.[3] Using this data, we first select a sample of individuals not in the armed forces aged 16 and over, as in our CPS sample. We then assign

[3]We note that we take the magnitude of excess retirements obtained from the CPS as our baseline estimate for two reasons. First, as discussed by Krusell et al. (2017) and Birinci and See (2023), monthly transition rates between employment statuses are underestimated when calculated from the SIPP compared with those obtained from the CPS. Second, the most recent available panel of the SIPP (i.e., the SIPP 2022) covers the reference period until December 2021, preventing us studying the aggregate retirement dynamics after 2021. However, the increase in the retired share between 2020 and 2021 is also observed in the SIPP, allowing us to draw conclusions from the underlying retirement patterns across the wealth distribution between pre and post COVID-19 episode.

Figure 2: Distribution of new retirees across the wealth distribution: 2019 vs 2020-2021



Note: This figure plots fractions of new retirees across wealth quintiles, separately for those who retire in 2019 and those who retire between 2020 and 2021 using data from the SIPP. Our sample consists of non-armed individuals aged 16 and over. Our measure of wealth is the total household-level of net worth.

these individuals into quintiles of the wealth distribution at each year, where our measure of wealth is the total household-level of net worth. Next, we identify individuals who, as in the CPS, report as retired in 2019 or between 2020 and 2021.[4] Figure 2 plots fractions of new retirees across wealth quintiles, separately for those who retire in 2019 and those who retire between 2020 and 2021. The comparison of retirement patterns across wealth quintiles between 2019 and 2020-2021 allows us to document whether the distribution of new retirements has changed before and after the COVID-19 pandemic.

Figure 2 shows that there is an increasing relationship between wealth quintiles and the share of new retirees in both episodes: higher quintiles cover a larger share of recent retirees. Our focus, however, is in the subtle yet important changes in this distribution between 2019 and 2020-2021. In particular, we find that the share of new retirees in lower quintiles increased by 0.7 pp in the first quintile and 2 pp in the third quintile between these two episodes, while the share of new retirees in the fifth quintile fell by 2.7 pp. As such, we conclude that, compared with the distribution of new retirees pre-COVID-19, the post-COVID-19 episode is not characterized by an increase in the fraction of new retirees with higher levels of wealth. If anything, 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.

[4]Appendix B.2 provides details on the construction of the data and these variables.

2.3 Post-2020 trends in the wealth distribution

Figure 2 is silent about changes in the level and distribution of wealth, however. This is particularly relevant due to significant fluctuations in realized asset returns after 2019. Thus, we now analyze changes in the level and the distribution. In terms of analyzing movements in various moments of the wealth distribution, we face a trade-off that is related to data availability. While the annual SIPP allows us to track movements in these moments over time, these data tend to understate levels and the extent of inequality in the wealth distribution. The Survey of Consumer Finances (SCF) conducted by the Board of Governors of the Federal Reserve corrects this issue at the cost of decreased frequency since it is typically a triennial survey. As such, we provide various moments of the wealth distribution from both of these two datasets.[5]

Panel (a) of Figure 3 presents deciles of the net worth distribution relative to the average level of net worth in 2019, separately using SIPP and SCF data. Overall, there is large inequality in the net worth distribution in both datasets, but the level of wealth and the extent of heterogeneity is relatively lower in the SIPP than in the SCF. For instance, the ratio of median to mean net worth is 0.53 in the SCF and 0.25 in the SIPP, while the ratio of the 90th percentile to mean net worth is 2.75 in the SCF and 2.47 in the SIPP.

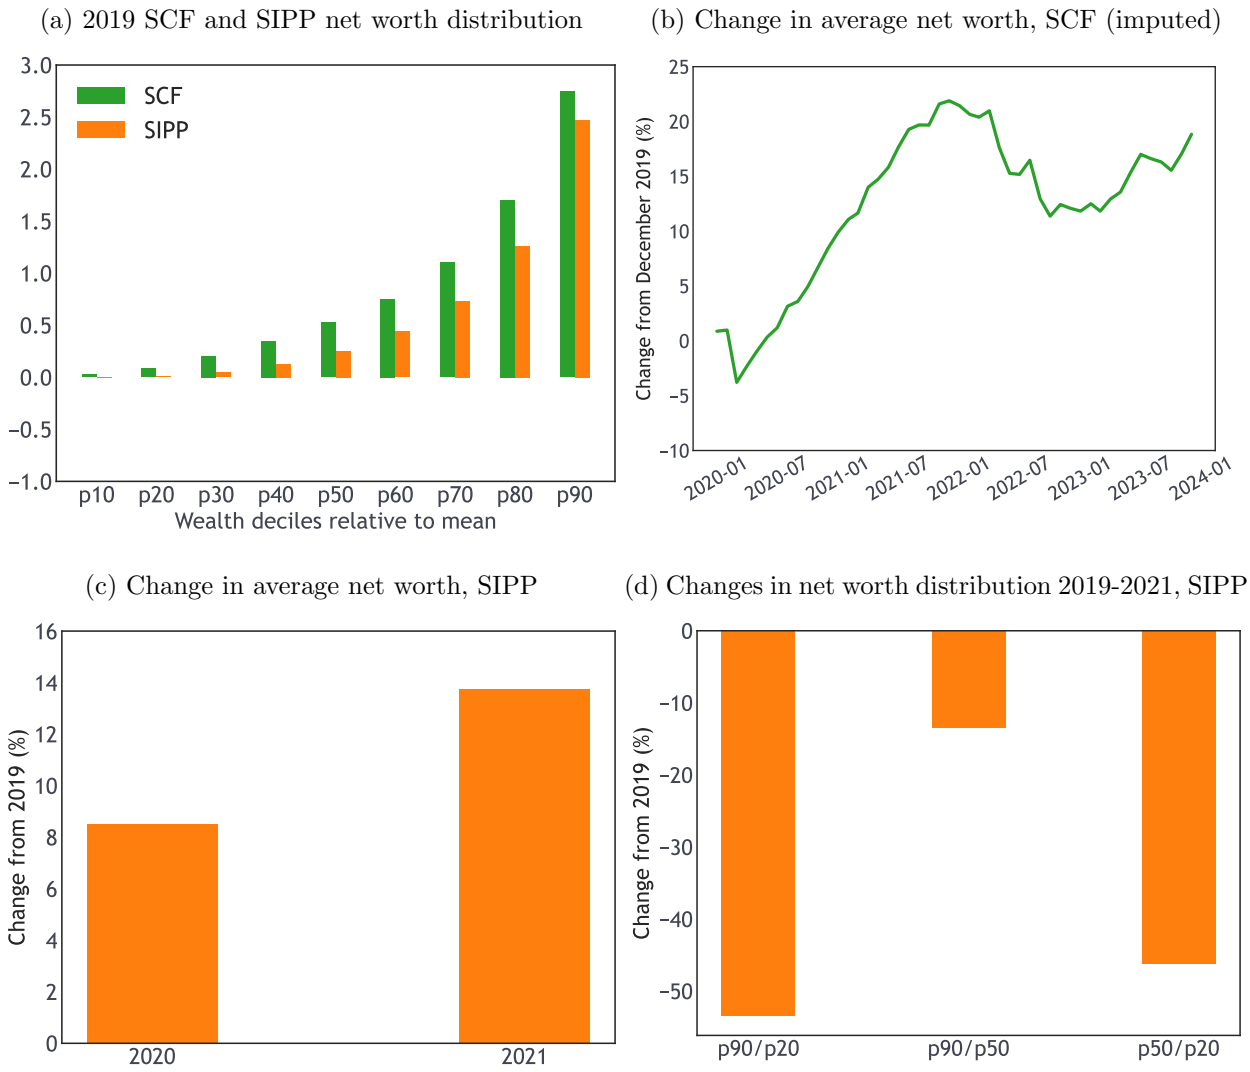
Panels (b) and (c) measure the percent change in the average net worth during the post-2020 period relative to that in 2019, with data from the SCF and SIPP, respectively. Panel (b) computes a measure of cumulative changes in the average net worth from the 2019 SCF that result from imputing realized returns to different asset classes after 2019. This procedure is described in detail in Section 4.2. Panel (b) shows that in spite of a slight decrease at the onset of the pandemic, the average net worth for U.S. households grew almost continuously through the end of 2021, at which point it had experienced 21.9% growth relative to December 2019. Panel (c) shows a similar result from the SIPP using the actual wealth changes, with the average net worth growing 8.5% in 2020 and 13.7% in 2021 relative to that in 2019.

Finally, Panel (d) presents percent changes in ratios of various percentiles of the net worth distribution from the SIPP between 2019 and 2021, a convenient way to summarize changes in the shape of the distribution over time. We look at changes in three percentile ratios: the p90/p20 ratio, which measures inequality between the top and bottom of the distribution, the p90/p50 ratio, which measures inequality between the top and the median, and the p50/p20 ratio that measures inequality between the median and the bottom.[6] Panel (d) shows a reduction in the inequality of net worth distribution across the board between 2019 and 2021, with the main driver being an increase of the 20th percentile relative to the 50th and the 90th percentiles. This

[5]Details of the SCF data and variable construction are provided in Appendix B.3.

[6]We note that the 10th percentile of the net worth distribution in 2021 is exactly zero. For this reason, while it is probably more natural to look at the p90/p10 ratio, we instead present the percent change in the p90/p20 ratio between 2019 and 2021. However, in terms of level changes, the increase in the 20th percentile between 2021 and 2019 is very similar to the increase in the 10th percentile between this episode.

Figure 3: Changes in the average and distribution of net worth since 2019



Note: Panel (a) presents deciles of the net worth distribution relative to the average level of net worth in 2019, separately using SIPP and SCF data. Panel (b) computes a measure of cumulative changes in the average net worth from the 2019 SCF that result from imputing realized returns to different asset classes after 2019. Panel (c) shows a similar result from the SIPP using the actual wealth changes in 2020 and 2021 relative to 2019. Panel (d) presents percent changes in ratios of various percentiles of the net worth distribution from the SIPP between 2019 and 2021. Our measure of wealth is the total household-level of net worth.

suggests that the distribution became relatively more compressed as a relatively larger mass moved from the left tail toward the median: both the p90/p20 and p50/p20 ratios decreased by more than 45% between 2019 and 2021, while the p90/p50 ratio declined by only 13%.

Next, we present a decision model of retirement that captures the joint distribution of retirement and wealth in Figure 2 in 2019 and the distribution of net worth across the population in Panel (a) of Figure 3 in 2019. Features of the data such as the increase in the retired share between 2020 and 2023 and movements in the wealth distribution or in the joint distribution of wealth and retirement after 2019 will not be explicitly targeted when disciplining this model and will instead be used to validate the predictions of our model.

3 Model

3.1 Overview

We combine a partial-equilibrium heterogeneous-agents incomplete markets overlapping generations model with a labor search framework to develop a model with the minimal set of ingredients that we believe allow us to quantify contributions of various sources behind the rise in the retired share between 2020 and 2023. Individuals start at age 16 and make consumption, savings, and labor supply decisions over their lifecycle in a frictional labor market. In particular, individuals observe their (i) wage—drawn from a distribution—in the labor market, (ii) job-finding and job-separation rates, (iii) real returns in their net worth, and (iv) mortality risk, and then decide on consumption, savings in a single composite asset, and labor supply at the extensive margin. We allow the exogenous job-separation rates to vary depending on the current wage (Birinci and See, 2023) and the mortality risk to vary with age. Labor supply decisions involve moving out of the labor force for employed and unemployed individuals, and moving into the labor force for non-participants. Importantly, we do not model retirement as a permanent state and allow existing retirees to unretire and search for new jobs (Nie and Yang, 2021; Hobijn and Şahin, 2022). Overall, individuals are heterogeneous with respect to their age, employment status (employed, unemployed, non-participant, or retired), wages, and the level of wealth.

3.2 Environment

Time is discrete and infinite. The economy is populated by a stationary mass of overlapping generations of agents. We focus on the decision problem of each of these agents and aggregate these decisions to compute macroeconomic and labor market variables. Agents are indexed by four states: age $j \in \{16, \ldots, 90\}$, wealth $a \in [-\underline{a}, \infty)$, employment status $\ell \in \{E, U, N, R\}$, and wage $w \in \mathbb{R}^+$. Agents enter their lifecycle at age 16, face an age-dependent probability of dying given by $1 - \pi(j)$ (where $\pi(j)$ is the survival probability of an agent aged j), and die with probability 1 once they reach the age of 90. In terms of the employment status, individuals can be employed E, unemployed U, out of the labor force N, or retired R.

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. All agents have access to a risk-free asset that pays an exogenously given return $\bar{r}$ on savings and $r^b > \bar{r}$ on borrowings. The return

function is thus given by[7]:

$$r(a) = \begin{cases} \bar{r}, & \text{if } a \geq 0 \\ r^b, & \text{otherwise.} \end{cases}$$

We now describe the problem of agents at each employment status.

Employed. Employed agents receive a constant wage w that is drawn upon receiving but before accepting a job offer (more on this below, when discussing the unemployed). They may exogenously separate from their job with probability $\delta(w)$, which depends on their wage level, or endogenously choose to quit from their jobs. Their problem is given by:

$$\begin{aligned} V^E(j,a,w) = \max_{c,a'} \quad & u(c,\ell=E,j) + \beta\pi(j)\delta(w)\max\{V^U(j+1,a'),V^N(j+1,a'),V^R(j+1,a')\} \\ & +\beta\pi(j)[1-\delta(w)]\max\{V^E(j+1,a',w),V^U(j+1,a'),V^N(j+1,a'),V^R(j+1,a')\} \\ \text{s.t. } & c + a' = w + [1+r(a)]a + T \\ & a' \geq -\underline{a}, \end{aligned}$$

where $\underline{a}$ is the borrowing constraint and T are government transfers. Government transfers T, which are separate from UI benefits b, are indented to capture three rounds of economic impact payments provided to individuals in the U.S. during the COVID-19 episode.[8] The dependence of the probability of separation on the wage, $\delta(w)$, reflects the fact that agents have different probabilities of separation across the wage distribution. If a separation occurs, the individual can choose to become unemployed, leave the labor force, or retire. If no exogenous separation takes place, she can choose to either stay in the current job or move to any of these other statuses.

Unemployed. Unemployed agents derive income from home production h and from unemployment benefits b. Their problem is given by:

$$\begin{aligned} V^U(j,a) = \max_{c,a'} \quad & u(c,\ell=U,j) + \beta\pi(j)(1-f)\max\{V^U(j+1,a'),V^N(j+1,a'),V^R(j+1,a')\} \\ & +\beta\pi(j)f\int_w \max\{V^E(j+1,a',w),V^U(j+1,a'),V^N(j+1,a'),V^R(j+1,a')\}\mathrm{d}G(w) \\ \text{s.t. } & c + a' = h + b + [1+r(a)]a + T \\ & a' \geq -\underline{a}. \end{aligned}$$

[7]Notice that our baseline model is a single-asset model whose rate of return does not depend on the level of wealth. Having a return function that depends on wealth is, in principle, an attractive way of capturing portfolio heterogeneity across wealth levels without having to include additional state and control variables. In our transition experiments, we allow returns to vary across individuals depending on their level of wealth. However, having this feature at the stationary state generates counterfactual predictions for labor supply as we discuss in greater detail in Section 7.

[8]We note that, in our transition experiments, these transfers will depend on the income level and age of individuals. We provide more details on disciplining these transfers later in Section 5.1.

The problem is similar to that of an employed agent, with the difference being that, at the beginning of the next period, the agent receives a job offer with probability f. When a job offer is received, the agent draws a wage from a distribution G and decides whether to accept the offer and become employed at that wage or choose to remain unemployed, leave the labor force, or retire. If the agent does not get a job offer, then she can still choose between unemployment, non-participation, and retirement.

Out of the labor force. Agents who are out of the labor force receive income from home production only. The other differences with respect to unemployment is that they do not suffer disutility from unemployment (i.e., job search) but have a lower probability of receiving offers. To capture the latter, we assume that only γ_N fraction of these individuals may receive a job offer, which again only happens with probability f.[9] The problem of this agent is given by:

$$\begin{aligned} V^N(j,a) = \max_{c,a'} \quad & u(c, \ell = N, j) + \beta\pi(j)(1-\gamma_N)\max\{V^U(j+1,a'), V^N(j+1,a'), V^R(j+1,a')\} \\ & + \beta\pi(j)\gamma_N f \int_w \max\{V^E(j+1,a',w), V^U(j+1,a'), V^N(j+1,a'), V^R(j+1,a')\}\mathrm{d}G(w) \\ & + \beta\pi(j)\gamma_N(1-f)\max\{V^U(j+1,a'), V^N(j+1,a'), V^R(j+1,a')\} \\ \text{s.t.} \quad & c + a' = h + [1 + r(a)]a + T \\ & a' \geq -\underline{a}. \end{aligned}$$

Retirement. Retirement is similar to being out of the labor force with three differences. First, only agents age 62 and over are allowed to retire. This is the minimum age at which agents become eligible for Social Security benefits and thus the earliest age at which retirement becomes meaningfully different from non-participation. Second, retirees earn Social Security benefits. We allow these these benefits to depend on age, as a parsimonious way of capturing the fact that these benefits depend on the age of retirement. We later show how we calibrate this function to match the U.S. Social Security Administration (SSA) formula that is used to compute retirement benefits. Third, the fraction of retired individuals who may receive a job offer while retired is γ_R instead of γ_N for those out of the labor force. The problem of retirees is given by:

$$\begin{aligned} V^R(j,a) = \max_{c,a'} \quad & u(c, \ell = R, j) + \beta\pi(j)(1-\gamma_R)\max\{V^U(j+1,a'), V^N(j+1,a'), V^R(j+1,a')\} \\ & + \beta\pi(j)\gamma_R f \int_w \max\{V^E(j+1,a',w), V^U(j+1,a'), V^N(j+1,a'), V^R(j+1,a')\}\mathrm{d}G(w) \\ & + \beta\pi(j)\gamma_R(1-f)\max\{V^U(j+1,a'), V^N(j+1,a'), V^R(j+1,a')\} \\ \text{s.t.} \quad & c + a' = h + y^{SS}(j) + [1 + r(a)]a + T \\ & a' \geq -\underline{a}. \end{aligned}$$

[9] This assumption allows us to capture direct transitions from non-participation to employment in the data.

Retirees are allowed to leave retirement by transitioning either to unemployment or non-participation or to employment if they receive an offer with probability $\gamma_R f$.

Death. At age $j = 91$, all agents die with probability 1 and obtain zero value, $V^\ell(j = 91, a, w) = 0, \forall(a, \ell, w)$. They are replaced with newborns. We abstract from bequests.

Birth. Agents enter the model at age $j = 16$, drawing their initial wealth from a distribution $Q(a)$. We assume that agents enter the model in the unemployment state.

3.3 Stationary distribution

We focus on macroeconomic variables that result from the aggregation of the individual decisions. There are no markets to clear in this model. Let $\lambda_t(j, a, w, \ell)$ denote the distribution over individual states. We calibrate the model to capture various data moments of the U.S. economy in 2019, which we interpret as the model's stationary state. At the stationary state, the distribution is such that it solves the fixed-point of the following equation:

$$\lambda(j, a, w, \ell) = \mathcal{T}[\lambda(j, a, w, \ell)],$$

where $\mathcal{T}$ is the transition function, which is described in more detail in Appendix A.1.

4 Calibration

We now describe the calibration as well as several important predictions of the model at the stationary state. The model period is a month. Our calibration strategy involves setting a few parameters externally and then internally calibrating most other parameters to match a series of moments related to labor market and demographic outcomes as well as income and wealth distributions. Since we will use our model to understand the labor market dynamics between 2020 and 2023, we interpret the stationary state of the model to be the U.S. economy at the end of 2019. For that reason, most calibration targets refer to moments of the U.S. economy in 2019.

4.1 Functional forms

We start by describing functional forms of some model elements. We assume that disutility functions for working at and searching for a job depend linearly on the individual's age:

$$\phi^E(j) = \phi_0^E + \phi_1^E \times j$$
$$\phi^U(j) = \phi_0^U + \phi_1^U \times j.$$

We assume that the distribution of wage offers $G(w)$ is log-normal with parameters (μ_w, σ_w). We also assume that the distribution of wealth for the newborn $Q(a)$ is log-normal with parameters (μ_a, σ_a). As previously explained, we assume that the separation rate varies with the wage

of the worker according to $\delta(w) = \bar{\delta} \times \exp[\eta_w^{\delta} \times (w - \mu_w)]$.[10] Finally, the Social Security benefits function $y^{SS}(j)$ is also parametrized and described in the next subsection.

4.2 Externally calibrated parameters

We set the coefficient of relative risk aversion σ to 2, a standard value in this class of models. We set μ_w, the mean of the wage distribution G, to be 0.5. The choice of the value for μ_w is a model normalization: through its data counterpart (the average wage in the U.S. economy in 2019), we are able to map data objects denominated in dollars to the model (excess returns, fiscal policies, etc.). Because government transfers (economic impact payments) other than the UI benefits T we have in mind are paid only during the COVID-19 episode, we set $T = 0$ at the stationary state. The other three key external inputs are the asset return function $r(a)$, the age-dependent survival probabilities $\pi(j)$, and the Social Security income function $y^{SS}(j)$.

Asset returns. We set the return on savings $\bar{r}$ to the median monthly return on net worth from the SCF in 2019. To obtain this value from the data, we follow the imputation process that combines data from the 2019 SCF with data on aggregate returns for different asset classes, which is described in detail in Faria-e-Castro and Jordan-Wood (2024). This imputation process relies on several assumptions: (i) the composition of asset portfolios that is observed in the 2019 SCF remains constant, including debt owed; and (ii) households are perfectly diversified within each asset class. Additionally, we compute returns only for changes in net worth that arise from asset classes for which we observe data on realized returns.

Formally, we define the net worth for household i at the beginning of 2019 as

$$NW_{i,2019} = \sum_{j \in J} A_i^j - B_i, \tag{1}$$

where A_i^j is the dollar value of assets of type j, and B_i is debt owed by the household in dollars. The asset classes j that we consider are stocks, real estate, corporate bonds, government bonds, private businesses, and other assets.[11] Given data on realized returns for each of these asset classes over some period τ, we can estimate net worth over this period as

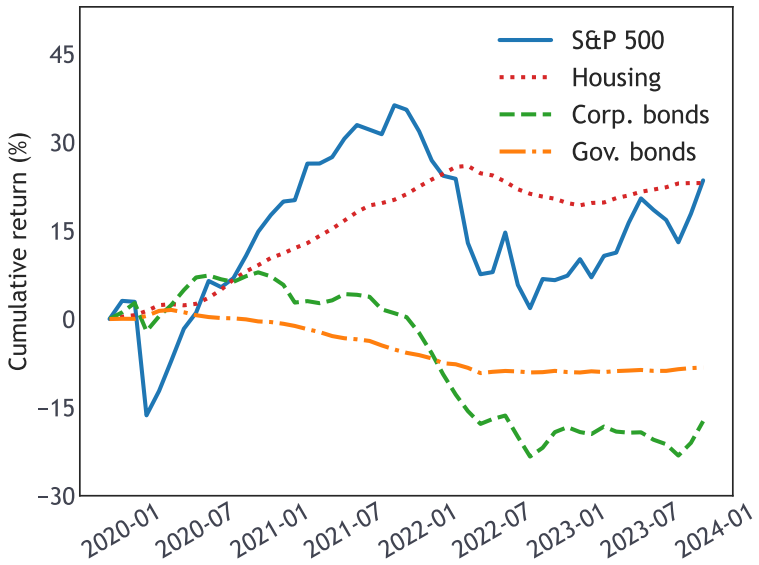
$$NW_{i,\tau} = \sum_{j \in J} R_\tau^j A_i^j - B_i,$$

where R_τ^j is the cumulative return on asset class j over the period τ. Given our assumption of perfect diversification within each asset class, we use publicly available data on composite real returns for each asset class as a proxy for R_τ^j. For stocks, we use the return on the S&P

[10] The functional form of $\delta(w)$ is inspired from Shimer (2005) who uses the same functional form when defining how the aggregate job separation rate changes with the aggregate productivity over time.

[11] Appendix B.3 provide details on how we construct these variables.

Figure 4: Cumulative real returns on selected asset classes



Note: This figure provides cumulative real returns on selected asset classes relative to 2019. Please refer to the main text for more details.

500 index; for housing, we use the Case-Shiller U.S. National Home Price Index; for corporate bonds, we use the ICE Bank of America U.S. Corporate Index; for government bonds we use the 10-year Treasury yield; and finally, we assume that the return on private businesses is equal to 50% of the return on the S&P 500. We assume zero real returns for all other asset classes. The cumulative returns for these asset classes are shown in Figure 4.[12]

This process allows us to compute the net return on net worth over the same period as

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

For calibration purposes, we look at the return on net worth between the beginning and the end of 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 given that the borrowing rate is distinct from the saving rate in our model, as well as the very wealthy given that our model is not expected to capture very high levels of wealth, and it encompasses over 82% of the households in the 2019 SCF. This procedure leads to a median real return of 1.45%, which implies $\bar{r} = 0.12\%$ at monthly frequency. For the borrowing rate, we assume that there is a 2% annualized spread over the saving rate (i.e., $r^b = \bar{r} + 0.17\%$), which is standard in this class of models.

Survival probabilities. To calibrate $\pi(j)$, we use the 2019 Actuarial Life Table from the SSA.[13] This table reports conditional death probabilities for males and females in each age group.

[12] Appendix B.4 contains a description of how these return series are constructed.
[13] Available at https://www.ssa.gov/oact/STATS/table4c6.html.

We compute an equally weighted average for men and women for each age group, and convert these annual conditional death probabilities into monthly conditional survival probabilities.

Social Security income. We approximate the current SSA formula for retirement benefits using a truncated linear function.[14] Upon retirement, Social Security benefits are scaled by a factor that depends on the distance of the age of the person at retirement and the "normal retirement age" (NRA). The NRA, as well as the scaling of the benefits, is a function of a person's birth year. To simplify the analysis, we calibrate the benefits function to that of someone born between the years of 1943 and 1954, which is likely to be the bulk of normal-age retirees for the period we are focusing on. For someone born on these dates, the NRA 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 the NRA. 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 NRA. We approximate these formulas with a linear function that captures the level and slope of the benefits. This exercise yields

$$y^{SS}(j) = \begin{cases} 0 & \text{if } j < 62 \\ \bar{y}^{SS} \times (-3.54 + 0.07 \times j) & \text{if } j \in [62, 70] \\ \bar{y}^{SS} \times (-3.54 + 0.07 \times 70) & \text{if } j > 70, \end{cases}$$

where $\bar{y}^{SS}$ is a constant that is internally calibrated.

Note that, unlike the data, Social Security benefits in the model depends on the current age j and not on the age of retirement. We do this simplification in the model for two reasons. First, it avoids the need to keep track of an additional state variable for the individual (age of retirement). Second, it avoids having to define a more complicated formula to account for instances where individuals move between retirement and employment after the age of 62.

4.3 Internally calibrated parameters

We internally calibrate the remaining parameters in the model. The full set of parameters, their values, and targeted data moments are summarized in Table 1. Most target moments are computed from either the 2019 SCF or the 2019 CPS, where we focus on a reference population that is 16 years of age and over and exclude those in the armed forces. All CPS target moments are based on averages across all months of 2019.

We start by describing the parameters that are specific to the consumption-savings problem.

[14] See Appendix 2 in Faria-e-Castro and Jordan-Wood (2024) for a more detailed description.

Table 1: Internally calibrated parameters

Parameter	Value	Moment	Source	Data	Model
		A. Individuals			
β	0.995	Fraction of population w/ NW ≤ 0	SCF	0.100	0.182
$\underline{a}$	-1.024	median $\frac{\text{credit limit}}{\text{quarterly labor income}}$ for $\ell = E$	SCF	0.740	0.632
μ_a	0.246	1st quartile of NW/Median NW	SCF	0.145	0.185
σ_a	0.329	3rd quartile of NW/Median NW	SCF	3.304	2.019
$\bar{y}^{SS}$	0.013	Retired share	CPS	0.183	0.150
h	0.218	$\frac{\text{median NW for age<62}}{\text{median NW for non-participant age<62}}$	SCF	0.437	0.597
b	0.240	$0.4 \times \mathbb{E}(w\|\ell = E)$	–	0.400	0.455
		B. Labor Market			
ϕ_0^E	0.033	Unemployment rate, all ages	CPS	0.037	0.022
ϕ_1^E	0.003	Unemployment rate, over 55	CPS	0.027	0.000
ϕ_0^U	0.432	LFPR, all ages	CPS	0.634	0.591
ϕ_1^U	0.109	LFPR, over 55	CPS	0.389	0.373
σ_w	0.084	p90/p50 ratio of log annual labor earnings	GRID	1.090	1.029
γ_N	0.096	Ratio of NE flows to total findings	CPS	0.171	0.146
γ_R	0.037	Ratio of RE flows to total findings	CPS	0.035	0.041
f	0.576	Job-finding rate	CPS	0.461	0.354
$\bar{\delta}$	0.032	Total separation rate	CPS	0.041	0.033
η_w^δ	-9.477	Q1/Q5 ratio of E to U or N or R rate	CPS	3.290	3.282

Note: SCF refers to the 2019 Survey of Consumer Finances. CPS refers to averages over the 12 months of 2019 for the Current Population Survey. GRID is the Global Repository of Income Dynamics (Guvenen et al., 2022). NW refers to net worth computed according to Equation (1) using the SCF. Please refer to main text for more detailed explanations of moments and parameters.

These parameters are summarized in Panel A of Table 1. The discount factor β is chosen to match the fraction of individuals with non-positive net worth in the SCF, which is equal to 9.66%. The borrowing limit is chosen to target the median value of the credit-limit-to-quarterly-income ratio among the employed, as reported in Kaplan and Violante (2014) using the SCF. The parameters governing the distribution of wealth among newborns, (μ_a, σ_a), are chosen to match the first and third quartiles of the distribution of net worth, relative to median net worth in the SCF. The level of Social Security benefits $\bar{y}^{SS}$ is chosen to target the average fraction of retirees in the population from the CPS. The monetary value of not working h is chosen to match the ratio of median net worth for individuals under the age of 62 to the median net worth of individuals in the same category who do not participate in the labor force. Finally, the level of UI benefits b is chosen to generate a 40% replacement rate out of the average wage.

Panel B reports the values of parameters and targeted data moments associated with the labor market block of the model. 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. The level and slope of the unemployment disutility function are chosen in a similar way, but to match the labor force participation rate of the general population and those aged 55 and over. The dispersion of the wage distribution σ_w is chosen to match the ratio of the 90th

to the 50th percentile of the log annual labor earnings distribution in the U.S. from GRID data (Guvenen et al., 2022). The likelihood of returning to the labor market for non-participants γ_N is chosen to match the ratio of flows from non-participation to employment relative to total flows to employment. The likelihood of returning to the labor market for retirees γ_R is chosen in an analogous manner. The finding rate 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 employment separation rate $\bar{\delta}$ is chosen to match flows out of employment in an analogous manner. Finally, the slope parameter of the employment separation rate η_w^δ is chosen to target the ratio of the employment separation rate for those in the first quintile of the wage distribution and for those in the fifth quintile.

4.4 Validation of model predictions at the stationary state

The last two columns of Table 1 show that the model does a reasonable job of matching targeted data moments. We now show that the model also matches a series of untargeted data moments in 2019 that are relevant for the economic forces that we seek to analyze. In particular, we show that the model does a good job in capturing the shape of the wealth distribution for the entire population, as well as wealth and labor income distributions among new retirees. This last point is especially important as it means that the model can adequately capture the financial characteristics of the agents who decide to retire.

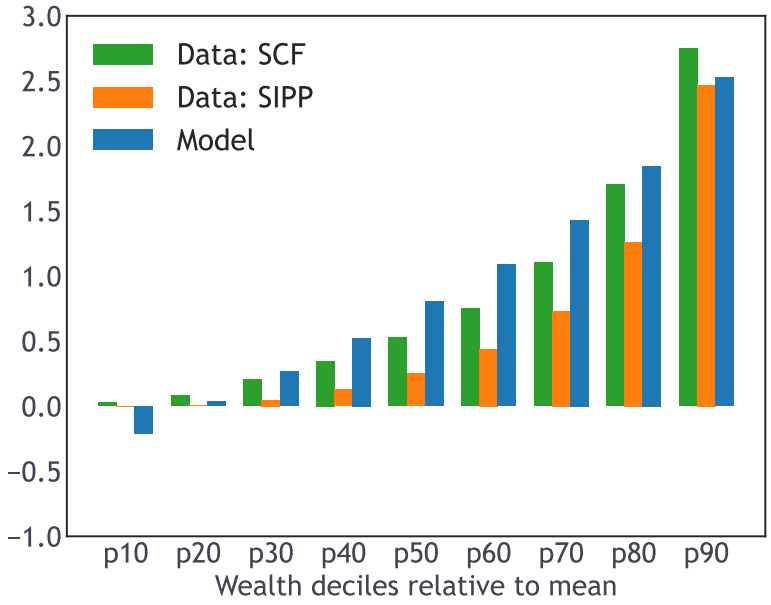
4.4.1 Unconditional wealth distribution

Figure 5 plots deciles of the economy-wide wealth distribution in the model's stationary state vs. deciles of the net worth distribution in the 2019 data from two sources, the SCF and the SIPP. The data series were previously plotted in Panel (a) of Figure 3. To ensure comparability between model numeraire units and dollars in the data, we report wealth deciles relative to mean wealth. We find that the model does a good job of matching the shape of the wealth distribution, especially relative to the SCF. The model does have slightly less wealth inequality than the data, which is expected as the model is not able to capture very high levels of wealth at the top.

4.4.2 Wealth and income distributions among new retirees

Since our analysis is focused on the drivers of retirement patterns between 2020 and 2023, it is particularly important that the model's stationary state generates the right patterns of retirement in 2019 in the data. If our model predicts that people who retire at the stationary state have similar characteristics vis-à-vis the data, this gives us confidence that the model predictions regarding new retirees in the main quantitative experiment covering the 2020-2023 episode will be sensible. Figure 6 plots fractions of new retirees across quintiles of wealth and income distributions in the model's stationary state vs. the 2019 SIPP data. Recall that Figure 2 already presented fractions of new retirees across wealth quintiles in the data, where we discussed

Figure 5: Wealth distribution: Model stationary state vs. SIPP and SCF 2019 data



Note: This figure presents deciles of the economy-wide wealth distribution in the model's stationary state vs. deciles of the net worth distribution in the 2019 data from two sources, the SCF and the SIPP. To ensure comparability between model numeraire units and dollars in the data, we report wealth deciles relative to mean wealth.

Figure 6: Fractions of new retirees across wealth and income quintiles: Model vs. data



Note: Panel (a) plots fractions of new retirees across wealth quintiles in the model's stationary state and in SIPP 2019. Panel (b) repeats the same results for the income quintiles. In the model, we identify people who retire in a given month, compute the wealth (or labor income) distribution in that month, and then assign these new retirees to one of the quintiles of that distribution. In the data, we implement the same calculations by looking at individuals who retire in 2019. Our measure of wealth (income) in the data is the total household-level of net worth (labor income).

how we obtained these results. We repeat the same calculations to obtain fractions of new retirees by labor income quintiles. In the model, we implement the same calculations as in the data: we identify people who retire in a given month, compute the wealth (or labor income) distribution in that month, and then assign these new retirees to one of the quintiles of that distribution.

We find that the model matches the qualitative patterns in the data: new retirees tend to be relatively wealthier and tend to have relatively lower labor income. The model matches almost exactly the positive dependence on wealth, while the relation of retirement with respect to income is slightly weaker in the model than in the data; but the qualitative pattern is mostly captured. These results indicate that the model is able to capture both the wealth effects of labor supply, with those who retire being more likely to be wealthy, and the substitution/opportunity cost effects, with those who retire being more likely to have lower labor income.

5 Aggregate Labor Market Dynamics during 2020-2023

Using the calibrated model, we now ask whether the model can generate the observed empirical changes in the aggregate labor market conditions between 2020 and 2023. To do so, in Section 5.1, we first describe how we measure and input the following five exogenous forces to the model: changes in (i) asset returns, (ii) labor market conditions, (iii) economic impact payments, (iv) UI benefits, and (v) mortality rates. Then, in Section 5.2, we present the results of our main experiment, where we feed in all these forces into our model and analyze whether the model is able to generate the empirical changes in the aggregate (i) retired share (i.e., excess retirement gap), (ii) unemployment rate, and (iii) employment-to-population ratio between 2020 and 2023. As none of these empirical patterns are targeted in our calibration exercise, the comparison between our model's predictions of these aggregate labor market outcomes and empirical changes in the data serve as yet another element of validation of the model. Overall, we show that the model is broadly able to capture these empirical patterns between 2020 and 2023.

5.1 Shocks

Starting from the stationary state, we feed five sequences of shocks to the model: a shock to the return on savings, which is heterogeneous across wealth and age; a shock to separation rates for employed, which is heterogeneous across wage levels; a shock to lump-sum transfers, which depends on age and income of the individual; a shock to UI benefits of unemployed; and a shock to age-dependent mortality rates. All these shocks are presented in Figure 7. We now describe in detail how we map each of these impulses from the data to the model.

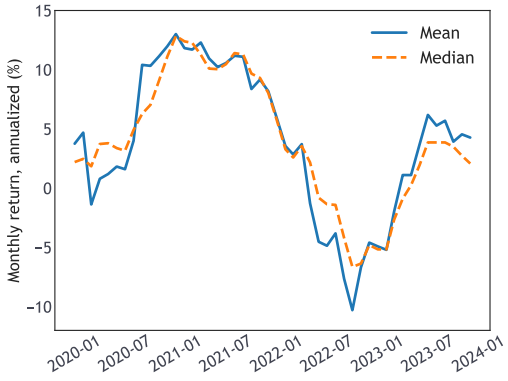
5.1.1 Asset returns

The first shock we feed to the model is the path of returns on net worth during 2020-2023. These elevated returns may have triggered wealth effects that led an above-average number of people to retire, as well as contributed to keeping people in retirement.

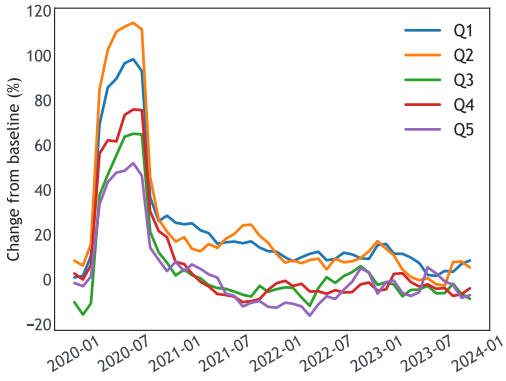
One important feature of the data we do not explicitly model is that agents at different levels of wealth and ages have different portfolios that may have earned different amounts of returns during this period. To capture this heterogeneity with our parsimonious environment,

Figure 7: Time series paths for exogenous shocks

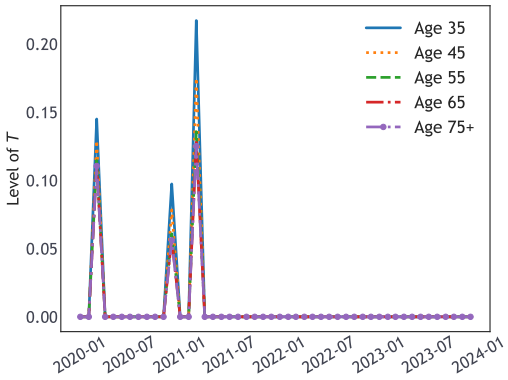
(a) Asset returns shock



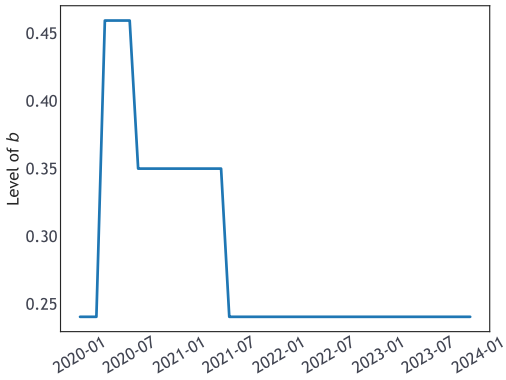
(b) Job separation rate shock



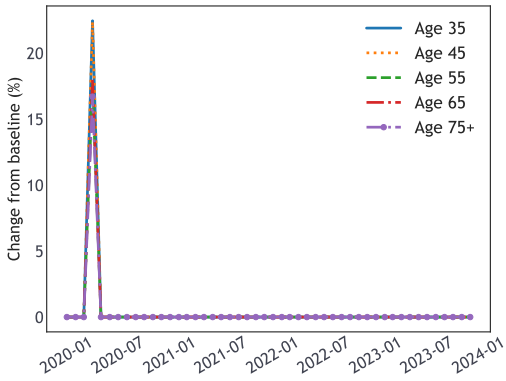
(c) Economic impact payments shock



(d) UI benefits shock



(e) Mortality rate shock



Note: Panel (a) plots the mean and median paths of the estimated monthly return (annualized) function $r_t(a, j)$. While this return function depends on the level of wealth and age, we only plot the mean and median values at each month for expositional purposes. Panel (b) plots shocks to the job separation rate at the stationary state $\delta(w)$. That is, for example, the job separation rate of those at the bottom two quintiles increase in mid 2020 by around 100% relative to their respective stationary state levels, while the job separation rate of those at the top quintile increase at that time by only around 50%. Panel (c) presents shocks to the economic impact payments T, where the payment amount depends on the age and the eligibility for these payments depend on the income level of the individual. These payments cover three rounds of economic impact payments made during the COVID-19 episode. Panel (d) plots the shocks to UI benefit amount during this episode, starting from the stationary-state level of $b = 0.24$ and returning back to this level after two rounds of additional payments. Finally, Panel (e) plots changes in mortality rates depends on the age of the individual. Due to significantly monthly variations, shocks in Panels (a) and (b) are smoothed by taking six-month moving averages.

we estimate the following regression specification:

$$r_{i,t}^{NW} = \beta_{0,t} + \beta_{1,t}\text{age}_i + \beta_{2,t}\text{age}_i^2 + \beta_{3,t}\text{age}_i^3 + \beta_{4,t}\left(NW_{i,2019}/\bar{w}\right) + \epsilon_{i,t}, \tag{2}$$

where $r_{i,t}^{NW}$ is the monthly return on net worth for household i in month t, computed using the same imputation procedure as described in Section 4.2, age_i is the age of the head of household, and $NW_{i,2019}/\bar{w}$ is the ratio of net worth in 2019 to average annual labor earnings in 2019 ($65,477 according to the CPS). This is a useful normalization for net worth, as it then allows us to parametrize returns in the model as a function of model-equivalent objects.[15] We estimate this regression for each month from January 2020 to December 2023.

We then use the estimated coefficients to specify a period return function $r_t(a, j)$ that yields heterogeneous returns for agents with different levels of wealth a and age j at each time t—a parsimonious way of accounting for the impact of portfolio heterogeneity on returns without having to explicitly model multiple assets. Due to significant month-to-month variation in returns, we use six-month moving averages of the estimated coefficients. We treat this return function as an exogenous shock that we feed into the model for our experiment. Panel (a) of Figure 7 plots the mean and median paths of the estimated monthly return (annualized) function: both mean and median increase in the early months of the pandemic, surpassing 10% in 2021. They then fall and become negative in 2022 and early 2023, but recover to positive levels later in 2023.

In terms of implementation, we replace the return function $\bar{r}$ in the budget constraint for each agent with positive wealth with $r_t(a, j)$. Importantly, these return shocks are unexpected and assumed to be transitory. Individuals expect the return on savings to be $\bar{r}$ in the following period at each point in time. This is therefore equivalent to a lump-sum windfall that does not distort the savings decision of each individual.[16] This is important since, if agents understand that returns increase considerably, it could trigger significant changes in consumption/savings behavior that could induce counterfactual expansions of labor supply through substitution effects. As agents perceive greater returns on their savings, they could actually work more in order to accumulate wealth faster and take advantage of these elevated returns. This is not only counterfactual but is conceptually at odds with the nature of this increase in returns in the data, which plausibly materialized in the form of unanticipated capital gains on housing and mutual funds. We discuss the counterfactual implications of such a setup in greater detail in Section 7.

5.1.2 Labor market conditions

The 2020-23 period was marked by a large increase in the aggregate job separation rate. In addition, the COVID-19 episode created a much larger increase in the job separation rate of

[15] We experimented with several specifications, and a cubic polynomial on age and linear in net worth (relative to mean income) seemed to provide the best combination of simplicity and explanatory power.

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

low-income workers, while those with relatively higher levels of income experienced a much smaller increase in the job separation rate. These increases in separations may have affected labor force participation by discouraging workers from participating given that the flow rate from unemployment to non-participation is much larger than the flow rate from employment to non-participation (Hobijn and Şahin, 2021). We capture both the magnitude and heterogeneity in separations by feeding exogenous paths of separation rates that vary by wage quintile. To this end, we use the CPS to calculate the percent change in the job-separation rate for each month in 2020-23 relative to the average job separation rate in 2019. We do this for each quintile of the wage distribution. Due to sizable fluctuations in monthly transition rates, we compute six-month moving averages of these changes. Panel (b) of Figure 7 plots the series that we feed to the model as period-by-period shocks to the job separation rate at the stationary state $\delta(w)$.[17] These series capture both the significant increase in separation rates and the substantial level of heterogeneity for workers in different wage quintiles, with those in lower quintiles being significantly more affected and experiencing a slower recovery back to 2019 levels.

5.1.3 Economic impact payments

The third exogenous shock we consider is economic impact payments during the COVID-19 episode. 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 $\{T_t\}$ in our model.

For each of the three rounds of transfers, households were eligible if their adjusted gross income (AGI) was lower than $75,000, and ineligible if their AGI exceeded $80,000, with payments phasing out linearly in between. Since these cutoffs are close to the 40th percentile of income in the data, we set this as the eligibility cutoff in the model.

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

[17] That is, for example, the job separation rate of those at the bottom two quintiles increase in mid 2020 by around 100% relative to their respective stationary state levels, while the job separation rate of those at the top quintile increase at that time by only around 50%.

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 contains information on the number of people under the age of 18 per age of the head. This allows us to compute a transfer modifier that depends on the age of the head: this modifier tends to be larger for younger household heads, who tend to live with more people under the age of 18, and is close to zero for those age 75 years and over.

Using the fact that the (ex-ante) average wage in the model $\mu_w = 0.5$ is equivalent to average monthly labor earnings of $\$65,477/12$ in the data, we convert the modified effective transfers to model units (appropriately deflated using the CPI). The procedure is explained in detail in Appendix A.2. The effective transfer series, as a function of age, is plotted in Panel (c).

5.1.4 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.[18] 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).

5.1.5 Mortality rates

The last shock we consider is a change in age-dependent mortality rates $\pi(j)$. The goal is not to match actual mortality patterns, but rather to shock agents' perceived mortality risk during 2020. This is potentially an important channel given that perceived increases in mortality operate as changes in the discount factor that may affect participation decisions especially for older agents.

We model the increase in mortality rates as a one-time shock in March 2020, when agents experience mortality rates switching from their baseline values to the empirical 2020 conditional mortality rates from the SSA life tables. The change in mortality rates depends on the age of the individual and is plotted in Panel (e).

5.2 Changes in aggregate labor market moments: Model vs data

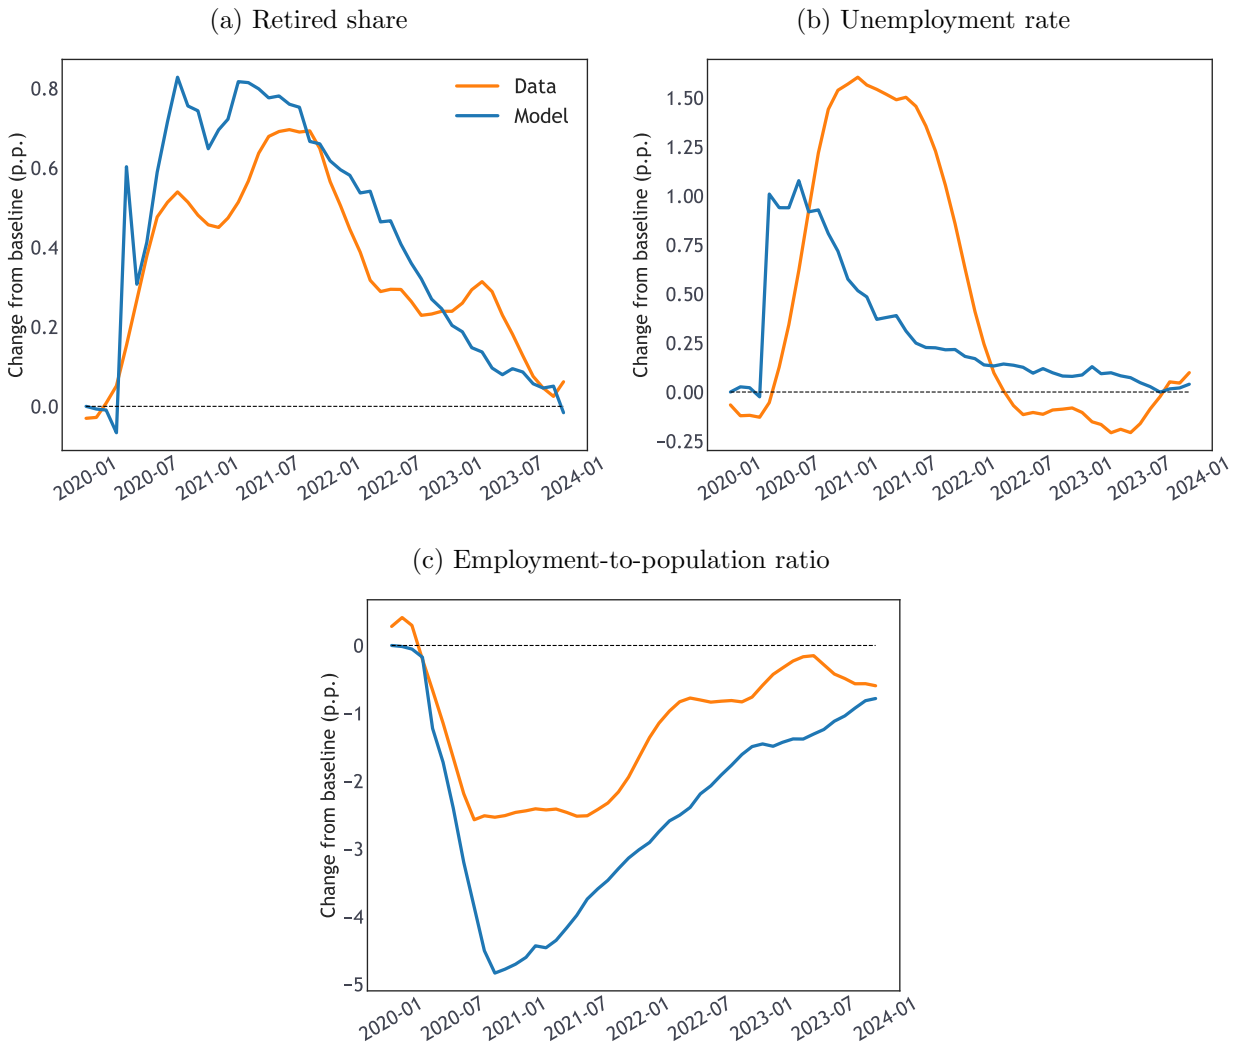
Next, we present the results of our experiment, where we feed in all shocks at the same time to the stationary state of the model and compare the resulting aggregate labor market dynamics along the transition to the observed empirical dynamics between 2020 and 2023.

Figure 8 plots the paths of the aggregate retired share (i.e., excess retirement gap) (Panel (a)), unemployment rate (Panel (b)), and employment-to-population ratio (Panel (c)) in the data and the model.

For the retired share in the data, we use the same definition as in Figure 1: the deviation

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

Figure 8: Changes in aggregate labor market moments: Model vs data



Note: This figure plots the paths of the aggregate retired share (i.e., excess retirement gap) (Panel (a)), unemployment rate (Panel (b)), and employment-to-population ratio (Panel (c)) in the data and the model. For the aggregate retired share in the data, we calculate the deviation of the actual retired share in the CPS relative to a linear trend. We take six-month moving averages both in the data and in the model, and plot the percentage point deviation from the 2019 average for the data and stationary state for the model. We also repeat the same procedure for the unemployment rate and the employment-to-population ratio. 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 for the unemployment rate is the unemployment rate net of temporary unemployment, as classified in the CPS.

of the actual retired share in the CPS relative to a linear trend. We take six-month moving averages both in the data and in the model, and plot the pp deviation from the 2019 average for the data (as in Panel (b) of Figure 1) and stationary state for the model. Note that the model matches both the magnitude and persistence of the increase in the retired share. The model predicts a slightly larger increase, peaking at 0.82 pp, while the data peak at 0.70 pp. Importantly, the model also matches the dynamics after this peak reasonably well. In particular, the model generates the slow decline in excess retirements through 2023, as in the data.

Similarly, for both the unemployment rate and the employment-to-population ratio, we take

six-month moving averages and plot the pp deviations from 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. With this caveat, the model slightly underestimates the increase in the unemployment rate by about 0.5 pp, but reasonably captures the dynamics.

Finally, the model overestimates the decline in the employment-to-population ratio by about 2 pp, but matches its slow recovery path in the data. In particular, both model and data are aligned with their prediction that the employment-to-population ratio is about 1 pp lower than the 2019 level at the end of 2023.

Taken together, these results suggest that the model does a satisfactory job in capturing untargeted aggregate dynamics between 2020 and 2023, even if it slightly overestimates movements out of the labor force and underestimates movements toward unemployment.

6 Decomposing the Retirement Boom

Having shown that the model does a good job in capturing the size and persistence of movements in terms of aggregate labor market variables, namely the rise in the retired share during 2020-2023, we now undertake a decomposition exercise where we quantify the relative importance of each of the five shocks in driving the increase in the retired share during this episode.

6.1 Decomposing the increase in retired share

Figure 9 offers 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. On the other hand, Panel (b) adds only one shock at a time, starting from the stationary state (without any shock).

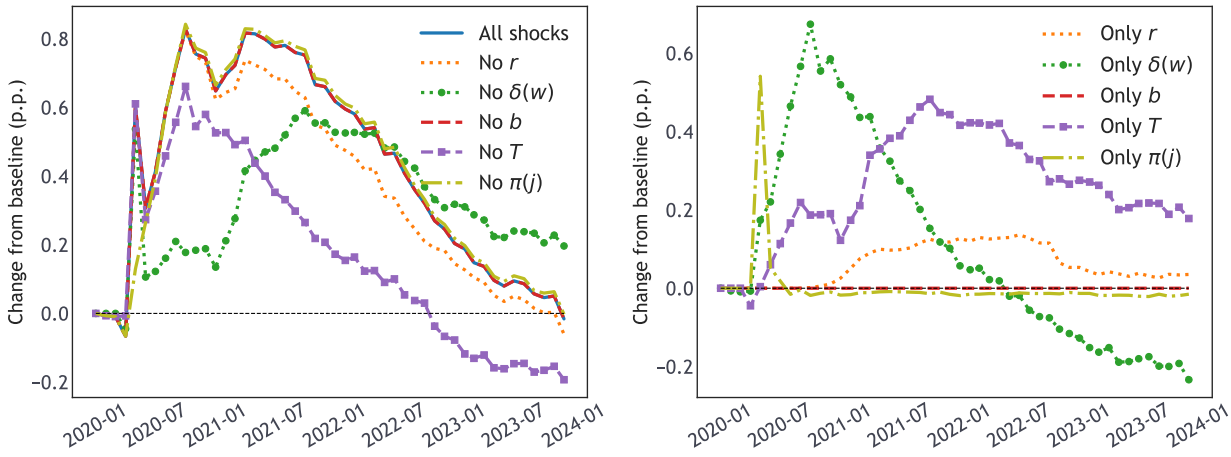
We show that labor market (job separation) shocks, as shown by green dotted lines in both panels, are most responsible for the increase in the retired share during 2020, while mortality shocks are only essential to explain the initial increase in the retired share during the first months of 2020, as shown by gold dashed-dotted lines. However, mortality and labor market shocks alone cannot explain the persistence of the rise in the retired share. As labor market conditions improve throughout 2021 and 2022, the retired share would have quickly fallen and even gone negative after the summer of 2022 (i.e., became lower than its level at the stationary state), as shown by green and gold lines in Panel (b).[19] Panel (b) also shows that the persistence of the rise in the retired share is explained primarily by economic impact payments (purple dashed line) and,

[19]To clarify, given that the change in the retired share in the economy with only mortality shocks (gold dashed-dotted line) in Panel (b) is close to zero after the summer of 2020, when we only have mortality and labor market shocks, retired share dynamics would be close to the green line in Panel (b). As a result, when we only have mortality and labor market shocks, the model cannot account for the persistence of the rise in the retired share.

Figure 9: Decomposing the increase in the retired share

(a) Removing one shock at a time

(b) Adding one shock at a time



Note: This figure provides 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. On the other hand, Panel (b) adds only one shock at a time, starting from the stationary state (without any shock).

to a lesser extent, by asset returns (orange dotted line). Importantly, these shocks, in isolation, would have predicted a much smaller increase during 2020-2021 and a more persistent increase in the retired share during 2023 than those observed in the data. Hence, we conclude that we need all of these four shocks to get the dynamics in the retired share right. On the other hand, changes in UI benefits play essentially no role in explaining excess retirements. This is because the flows between retirement and unemployment are negligible in our model.

Table 2 offers a formal decomposition to quantify the contribution of all these five shocks on the rise in the retired share separately 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 discussions in the previous paragraph: 65% of the increase in the retired share in 2020 is accounted for by changes in labor market conditions. This share drops to 41% in 2021 and becomes negative thereafter. Economic impact payments explain 48% in 2021 and 83% in 2022. Changes in asset returns explain 26% in 2022 and 19% in 2023. Note that the contribution of economic impact payments exceeds 100% in 2023, which again confirms that this force in isolation is unable to correctly account for the dynamics of aggregate excess retirements, and the offsetting effects of improving labor market conditions are important during this period.

In sum, our quantitative exercise indicates that deteriorating labor market conditions were a key driver of excess retirements during the early pandemic period, 2020-21, while fiscal policy measures and elevated asset returns contributed to keeping excess retirements high in the latter period, 2022-23. The importance of labor market conditions and fiscal transfers is suggestive that the increase in retirements may have been driven primarily by those at the bottom of the income

Table 2: Formal decomposition of the increase in the retired share

	Asset returns	Labor market	UI benefits	Economic impact payments	Mortality
2020	0.3%	71.9%	0.0%	22.4%	12.8%
2021	12.5%	41.3%	0.0%	47.5%	-1.7%
2022	26.0%	-9.8%	0.0%	82.6%	-3.3%
2023	18.7%	-73.1%	0.0%	137.9%	-11.5%

Note: This table provides changes in the retired share from the stationary state which are obtained by feeding each shock one at a time. Values present the percentage of the change obtained by feeding these shocks, averaged over each year. Due to interactions and averaging, values may not sum up to 100%.

and wealth distribution, who faced relatively worse labor market prospects and were eligible and more sensitive to income effects arising from fiscal transfers. We study the composition of new retirees in more detail in Section 6.3.

6.2 Decomposing the increase in unemployment rate

Figure 10 repeats the same exercise for the unemployment rate, with Panel (a) removing one shock at a time and Panel (b) adding one shock at a time. There are three key takeaways. First, the increase in UI benefits and labor market shocks both contribute to explaining the increase in the unemployment rate in 2020. Second, while the UI expansion plays an equally important role as labor market shocks in generating this increase in 2020, labor market shocks are more important in accounting for the persistence of the unemployment dynamics, as in the data. Third, mortality shocks, economic impact payments, and changes in asset returns play a relatively minor role in driving unemployment dynamics. The formal decomposition that quantifies these results is presented in Table 3.[20]

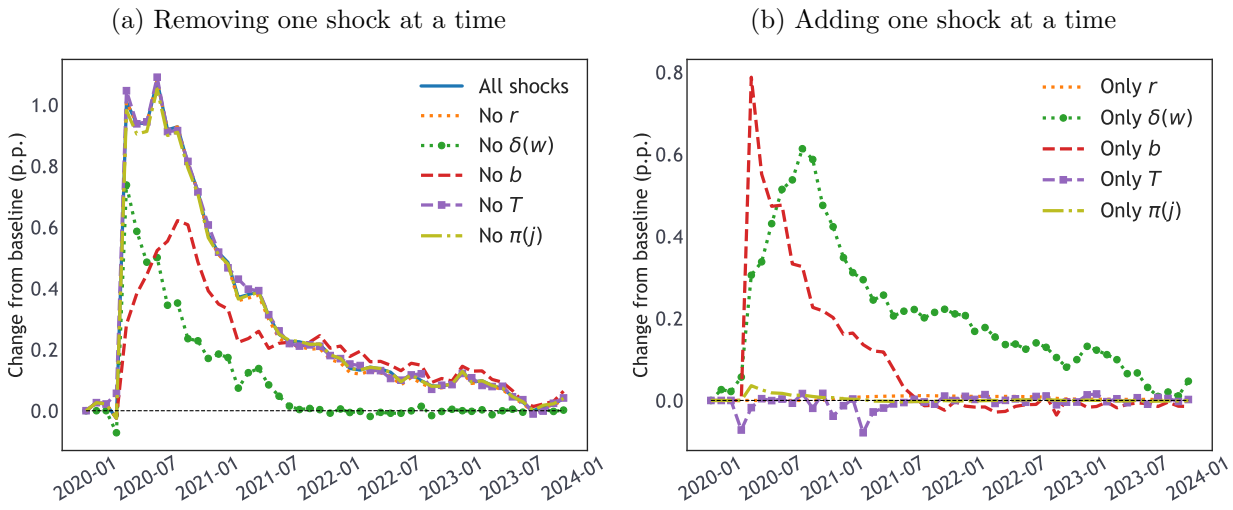
Table 3: Formal decomposition of the increase in the unemployment rate

	Asset returns	Labor market	UI benefits	Economic impact payments	Mortality
2020	0.02%	52.92%	45.14%	0.02%	1.85%
2021	3.23%	79.78%	18.11%	-3.86%	-0.27%
2022	6.98%	118.23%	-10.93%	1.10%	0.23%
2023	-23.01%	-19.95%	163.09%	-86.39%	19.68%

Note: This table provides changes in the unemployment rate from the stationary state which are obtained by feeding each shock one at a time. Values present the percentage of the change obtained by feeding these shocks, averaged over each year. Due to interactions and averaging, values may not sum up to 100%.

[20]Because the unemployment rate fully recovers in 2023 to its stationary state level in the model, as in the data, values presented in Table 3 for 2023 are not very informative as they represent the contribution of each shock to a very small change in the aggregate. However, we still present these numbers with this caveat for completeness.

Figure 10: Decomposing the increase in the unemployment rate



Note: This figure provides two alternative decompositions that shed light in the importance of each exogenous force at each point in time on the increase in the unemployment rate. Panel (a) plots the baseline (with all shocks included) and removes one shock at a time. On the other hand, Panel (b) adds only one shock at a time, starting from the stationary state (without any shock).

6.3 Validation of model predictions 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 micro data that are relevant for the mechanisms of interest? Comparing the outcomes from the model along the transition against the outcomes from the micro data also renders credibility to our quantitative decomposition on the sources of changes in aggregate variables. For this reason, in this section, we show that the model delivers three predictions that are broadly in line with the predictions of micro data presented in Section 2. In particular, the model broadly matches changes in (i) the average net worth, (ii) the distribution of net worth, and (iii) the distribution of new retirees across the net worth distribution in the data between 2020 and 2023 relative to 2019.

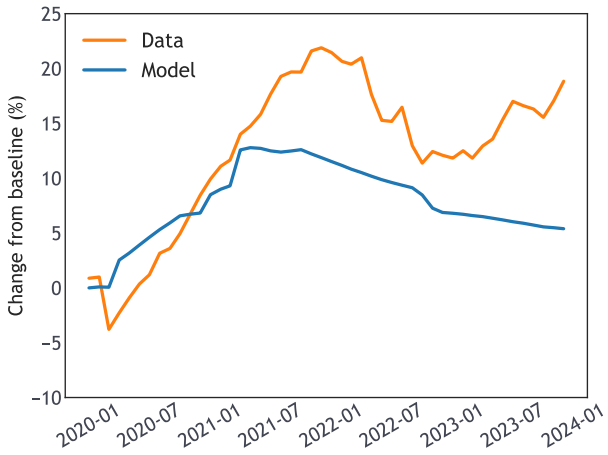
6.3.1 Changes in the average net worth

Panel (a) in Figure 11 plots the average net wealth in the model versus the average net worth in the SCF for the period in analysis, computed using the imputation procedure described in Section 4.2. 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: an increase throughout 2020 that stabilizes in 2021 and then turns to a gradual decline as monthly returns become negative in 2022. Compared with the SCF, the model underestimates the rise in the average net worth: the peak in the model is 13% while the peak in the data is 22%.

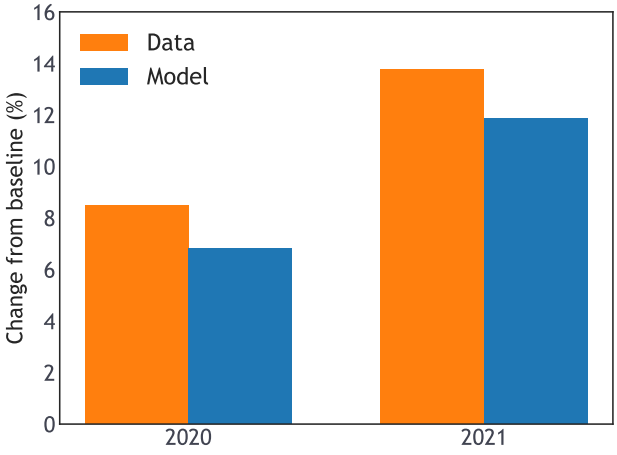
The advantage of the SCF imputation is that it generates a full time series of the net worth distribution. Recall, however, that the SCF imputation procedure requires two assumptions that

Figure 11: Model vs micro data: Net worth dynamics

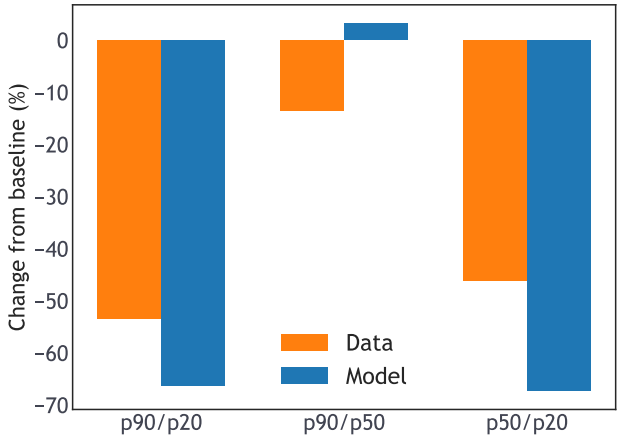
(a) Average net worth: Model vs SCF imputation



(b) Average net worth: Model vs SIPP



(c) Distribution of net worth: Model vs SIPP



Note: Panel (a) plots the average net wealth in the model versus the average net worth in the SCF for the period in analysis, computed using the imputation procedure described in Section 4.2. 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. Panel (b) plots the same by comparing the model with the SIPP data in the two years for which the SIPP is available as of this writing, 2020 and 2021. Panel (c) plots percent changes in different percentile ratios of the wealth distribution between 2019 and 2021 in the SIPP and the model.

may be strong over longer periods of analysis: households do not adjust their portfolios and are perfectly diversified within asset classes. To address this issue, we next use the SIPP data that also contain information on net worth. The main disadvantage is that the SIPP provides only a yearly snapshot. Panel (b) plots the percent change in the average net worth for the SIPP versus the model in the two years for which the SIPP is available as of this writing, 2020 and 2021. The model does an almost perfect job in capturing the average increase in net worth in both years relative to 2019. Importantly, the SIPP wealth data are not subject to the two caveats of the SCF imputation and are therefore more likely to reflect actual net worth changes over time.

6.3.2 Changes in the distribution of net worth

We have shown that the model does a good job in matching the changes in the average of net worth—but what about changes in the distribution of net worth? Panel (c) uses SIPP data to quantify how different percentile ratios of the wealth distribution changed between 2019 and 2021, as presented and explained in Section 2. This figure shows that the model does a good job in explaining the compression in the wealth distribution that is primarily driven by changes in its left tail. Both in the model and in the data, the p90/p20 ratio and the p50/p20 ratio decline largely. On the other hand, the model predicts only a small increase in the p90/p50 ratio, while this ratio fell slightly in the data.

Overall, the model reproduces the overall dynamics of the wealth distribution between 2019 and 2021, which involved an increase in the average net worth and a reduction of inequality in net worth driven primarily by faster growth of the bottom quantile. The fact that the model matches these empirical patterns is important if we ever expect strong wealth effects on labor supply during this episode. By showing that the model matches the increase in the average net worth and the increase in net worth especially for wealth-poor individuals, we are giving this mechanism a fair chance in explaining aggregate participation dynamics during this period.

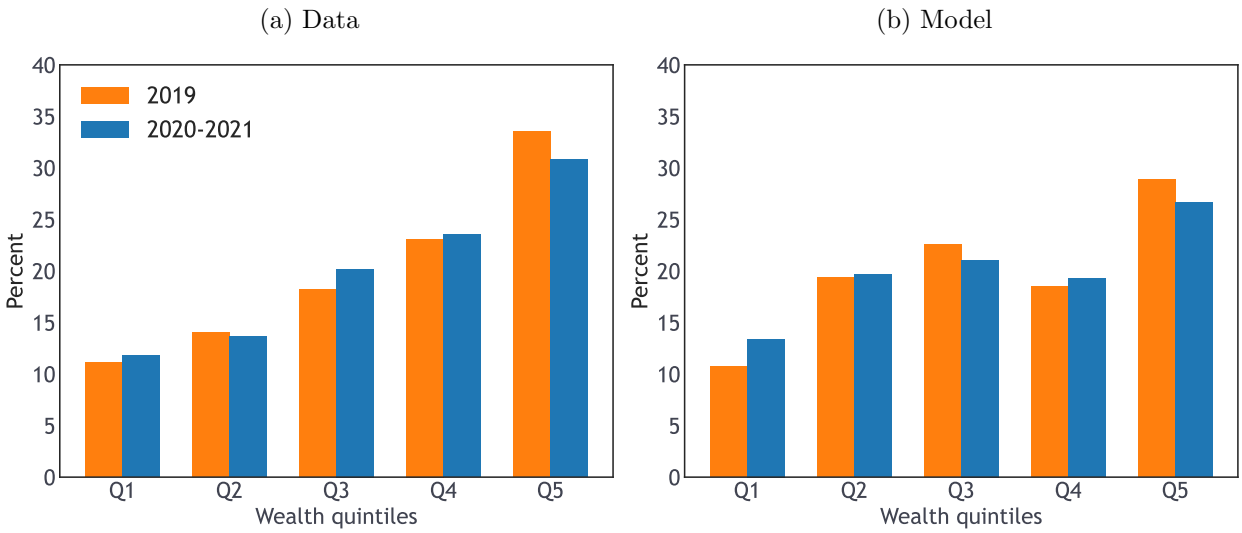
6.3.3 Changes in the distribution of new retirees across the net worth distribution

Figure 12 plots fractions of new retirees across wealth quintiles, separately for those who retire in 2019 and those who retire between 2020 and 2021, using data from the SIPP (Panel (a)) and from the model (Panel (b)). Calculation of these outcomes in the data is already explained in Section 2. We repeat the same calculations in the model.

As discussed in Section 2, Panel (a) reveals that the post-COVID-19 episode is not characterized by an increase in the fraction of new retirees with higher levels of wealth. If anything, retirements during 2020-2021 were slightly tilted toward people with lower 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 qualitative pattern: a slight increase in the share of new retirees in the first quintile and a slight decline in the fifth quintile, indicating that retirements during 2020-2021 were slightly tilted toward wealth-poor individuals. This result also makes sense in light of our findings regarding the decomposition of excess retirements: most of these retirements were driven by labor market conditions in the earlier stages of the crisis and by economic impact payments at a later stage. Agents with lower incomes were more likely to be affected by the negative labor market conditions (as their separation rates were disproportionately higher) and were also more susceptible to income effects from economic impact payments due to the eligibility limits on these transfers and their size relative to their usual income.

In summary, the model not only matches the aggregate increase in the retired share during

Figure 12: Model vs micro data: Distribution of new retirees by net worth



Note: This figure plots fractions of new retirees across wealth quintiles, separately for those who retire in 2019 and those who retire between 2020 and 2021 using data from the SIPP (Panel (a)) and from the model (Panel (b)).

this episode but also generates the distribution of new retirees by wealth groups.

7 Heterogeneous Returns to Wealth and Labor Supply

Our implementation of the asset returns shock described in Section 5.1.1 is based on assumptions that may seem strong. In particular, we assume that the dependence of returns on wealth and age is completely unexpected every period and that agents expect returns on savings to return to their stationary state value $\bar{r}$ in following periods. This allows us to effectively implement returns on wealth that are heterogeneous across the wealth distribution without this heterogeneity distorting any savings decisions directly through substitution effects. However, we note that this baseline experiment preserves distortion of decisions through wealth effects (as it is intended).

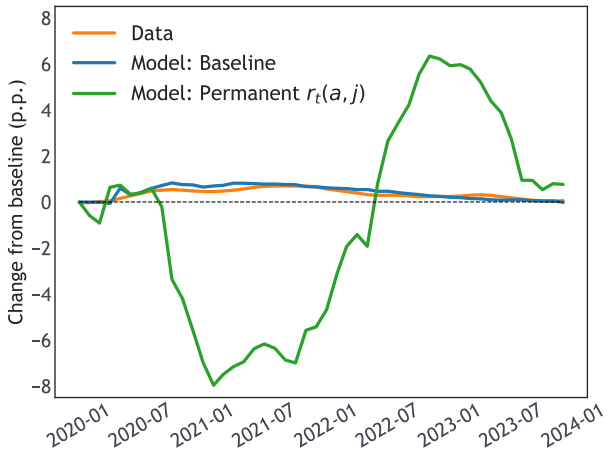
7.1 Counterfactual experiment

We now ask what are the consequences of relaxing this assumption. In particular, what happens if agents understand, both along the transition and at the stationary state, that the return function does not revert back to its stationary state level in future periods $\bar{r}$ but rather remains constant at $r_t(a, j)$? That is, we now assume that in period t agents observe $r_t(a, j)$ and expect future returns on savings to be governed by that function permanently. Note that the agents are still surprised, as the function is indexed by time and changes every period. Relaxing our assumption in the baseline experiment allows us to understand the consequences of introducing substitution effects on top of the already present wealth effects.

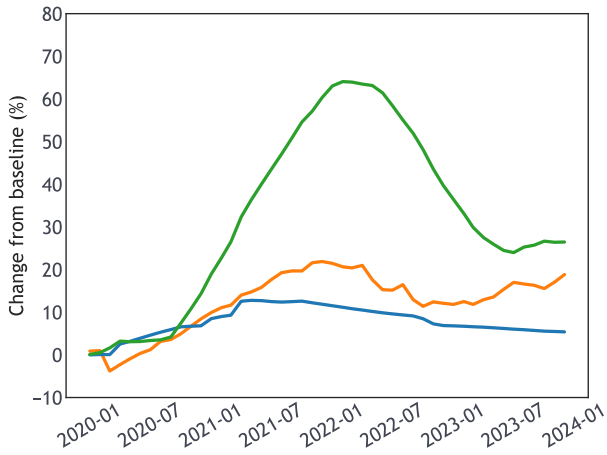
Figure 13 presents the paths of the excess retirement share and the average net worth that we obtain by running our baseline experiment under this different assumption regarding agent

Figure 13: Counterfactual experiment: Agents expect $r_t(a, j)$ to remain permanently

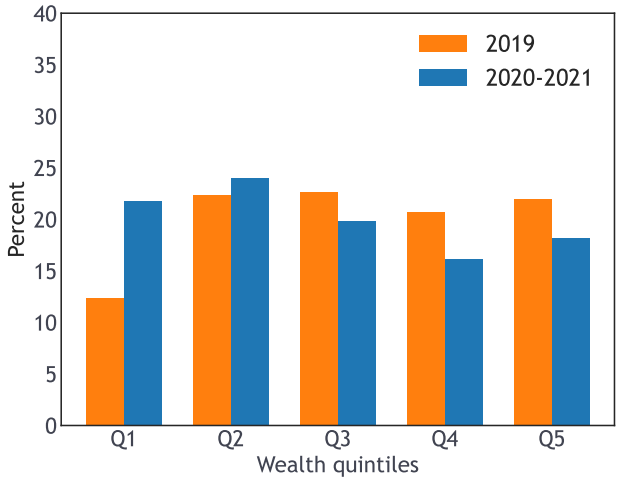
(a) Retired share



(b) Average net worth



(c) Distribution of new retirees by net worth



Note: This figure presents results from a counterfactual experiment where we now assume that, in each period t, agents observe the heterogeneous return function on savings $r_t(a, j)$ and expect future returns on savings to be governed by that function permanently. Panel (a) plots the paths of the aggregate retired share (i.e., excess retirement gap) and Panel (b) plots the average net wealth in the SCF data (orange line), baseline model (blue line), and the baseline model under this counterfactual experiment (green line). Panel (c) plots fractions of new retirees across wealth quintiles, separately for those who retire in 2019 and those who retire between 2020 and 2021 in the baseline model under this counterfactual experiment.

expectations (green line). The implications for the changes in the aggregate retired share are wildly counterfactual when compared with the data (orange line) and the baseline model (blue line): we observe an 8 pp decline at the start of 2021 from the stationary state, which is followed by a 14 pp increase between 2021 and 2022 (or, equivalently, a 6 pp increase from the stationary state). Why is this happening? Since we find that returns typically increase in wealth, consistent with the literature, when agents internalize this, they have an incentive to participate in the labor force to earn more income and accumulate wealth. This effect turns out to be much stronger than the dissuading effects caused by poor labor market conditions and expansion of fiscal transfers. Panel (b) shows that these incentives for wealth accumulation result in a large increase in the

average wealth, which peaks at 64% growth during late 2021 and early 2022 instead of at 22% growth in the data. At this point, as agents are much wealthier, wealth effects start to overcome these incentives to work, and the model generates a large increase in the retired share in late 2022 and throughout 2023, which also coincides with a sizable deaccumulation of average wealth.

Panel (c) plots the distribution of new retirees by net worth in the model separately for 2019 and 2020-2021 under this counterfactual assumption. We now see that shares of new retirees across wealth quintiles in 2020-2021 are mostly equal, which is at odds with the empirical evidence (Panel (a) in Figure 12) and predictions of the baseline model (Panel (b) in Figure 12). In the counterfactual experiment, compared with the 2019 distribution, the share of new retirees in the bottom quintile increases significantly in 2020-2021, while the share in the top quintile declines. When agents expect the heterogeneous return function to stay permanent, because expected returns are much higher among wealthier agents, it becomes costly for them to retire. Poorer agents, however, do not benefit much from the return function and thus prefer to retire instead.

7.2 Implications for recent literature on heterogeneous returns

These results have implications for the burgeoning literature on household financial and inequality that documents the heterogeneity in returns to wealth in a variety of settings (Bach et al., 2020; Fagereng et al., 2020; Ozkan et al., 2023; Smith et al., 2022). The empirical literature overwhelmingly documents that returns increase with the level of wealth. We show that introducing an explicit dependence in line with these empirical patterns can significantly distort households' labor supply decisions. Model-based treatments of heterogeneous returns either do not explicitly model such dependence or do not study the impact of such dependence in labor supply decisions (Benhabib and Bisin, 2018; Benhabib et al., 2019; Xavier, 2021). We show that this dependence, when internalized by agents and combined with endogenous labor supply decisions, may lead to counterfactual predictions in a standard household decision model.

8 Conclusion

In this paper, we developed a quantitative model that allows us to understand and quantify the key potential drivers behind the rise in retirements experienced in the U.S. economy since 2020. Our framework is an incomplete markets, overlapping generations model combined with a frictional labor market. We analyze the ability of five different channels to explain excess retirements during 2020-2023 using this model: elevated asset returns, changes in labor market conditions, 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 our model is able to match the magnitude and persistence of excess retirements when all these forces are active. In our decomposition exercise, we show that labor market disruptions through increased job separation rates were the primary factor driving the increase in the retired

share during 2020, while increased mortality rates were essential only in explaining the initial increase during the first months of 2020. We also find that the persistence of the rise in the retired share is explained by the provision of economic impact payments and, to a lesser extent, elevated returns on assets, in spite of improving labor market conditions post-2021.

The fact that labor market conditions and fiscal transfers conditional on income explain the bulk of excess retirements suggests that these excess retirements were mostly caused by lower-income and -wealth households. We show that this prediction of the model is corroborated by the evidence in the micro data by showing that the model matches the changes in fractions of new retirees across wealth and income distributions. In particular, the model predicts that the share of new retirees at the bottom quintile slightly increases, while the share of new retirees at the top quintile slightly falls, leading to less inequality in fractions of new retirees across wealth quintiles relative to that in 2019, as in the data.

We conclude our analysis by showing that environments with endogenous labor supply and heterogeneous returns to wealth that are internalized by agents would have generated counterfactual predictions for changes in the retired share, average net worth, and fractions of new retirees across the wealth distribution. As such, our findings reveal shortcomings of predictions of a model with heterogeneous returns when labor supply decisions are endogenous. Since the recent literature has documented ample evidence indicating that returns on wealth increase in wealth, our results are potentially relevant for future research in this literature.

References

Bach, L., L. E. Calvet, and P. Sodini (2020): "Rich pickings? Risk, return, and skill in household wealth," *American Economic Review*, 110, 2703–47.

Benhabib, J. and A. Bisin (2018): "Skewed wealth distributions: Theory and empirics," *Journal of Economic Literature*, 56, 1261–91.

Benhabib, J., A. Bisin, and M. Luo (2019): "Wealth distribution and social mobility in the US: A quantitative approach," *American Economic Review*, 109, 1623–1647.

Benson, D. and E. French (2011): "How do sudden large losses in wealth affect labor force participation?" *Chicago Fed Letter*, 282.

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

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

Coronado, J. L. and M. G. Perozek (2003): "Wealth effects and the consumption of leisure: Retirement decisions during the stock market boom of the 1900s," Finance and Economics Discussion Series 2003-20, Board of Governors of the Federal Reserve System (U.S.).

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

——— (2016): "Savings after retirement: A survey," *Annual Review of Economics*, 8, 177–204.

Fagereng, A., L. Guiso, D. Malacrino, and L. Pistaferri (2020): "Heterogeneity and persistence in returns to wealth," *Econometrica*, 88, 115–170.

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

——— (2021b): "The COVID retirement boom," *Economic Synopses*, 1–2.

Faria-e-Castro, M. and S. Jordan-Wood (2024): "Pandemic labor force participation and net worth fluctuations," *Federal Reserve Bank of St. Louis Review*, 106, 40–58.

Guvenen, F., L. Pistaferri, and G. L. Violante (2022): "Global trends in income inequality and income dynamics: New insights from GRID," *Quantitative Economics*, 13, 1321–1360.

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

Hobijn, B. and A. Şahin (2022): "Missing workers and missing jobs since the pandemic," Tech. rep., National Bureau of Economic Research.

Kaplan, G. and G. L. Violante (2014): "A model of the consumption response to fiscal stimulus payments," *Econometrica*, 82, 1199–1239.

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

Montes, J., C. Smith, and J. Dajon (2022): ""The great retirement boom": The pandemic-era surge in retirements and implications for future labor force participation," Finance and Economics Discussion Series 2022-081, Board of Governors of the Federal Reserve System (U.S.).

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

Ozkan, S., J. Hubmer, S. Salgado, and E. Halvorsen (2023): "Why are the wealthiest so wealthy? A longitudinal empirical investigation," Working Papers 2023-004, Federal Reserve Bank of St. Louis.

Powell, J. H. (2022): "Inflation and the labor market," Speech at the hutchins center on fiscal and monetary policy, brookings institution, Board of Governors of the Federal Reserve System.

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

Smith, M., O. Zidar, and E. Zwick (2022): "Top wealth in America: New estimates under heterogeneous returns*," *The Quarterly Journal of Economics*, 138, 515–573.

Xavier, I. (2021): "Wealth inequality in the US: The role of heterogeneous returns," Working paper, Federal Reserve Board of Governors.

Online Appendix

A Model Appendix

This Appendix first provides details on the construction of the transition function for the stationary distribution discussed in Section 3.3. Next, we discuss in detail how we implement the economic impact payments in our model to supplement our discussions in Section 5.1.3.

A.1 Transition function for the stationary distribution

We now write down the law of motion for the distribution that defines the transition function $\mathcal{T}$. Let $a^p(j, a, w, \ell)$ be the policy function for savings of an agent with states (j, a, w, ℓ), and $\ell^p(j+1, a, w', \ell)$ be the policy function for employment status for an agent who has received an offer w' (recall that participation decisions are taken at the beginning of the following period, hence the discrepancy in age). For ease of notation, let us set $w = 0$ if the agent is not employed.

To describe the law of motion of the distribution, it is easier to describe the law of motion of each employment state separately. The law of motion for employed agents is:

$$\begin{aligned}\lambda(j+1, a', w', \ell' = E) = &\sum_a \sum_w \pi(j)\lambda(j, a, w, E)\mathbb{I}[a^p(j, a, w, E) = a']\mathbb{I}[\ell^p(j+1, a', w', E) = E]\mathbb{I}[w' = w][1 - \delta(w)] \\ &+ \sum_a \pi(j)\lambda(j, a, 0, U)\mathbb{I}[a^p(j, a, 0, U) = a']\mathbb{I}[\ell^p(j+1, a', w', U) = E] f g(w') \\ &+ \sum_a \pi(j)\lambda(j, a, 0, N)\mathbb{I}[a^p(j, a, 0, N) = a']\mathbb{I}[\ell^p(j+1, a', w', N) = E]\gamma_N f g(w') \\ &+ \sum_a \pi(j)\lambda(j, a, 0, R)\mathbb{I}[a^p(j, a, 0, R) = a']\mathbb{I}[\ell^p(j+1, a', w', R) = E]\gamma_R f g(w').\end{aligned}$$

For the employed at a given level of wealth and wage, we have all those who were already employed earning that wage who were not separated, as well as all those who were not employed and accepted a job offer at that given wage.

The law of motion for unemployed agents is:

$$\begin{aligned}
\lambda(j+1,a',0,\ell'=U) = &\sum_a\sum_w \pi(j)\lambda(j,a,w,E)\mathbb{I}[a^p(j,a,w,E)=a']\mathbb{I}[\ell^p(j+1,a',w,E)=U][1-\delta(w)]\\
&+\sum_a\sum_w \pi(j)\lambda(j,a,w,E)\mathbb{I}[a^p(j,a,w,E)=a']\mathbb{I}[\ell^p(j+1,a',0,U)=U]\delta(w)\\
&+\sum_a\sum_{w'} \pi(j)\lambda(j,a,0,U)\mathbb{I}[a^p(j,a,0,U)=a']\mathbb{I}[\ell^p(j+1,a',w',U)=U]fg(w')\\
&+\sum_a \pi(j)\lambda(j,a,0,U)\mathbb{I}[a^p(j,a,0,U)=a']\mathbb{I}[\ell^p(j+1,a',0,U)=U](1-f)\\
&+\sum_a\sum_{w'} \pi(j)\lambda(j,a,0,N)\mathbb{I}[a^p(j,a,0,N)=a']\mathbb{I}[\ell^p(j+1,a',w',N)=U]f\gamma_N g(w')\\
&+\sum_a \pi(j)\lambda(j,a,0,N)\mathbb{I}[a^p(j,a,0,N)=a']\mathbb{I}[\ell^p(j+1,a',0,N)=U][1-\gamma_N+\gamma_N(1-f)]\\
&+\sum_a\sum_{w'} \pi(j)\lambda(j,a,0,R)\mathbb{I}[a^p(j,a,0,R)=a']\mathbb{I}[\ell^p(j+1,a',w',R)=U]\gamma_R fg(w')\\
&+\sum_a \pi(j)\lambda(j,a,0,R)\mathbb{I}[a^p(j,a,0,R)=a']\mathbb{I}[\ell^p(j+1,a',0,R)=U][1-\gamma_R+\gamma_R(1-f)]\\
&+m\pi(j)q(a')\mathbb{I}[j=16].
\end{aligned}$$

For the unemployed, we have all those employed who move to unemployment either voluntarily or after job loss. For each other category, we have those who draw a wage and choose to remain unemployed and those who do not have a chance of drawing a wage and choose to move to (or stay in) unemployment. Finally, the term m is the mass of newborn agents who draw initial wealth a' with density $q(a')$ and enter directly into unemployment. This is equal to the total mass of agents who die each period to keep a constant population, i.e., $m = \sum_{j=16}^{90}\prod_{i=16}^{j-1}[1-\pi(i)]\pi(j)$.

The law of motion for non-participants is:

$$\begin{aligned}
\lambda(j+1,a',0,\ell'=N) = &\sum_a\sum_w \pi(j)\lambda(j,a,w,E)\mathbb{I}[a^p(j,a,w,E)=a']\mathbb{I}[\ell^p(j+1,a',w,E)=N][1-\delta(w)]\\
&+\sum_a\sum_w \pi(j)\lambda(j,a,w,E)\mathbb{I}[a^p(j,a,w,E)=a']\mathbb{I}[\ell^p(j+1,a',0,U)=N]\delta(w)\\
&+\sum_a\sum_{w'} \pi(j)\lambda(j,a,0,U)\mathbb{I}[a^p(j,a,0,U)=a']\mathbb{I}[\ell^p(j+1,a',w',U)=N]fg(w')\\
&+\sum_a \pi(j)\lambda(j,a,0,U)\mathbb{I}[a^p(j,a,0,U)=a']\mathbb{I}[\ell^p(j+1,a',0,U)=N](1-f)\\
&+\sum_a\sum_{w'} \pi(j)\lambda(j,a,0,N)\mathbb{I}[a^p(j,a,0,N)=a']\mathbb{I}[\ell^p(j+1,a',w',N)=N]\gamma_N fg(w')\\
&+\sum_a \pi(j)\lambda(j,a,0,N)\mathbb{I}[a^p(j,a,0,N)=a']\mathbb{I}[\ell^p(j+1,a',0,N)=N][1-\gamma_N+\gamma_N(1-f)]\\
&+\sum_a\sum_{w'} \pi(j)\lambda(j,a,0,R)\mathbb{I}[a^p(j,a,0,R)=a']\mathbb{I}[\ell^p(j+1,a',w',R)=N]\gamma_R fg(w')\\
&+\sum_a \pi(j)\lambda(j,a,0,R)\mathbb{I}[a^p(j,a,0,R)=a']\mathbb{I}[\ell^p(j+1,a',0,R)=N][1-\gamma_R+\gamma_R(1-f)].
\end{aligned}$$

For the non-participants, the law of motion is similar to that of unemployment.

Finally, the law of motion for retirees is very similar to that of non-participation with the exception that only agents age 62 and over can retire:

$$\begin{aligned}
\lambda(j+1,a',0,\ell'=R) = &\sum_a \sum_w \pi(j)\lambda(j,a,w,E)\mathbb{I}[a^p(j,a,w,E)=a']\mathbb{I}[\ell^p(j+1,a',w,E)=R][1-\delta(w)]\mathbb{I}[j\geq 62] \\
&+\sum_a \sum_w \pi(j)\lambda(j,a,w,E)\mathbb{I}[a^p(j,a,w,E)=a']\mathbb{I}[\ell^p(j+1,a',0,U)=R]\delta(w)\mathbb{I}[j\geq 62] \\
&+\sum_a \sum_{w'} \pi(j)\lambda(j,a,0,U)\mathbb{I}[a^p(j,a,0,U)=a']\mathbb{I}[\ell^p(j+1,a',w',U)=R]fg(w')\mathbb{I}[j\geq 62] \\
&+\sum_a \pi(j)\lambda(j,a,0,U)\mathbb{I}[a^p(j,a,0,U)=a']\mathbb{I}[\ell^p(j+1,a',0,U)=R](1-f)\mathbb{I}[j\geq 62] \\
&+\sum_a \sum_{w'} \pi(j)\lambda(j,a,0,N)\mathbb{I}[a^p(j,a,0,N)=a']\mathbb{I}[\ell^p(j+1,a',w',N)=R]\gamma_N fg(w')\mathbb{I}[j\geq 62] \\
&+\sum_a \pi(j)\lambda(j,a,0,N)\mathbb{I}[a^p(j,a,0,N)=a']\mathbb{I}[\ell^p(j+1,a',0,N)=R][1-\gamma_N+\gamma_N(1-f)]\mathbb{I}[j\geq 62] \\
&+\sum_a \sum_{w'} \pi(j)\lambda(j,a,0,R)\mathbb{I}[a^p(j,a,0,R)=a']\mathbb{I}[\ell^p(j+1,a',w',R)=R]\gamma_R fg(w')\mathbb{I}[j\geq 62] \\
&+\sum_a \pi(j)\lambda(j,a,0,R)\mathbb{I}[a^p(j,a,0,R)=a']\mathbb{I}[\ell^p(j+1,a',0,R)=R][1-\gamma_R+\gamma_R(1-f)]\mathbb{I}[j\geq 62].
\end{aligned}$$

A.2 Economic impact payments

Next, we describe in detail the mapping of economic impact payments during the 2020-2021 episode in the U.S. to the model. There were two features of these transfers that required detailed consideration when mapping to the model. First, not all households were eligible for these transfers, and the eligibility depended on income level. Second, the effective amount of transfer payments depended on household characteristics, especially the number of dependents.

For each of the three rounds of impact payments, households were eligible if their adjusted gross income (AGI) was lower than $75,000 and ineligible if their AGI exceeded $99,000 for the first two rounds and $80,000 for the third round, with payments phasing out linearly between these thresholds. These values roughly correspond to the 40th percentile of the income distribution. Thus, we set the 40th percentile of the stationary state income distribution as the eligibility cutoff in the model.

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 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,

Table A.1: Effective transfers for each age group of householder

Age of householder	Modifier	1st round	2nd round	3rd round
Under 20 years	0.42	0.13	0.08	0.17
20-24 years	0.18	0.12	0.06	0.14
25-29 years	0.34	0.12	0.07	0.16
30-34 years	0.61	0.14	0.09	0.20
35-39 years	0.78	0.15	0.10	0.22
40-44 years	0.64	0.14	0.09	0.20
45-49 years	0.43	0.13	0.08	0.17
50-54 years	0.22	0.12	0.07	0.15
55-59 years	0.11	0.11	0.06	0.14
60-64 years	0.08	0.11	0.06	0.13
65-74 years	0.05	0.11	0.06	0.13
75 years and over	0.03	0.11	0.06	0.13

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.

by the age of householder, for 2019.[21] 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 should receive. The 2019 dependent modifiers are presented in Table A.1. The effective transfer per eligible individual is then the adult transfer plus supplement for the dependent times the modifier for that individual's age.

First round. The first round of transfers was associated with the Coronavirus Aid, Relief, and Economic Security (CARES) Act and took place in March 2020. These transfers consisted of \$1,200 per person plus \$500 per child under the age of 17. We convert these amounts to the model by using the fact that the ex-ante average wage in the model is $\mu_w = 0.5$ and its data equivalent are average monthly earnings of around $\$65,000/12$, as per the 2019 CPS. Using CPI deflators $P^{2019}_{2020} = 1.012$ and $P^{2019}_{2021} = 1.059$, we arrive at the following transfer values for adults and children:

$$T^{adult}_{2020m3} = \frac{1200/1.012 \times (12 \times 0.5)}{65000} = 0.10946$$
$$T^{child}_{2020m3} = \frac{500/1.012 \times (12 \times 0.5)}{65000} = 0.04561.$$

The effective transfer is then computed as the adult transfer plus the relevant modifier times the dependent transfer. For example, for a household between 20-24 years of age, the effective

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

transfer amounts from the first round is computed as $0.10946 + 0.04561 \times 0.18 = 0.12$, which is shown in the second column of Table A.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:

$$T^{adult}_{2020m12} = \frac{600/1.012 \times (12 \times 0.5)}{65000} = 0.05473$$
$$T^{child}_{2020m12} = T^{adult}_{2020m12}.$$

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

$$T^{adult}_{2021m3} = \frac{1400/1.059 \times (12 \times 0.5)}{65000} = 0.12203$$
$$T^{child}_{2021m3} = T^{adult}_{2021m3}.$$

B Data Appendix

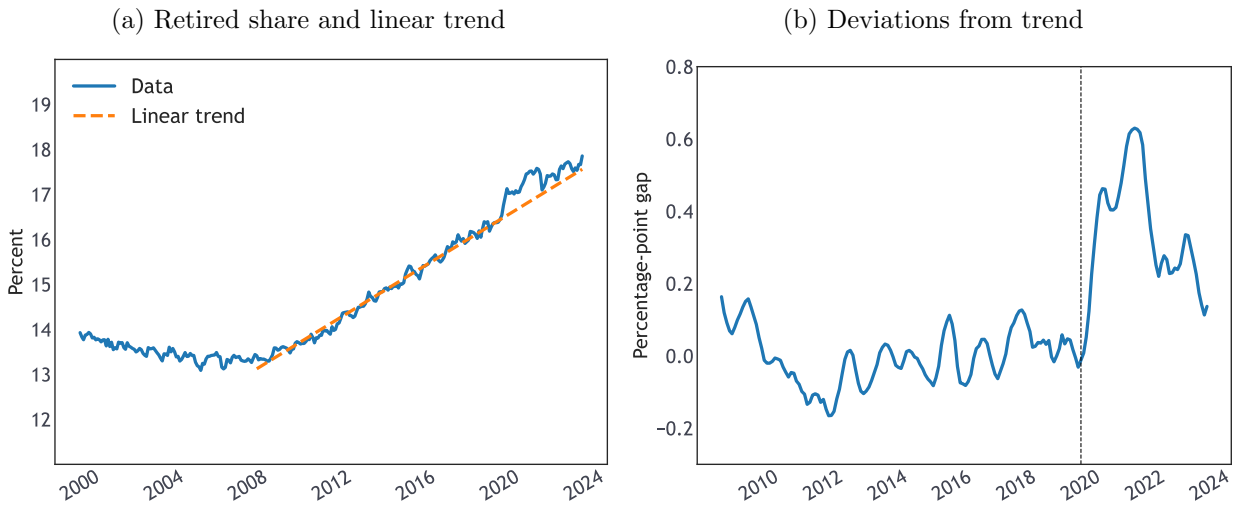
In this Appendix, we provide details on our empirical analysis to supplement the discussions in the main text and provide additional results from the data.

B.1 Current Population Survey

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.

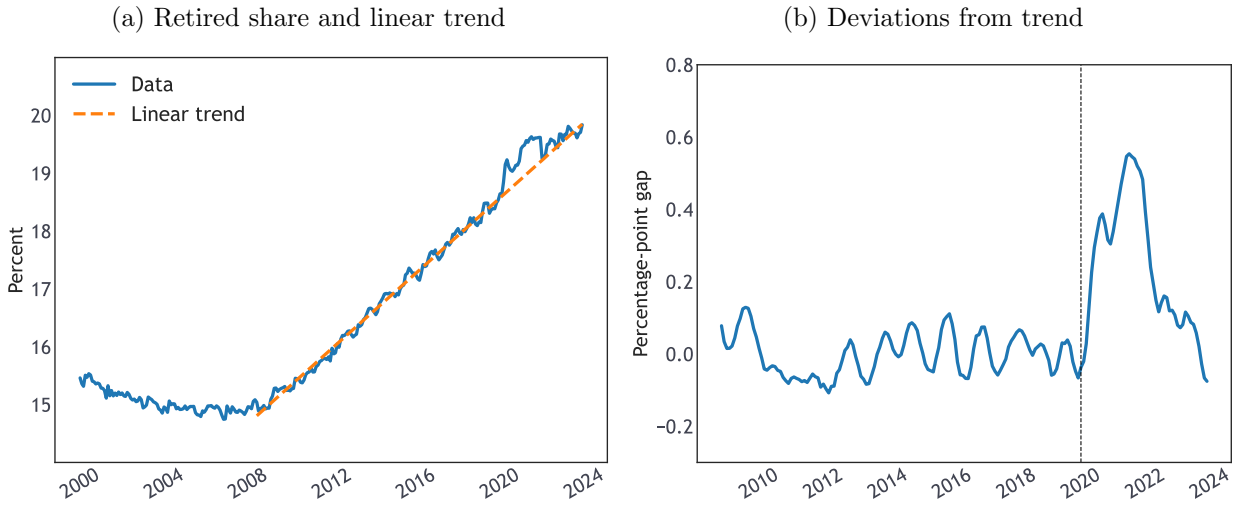
We have also experimented with alternative definitions of retirement. Figures B.1 and B.2 replicate Figure 1 for two such alternative definitions. Figure B.1 considers a stricter definition where a person is considered retired if `EMPSTAT` is equal to 36 and age is equal to or greater than 62. This is 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 B.2, on the other hand, considers a slightly broader definition of retirement: `EMPSTAT` is equal to or greater than 30 and age is equal to or greater than 62. This means that we define the retired group as all people who are at least 62 years old and are out of the labor force. Figures B.1 and B.2 show that our measure of the retired share (i.e., excess retirement share) is robust to alternative definitions of retirement.

Figure B.1: Alternative retirement definition: Self-identified retirees who are min 62 years old



Note: Panel (a) plots the retired share in the U.S. economy which we calculate 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.

Figure B.2: Alternative retirement definition: Non-participants who are min 62 years old



Note: Panel (a) plots the retired share in the U.S. economy which we calculate 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.

B.2 Survey of Income and Program Participation

We use the SIPP data for two purposes. First, we calculate the wealth distribution for each year between 2019 and 2021. These results are presented in Figure 3. Second, we calculate the distribution of new retirees across the wealth distribution, separately for those who retire in 2019 and those who retire between 2020 and 2021. These results are presented in Figure 2. We now provide details on these calculations.

For these calculations, we use SIPP 2020, 2021, and 2022 panels covering data from the start of 2019 to the end of 2021.[22] Our sample consists of all individuals (excluding those in armed forces) aged 16 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 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 given then 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. Then, for each year, we calculate the average and various percentiles of the net worth distribution using weights.

Distribution of new retirees across 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, $\texttt{RWKESR}j = 5$ 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, $\texttt{RWKESR}j \leq 2 \, \forall j \in \{1, 2, 3, 4, 5\}$. 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.[23] Then, we assign each new retiree in 2019 into quintiles of the economy-wide net worth distribution (as calculated above) in 2019

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

[23] 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.

using their own household-level net worth. This allows us to calculate the fraction of new retirees in 2019 at each net worth quintile among all new retirees in 2019. We repeat the same procedure to calculate the same moments for new retirees between 2020 and 2021.

Distribution of new retirees across income quintiles. Recall that we also present the empirical distribution of new retirees across labor income quintiles in 2019 in Figure 6. We obtain this result following the same procedure as above except that we use total labor earnings (instead of net worth) to classify individuals into quintiles of the labor earnings distribution. We measure labor earnings 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.[24]

B.3 Survey of Consumer Finances

We use the 2019 wave of the Survey of Consumer Finances, 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 in our regressions.

B.4 Asset returns

Finally, we explain how we obtain time series of cumulative real returns on selected asset classes relative to 2019, as shown in Figure 4. All monthly series for these asset returns are taken from FRED, from where we report the mnemonics. For stocks 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.

[24] For the first job, weekly earnings are given by `TJB1_WKSUM1` to `TJB1_WKSUM5`. For the second job, weekly earnings 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.