

NBER WORKING PAPER SERIES

WHO BENEFITS FROM RETIREMENT SAVING INCENTIVES IN THE U.S.?
EVIDENCE ON GAPS IN RETIREMENT WEALTH ACCUMULATION
BY RACE AND PARENTAL INCOME

Taha Choukhmane
Jorge Colmenares
Cormac O'Dea
Jonathan L. Rothbaum
Lawrence D.W. Schmidt

Working Paper 32843
<http://www.nber.org/papers/w32843>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2024, Revised June 2026

We are very grateful to Jackson Howell and Cashua Wang who provided outstanding research assistance. We thank Franca Glenzer, Luigi Guiso, Pascal Noel, and Natasha Sarin for illuminating discussions of this paper. We also thank James Choi, Amir Kermani, Debbie Lucas, Olivia S. Mitchell, Jim Poterba, John Sabelhaus, Jason Seligman, Angelino Viceisza, and Stephen Zeldes for extremely helpful conversations and comments. We thank Cristina Tello Trillo for her assistance in building the bridge from the Form 5500 data to the individual-level data. We are grateful to numerous seminar and conference participants for their helpful comments and suggestions. This report is released to inform interested parties of ongoing research and to encourage discussion. Any views expressed on statistical, methodological, technical, or operational issues are those of the authors and not necessarily those of the U.S. Census Bureau. The data in this paper have been cleared by the Census Bureau's Disclosure Review Board release authorization numbers CBDRB-FY22-SEHSD003-001, CBDRB-FY22-SEHSD003-017, CBDRB-FY22-SEHSD003-033, CBDRB-FY23-SEHSD003-043, CBDRB-FY23-0494, CBDRB-FY24-0151, CBDRB-FY25-0400, and CBDRB-FY26-0106. The research reported herein was performed pursuant to a grant from the U.S. Social Security Administration (SSA) funded as part of the Retirement and Disability Research Consortium through the Michigan Retirement and Disability Research Center Award RDR18000002. The opinions and conclusions expressed are solely those of the author(s) and do not represent the opinions or policy of SSA, any agency of the Federal Government, or the National Bureau of Economic Research. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or

usefulness of the contents of this report. Reference herein to any specific commercial product, process or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply endorsement, recommendation or favoring by the United States Government or any agency thereof. We are very grateful to the MIT Institute Community & Equity Office, the MIT Sloan School of Management, the Yale University Tobin Center for Economic Policy, and the Yale University Arthur M. Okun Research and Policy Fund for co-funding this work.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2024 by Taha Choukhmane, Jorge Colmenares, Cormac O'Dea, Jonathan L. Rothbaum, and Lawrence D.W. Schmidt. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Who Benefits from Retirement Saving Incentives in the U.S.? Evidence on Gaps in Retirement Wealth Accumulation by Race and Parental Income

Taha Choukhmane, Jorge Colmenares, Cormac O'Dea, Jonathan L. Rothbaum, and Lawrence D.W. Schmidt

NBER Working Paper No. 32843

August 2024, Revised June 2026

JEL No. D14, G5, H2, H31, J15, J32

ABSTRACT

U.S. employers and the federal government devote the equivalent of 1.5% of GDP annually toward promoting defined contribution (DC) retirement savings. Using a new employer-employee linked dataset covering millions of Americans, we show that tax and employer matching incentives disproportionately benefit White and Asian workers compared to their similar-income Hispanic, Black, and American Indian or Alaska Native coworkers. Similarly, these incentives disproportionately benefit those with richer parents compared to those from lower-income families. Breaking the link between contribution choices and saving subsidies through revenue-neutral reforms could close up to one-third of the DC wealth gaps by race and parental income.

Taha Choukhmane
Massachusetts Institute of Technology
MIT Sloan School of Management
and NBER
tahac@mit.edu

Jorge Colmenares
Harvard University
jcolmenares@fas.harvard.edu

Cormac O'Dea
Yale University
Department of Economics
and NBER
cormac.odea@yale.edu

Jonathan L. Rothbaum
U.S. Census Bureau
jonathan.l.rothbaum@census.gov

Lawrence D.W. Schmidt
Massachusetts Institute of Technology
and NBER
ldws@mit.edu

1 Introduction

Every year, the equivalent of 1.5% of U.S. GDP is devoted to encouraging contributions to retirement savings plans such as 401(k) accounts.¹ Around 100 million Americans have access to such defined contribution (DC) plans through their employers, and these accounts offer an attractive vehicle for long-term saving. Contributions are taxed favorably, and over 80% of employers further subsidize savers by matching their employees’ contributions (Arnoud et al., 2021). This institutional design, therefore, rewards those who can, and do, save more for retirement. Employees who do not contribute receive neither tax benefits nor employer-matching contributions.

In this paper, we ask how much retirement savings incentives contribute to racial, ethnic, and inter-generational wealth inequality in the United States.² To address this question, we create a unique dataset linking (1) the retirement contributions and withdrawals of millions of Americans from their federal tax filings, (2) demographic information from survey responses to 10 years of the American Community Survey (ACS), and (3) hand-collected data on the characteristics of over 6,000 DC retirement plans, covering approximately 35 million employees from employers’ Form 5500 filings. We use this new dataset to assess the distributional impacts of retirement saving incentives by race and parental income.

To do so, we first measure contributions to DC accounts (as groups that contribute more receive more tax and matching benefits) and measure withdrawals (as groups that take more early withdrawals forgo long-term tax benefits and face tax penalties). We then build a model to translate the observed differences in access and contributions to DC plans into differences in lifetime wealth accumulation and to study the mechanical impact of budget-neutral reforms to tax and employer saving subsidies on wealth inequality. The model also allows us to study retirement saving incentives in the context of the broader tax and retirement system and, importantly, takes into account Social Security benefits.

We divide our analysis into three parts. First, we document significant differences in retirement contributions by race and parental income among workers with similar ages and earnings. Black, Hispanic, and American Indian and Alaska Native (AIAN) workers whose employers offer a DC plan contribute significantly less than White workers with comparable earnings, while Asian workers contribute significantly more. Similarly, workers who grew up

¹In 2022, federal tax expenditures on defined contribution (DC) retirement accounts amounted to \$185 billion (US Department of the Treasury, 2024). In 2022, private sector employers contributed more than \$212 billion into these accounts (Department of Labor, 2024)—mainly in the form of matching contributions.

²We focus on the five largest racial and ethnic groups: non-Hispanic White, Hispanic, non-Hispanic Black, non-Hispanic Asian, and non-Hispanic American Indian and Alaska Native, who together make up 98.2% of the individuals in our sample. We will often use “race” to refer to both race and ethnicity, “White” to refer to non-Hispanic White, “Black” to refer to non-Hispanic Black, “Asian” to refer to non-Hispanic Asian, and “AIAN” to refer to non-Hispanic American Indian and Alaska Native.

in lower-income families contribute less than those with higher-income parents. While access to employer-sponsored retirement plans also varies by race and parental income, these differences are small when comparing workers with similar characteristics. In contrast, sizable contribution gaps remain among workers at employers with DC plans, even when comparing those with similar education, occupation, tenure, and employer. Employer matching amplifies these savings differences by channeling more employer contributions to groups that save more. Notably, among workers with similar ages and incomes, employer matching widens the Black-White contribution gap by 28%, the Hispanic-White gap by 33%, and the gap between workers from the highest and lowest parental income quintiles by 34%. Comparing workers with similar characteristics, and in particular age and income, is important because these workers accrue similar Social Security entitlements. Therefore, residual differences in subsidies are not directly mitigated by other aspects of the broader retirement system.

In the second part of the paper, we turn to early withdrawals. We find that groups with lower contributions withdraw more frequently from their retirement accounts. Among savers with at least \$1,000 in recent contributions and similar ages and incomes, pre-retirement distributions vary significantly by race: Black savers are 86% more likely to take an early withdrawal than White workers, Hispanic and AIAN savers are 17% and 32% more likely, respectively, while Asian savers are 30% less likely. Similarly, savers with parents in the bottom income quintile are 57% more likely to take an early withdrawal than comparable savers with parents in the top quintile. Because groups already contributing less also withdraw more frequently, early withdrawals amplify existing gaps by further depleting smaller account balances. However, this amplification effect is partially offset by a countervailing force: groups with higher contributions have larger balances available to withdraw from, allowing for larger withdrawal amounts. The net effect of these two forces varies across groups. Accounting for withdrawals amplifies the Black-White, Asian-White, and top-bottom parental income *net* contribution gaps, while attenuating the Hispanic-White gap. Beyond their direct impact on wealth accumulation, early withdrawals also shape the lifetime distributional impact of DC tax subsidies, as savers often forgo tax benefits and can face tax penalties.

In the third part of the paper, we develop a micro-simulation model that uses our data on access to DC plans and flows of earnings, employee contributions, employer matches, and early withdrawals to compute measures of retirement wealth. We use this model for three purposes. First, we evaluate the lifetime distribution of savings subsidies. The value of these for any individual, especially the benefits from favorable taxation, depend on the entire trajectory of their incomes, contributions, and withdrawals, and so cannot be calculated directly from the data. Second, the model allows us to incorporate taxation and Social Security benefit accrual explicitly and, therefore, evaluate the distributional effect of

retirement incentives in the context of the broader U.S. retirement system. Finally, we use the model to assess the mechanical effect of revenue-neutral reforms that would redistribute employer-match dollars within each firm and federal tax expenditures across the population so that they are distributed 1) proportionally to earnings but 2) independently of workers' own contribution choices. We estimate that, for median earners, these reforms could close each of the Black-White, Hispanic-White, and Asian-White DC wealth gaps by a quarter to a third, and the gap between those with the richest and poorest parents by close to a third. This estimated change in the gaps would be (i) similar if workers reduced their contributions in response to removing match and tax incentives (assuming an elasticity calibrated to the upper end of empirical estimates), and (ii) only modestly smaller if the match redistribution happened across employees in the same income bin within each firm, rather than across all employees in that firm.

The first branch of the literature to which we contribute is that on the design and impact of retirement savings incentives. A large literature has studied the effect of tax and employer matching incentives on private savings and finds small to insignificant behavioral responses to these incentives (see Choi (2015) and Friedman (2015) for reviews). Retirement incentives seem to mainly shift the location of existing savings rather than creating new savings (Chetty et al., 2014; Choukhmane and Palmer, 2025). In contrast to this large literature, the mechanical effect of these incentives has received less attention. While the distributional impact of the federal tax expenditure by income has been studied (Burman et al., 2004; Brown et al., 2022), quantitative evidence on the full effect of retirement saving subsidies (including employer subsidies) has been limited due to a lack of systematic data on retirement plan characteristics. The fact that we directly observe the match schedule means that we can precisely measure the contribution of employer subsidies to retirement wealth. In doing so, we are able to quantify the total contribution of these savings supports to retirement wealth. We find their contribution to be large: at the bottom of the lifetime earnings distribution, they account for 40% of DC wealth, while at the top, they account for almost 50% of (a much larger base amount of) DC wealth. In addition to improving the measurement of saving subsidies, we contribute to this literature by bringing in richer demographic characteristics. We show how the receipt of subsidies differs by race and parental income for individuals who have the same lifetime income and thus likely face similar tax incentives and expect similar Social Security benefits. The disparities in subsidy receipt among co-workers of similar ages, incomes, and geographic locations are less likely than disparities by income to be mitigated by other features of the broader retirement system and Social Security formula.

The second branch of the literature to which we contribute is that concerned with race/ethnicity, earnings, and wealth in the U.S. The gap between Black and White wealth

is large (Oliver and Shapiro, 1989; Darity and Nicholson, 2005), stable since the 1980s (Derenoncourt et al., 2022), and cannot be fully accounted for by earnings differences (Blau and Graham, 1990; Barsky et al., 2002; Altonji and Doraszelski, 2005; Kuhn et al., 2020; Derenoncourt et al., 2023; Catherine et al., 2025). There is also a large wealth gap between Hispanic and non-Hispanic White individuals, with the latter being approximately four times wealthier (Sabelhaus and Thompson, 2021). Our contribution to this literature is to study one channel that contributes to both earnings and wealth inequality by race and ethnicity. On the earnings side, our contribution is to measure an often-unmeasured component of earnings—the employer match—which gives a compensation premium to those who save more. In doing so, we contribute to a rich literature on racial earnings inequality.³ On the wealth side, Derenoncourt et al. (2022) emphasize that differences in rates of return are the dominant factor shaping the lack of racial wealth convergence over the past 30 years, the period in which DC accounts have emerged as the main vehicle for private retirement savings. Our results shed light on an important mechanism generating such differences in rates of return across racial groups, even holding portfolio risk constant: differences in the take-up of employer match and tax incentives.

The differences by race in broad measures of wealth referenced above are also seen in retirement wealth (Hou and Sanzenbacher, 2021; Francis and Weller, 2021; Viceisza et al., 2022; Wolff, 2023). Closer to our paper, a number of studies have examined racial differences in savings rates. While studies using administrative retirement record-keeper data from one firm (Kuan et al., 2015) or a small number of firms (Ariel/AON Hewitt, 2009) find differences in saving rates by race, the literature using representative survey data has typically found that these patterns no longer hold after accounting for income differences.⁴ However, surveys can suffer from significant measurement error, as shown by studies comparing self-reported measures of access and contributions to DC plans with the respondents’ tax records.⁵ Administrative tax data combines the strength of both data sources: they have better coverage and are more representative than record-keeper data while suffering from less

³Altonji and Blank (1999) offer a comprehensive review of studies to that date, and Bayer and Charles (2018), Chetty et al. (2020), and Derenoncourt and Montialoux (2021) provide more recent evidence.

⁴A review of the evidence by Darity Jr et al. (2018) concludes that “the finding advanced in peer-reviewed articles in economic journals is clear: there is no evidence that black Americans have a lower savings rate than white Americans once household income is taken into account”. We show a similar result using data from the Survey of Consumer Finances (SCF), albeit with wide confidence intervals.

⁵Dushi and Iams (2010) find that 24% of private-sector and 36% of public-sector respondents to the 2006 Survey of Income and Program Participation (SIPP) misreport making a tax-deferred DC contribution relative to their W2 records (with both false positives and negatives being common). Similarly, Dushi and Honig (2015) find that the average absolute difference between annual DC contributions reported in the Health and Retirement Study (HRS) and respondents’ W2 records was approximately 1.5 times larger than the mean DC contribution in the W2s. Likewise, Bee and Mitchell (2017) show that survey respondents vastly underreport their DC plan withdrawals.

measurement and misreporting error than surveys. Using such data yields a different result than that seen in prior studies: racial gaps in contributions persist even among co-workers and after controlling for individual and household incomes. Additionally, the scale of our data allows us to measure differences in saving and wealth accumulation for racial groups with smaller population shares, such as Asian and AIAN individuals, for whom evidence on wealth accumulation and saving behavior has been more limited due to small sample sizes in most surveys. Other recent papers also use tax data to study differences in participation in retirement plans (Yogo et al., 2023) and penalized early withdrawals (Coyne et al., 2022). We add to this literature, first by measuring contributions (in addition to participation and withdrawals), second by observing racial self-identification directly rather than imputing it, and third by bringing in detailed information on plan features—which are important determinants of differences in contribution rates.

Our paper contributes to a third, closely related literature that evaluates the disparate impact of policies by race. This includes research examining racial disparities in welfare programs (Darity and Myers, 1983, 1987), unemployment insurance (Kuka and Stuart, 2021; Skandalis et al., 2022), credit and mortgage access (Myers Jr, 1995; Ross and Yinger, 2002; Bhutta and Hizmo, 2021; Bakker et al., 2025), housing returns (Kermani and Wong, 2021), property tax assessments (Avenancio-León and Howard, 2022), and financial aid for college (Levine and Ritter, 2022). In the context of retirement plans, Brown (2021) argues that the design of retirement incentives favors activities that are more likely to be carried out by White Americans (retirement saving) and penalizes activities that are more likely to be carried out by Black Americans (early withdrawals). More broadly, Hamilton and Darity (2017) argue that “if the existing federal asset-promotion budget were allocated in a more progressive manner, federal policies would go a long way toward eliminating racial disparities and building an inclusive economy for all Americans.” We quantify how much changing a major component of the U.S. asset-promotion budget, namely the design of retirement savings subsidies, could affect racial and intergenerational wealth inequality.

The fourth branch of literature to which we contribute is that on intergenerational persistence in wealth. The correlation in wealth across generations has been well documented (Charles and Hurst, 2003). Recent work emphasizes the importance of heterogeneity in rates of return for cross-sectional wealth inequality (Fagereng et al., 2020). Our paper draws a link between these two phenomena. While it has long been known that the rich save more (Dynan et al., 2004), we show that the *children* of the rich save more conditional on their own earnings. The saving in question here is, by virtue of matching, one with an extraordinary rate of return. This correlation between the resources of one generation and the rates of return of the next will directly contribute to intergenerational persistence in wealth. This

channel also relates to a theme that has been emphasized in the literature on wealth gaps by race in the U.S. Chiteji and Hamilton (2002) and Charles and Hurst (2002) highlight the role of the family in savings decisions and the direction of intergenerational transfers: Black individuals are both more likely to provide financial support to their parents and less likely to receive support from their parents than White individuals. This is consistent with our finding that, even among those with similar *individual* characteristics, accounting for differences in parental and household resources reduces gaps in contributions by race.

We proceed as follows. Section 2 discusses the institutional background. Section 3 introduces our data. Section 4 gives our results on gaps in retirement saving rates by race and parental income. Section 5 turns to early withdrawals. Section 6 uses our data and a microsimulation model to study the implications of the saving patterns we document for the distribution of wealth at retirement, and to simulate alternative retirement savings policies. Section 7 concludes.

2 Institutional Background

DC plans have become the dominant vehicle through which Americans save for retirement. Sixty percent of U.S. civilian workers now have access to an employer-sponsored DC plan (Myers and Topoleski, 2020). Participants in these plans can make pretax contributions to their accounts (up to an annual maximum employee contribution of \$23,000 in 2024), thereby deferring income taxes to when they retire and when they will (likely) face lower tax rates. In addition to the advantages that this deferral brings, dividends and capital gains are untaxed provided that they remain in the account. Wealth held in DC plans is illiquid: participants generally face tax penalties on withdrawals made before the age of 59.5, though some plans permit borrowing against existing DC balances.

DC plans provide substantial flexibility and discretion to participants in deciding how much to save. This structure contrasts substantially with defined benefit (DB) plans, in which employees typically only choose whether to participate, and *employer* contributions do not depend on any choice that the *employee* makes. The shift away from DB toward DC plans in recent years moves considerable risk related to financing retirement income from employers to employees.⁶ Whereas traditional pension plans insure against mortality risk and the lion’s share of risks associated with fluctuations in investment returns, DC plans force households to self-insure against these risks. In the vast majority of DC plans,

⁶Only a quarter of civilian workers now have access to a DB pension (Myers and Topoleski, 2020), a share that continues to fall. DC plans are becoming, alongside Social Security, one of the largest sources of income in retirement. Devlin-Foltz et al. (2016) show that over the past 30 years, the dynamics of retirement wealth have had a moderating impact on overall wealth inequality. They also find, however, that DC wealth is more concentrated than DB wealth.

employers match employee contributions at some rate up to a cap, meaning that the amount contributed by the employer depends on how much the employee chooses to save.

In contrast, the rules governing both employee and employer Social Security contributions are more rigid. Social Security payments are financed via non-discretionary FICA payroll tax contributions from both employers and employees on each dollar of labor earnings up to a taxable maximum. Social Security benefits to workers are then computed as a function of the worker’s earnings history. These benefit amounts are progressive, implying that low-income workers generally receive larger benefit payments per dollar of payroll tax contributions than higher-income workers in the same cohorts. The progressivity of Social Security means that measures of wealth that include it display less inequality and narrower racial gaps than measures without it (Catherine and Sarin (2023), Sabelhaus (2023)).

As DC plans become more dominant and DB coverage recedes, individuals’ decisions play a larger role in determining retirement wealth, and employer plan design can amplify the implications of these decisions for wealth inequality. Endogenous DC participation also implies that the benefits paid to employees in the form of matching contributions will not be equally distributed across workers, even among those with identical earnings (and Social Security entitlements). To study the interplay between individual saving decisions and firm matches, we therefore need data that contain both the saving decisions made by individuals and the full match schedules offered by their employers.

3 Data

We build a new dataset linking administrative data on retirement savings and the demographics of a large sample of U.S. employees with a newly constructed data set on employer-sponsored retirement plan characteristics.

3.1 Employee data

For the analysis, we start with all individuals in the 2008 to 2017 waves of the American Community Survey (ACS).⁷ We link ACS respondents to other administrative data using protected identification keys (PIKs).⁸ With this method, 90–94% of ACS respondents are

⁷From 2005 to 2019, the ACS averaged over 3.2 million individuals surveyed, including a sample expansion from 2010 to 2012. Refer to <https://www.census.gov/acs/www/methodology/sample-size-and-data-quality/> for more information on the ACS sample and response rates over time; accessed 7/12/2024.

⁸PIKs are assigned by a probabilistic matching algorithm that compares the characteristics of records in Census, survey, and administrative data to those in a reference file constructed from the Social Security Administration Numerical Identification System and other federal administrative data. PIKs correspond one-to-one with SSNs and so allow us to link individuals over time and across data sources. For more information, see Wagner and Layne (2014).

successfully assigned a PIK in any given year (Ferrie et al., 2021).⁹ Next, we link ACS respondents with their 1040, W-2, and 1099-R tax filings. The ACS provides individuals’ race, year, age, education, gender, occupation, and location at the time of the survey. The 1040 and W-2 filings provide other socioeconomic and demographic indicators, including wages, deferred compensation, family structure, employer identification number (EIN), employment tenure, spousal income, and intergenerational linkages (through which we can observe parental income). Appendices A.1 and A.2 provide detailed overviews of our data build and variable construction, respectively.

For our primary analysis, we report differences in contributions and withdrawals between the five largest racial and ethnic groups in the US, which cover 98.2% of the sample: non-Hispanic White, non-Hispanic Black, non-Hispanic Asian, non-Hispanic AIAN, and Hispanic. Following Chetty et al. (2020), these groups are constructed using the answers to two questions asked in the ACS about race and Hispanic origin. For brevity, throughout the paper, we use the term race to refer to both race and ethnicity. The “Hispanic” group includes any individuals who answer yes to the question about Hispanic origin. We let “AIAN”, “Asian”, “Black”, and “White” refer to non-Hispanic AIAN, Asian, Black, and White individuals, respectively.¹⁰ As is well established in the literature, each of these groups includes many heterogeneous subgroups, and the results we document might not apply equally to all of these subgroups.

3.2 Employer retirement plan data

All employers must submit an annual regulatory form (Form 5500) on their U.S. retirement plans to the federal government. Plans with over 100 participants provide narrative descriptions of plan characteristics. We create a data set by extracting these descriptions from the original free-form text.¹¹ We set out to do this for the largest 4,730 plans in the U.S. and a random sample of 1,471 smaller plans. These employers cover a substantial portion of the U.S. population; in 2017, 37 million employees were eligible for these large plans, constituting about half of employees with access to private and nonprofit-sector DC retirement plans.¹² Appendix A.1.3 provides further details. These plan-level data, further detailed in

⁹As noted in Bond et al. (2014), there is some selection into linkage, for example by age, race, and citizenship status. However, we do not believe that the magnitudes of these differences will bias our estimates substantially. For example, in 2010, the linkage rate for Black ACS respondents was 91.4%, compared to 93.5% for White respondents.

¹⁰Every regression estimate in the paper includes all racial groups listed in the ACS, including the remaining 1.8% of our primary sample: Native Hawaiian and Other Pacific Islander (0.15%), Some other race (0.15%), and Two or more (1.5%), which is anyone who identifies as more than one non-Hispanic group.

¹¹Refer to <https://www.dol.gov/agencies/ebsa/about-ebsa/our-activities/public-disclosure/foia/form-5500-datasets>.

¹²Note that while employers include firms, hospitals, non-profits, and other non-firm employers, “firm”

Arnoud et al. (2021) and Choukhmane et al. (2025), include information on vesting schedules, auto-enrollment, and, crucially for our question, employer match schedules. These match schedules are typically concave functions of employee contribution rates, often linear up to a threshold.

3.3 Samples

We match the retirement plan data with our employee data using numeric identifiers such as the EIN, telephone number, name, and address fields. The two main resulting samples are: the “Form 5500 sample,” which links ACS respondents to employer retirement plan data, and the “Parent-Form 5500 sample,” a subset of the former that can be linked to parents. Both are restricted to: i) ages 25–59.5 (though, all individuals in the Parent-Form 5500 sample are under 40) and ii) individuals with annual labor income in the year they respond to the ACS above \$8,000 (Box 1 wages plus deferred compensation) in nominal terms, which corresponds to roughly 20 hours per week at the current federal minimum wage. We exclude observations with zero Box 1 wages or missing ACS demographics.

We impose two restrictions on which employers we retain in the linked data. First, so that a link to a firm can serve as a link to a particular match formula, the employer must use the same match formula for all employees in the plan. This is the case for about 70% of the plans in our linked Form 5500 sample. Second, to limit incorrect matches or inaccurate formulas, we filter out plans for which there is too large of a discrepancy between the ratio of employer to total contributions from Form 5500 and an analogous measure computed from applying our matching rules to the linked W-2 data. Around 75% of the remaining plans satisfy this test and are retained in our final sample. See Appendix A.3 for further details on these two restrictions, and Table A.1 on how the sample size changes sequentially with the selection criteria.

The final Form 5500 sample includes about 1.58 million individuals, while the Parent-Form 5500 sample, restricted to the 1978–1994 birth cohorts (due to data limitations on when parent-child links are observable), includes 435,000 younger individuals. We weight estimates to be nationally representative of the population of US workers with employers offering retirement plans with at least 100 participants.

To assess the representativeness of our Form 5500 linked employer-employee sample, we compare it to a broader sample of all ACS respondents at employers offering a DC plan (approximately 9.7 million individuals). Observable characteristics (Appendix Table A.2) and savings behavior across groups (Appendix Figure A.7) are broadly similar between our Form 5500 sample and the full ACS sample of workers at DC-offering employers. Appendix and “employer” are used interchangeably throughout this paper.

A.3 provides additional details on sample construction and weighting, while Appendix A.4 further discusses representativeness.

3.4 Retirement savings outcomes

Our six primary measures of retirement saving and withdrawals are: i) **Employee contributions**: deferred compensation reported in Box 12 of the W-2, corresponding to contributions to an employer-sponsored plan (e.g., 401(k)). The contribution *rate* is defined as a percentage of salary, i.e., Box 12 divided by the sum of Box 1 wages and Box 12.¹³ ii) **Employee plus employer contributions**: the sum of employee contributions and employer match, calculated using match formulas from Form 5500 data. iii) **Participation**: indicator equal to one if a worker makes a positive contribution. iv) **Early withdrawal**: dummy equal to one if an individual aged 25–54 reports a distribution on Form 1099-R.¹⁴ v) **Withdrawal amount**: withdrawals in year $t + 1$ as a share of salary in year t . vi) **Net contribution rate**: employee plus employer contribution rate in year t minus early withdrawals in $t + 1$, relative to year t income. Unless otherwise specified, all dollar-denominated variables, including those used to compute rates, are deflated to base year 2017 using the Consumer Price Index provided by the Bureau of Labor Statistics. Appendix A.2 provides further details on variable construction.

4 The Distribution of Retirement Contributions

Groups that contribute more to retirement accounts receive more employer matching and tax benefits. In this section, we document differences by race and parental income in retirement contribution rates and the interplay between these gaps and saving incentives. First, we find that employee contribution gaps are large and that they are amplified by employer matching, where sizable differences remain even among workers in the same firm with similar individual-level characteristics. Second, given the notable differences in parental income by race, we consider the interplay between contribution gaps by race and parental income. Finally, we conclude by discussing the role of other plausible correlates of contribution gaps.

4.1 Contribution gaps by race and parental income at firms offering a DC plan

Raw contribution gaps are large and amplified by employer matches. Before presenting our main analysis, which compares savings rates across groups with similar incomes,

¹³We do not observe direct contributions to IRA accounts. However, in 2017, over 96% of inflows to traditional IRAs came from rollovers rather than direct contributions, and these rollover assets would have been captured in our data when originally contributed to employer-sponsored retirement plans (Myers, 2024).

¹⁴After this age, withdrawals upon separation from an employer do not incur a tax penalty.

we briefly discuss raw differences between groups. Among employees at firms offering DC plans, contribution rates vary significantly by race and ethnicity. White workers contribute an average of 4.2% of salary, while Black (2.4%), Hispanic (2.6%), and AIAN (2.8%) workers contribute substantially less—between 33% and 43% below the White worker average. Asian workers have the highest average contribution rate at 5.4% of salary, consistent with prior research findings (Gindelsky and Martin (2024)).

We observe similar patterns when examining differences by parental income, though this variable is available only for the younger cohorts who are in our Parent-Form 5500 sample. Given that younger workers typically contribute less, contribution rates are lower in this subsample: averaging 2.8% compared to 3.8% in the full sample. Table 1 shows that workers with parents in the top income quintile contribute nearly twice as much as those whose parents are in the bottom quintile.

Employer matching amplifies the effect of these differences in employee contribution rates. The average matching contribution rates in Table 1 imply that employer matching widens racial contribution gaps (as a share of salary) by 35-39%, except for the White-Asian gap, which is only modestly affected. Similarly, the gap between workers from the top and bottom parental income quintiles widens by 42% after accounting for employer matching.

Which mediating factors *should* we include? A natural question is the extent to which these sizable gaps persist conditional on observable characteristics. In an analysis such as ours, whether to partial out the influence of a given mediating factor depends on the question at hand.¹⁵ Our objective is to study the distributional effects of retirement saving subsidies. These subsidies do not exist in isolation, and a challenge to interpreting raw differences across groups is that this mechanical impact of employer matching and retirement tax benefits could be offset by other aspects of the broader U.S. retirement system. For instance, those receiving less in employer matching and tax benefits may benefit more from Social Security’s progressive benefit formula, which provides higher returns per dollar of payroll taxes for lower earners.

We present results using two sets of mediating variables to address different questions. First, we include calendar time, age, and labor income as mediators, which are the cross-sectional versions of the key determinants of Social Security benefit accrual. This isolates retirement saving gaps that are unlikely to be offset by Social Security’s progressive structure. Second, we add mediators for location, education, occupation, gender, tenure, and firm characteristics. This second analysis examines what explains the observed saving differences across groups and, of independent interest, reveals the extent to which gaps (conditional on

¹⁵See, for example, Neal and Johnson (1996), Lang and Manove (2011), and Carneiro et al. (2005) for discussions of test score and education controls in racial wage gap regressions.

age, time, and income) can be statistically accounted for by differences along these additional dimensions.

Quantifying gaps conditional on characteristics. To quantify gaps conditional on characteristics, we estimate weighted linear models for worker i of the form:

$$y_i = \alpha + \beta group_i + X_i' \delta + \epsilon_i. \quad (1)$$

Since optimal retirement savings can depend on age and income in a highly nonlinear way, all specifications follow a nonparametric reweighting approach that flexibly controls for age, calendar time, and income, similar to Bell et al. (2019). Specifically, we reweight observations in each cell defined by five-year age bin $\times$ income deciles (within each age bin) $\times$ year, such that each group has the same age-income-year distribution as the omitted category (e.g., White workers in our analysis by race).

Equation (1) is fairly standard in the literature on wage gaps (see, e.g., Cahuc et al., 2014, Ch. 8). When studying racial gaps, $group_i$ is a race indicator for the groups listed in Section 3.1 and the omitted category is White workers. Our analysis restricts attention to the five largest racial groups in the U.S.; therefore, the coefficients we report are those on the dummies in β for Hispanic, Black, Asian, and AIAN group membership. For the intergenerational analysis, $group_i$ is an indicator for the parental income quintile bins, and the top bin is the omitted category.

Motivated by the discussion in the previous subsection, we begin with equation (1) without controls ($X_i = \emptyset$), which shows differences after nonparametrically reweighting to balance age, income, and calendar time distributions across groups. We then progressively add linear controls in a “regression cascade,” where each additional variable reduces gaps only if it both predicts the outcome (after reweighting) and correlates with group membership.

Age and earnings account for substantial portions of the contribution gaps. The first set of bars (i.e., “Year + Age + Income”) in Figures 1(a) and 1(b) show contribution gaps after nonparametrically reweighting distributions by age, income, and calendar year to match that of White workers (left panel) and workers with parents in the top income quintile (right panel). Gaps are expressed as differences (y-axis) and percentages (above or below each bar) relative to the omitted category.

These observable worker characteristics account for roughly half of the raw racial and parental income gaps shown in Table 1. The one notable exception is the White-Asian gap, which remains virtually unchanged. Income is the primary factor driving the attenuation of

the gap.¹⁶ This result reflects two well-established patterns: (i) sizable income gaps exist across racial groups and between individuals with different levels of parental income, and (ii) high earners contribute more to DC accounts (Dynan et al., 2004).

The gaps by race and parental income persist across the income distribution and throughout workers' lifecycles, with gaps actually widening in percentage point terms at older ages and at higher earnings levels (Appendix Figure A.9). Even after accounting for age and earnings differences, total contribution gaps remain quantitatively large: Black workers contribute 1.2 p.p. less of earnings than White workers, Hispanic workers 1.0 p.p. less, and AIAN workers 0.7 p.p. less. Similarly, those from the bottom parental income quintile contribute 1.1 p.p. less than those from the top quintile.

Gaps remain even among co-workers with similar earnings, gender, occupation, tenure, education, and location. The final bars in Figures 1(a) and 1(b) report residual gaps after accounting for a number of additional differences in the individual characteristics of the workers in our sample. In particular, we include indicators for gender, four different educational attainment levels, occupation codes, and four different employment tenure levels.¹⁷ In addition, we absorb fixed effects for county of residence and the EIN of the employer. These EIN fixed effects are identified off of coworkers and absorb unobserved drivers of average contribution rates across firms, which include plan characteristics such as the match schedule. We discuss potential rationales for how each characteristic could impact savings rates in Appendix B.1.

Adding these individual characteristics further narrows most estimated gaps: by roughly half for Hispanic-White differences, by roughly 40% for both the AIAN-White and the top-bottom parental income quintile gaps, and by 13% for the Black-White gap. The Asian-White gap increases slightly by less than 4%.

Differences in contributions remain quantitatively large even after accounting for our broad range of characteristics (e.g., income, education, occupation, and employer). Black workers still have total contribution rates about 1 p.p. of earnings less than comparable White workers, Hispanic workers about 0.5 p.p., and those from the bottom parental income quintile contribute 0.75 p.p. less than those from top-quintile families.¹⁸ As a point of reference, the 1 p.p. Black-White gap is comparable in magnitude to the 0.9 p.p. effect of going from less than a high school degree to a college degree, as estimated in our regression controlling for individual characteristics.

¹⁶The distribution of age and year across groups is balanced, so reweighting by these will have little impact on the estimated gaps.

¹⁷In these graphs, we show results adding all regressors together; we detail the incremental impact on gaps from sequentially adding each individual characteristic in Appendix Figure A.3.

¹⁸We additionally report the estimates for all racial groups in Appendix Figure A.2(a).

Both intensive and extensive margins matter for contribution gaps, with the intensive margin playing a larger role when comparing otherwise similar workers. We next examine the roles of the extensive and intensive margins in driving contribution gaps. Panels (c) and (d) of Figure 1 show gaps in the probability of participating; panels (e) and (f) show gaps in contributions among those who participate. Both margins contribute to the contribution gaps, but after accounting for individual characteristics, the remaining gaps are driven primarily by the intensive margin. In contrast, participation gaps among workers with similar individual characteristics are modest, ranging from 1.7 to 2.7 p.p. across racial and parental income groups.

Taking stock and comparing our results with survey data. We document significant differences in DC retirement contributions, even for workers with similar ages, earnings, and other individual characteristics. In Appendix C, we reproduce our baseline analysis on contribution differences for the three largest racial groups using data from the Survey of Consumer Finances, the gold standard source of survey information on wealth in the U.S. Consistent with findings from previous research (see Darby Jr et al. (2018) for a review), we find estimates of gaps are imprecisely estimated. It is hard to make precise statements about the Black-White contribution difference once we account for differences in income. Confidence intervals are large and not only overlap with zero but also come close to and even overlap with some of our estimates using administrative data. This finding suggests that survey data that has been typically used to study this question might be under-powered to detect (even sizeable) differences in retirement contribution by race.¹⁹ This exercise highlights the value of large-scale administrative data, which not only enables precise measurement of racial contribution gaps but also allows us to examine gaps by parental income and study the interplay between race and family background in shaping savings behavior.

4.2 Household characteristics and the interplay between race and parental income

Our analysis thus far has analyzed gaps by race and parental income separately and focused solely on the mediating role of *individual* characteristics. Next, we investigate the role of household factors and draw out connections between race and parental income.

Accounting for differences in family structure and spousal income shrinks contribution gaps across most racial groups. Figure 2 presents racial contribution gaps for the subsample with available parental income data. As noted in Section 3, this subsample

¹⁹For example, in the 2013 wave of the SCF, only 167 Black and 94 Hispanic respondents reported having access to a DC plan through their employer; in contrast, our sample includes an average of roughly 22,000 Black and Hispanic workers every year whose employer sponsors a DC plan.

is younger than the overall sample we used in the previous section to study racial gaps in isolation. The first set of bars reports residual gaps by race after including all individual-level variables from Figure 1(a). While gaps remain similar in percentage terms across most groups, the absolute gaps (measured as percentage points of earnings) are smaller in this sample of younger workers who have lower average saving rates. The exception is the Asian-White gap, which is only half as large in percentage terms in this subsample.

The second set of bars in Figure 2 shows racial gaps once we include variables capturing the structure of the household (the individual’s marital status and the number of dependent children they list on form 1040) and spousal income (deciles if the spouse is working, plus an indicator for having a spouse with no W-2 income). Incorporating these household factors has very little effect on the Asian-White gap but reduces the estimated Black-White, Hispanic-White, and AIAN-White gaps by 10% to 20%. These incremental reductions are obtained *after* accounting for the mediating effects of individual-level characteristics (including income, education, firm, and occupation). This effect arises as, relative to their White co-workers, Hispanic, Black, and AIAN workers are much more likely to be single parents and less likely to have high-income spouses, both of which are correlated with contributing less.

Accounting for differences in parental income reduces the residual Hispanic-White gap by roughly 40% and both the Black-White and AIAN-White gaps by 20%.

We document substantial differences in contribution rates for workers with different levels of parental income, even conditional on a rich set of individual characteristics; that is, the children of the rich save more. Even after accounting for the host of individual- and household-level characteristics, children of parents in the top decile contribute 0.78 p.p. more than those with parents in the bottom decile (Appendix Figure A.4(a)). The strong association between parental income and own saving has relevance for differences in saving by race given that there are very large gaps in parental income by race. These gaps are illustrated in Appendix Figure A.4(b), which shows that Hispanic, Black, and AIAN workers are over 50% more likely to have parents in the bottom three parental income deciles, while White workers are significantly more likely to have higher-income parents. The distribution of parental income for Asian workers is roughly similar to that in the overall population.

Taken together, these two facts—that children with high-income parents save more, and that parents of White workers have much higher average incomes than those of Hispanic, Black, and AIAN workers—imply that parental income can play a mediating role in the racial savings gap. We quantify this effect in Figure 2. Including indicators for each decile of parental income at age 16 reduces the contribution gaps that remain (after accounting for

the part mediated by individual- and family-level characteristics) by roughly 15% to 40%: from 0.35 p.p. to 0.22 p.p. of earnings for the Hispanic-White gap, from 0.70 p.p. to 0.55 p.p. for the Black-White gap, and from 0.48 p.p. to 0.40 p.p. for the AIAN-White gap.²⁰

Before turning to differences in early withdrawals by race and parental income, we report several results on other potential drivers of contribution gaps.

4.3 Difference in the availability of a DC or DB plan

Differences in access to a retirement plan by race and parental income are largely accounted for by firm size. Our main empirical analysis focuses on workers whose employer offers a DC retirement plan. Yet, one third of private sector employees do not have access to such an employer-sponsored plan (a dimension we incorporate in our microsimulation model in Section 6). To the extent that access to such plans varies by race and parental income, differences in access could also affect differences in receipt of subsidies.

To measure the availability of a DC plan at a given employer, we classify a firm as offering a DC plan in a given year if at least 5% (or, as a robustness check, 25%) of its employees made a positive DC contribution. This approach has the advantage of relying solely on tax filings, without requiring a link to Form 5500 data.

An employer might choose not to sponsor a DC plan if they already offer a DB plan. To measure the availability of a DB plan, we link data from Form 5500 filings at the EIN level and classify a firm as offering a DB plan if (i) it has filed a Form 5500 for a DB plan and (ii) the plan is open to new participants.²¹

We find that, holding income constant, differences in access by parental income are modest, while those by race and ethnicity are larger. Black and AIAN workers are, respectively, 6.7 p.p. and 1.7 p.p. more likely than White workers to be employed by firms offering a DC or DB plan, while Hispanic and Asian workers are, respectively, 3.2 and 4.1 p.p. less likely to work for such firms (Figure 3(a)).²² These patterns suggest that, once we account for differential access, the Black-White, Asian-White, and AIAN-White gaps in receipt of savings subsidies we document among workers with access to a plan might be (modestly) smaller in the broader worker population, and the Hispanic-White gaps might be (modestly) bigger. In our micro-simulation model, introduced in Section 6, we account for these differences in

²⁰We also evaluate the importance of including a dummy for parents having contributed to a DC account, a proxy for familiarity with and exposure to these accounts. This does not affect the size of the residual contribution gap conditional on the mediating factors included in Figure 2.

²¹In Appendix Figure A.8, we show that racial contribution gaps are similar for the subsample of workers whose employer does not sponsor a DB plan.

²²Although DC plans are the most common type of retirement plans in our data, we use a broader measure of access that also includes DB plans, since some employers might offer a DB plan in place of a DC plan.

access in our measurement of lifetime employer matching and tax benefits.

A natural question is what drives the fact that Black workers, conditional on income, are more likely to have access to a retirement account. Employer size turns out to be relevant. Appendix Figure A.6(b) shows that larger employers are more likely to offer a DC account, while Figure A.6(c) shows that Black workers are most likely to work in the largest employers. Consistent with this, once we include indicator variables for employer size bins in the regression depicted in Figure 3(a), the Black-White and AIAN-White gaps in access flip sign: Black and AIAN workers are less likely than White workers to be at employers that sponsor a DC or DB plan after controlling for size. The Hispanic-White gap, already negative, widens further.

4.4 Other potential drivers of contribution differences

While the goal of this paper is not to fully decompose the differences in contributions by race and parental income, we do explore other plausible correlates of contribution gaps.

The role of vesting. Not all employer contributions are immediately vested. When vesting is not immediate, differences in earnings risk could generate different incentives to contribute by race. To investigate this channel, Appendix Figure A.8 re-runs our baseline regressions for racial gaps on a sample that includes only employees who are fully vested. The patterns in this sample (both the gaps after reweighting by age and income and how they change as regressors are added) are very similar to our baseline results. From this, we conclude that vesting has a limited effect on contribution differences by race.

The role of auto-enrollment. Employer-sponsored savings plans are increasingly moving from an opt-in to an automatic enrollment regime. Earlier evidence has shown that, in the short-run, the positive effect of automatic enrollment on participation rates is larger for Black and Hispanic workers relative to their White counterparts (Madrian and Shea, 2001). However, in the medium-run, the savings gains from auto-enrollment are largely attenuated (Choukhmane, 2025). Comparing employees hired before and after the adoption of auto-enrollment, we find that automatic enrollment does not reduce residual contribution gaps. This evidence is consistent with our finding that residual contribution gaps largely reflect differences in intensive-margin saving rather than differences in extensive-margin participation (the latter being more likely to be affected by auto-enrollment).

The role of life expectancy. Those who have lower life expectancy have less of an incentive to save for retirement. Differences in life expectancy by race and parental income (see, e.g., Schwandt et al., 2021, for evidence on the Black-White gap) could therefore account for some of the contribution differences we document. We make two remarks on how this

could relate to our findings. First, differences in life expectancy, conditional on income and other individual variables such as education and occupation, are likely to be smaller than unconditional differences. They are, therefore, less likely to fully account for the residual contribution gaps (after conditioning on individual characteristics) that we emphasize.²³

Second, and most importantly, differences in life expectancy do not directly impact our study’s primary focus, which is the distribution of employer matching subsidies and tax benefits. Unlike in DB plans, where life expectancy determines the present value of benefits, the value of a DC account is not directly linked to survival. Savings can often be accessed (albeit subject to penalties) before age 59.5, after which the full balance is liquid. These assets can also be bequeathed.

5 The Distribution of Early Withdrawals and their Impact on Net Retirement Contributions

The previous section documented racial and intergenerational gaps in contribution *inflows* into DC accounts. However, the distributional incidence of the federal tax expenditure, as well as wealth at retirement, depends on *net* contributions: inflows minus pre-retirement *outflows*. In this section, we examine the role of early withdrawals. Our analysis has two parts. First, we document differences in early withdrawals among past DC savers, defined as those who have contributed at least \$1,000 (in nominal terms) over the past four years and thus have balances from which to withdraw. We find that groups with lower contributions withdraw more frequently but, conditional on withdrawing, tend to take out smaller amounts. Second, we combine these extensive and intensive margin gaps with contribution differences to measure differences in *net* contribution rates in the broader sample, including both past savers and those who have not contributed in the past. We find that accounting for withdrawals amplifies the Black-White and Asian-White *net* contribution gaps, narrows the Hispanic-White gap, and amplifies the gap by parental income.

5.1 Early Withdrawals Differences among Past DC Savers

Institutional background, sample, and variable definitions. While employer-sponsored retirement plans are designed to finance retirement consumption, savers can access these resources early subject to potential tax penalties. Withdrawals before age 59.5 typically incur a 10% penalty unless they qualify for an exception.²⁴ Despite these restrictions, early with-

²³Furthermore, to the extent that saving decisions are determined by subjective survival probabilities, there is evidence that Black Americans have higher subjective life-expectancy than White Americans. See Palloni and Novak (2016) for related results and a discussion of the literature.

²⁴Before 2020, early withdrawals also triggered a six-month contribution suspension.

drawals are common: Goodman et al. (2021) find that outflows from DC plans and IRAs over 12 years exceed 20% of inflows. In our analysis, due to data limitations, we focus on all withdrawals (both penalized and unpenalized) but only count those of at least \$1,000 (in nominal terms) to exclude automatic cash-outs that might not reflect an actual decision from savers.²⁵

We make three decisions regarding the sample and variable definitions in our analysis of early withdrawals that are worth highlighting. First, we focus on individuals under age 55, since those 55 and older face fewer withdrawal restrictions.²⁶ Second, because our sample consists of individuals who are employed in the year they respond to the ACS (which we denote as year t), we examine withdrawals taken in year $t + 1$, regardless of employment status, and express amounts as a share of year t income. This allows us to capture withdrawals among those who become unemployed in $t + 1$. Third, since those who have not contributed cannot take early distributions, in this subsection we further restrict the sample to past savers, defined as those who have made cumulative retirement contributions of at least \$1,000 (in nominal terms) over the preceding 4 years. This allows us to examine differences in early withdrawal behavior across groups while holding constant the ability to withdraw, and we relax this assumption in the next subsection.

Early withdrawals are common, especially among Black savers and those with lower parental income. In our main sample, 13.6% of past DC savers withdraw at least \$1,000 (in nominal terms) annually, with withdrawal amounts averaging 28% of prior-year income (Table 1). The propensity to take an early withdrawal varies substantially by race. It is lowest, at 7.8%, for Asian past savers and is higher at 12.3% for White, 14.7% for Hispanic, 16.0% for AIAN past savers. The highest rate is for Black past savers: 23.4% of whom take an early withdrawal, nearly double the White rate.

In the sample of younger past savers, for whom we have data on parental income, early withdrawals of at least \$1,000 (in nominal terms) are nearly twice as common among those whose parents are in the bottom income quintile (16.0%) relative to those whose parents are in the top income quintile (9.1%). Among these younger cohorts which tend to have smaller balances, average withdrawal amounts are smaller, around 20% of past-year income.

Among savers with similar age and income, groups with lower contributions withdraw more frequently but, conditional on withdrawing, take smaller amounts.

Figure 4 shows regression cascades following the format of Figure 1. For each outcome,

²⁵Employers can automatically cash out balances under \$1,000 for terminated employees, so smaller withdrawals may not reflect active choice.

²⁶Individuals separating from employers at or after age 55 can withdraw from their current 401(k) without penalty.

we show two specifications: one reweighting based on year, age, and income, and another including our full set of individual mediators (education, gender, occupation, county of residence, tenure, and firm). The first row shows early withdrawal amounts as a proportion of prior-year income among all past savers.²⁷ The second and third rows decompose these overall effects into extensive and intensive margins.

At the extensive margin, Panels (c)-(d) show that groups with lower contribution rates withdraw more frequently. Black, Hispanic, and AIAN savers are, respectively, 86%, 17%, and 32% more likely to take an early withdrawal than White workers with similar age and income, while Asian savers are 30% less likely to do so. Similarly, savers with parents in the bottom income quintile are 57% more likely to take an early withdrawal than comparable savers with parents in the top quintile. These large racial and parental withdrawal gaps remain even among co-workers with similar earnings, gender, occupation, tenure, education, and location. However, intensive margin effects partially offset these extensive margin differences. Panels (e)-(f) show that groups with lower contributions (and presumably fewer resources to draw from) take smaller withdrawals as a share of prior-year income. Black, Hispanic, and lowest parental income quintile savers take withdrawals that are, respectively, 24%, 13%, and 10% smaller than White or top parental income quintile savers of similar age and income. These intensive margin differences attenuate (and in the case of Hispanic savers, fully offset) the effect of more frequent withdrawals.

Combining both margins, Panels (a)-(b) show that we cannot reject that Hispanic and AIAN savers withdraw similar amounts (as a share of income) as White workers with comparable characteristics. In contrast, large Black-White, Asian-White, and (to a lesser extent) top-bottom parental income withdrawal gaps persist among those with similar age and income, and even when including our full set of individual characteristics, reinforcing the contribution differences documented previously.

Early withdrawals are most common after large income declines, with wider racial and parental income gaps. Early withdrawals are known to be more common when income declines, which is often when liquidity demand is elevated (Coyne et al., 2022). To examine how this relationship varies by race and parental income, we sort workers into ventiles based on income growth between year t (the year in which the respondent fills out the ACS and meets a minimum earnings requirement) and year $t + 1$. Figure 5 shows three patterns. First, the prevalence of early withdrawals rises sharply following large income declines. Second, at nearly all income growth levels, Black, Hispanic, and AIAN savers withdraw more frequently than White savers, while Asian savers withdraw less; similarly,

²⁷We report the estimates for all racial groups in Appendix Figure A.2(b).

withdrawal probabilities decrease monotonically with parental income. Third, these racial and parental income gaps widen following severe income drops. Among workers in the bottom income growth ventile, 51% of Black and 42% of Hispanic savers withdraw at least \$1,000 (in nominal terms), compared to 35% of White and 21% of Asian savers. By parental income, the propensity to take a withdrawal in this bottom ventile ranges from 42% for those with parents in the lowest quintile to 25% for those with parents in the highest quintile.

What potential channels drive differences in early withdrawals among past savers?

Large differences in early withdrawal rates among past DC savers are consistent with Black savers and those with lower-income parents having stronger liquidity needs and less access to alternative liquidity sources. Coyne et al. (2022) highlight that early withdrawal propensity, despite tax penalties, can serve as a measure of differences in liquidity valuation. Supporting this interpretation, Ganong et al. (2020) find that Black and Hispanic households cut consumption more than White households following similar income shocks, suggesting racial differences in liquidity constraints. The same mechanism could also explain differences by parental income: those with wealthier parents may benefit from both familial financial support and less need to provide assistance to family members. Indeed, there is evidence that richer parents support their children financially and insure them against shocks (Andersen et al., 2020; Mukherjee, 2022; Fagereng et al., 2023), while poorer parents rely more on their adult children for financial assistance.²⁸ Finally, the institutional design of DC loan may also contribute to these patterns. Plans often require loans to be fully repaid at job separation, with outstanding balances treated as penalized early distributions.²⁹ Since income declines often coincide with job separations, workers must repay loans precisely when liquidity needs are highest. This could, in turn, reduce workers' incentive to contribute in the first place (Mitchell et al., 2007; Briere et al., 2022).

5.2 The Impact of Early Withdrawals on Net Contribution Gaps

Section 4 documented gaps in contributions to DC accounts, while the previous subsection documented gaps in early withdrawals. This subsection brings these two components together to examine *net* contribution rates, which account for both inflows into and outflows from retirement accounts. To do so, we extend our sample to all workers under the age of 55 (without restricting based on past savings behavior).

²⁸For example, a Pew report found that lower income children are more likely to provide financial support to their parents (Parker and Patten, 2013).

²⁹Among Vanguard plans in 2022, 82% offered loans and 61% required repayment at separation (Vanguard, 2023). Using Vanguard plan data, Beshears et al. (2024) find that within two years of loan issuance, default rates were 4% for participants who remained employed versus 71% for those who separated.

Early withdrawals amplify net contribution gaps for Black, Asian, and low-parental income workers and attenuate them for Hispanic workers. Figure 6 shows differences in net contribution rates, defined as year t contributions minus year $t + 1$ withdrawals divided by year t income. These net rates show how contribution gaps and early withdrawals together shape wealth accumulation gaps. For Black and Asian workers, accounting for withdrawals substantially amplifies gaps. Black workers' net contributions are 1.77 p.p. of income lower than White workers of similar age and income (a 63% difference), while Asian workers' net contributions are 2.24 p.p. higher (a 80% difference). The magnitude of these differences remains largely unchanged when comparing co-workers with the same education, tenure, occupation, and location. Parental income gaps also widen when accounting for withdrawals, though less dramatically. Those with parents in the bottom income quintile have net contributions 1.37 p.p. lower than those with top-quintile parents, a 31% difference. However, our comprehensive set of mediators can account for roughly half of this gap, suggesting that differences in these individual characteristics account for a substantial share of the parental income difference. The Hispanic-White gap shows the opposite pattern: withdrawals actually narrow the net contribution gap to 0.27 p.p. of income (a 10% difference), and it becomes statistically insignificant once our full set of individual characteristics is included.

Accounting for household and parental characteristics reduces net contribution gaps. Next, we explore how household structure and parental background interact with race to shape net contribution patterns. As in Figure 2, we focus here on our subsample of younger workers, for whom we have data on parents' characteristics. It is worth noting that net contribution patterns differ in this younger subsample.³⁰ Figure 6(c) shows that differences in household and parental characteristics account for approximately one quarter of the Black-White and AIAN-White net contribution gaps that remain after accounting for individual factors, and around half of the Hispanic-White gap. In contrast, the Asian-White net contribution gap *increases* once we account for parental income.

To assess how these differences in net contributions shape the mechanical impact of employer matching and tax subsidies, the next section develops a microsimulation model.

³⁰For instance, the Hispanic-White gap is insignificant in the full sample after accounting for individual characteristics, but remains significant in the younger sample. This difference across samples likely reflects the age-specific withdrawal patterns shown in Appendix Figure A.10(a): younger Hispanic savers withdraw more frequently than their White counterparts, but this difference disappears for older savers.

6 The Lifetime Mechanical Impact of Current and Alternative Retirement Savings Policies

We have documented substantial heterogeneity in annual contributions to and withdrawals from DC accounts. In this section, we combine our data with a micro-simulation model to examine the mechanical impact of retirement saving subsidies on wealth at retirement. The model, which is described in full in Appendix D, simulates data on wealth in retirement by bringing together i) our data on access to flows in and out of DC funds, ii) a specification of the federal income tax code (from NBER TAXSIM), iii) Social Security rules, and iv) assumptions about portfolio composition, asset returns, and the draw-down of wealth in retirement.

We use this model for three purposes. First, we examine the lifetime distribution of both employer matching contributions and the federal tax expenditure. This analysis depends on the whole life-cycle trajectory of income, contributions, and withdrawals. Second, we use the model to place our distributional analysis in the context of the broader U.S. retirement system by comparing differences in the allocation of subsidies between individuals with similar Social Security benefits. Lastly, we study the mechanical impact of revenue-neutral reforms that would redistribute employer-match dollars within each firm and federal tax expenditures across the population to be allocated 1) proportionally to earnings but 2) independently of workers' own contribution choices.

6.1 Micro-simulation model

6.1.1 The components of wealth at retirement

The full model is outlined in Appendix D. Here, we summarize key inputs and outputs. The model inputs are data on earnings over the life cycle, employee and employer matching contributions to employer-sponsored DC accounts, and withdrawals from those accounts over the working life. We do not observe *full* life cycle paths (we have at most 13 years of data for any one individual), so we construct simulated full-life cycle paths using these data and a hot-deck imputation procedure. This procedure, described in Appendix D.2, initializes the population with individuals aged 25-26 in our data, and then builds a panel by matching these individuals to those aged 27-28 with similar characteristics.³¹ Because the simulation procedure requires us to match on multiple characteristics within race it requires a large sample size, and we thus only simulate lifecycles for the four largest race groups (White,

³¹An implication of this is that the earnings of groups inherit the distributional characteristics of the initial 25-26 year-old population. For example, the fact that, among younger cohorts, Asian workers earn more than each of the other racial groups, means that they have higher earnings over the whole lifecycle. See Figure A.19 for a diagram of how lifetime earnings are imputed from shorter windows in the model.

Hispanic, Black, and Asian). Since parent income is only observable for younger cohorts, it is not feasible to condition on parental income for all cohorts. Therefore, our simulations only allow parental income to influence future outcomes via its correlation with the initial observable variables on which we condition.³²

The key model outputs are:

- **DC wealth** (A_i^{DC}): is defined as the discounted value of after-tax withdrawals from the simulated DC account balance. We assume that savers employ a draw-down rule that keeps withdrawals constant in retirement. We divide DC wealth into three components:
 - **Lifetime tax expenditure:** A_i^T is the part of DC wealth arising from its favorable tax treatment. We define this as the difference between A_i^{DC} and the discounted value of withdrawals that worker i would have received if she had instead saved the same proportion of salary each year in a taxable account. A_i^T therefore represents the tax advantage obtained from deferring the taxation of contributions, having tax-free growth of assets, and being exempt from capital gains, net of tax penalties on early withdrawals from DC accounts.³³
 - **Employee contributions:** A_i^{EE} is the portion of DC wealth, exclusive of tax benefits (i.e., $A_i^{DC} - A_i^T$), accruing from *employee* contributions.
 - **Employer contributions:** A_i^{ER} is the portion of DC wealth, exclusive of tax benefits (i.e., $A_i^{DC} - A_i^T$), accruing from *employer* contributions.
- **Social Security Wealth:** Social Security wealth (SS_i) is measured as the discounted stream of benefits an individual will receive through retirement. We calculate this benefit by applying the Social Security formula to each simulated history of earnings. We assume everyone claims on turning 66.
- **Broad retirement wealth:** Broad retirement wealth (A_i^{BR}) is the sum of DC wealth (A_i^{DC}) and the discounted value of Social Security payments (SS_i). This measure of broad retirement wealth omits DB pension wealth and non-tax-favored savings.³⁴

Appendix Figures A.12 and A.13 illustrate model outputs (earnings, DC wealth, and

³²Specifically, we initiate our microsimulations using the empirical distribution of parental income and child characteristics at age 25, including income and DC plan participation. To the extent that future income and plan participation are also positively correlated, but not perfectly so, with parental income, our estimates of the relationship between parental income and lifetime outcomes of the children will be attenuated.

³³Appendix D.8.2 compares our estimates of aggregate tax expenditure to DC savings with official Treasury Department figures, while Appendix D.9 discusses features that are not included in the model (including inflation, state income taxes, and heterogeneity in asset allocations).

³⁴We expect both omitted components to have a limited impact on our main findings. DB coverage has declined sharply and is concentrated in older cohorts and a few sectors, while non-tax-favored retirement savings are concentrated at the top of the income distribution and are small relative to DC and Social Security wealth for the bulk of households we study.

Social Security) by race and parental income. In these figures, and in our analysis in this section, we divide the population into six groups based on their lifetime earnings: the bottom four income quintiles and the top two income deciles. We split the top income quintile in two as those in the top income decile are much more likely to be constrained by the annual federal contribution limit. Our microsimulations generate significant dispersion in DC wealth at retirement within each of these lifetime income groups. By contrast, there is very little variation in average Social Security wealth SS_i by race and parental income within each bin because Social Security payouts are calculated based on lifetime earnings.³⁵

6.2 The distribution of current retirement savings subsidies

Employer and tax subsidies account for over 40% of DC balances at retirement.

We decompose wealth at retirement into three components: those arising from employee saving elections, employer matches, and favorable tax treatment. Figure 7(a) reports the share of retirement wealth coming from each component across the lifetime earnings distribution.

Employee contributions account for the largest share of wealth at retirement—approximately three-fifths of DC wealth among those in the bottom quintile of the lifetime earnings distribution and one-half in the top decile. The remainder of DC wealth comes from a combination of employer matches and the tax expenditure. Matches account for approximately one-quarter of DC wealth across all groups, while the federal tax expenditure contributes between 14% at the bottom of the lifetime earnings distribution and 22% of DC wealth at the top.

While the *shares* of wealth arising from each source differ only modestly by lifetime earnings, the *level* of each component scales, approximately, with the saving done by those in each group. Figure 7(b) gives the level of each component expressed as a share of average annual lifetime earnings. Given that savings rates increase in earnings, the subsidies, as a share of earnings, also increase with earnings. In the bottom quintile of lifetime earnings, cumulative tax and employer subsidies are worth less than 70% of annual lifetime earnings. By comparison, these subsidies are worth more than 250% of annual lifetime earnings for individuals in the top decile of earnings. Differences are larger still in dollar terms (as shown in Table 2).

Next, we extend the distributional analysis beyond lifetime earnings, showing how racial and parental income differences in contributions and withdrawals conditional on earnings (as emphasized in Sections 4 and 5) affect the receipt of saving subsidies over the lifecycle.

³⁵In reality, heterogeneity in Social Security wealth among those with the same income could arise if individuals differ in claiming ages and if the benefit adjustments are actuarially unfair (Munnell and Chen (2019)) or if life expectancies differ.

Savings subsidies amplify racial and intergenerational retirement wealth inequality, even among those with similar earnings and Social Security. Table 2 shows differences in the receipt of lifetime match and tax subsidies by race (panel (a)) and parental income (panel (b)) in dollar terms. Earnings bins are defined in the population, so observations within each column have similar incomes. Appendix Figure A.14 expresses the value of subsidies as a percentage of average annual lifetime earnings

Differences between groups mirror the differences in saving rates documented in Section 4. In each of the bottom five income groups, Asian individuals receive the highest combined tax and matching subsidies, followed by White, then Hispanic, and finally Black individuals, all with similar lifetime earnings. There are differences in all income groups, but they are largest as a proportion of lifetime earnings in the middle of the distribution. Asian individuals in the middle quintile receive combined tax and matching subsidies worth \$91,100 (218% of average annual earnings), compared to \$63,700 (161%) for White, \$58,000 (149%) for Hispanic, and \$44,400 (115%) for Black individuals.

Table 2(b) shows the value of DC subsidies by parental income (with the values as a percentage of average annual earnings given in Appendix Figure A.14). Each successive set of columns again represents a group based on an individual’s lifetime earnings. Because the children of the rich save more, even conditional on their own earnings (see Figure 1(b)), subsidies for savers advantage them. In the middle of the lifetime earnings distribution, children of parents in the bottom quintile receive combined subsidies worth \$58,900 (146% of average annual earnings), while children from the top quintile receive \$65,800 (172%).

6.3 Simulating alternative retirement savings policies

In this section, we use the micro-simulation model to evaluate the mechanical impact of budget-neutral reforms that would break the link between private saving and the amount of employer matching benefits and tax subsidies that individuals receive. The counterfactual we consider is an environment in which a) all employees in each firm receive a contribution that is the same percentage of their earnings and b) all workers in the economy get a share of the tax expenditure that is the same proportion of their lifetime earnings. Figure 8 illustrates the changes in DC wealth under this counterfactual policy and an alternative we discuss later in the subsection.

6.3.1 Description of the counterfactual policy exercise

The counterfactual exercise we discuss below changes the allocation of both employer and tax subsidies for retirement saving. A full description of the exercise is given in Appendixes D.10-D.12; here we provide a summary. In both our reforms of employer matches and the

tax subsidy, it is important to note that, while we break the link between savings choice and subsidies, we impose that subsidies remain proportional to earnings (which limits the extent of redistribution along the income distribution).

Employer contributions. We first redistribute the employer matching contributions within each firm. That is, we calculate the aggregate employer matching contribution made by each employer, and we divide these contributions by aggregate compensation. This gives a counterfactual proportion of salary that, if given to all employees regardless of how much they elected to contribute, would cost the same as the status quo.³⁶ We assume that wages are unaffected by this reform. This implies that the incidence is borne fully by workers, as we discuss further below.

Tax expenditures. Next, we calculate the aggregate tax expenditure on DC retirement savings. We redistribute this tax expenditure such that every individual receives a direct government contribution to their retirement account calculated as a proportion of lifetime earnings. This proportion is uniform across individuals and chosen in order to keep the aggregate tax expenditure constant.

Incidence assumptions. We assume that firms do not adjust compensation in response to the tax reform or to changes in the distribution of employer contributions. In particular, workers who lose match or tax subsidies are not compensated through higher wages, and those who gain subsidies do not see offsetting wage reductions. Under this assumption, the reforms are fully incident on employees.

Alternative counterfactual. In panels (b) and (d) of Figure 8, we consider an alternative scenario that relaxes this incidence assumption for the counterfactual distribution of employer contributions (without changing the assumptions of the tax counterfactual). Specifically, for firms with more than 200 workers we hold the total compensation package constant within each firm and earnings decile, but equalize employer contribution rates (as a proportion of earnings) within each group.³⁷ This approach therefore only redistributes matching contributions among co-workers with similar incomes and is therefore even less likely to induce changes in wages.

Behavioral responses. In our baseline counterfactual exercise, we assume that indi-

³⁶While employer-matching contribution formulas are chosen by employers, the government can encourage employers to adopt specific contribution formulas. Arnoud et al. (2021) estimate that a majority of employees are covered by plans with a safe-harbor matching formula. Our counterfactual can be thought of as a change in safe harbor rules that shifts all employers away from offering matching contributions and toward non-elective contributions.

³⁷For firms with fewer than 200 employees, accurately estimating an income gradient is difficult given the small sample sizes. For these firms, we thus retain our previous approach of redistributing and equalizing employer contribution rates (as a proportion of earnings) at the firm level. In Form 5500 data, plans with more than 200 eligible participants account for more than 93% of aggregate employer dollars in 2017.

vidual saving rates are unchanged across the different counterfactual exercises. Removing employer matching and tax incentives could induce individuals to save less for retirement, which may affect our distributional conclusions as we discuss in Appendix D.13. However, there is no consensus in the literature on how much private saving responds to employer matching and tax incentives.³⁸ As a result, our baseline assumption of no behavioral response in private saving will likely be a reasonable approximation. In an extension (discussed in Appendix D.13), we recalculate the results assuming that each dollar of employer matching or tax subsidies generates either 10 cents (which corresponds to the upper bound of the 95% confidence interval in Chetty et al. (2014)) or 30 cents of additional employee savings. Even under these assumptions, as shown in Appendix Figure A.16, the relative changes in DC wealth gaps by race and parental income remain quantitatively similar to the baseline.

6.3.2 The reform would reduce inequality in retirement wealth accumulation by own earnings, race, and parental income

The left panels of Figure 8 summarize the effect of our baseline counterfactual exercise on DC wealth by race (panel (a)) and parental income (panel (c)). We express the change as a proportion of broad retirement wealth (the sum of DC wealth and Social Security wealth).³⁹ Table 3 complements this analysis by expressing the changes outcomes in dollar values.

Results by own lifetime earnings. The take-up of matching and tax incentives is increasing in earnings; therefore, allocating tax expenditures and employer matching contributions in proportion to lifetime earnings (rather than in proportion to saving) would redistribute resources toward lower-income workers. We find that our revenue-neutral reform would significantly raise tax and employer transfers to the bottom 80% of the lifetime earnings distribution, resulting in around \$20,000 in additional DC wealth at retirement on average (Table 3). As shown in Figure 8, the relative gains are largest for the bottom 20% of lifetime earners (with gains of around 7.5% of broad retirement wealth) and remain sizeable in the middle of the lifetime earnings distribution (with gains of around 4.5% of broad retirement wealth). These gains at the bottom and middle of the earnings distribution come at the expense of lower resources allocated to top earners (especially in the top decile of lifetime earnings). The loss in DC retirement wealth for the top decile is worth about 3%

³⁸Choi (2015) reviews the literature on matching and finds that it is associated with a small positive effect on participation and an ambiguous effect on average contribution rates. Regarding tax incentives, a review by Friedman (2015) notes that “tax subsidies appear to primarily affect the allocation of savings across accounts, rather than the total amount of savings”. Related empirical results, discussed further in Appendix D.13, include Engen et al. (1996), Poterba et al. (1996), Chetty et al. (2014), Duflo et al. (2006), Engelhardt and Kumar (2007), Ramnath (2013), and Choukhmane and Palmer (2025).

³⁹Appendix Figure A.17 expresses the changes as a proportion of the average annual lifetime earnings of each group, and it has the same pattern as in Figure 8. Meanwhile, Appendix Figure A.18 gives results for our baseline where we separately, rather than simultaneously, redistribute tax and matching subsidies.

of broad retirement wealth. Because the losses are concentrated among those with higher lifetime income and the gains are concentrated among those with lower lifetime income, the relative gains (in percentage terms) from this counterfactual policy are much larger than the relative losses.⁴⁰

The right panels of Figure 8 illustrate similar calculations for the alternative counterfactual which redistributes tax benefits across the population but redistributes employer matches within each income decile of each firm. Since lower-paid workers contribute less on average, changing our incidence assumption to eliminate redistribution across income deciles within the firm modestly attenuates the impacts of the reform. Compared to the baseline, the increase in broad retirement wealth in this alternative counterfactual is between 1.7 p.p. and 2.3 p.p. smaller for low-income earners and between 0.9 p.p. and 1.8 p.p. smaller for middle-income earners. This represents around a 20% smaller increase in broad retirement wealth at the middle and bottom of the earnings distribution for most racial and parental income groups.

Results by race. Among those with the same lifetime earnings, the counterfactual policy would redistribute more to those with lower saving rates (Black and Hispanic individuals) and less to those with higher saving rates (Asian and White individuals). The differences by race are largest in the middle of the lifetime earnings distribution. The counterfactual reform increases broad retirement wealth by about 8% for Black individuals, just under 5% for Hispanic, about 4% for White and around 1% for Asian individuals. At the top of the lifetime earnings distributions, all groups see a decrease in broad wealth, but Asian individuals (who have very high saving rates) lose substantially more than all other groups.

We can quantify how much such a policy could change racial gaps in DC wealth. Table 3 shows that the mechanical impact of the reform would increase DC wealth among White, Hispanic, Black, and Asian individuals in the middle of the lifetime earnings distribution by about \$22,500, \$25,200, \$37,000, and \$8,200 respectively. This would reduce the gap between the DC wealth of Black and White individuals in this lifetime earnings group from 38% to 25%, and that between Hispanic and White individuals from 14% to 11%.

Results by parents' income. Figure 8(c) shows the effect of our reform by both own and parental income. Across all lifetime earnings groups, those with lower-income parents benefit more from the reform than those with richer parents, and these differences by parental income (conditional on own earnings) are larger for those with above-median lifetime earnings.

⁴⁰While these reforms are designed to be revenue neutral for the government and aggregate-compensation neutral for the firms, they lead to a net increase in wealth on retirement as matching resources are transferred from older workers to younger workers, who have more time to retirement to benefit from asset returns.

For instance, Table 3 shows that those in the bottom-income group with the lowest and highest-income parents have similar gains from the reform (around \$17,650 in additional DC wealth). In contrast, in the fourth quintile of earnings, those with the lowest-income parents gain an additional \$18,100 while those with the richest parents gain an additional \$7,000. Losses are concentrated among those with both high earnings and high parental income; in the top earnings decile, those with the poorest parents experience a \$33,600 drop in DC wealth, whereas those with parents in the top income quintile experience a \$80,100 drop in DC wealth. While these wealth losses are large in absolute terms, they represent less than a 5% reduction in broad retirement wealth for top earners with the richest parents. Given that those with lower-income parents benefit more from the reform, we estimate that the reform could close the gap in DC wealth accumulation by parental income by about one third.⁴¹

7 Conclusion

Since the introduction of the permanent income tax system in 1913, the U.S. has promoted retirement saving with tax subsidies and employer contributions. A long-standing concern is that these subsidies are regressive and largely favor higher-income individuals. This worry has sparked a long tradition of economics research studying the distributional effects and optimal design of the retirement system (Diamond, 1977; Kotlikoff et al., 1982; Geanakoplos et al., 1999; Moser and Olea de Souza e Silva, 2019). This concern is also reflected in the regulatory framework; since 1942, U.S. pension plans have been required to pass an annual nondiscrimination test to ensure that the benefits of the plan do not disproportionately accrue to highly compensated employees.⁴² In addition to income-based nondiscrimination tests, the Social Security formula is progressive and offers higher replacement rates for individuals with lower lifetime earnings. One view is that these more progressive aspects of the U.S. retirement system balance the income-regressive nature of retirement saving subsidies.

In this paper, we challenge this view by studying the distributional properties of retirement saving subsidies across individuals who have similar incomes but differ along other dimensions (with a focus on differences by race, ethnicity, and parental income). We find that the current system channels more tax and employer resources toward workers who are White or Asian and have richer parents than toward their similar-income coworkers who are

⁴¹Appendix D.13 reports analogous results under alternative assumptions about behavioral responses. We find that, even when assuming elasticities of private savings to financial incentives that exceed most empirical estimates, our baseline counterfactual policy continues to raise retirement wealth accumulation in the bottom half of the lifetime earnings distribution.

⁴²To pass the nondiscrimination test, the employer must show that differences between the average employee and employer contribution rates for highly compensated and non-highly compensated employees are sufficiently small. Employers can avoid these annual tests by adopting a set of plan features that qualify a plan as a safe harbor plan.

Hispanic, Black, or AIAN and have lower-income parents. Analogous distributional comparisons could be made by other characteristics, which are important for saving. While we do not emphasize them in our paper, our regression estimates indicate that conditional on income, those with more education save more than those with less education, while single parents save less than couples with children. The consequent effects on wealth accumulation are large and are not directly offset by other aspects of the retirement system. The Social Security formula does not vary by race, education, or parental background, and employer nondiscrimination tests only consider one’s compensation. Our results thus suggest that future research on the distributional impact and optimal design of retirement systems would benefit from looking beyond income.

References

- Altonji, J. G. and R. M. Blank (1999). Race and gender in the labor market. *Handbook of Labor Economics* 3, 3143–3259.
- Altonji, J. G. and U. Doraszelski (2005). The role of permanent income and demographics in black/white differences in wealth. *Journal of Human Resources* 40(1), 1–30.
- Andersen, A. L., N. Johannesen, and A. Sheridan (2020). Bailing out the kids: new evidence on informal insurance from one billion bank transfers.
- Arapakis, K., G. Wettstein, and Y. Yin (2023). What is the insurance value of social security by race and socioeconomic status? Technical report, Center for Retirement Research.
- Ariel/AON Hewitt (2009). 401(k) Plans in Living Color.
- Arnoud, A., T. Choukhmane, J. Colmenares, C. O’Dea, and A. Parvathaneni (2021). The Evolution of U.S. Firms’ Retirement Plan Offerings. Evidence from a New Panel Data Set. Technical report, NBER.
- Avenancio-León, C. F. and T. Howard (2022). The assessment gap: Racial inequalities in property taxation. *The Quarterly Journal of Economics* 137(3), 1383–1434.
- Bakker, T. J., S. DeLuca, E. A. English, J. S. Fogel, N. Hendren, and D. Herbst (2025). Credit access in the united states. Technical report, National Bureau of Economic Research.
- Barber, B. M. and T. Odean (2001). Boys will be boys: Gender, overconfidence, and common stock investment. *The Quarterly Journal of Economics* 116(1), 261–292.
- Barsky, R., J. Bound, K. K. Charles, and J. P. Lupton (2002). Accounting for the black–white wealth gap: a nonparametric approach. *Journal of the American Statistical Association* 97(459), 663–673.
- Bayer, P. and K. K. Charles (2018). Divergent paths: A new perspective on earnings differences between black and white men since 1940. *The Quarterly Journal of Economics* 133(3), 1459–1501.
- Bee, A. and J. Mitchell (2017). Do Older Americans Have More Income Than We Think? *U.S. Census Bureau SEHSD Working Paper #2017-39*.
- Bell, A., R. Chetty, X. Jaravel, N. Petkova, and J. Van Reenen (2019). Who becomes an

- inventor in america? the importance of exposure to innovation. *The Quarterly Journal of Economics* 134(2), 647–713.
- Beshears, J., J. J. Choi, J. Dickson, A. Goodman, F. Greig, and D. Laibson (2024). Does 401 (k) loan repayment crowd out retirement saving? implications for plan design. *Implications for Plan Design (October 03, 2024)*. Wharton Pension Research Council Working Paper (2024-22).
- Bhutta, N. and A. Hizmo (2021). Do minorities pay more for mortgages? *The Review of Financial Studies* 34(2), 763–789.
- Blau, F. D. and J. W. Graham (1990). Black-white differences in wealth and asset composition. *The Quarterly Journal of Economics* 105(2), 321–339.
- Bond, B., J. D. Brown, A. Luque, A. O’Hara, et al. (2014). The nature of the bias when studying only linkable person records: Evidence from the American Community Survey. *Center for Administrative Records Research and Applications Working Paper* 8.
- Briere, M., J. Poterba, and A. Szafarz (2022, May). Precautionary Liquidity and Retirement Saving. *AEA Papers and Proceedings* 112, 147–150.
- Brown, D. A. (2021). *The Whiteness of Wealth: How the Tax System Impoverishes Black Americans—And How We Can Fix It*. Crown Publishing Group (NY).
- Brown, J. (2002). Differential mortality and the value of individual account retirement annuities. In *The distributional aspects of social security and social security reform*, pp. 401–446. University of Chicago Press.
- Brown, J. R., J. M. Poterba, and D. P. Richardson (2022). Trends in Retirement and Retirement Income Choices by TIAA Participants: 2000–2018. Technical report, National Bureau of Economic Research.
- Burman, L. E., W. G. Gale, M. Hall, and P. Orszag (2004). Distributional Effects of Defined Contribution Plans and Individual Retirement Accounts.
- Cahuc, P., S. Carcillo, and A. Zylberberg (2014). *Labor economics*. MIT press.
- Carneiro, P., J. J. Heckman, and D. V. Masterov (2005). Labor market discrimination and racial differences in premarket factors. *The Journal of Law and Economics* 48(1), 1–39.
- Catherine, S., E. J. Lu, and J. D. Paron (2025). A dynamic model of the racial wealth gap. *Available at SSRN* 5340625.
- Catherine, S. and N. Sarin (2023). Social Security and the Racial Wealth Gap. Technical report.
- Charles, K. K. and E. Hurst (2002). The transition to home ownership and the black-white wealth gap. *Review of Economics and Statistics* 84(2), 281–297.
- Charles, K. K. and E. Hurst (2003). The correlation of wealth across generations. *Journal of Political Economy* 111(6), 1155–1182.
- Chen, A. and A. H. Munnell (2020, November). How Much Taxes Will Retirees Owe on Their Retirement Income? Technical Report Working Paper 2020-16, Center for Retirement Research at Boston College, Chestnut Hill, MA.
- Chetty, R., J. N. Friedman, S. Leth-Petersen, T. H. Nielsen, and T. Olsen (2014). Active vs. passive decisions and crowd-out in retirement savings accounts: Evidence from Denmark. *The Quarterly Journal of Economics* 129(3), 1141–1219.
- Chetty, R., N. Hendren, M. R. Jones, and S. R. Porter (2020). Race and economic opportunity in the United States: An intergenerational perspective. *The Quarterly Journal of Economics* 135(2), 711–783.

- Chiteji, N. S. and D. Hamilton (2002). Family connections and the black-white wealth gap among middle-class families. *The Review of Black Political Economy* 30(1), 9–28.
- Choi, J. J. (2015). Contributions to defined contribution pension plans. *Annual Review of Financial Economics* 7, 161–178.
- Choukhmane, T. (2025). Default options and retirement saving dynamics. *American Economic Review* 115(11), 3749–3787.
- Choukhmane, T., L. Goodman, and C. O’Dea (2025). Efficiency in household decision-making: Evidence from the retirement savings of us couples. *American Economic Review* 115(5), 1485–1519.
- Choukhmane, T. and C. Palmer (2025). The effect of increasing retirement saving on consumption, balance sheets, and welfare. Technical report.
- Congressional Budget Office (2021). The Distribution of Major Tax Expenditures in 2019. Technical report, Congressional Budget Office.
- Coyne, D., I. Fadlon, and T. Porzio (2022). Measuring Valuation of Liquidity with Penalized Withdrawals. Technical report, National Bureau of Economic Research.
- Darity, W. and S. L. Myers (1983). Changes in Black Family Structure: Implications for Welfare Dependency. *American Economic Review* 73(2), 59–64. Publisher: American Economic Association.
- Darity, W. and M. J. Nicholson (2005). Racial wealth inequality and the Black family. *African American Family Life: Ecological and Cultural Diversity* 78, 85.
- Darity, W. A. and S. L. Myers (1987). Do Transfer Payments Keep the Poor in Poverty? *American Economic Review* 77(2), 216–222. Publisher: American Economic Association.
- Darity Jr, W., D. Hamilton, M. Paul, A. Aja, A. Price, A. Moore, and C. Chiopris (2018). What we get wrong about closing the racial wealth gap. *Samuel DuBois Cook Center on Social Equity and Insight Center for Community Economic Development* 1(1), 1–67.
- Department of Labor (2024). Private Pension Plan Bulletin. Technical report, Department of Labor.
- Derenoncourt, E., C. H. Kim, M. Kuhn, and M. Schularick (2022, June). Wealth of Two Nations: The U.S. Racial Wealth Gap, 1860-2020. Working Paper 30101, National Bureau of Economic Research.
- Derenoncourt, E., C. H. Kim, M. Kuhn, and M. Schularick (2023). *Unemployment risk, portfolio choice, and the racial wealth gap*. ECONtribute.
- Derenoncourt, E. and C. Montialoux (2021). Minimum wages and racial inequality. *The Quarterly Journal of Economics* 136(1), 169–228.
- Devlin-Foltz, S., A. Henriques, and J. Sabelhaus (2016). Is the US retirement system contributing to rising wealth inequality? *RSF: The Russell Sage Foundation Journal of the Social Sciences* 2(6), 59–85.
- Diamond, P. (1977). A framework for social security analysis. *Journal of Public Economics* 8(3), 275–298.
- Dong, Y. and C.-Y. J. Peng (2013). Principled missing data methods for researchers. *SpringerPlus* 2(1), 222.
- Dufo, E., W. Gale, J. Liebman, P. Orszag, and E. Saez (2006). Saving incentives for low-and middle-income families: Evidence from a field experiment with H&R Block. *The Quarterly Journal of Economics* 121(4), 1311–1346.

- Dushi, I. and M. Honig (2015). How much do respondents in the health and retirement study know about their contributions to tax-deferred contribution plans? a cross-cohort comparison. *Journal of Pension Economics & Finance* 14(3), 203–239.
- Dushi, I. and H. M. Iams (2010). The impact of response error on participation rates and contributions to defined contribution pension plans. *Soc. Sec. Bull.* 70, 45.
- Dynan, K. E., J. Skinner, and S. P. Zeldes (2004). Do the rich save more? *Journal of Political Economy* 112(2), 397–444.
- Engelhardt, G. V. and A. Kumar (2007). Employer matching and 401 (k) saving: Evidence from the health and retirement study. *Journal of Public Economics* 91(10), 1920–1943.
- Engen, E. M., W. G. Gale, and J. K. Scholz (1996). The illusory effects of saving incentives on saving. *Journal of Economic Perspectives* 10(4), 113–138.
- Fagereng, A., L. Guiso, D. Malacrino, and L. Pistaferri (2020). Heterogeneity and persistence in returns to wealth. *Econometrica* 88(1), 115–170.
- Fagereng, A., L. Guiso, L. Pistaferri, and M. A. Ring (2023). Insuring Labor Income Shocks: The Role of the Dynasty.
- Farber, H. S. (1994). The Analysis of Interfirm Worker Mobility. *Journal of Labor Economics* 12(4), 554–593.
- Feenberg, D. and E. Coutts (1993). An Introduction to the TAXSIM Model. *Journal of Policy Analysis and Management* 12(1), 189–194.
- Ferrie, J., C. Massey, and J. Rothbaum (2021). Do Grandparents Matter? Multigenerational Mobility in the United States, 1940–2015. *Journal of Labor Economics* 39(3).
- Fidelity (2023). A Lifetime Investment Approach in a Single Strategy.
- Fidelity (2023). Fidelity Freedom Fund, Selected Years.
- Francis, D. V. and C. E. Weller (2021). Retirement inequality by race and ethnicity. *Public Policy & Aging Report* 31(3), 83–88.
- Friedman, J. (2015). Tax Policy and Retirement Savings. *Paper prepared for the Economics of Tax Policy Conference at the Brookings Institution.*
- Ganong, P., D. Jones, P. J. Noel, F. E. Greig, D. Farrell, and C. Wheat (2020). Wealth, race, and consumption smoothing of typical income shocks. Technical report, National Bureau of Economic Research.
- Geanakoplos, J., O. S. Mitchell, and S. Zeldes (1999). Social Security Money’s Worth. In O. S. Mitchell, R. Myers, and H. Young (Eds.), *Prospects for Social Security Reform*, pp. 79–151. Philadelphia, PA: Univ. of Pennsylvania Press.
- Gindelsky, M. and R. Martin (2024). Bls working papers.
- Goldin, C. (2021). *Career and family: Women’s century-long journey toward equity*. Princeton University Press.
- Goodman, L., J. Mortenson, K. Mackie, and H. R. Schramm (2021). Leakage From Retirement Savings Accounts In The United States. *National Tax Journal* 74(3), 689–719.
- Gourinchas, P.-O. and J. A. Parker (2002). Consumption over the life cycle. *Econometrica* 70(1), 47–89.
- Güvenen, F., S. Ozkan, and J. Song (2014). The nature of countercyclical income risk. *Journal of Political Economy* 122(3), 621–660.
- Hamilton, D. and W. A. Darity (2017). The Political Economy of Education, Financial Literacy, and the Racial Wealth Gap. *Federal Reserve Bank of St. Louis Review* 99(1), 59–76.


- Hou, W. and G. T. Sanzenbacher (2021). Measuring racial/ethnic retirement wealth inequality. *The Journal of Retirement* 8(3), 12–28.
- Internal Revenue Service (2023, September). SOI Tax Stats - Individual Information Return Form W-2 Statistics.
- Joint Committee on Taxation (2019, July). Background on Revenue Sources for the Social Security Trust Funds. Technical Report JCX-41-19.
- Jordà, O., K. Knoll, D. Kuvshinov, M. Schularick, and A. M. Taylor (2019, 04). The Rate of Return on Everything, 1870–2015. *The Quarterly Journal of Economics* 134(3), 1225–1298.
- Kermani, A. and F. Wong (2021). Racial Disparities in Housing Returns. Technical report, National Bureau of Economic Research.
- Kotlikoff, L. J., A. Spivak, and L. H. Summers (1982). The Adequacy of Savings. *American Economic Review* 72(5), 1056–1069.
- Krogstad, J. M. and C. Im (2025, May). Key facts about asian americans, a diverse and growing population.
- Kuan, K. Y., M. R. Cullen, and S. Modrek (2015). Racial disparities in savings behavior for a continuously employed cohort. Technical report, National Bureau of Economic Research.
- Kuhn, M., M. Schularick, and U. I. Steins (2020). Income and wealth inequality in America, 1949–2016. *Journal of Political Economy* 128(9), 3469–3519.
- Kuka, E. and B. A. Stuart (2021). Racial Inequality in Unemployment Insurance Receipt and Take-Up. Technical report, National Bureau of Economic Research.
- Lang, K. and M. Manove (2011). Education and labor market discrimination. *American Economic Review* 101(4), 1467–1496.
- Levine, P. B. and D. Ritter (2022). The Racial Wealth Gap, Financial Aid, and College Access. https://conference.nber.org/conf_papers/f162246.pdf.
- Lusardi, A., P.-C. Michaud, and O. S. Mitchell (2017). Optimal financial knowledge and wealth inequality. *Journal of Political Economy* 125(2), 431–477.
- Lusardi, A. and O. S. Mitchell (2014). The economic importance of financial literacy: Theory and evidence. *American Economic Journal: Journal of Economic Literature* 52(1), 5–44.
- Madrian, B. C. and D. F. Shea (2001). The power of suggestion: Inertia in 401 (k) participation and savings behavior. *The Quarterly Journal of Economics* 116(4), 1149–1187.
- Mitchell, O. S., S. P. Utkus, and T. Yang (2007). Turning workers into savers? Incentives, liquidity, and choice in 401 (k) plan design. *National Tax Journal* 60(3), 469–489.
- Moser, C. and P. Olea de Souza e Silva (2019). Optimal paternalistic savings policies. *Columbia Business School Research Paper* (17-51).
- Mukherjee, A. (2022). Intergenerational altruism and retirement transfers: Evidence from the social security notch. *Journal of Human Resources* 57(5), 1466–1497.
- Munnell, A. H. and A. Chen (2019). Are social security’s actuarial adjustments still correct?
- Myers, E. and J. Topoleski (2020). Worker Participation in Employer-Sponsored Pensions: Data in Brief. Technical report.
- Myers, E. A. (2024). Traditional, roth, and rollover individual retirement account (ira) ownership in 2022. Technical Report R48456, Congressional Research Service. Accessed October 16, 2025.
- Myers Jr, S. L. (1995). Racial discrimination in housing markets: Accounting for credit risk. *Social Science Quarterly*, 543–561.

- Neal, D. A. and W. R. Johnson (1996). The role of premarket factors in black-white wage differences. *Journal of Political Economy* 104(5), 869–895.
- Oliver, M. L. and T. M. Shapiro (1989). Race and wealth. *The Review of Black Political Economy* 17(4), 5–25.
- Palloni, A. and B. Novak (2016). Subjective survival expectations and observed survival: How consistent are they? *Vienna Yearbook of Population Research* 14, 187.
- Parker, K. and E. Patten (2013). Financial support across generations. Technical report, Pew Research Center. Social & Demographic Trends Project.
- Poterba, J. M., S. F. Venti, and D. A. Wise (1996). How retirement saving programs increase saving. *Journal of Economic Perspectives* 10(4), 91–112.
- Ramnath, S. (2013). Taxpayers’ responses to tax-based incentives for retirement savings: Evidence from the Saver’s Credit notch. *Journal of Public Economics* 101, 77–93.
- Ross, S. L. and J. Yinger (2002). *The color of credit: Mortgage discrimination, research methodology, and fair-lending enforcement*. MIT press.
- Rubin, D. (1987). *Multiple Imputation for Nonresponse in Surveys*. Wiley.
- Sabelhaus, J. (2023). How Does Social Security Affect the Racial Wealth Gap? *Ann Arbor, MI. University of Michigan Retirement and Disability Research Center (MRDRC) Working Paper*.
- Sabelhaus, J. E. and J. P. Thompson (2021). Racial Wealth Disparities: Reconsidering the Roles of Human Capital and Inheritance.
- Schafer, J. L. (1999). Multiple imputation: a primer. *Statistical methods in medical research* 8(1), 3–15.
- Schwandt, H., J. Currie, M. Bär, J. Banks, P. Bertoli, A. Bütikofer, S. Cattan, B. Z.-Y. Chao, C. Costa, L. González, et al. (2021). Inequality in mortality between Black and White Americans by age, place, and cause and in comparison to Europe, 1990 to 2018. *Proceedings of the National Academy of Sciences* 118(40), e2104684118.
- Sialm, C. and H. Zhang (2020). Tax-Efficient Asset Management: Evidence from Equity Mutual Funds. *The Journal of Finance* 75(2), 735–777.
- Skandalis, D., I. E. Marinescu, and M. Massenkoff (2022). Racial inequality in the us unemployment insurance system. *Available at SSRN 4097104*.
- Social Security Administration (2023a, September). Benefit Formula Bend Points.
- Social Security Administration (2023b, September). Maximum Taxable Earnings.
- Topoleski, J. J. (2018). Worker participation in employer-sponsored pensions: A fact sheet.
- US Department of the Treasury (2024). Tax Expenditures, FY2024. Technical report.
- U.S. Department of the Treasury (2025). Tax expenditures. <https://home.treasury.gov/policy-issues/tax-policy/tax-expenditures>. Accessed: 2025-10-18.
- Vanguard (2023). How america saves 2023. Technical report.
- Viceisza, A., A. Calhoun, and G. Lee (2022). Racial and Ethnic Disparities in Retirement Outcomes: Impacts of Outreach. *The Review of Black Political Economy*, 00346446231182343.
- Wagner, D. and M. Layne (2014). The Person Identification Validation System: Applying the Center for Administrative Records and Research and Applications’ Record Linkage Software. *Center for Administrative Records Research and Applications Working Paper*.
- Watson, J. and M. McNaughton (2007). Gender differences in risk aversion and expected retirement benefits. *Financial Analysts Journal* 63(4), 52–62.


- Wolff, E. N. (2023). Trends in the Retirement Preparedness of Black and Hispanic Households in the US. Technical report, National Bureau of Economic Research.
- Yogo, M., N. Cox, and A. Whitten (2023). WI23-12: Financial Inclusion Across the United States.

Figure 1: Differences in contribution rates by race and parental income


(a) Employee + Match DC Contribution Rate, by race




(b) Employee + Match DC Contribution Rate, by parental income




(c) Participation Rate, by race




(d) Participation Rate, by parental income



(e) Employee + Match DC Contribution Rate (contrib. > 0), by race




(f) Employee + Match DC Contribution Rate (contrib. > 0), by parental income



Notes: The right panels report gaps between White and workers of other racial/ethnic groups, specifically Asian, Black, Hispanic, and AIAN workers. The left panels report the gaps between workers with parental incomes in quintile 5 and quintile 1 to 4. All coefficients are estimated according to Equation (1), where White (Quintile 5) is absorbed as the omitted category. Therefore, the coefficients on the race (parent) terms, plotted in the figures measure the average gap between White and Asian, Black, Hispanic, or AIAN (Quintile 5 and Quintile 1 to 4). As described in 4.1, we nonparametrically reweight to balance age, income, and calendar time distributions across groups. We then progressively add linear individual characteristics, i.e., “Indiv.,” which includes gender, education, tenure, county, occupation, and employer. Please see Figure A.3 for a version of panel (a) with bars for each specification. In panels (a) and (b), the darker shaded regions represent the employee DC contribution rate gaps, while the lighter regions are the employer match gaps. The percentages printed above the bars represent the percentage difference relative to the average level for the omitted category.


Figure 2: The role of household and parental characteristics in racial contribution gaps

Combining race and parental income (younger cohorts)




Notes: Following the structure of Figure 1, the outcome variables are employee and total (employee + employer match) contribution rates. “+Indiv.” includes the previously defined individual characteristics. “+HH” adds household family structure and spousal income, and “+Parental Income” controls for parental income decile. The sample is restricted to the 1978-1994 birth cohort, which explains differences in “+Indiv.” estimates compared to Figure 1. See Section 3 and Appendix A.3.3 for additional details about the samples.

Figure 3: Employment at a firm offering a Defined Contribution or Defined Benefit Plan




Notes: Following the structure of Figure 1, we report gaps in access to a DC or DB plans. “Num Emp” indicates employer size bins. Employment at a firm offering a DC plan is measured from W2 data (i.e., we define an employer as offering a DC plan if at least 5% of its employees have strictly positive DC contributions), and we measure DB access from Form 5500 filings. As a robustness check for the 5% threshold, we re-run the analysis with a higher threshold (25%), shown in Appendix Figure A.5.

Figure 4: Differences in early withdrawals $\geq \$1,000$ among past DC savers




Notes: All panels follow the structure of Figure 1. The sample includes past DC savers (i.e., cumulative contributions $\geq \$1,000$ in nominal terms over the prior 4 years) under age 55 who were employed in survey year t . Early withdrawals are defined as distributions $\geq \$1,000$ (in nominal terms) taken in year $t + 1$, regardless of employment status at that time. Panels (a) and (b) show the withdrawal amounts as a share of year t income, panels (c) and (d) show the probability of taking an early withdrawal, and panels (e) and (f) are amounts like the first row but conditional on taking a withdrawal. See Appendix A.2.1 for additional details.

Figure 5: Early withdrawals $\geq \$1,000$ by ventile of earnings growth among past DC savers




Notes: These figures show the probability of early withdrawals of at least \$1,000 (in nominal terms) in year $t + 1$ by income growth ventile (comparing year $t + 1$ to year t income), separately by race and parental income quintile. The sample includes past DC savers (i.e., those who contributed at least \$1,000 (in nominal terms) over the prior 4 years), who are under age 55, and observed in the ACS in year t . Sample restrictions are the same as in Figure 4.

Figure 6: Differences in net contribution rates: $(\text{contribution}_t - \text{withdrawal}_{t+1})/\text{income}_t$





Notes: All panels follow the structure of Figure 1. Panel (a) includes all workers age < 55 in the primary sample. Panels (b) and (c) use the 1978-1994 birth cohort subset, which explains differences in the “+Indiv.” estimates between panels (a) and (c). Panel (c) additionally includes “+HH Charac.” (household family structure and spousal income) and “+Parental Income” (controls for parental income decile).

Figure 7: Components of DC wealth: employee contributions and subsidies



Notes: These figures decompose DC wealth at retirement into employee contributions, employer matching contributions, and federal tax subsidies. Panel (a) shows the share of DC wealth from each source (may not sum to 100 due to rounding). Panel (b) shows each component as a proportion of average annual lifetime earnings. Both panels break down results by the first four lifetime earnings quintiles (Q1-Q4) and the top two deciles (D9-D10).

Figure 8: Changes in wealth in the baseline and alternative counterfactuals
By race/ethnicity



Notes: This figure shows counterfactual changes in retirement wealth as a share of broad retirement wealth (i.e., the sum of DC and social security wealth) by own earnings level and by race (panels a-b) or parental income (panels c-d). Panels (a) and (c) use the baseline counterfactual redistributing tax benefits across all workers and matching within each firm. Panels (b) and (d) use the alternative counterfactual that redistributes matching within firm $\times$ income bins. Both panels break down results by the first four lifetime own earnings quintiles (Q1-Q4) and the top two deciles (D9-D10). Error bars show 95% confidence intervals.

Table 1: Worker-level summary statistics, by respondent race and parental income

	Aggregate F5500		Race					Parent Inc. Quintile		
	All	Parent	White	Hispanic	Black	Asian	AIAN	Q1	Q3	Q5
Age	41.2 (9.9)	30.0 (3.6)	42.1 (10.0)	39.4 (9.6)	40.1 (9.7)	39.7 (9.2)	40.8 (9.9)	30.1 (3.6)	30.0 (3.6)	30.1 (3.6)
W-2 total comp. (\$)	72,070 (210,600)	49,410 (69,830)	80,750 (247,100)	50,770 (119,900)	46,230 (81,680)	92,880 (175,800)	48,690 (68,680)	38,510 (33,110)	46,070 (41,060)	68,480 (103,000)
Parental AGI (\$)	0 (0)	90,110 (261,100)	106,800 (272,300)	56,160 (188,200)	50,880 (55,600)	86,950 (241,600)	60,350 (54,110)	14,180 (71,540)	66,460 (11,860)	233,500 (556,400)
Participation dummy (%)	65.1 (47.7)	59.2 (49.1)	68.8 (46.3)	54.5 (49.8)	56.5 (49.6)	71.3 (45.3)	55.9 (49.7)	51.0 (50.0)	58.5 (49.3)	69.0 (46.3)
Employee contrib. rate (% of comp.)	3.8 (4.6)	2.8 (3.5)	4.2 (4.8)	2.6 (3.6)	2.4 (3.3)	5.4 (5.9)	2.8 (3.8)	2.0 (2.9)	2.6 (3.3)	3.9 (4.2)
Employer match contrib. (% of comp.)	1.9 (2.0)	1.6 (1.9)	2.1 (2.0)	1.5 (1.9)	1.4 (1.8)	2.3 (2.1)	1.6 (1.9)	1.3 (1.7)	1.6 (1.9)	2.1 (2.0)
Employer match (\$)	1,690 (2,808)	1,063 (1,806)	1,965 (3,069)	988 (1,891)	857 (1,608)	2,451 (3,242)	1,007 (1,871)	671 (1,329)	946 (1,587)	1,725 (2,436)
Foregone matching (% of comp.)	1.7 (1.9)	2.0 (1.9)	1.5 (1.8)	2.0 (1.9)	2.1 (1.9)	1.2 (1.7)	2.2 (2.0)	2.3 (1.9)	2.0 (1.9)	1.6 (1.8)
Total contrib. > 0 (%)	8.8 (5.5)	7.4 (4.6)	9.1 (5.6)	7.6 (4.8)	6.8 (4.5)	10.9 (6.5)	7.7 (5.0)	6.4 (4.2)	7.2 (4.4)	8.7 (5.0)
Early withdrawal subsample: Age<55 Any status in t+1, working in t										
Withdrawal _{t+1} >\$1000 dummy (%)	10.3 (30.4)	8.1 (27.2)	10.1 (30.1)	9.0 (28.7)	14.6 (35.3)	6.3 (24.3)	10.6 (30.8)	8.6 (28.1)	8.2 (27.4)	7.3 (26.0)
... past DC saver	13.6 (34.3)	12.1 (32.6)	12.3 (32.8)	14.7 (35.4)	23.4 (42.3)	7.8 (26.9)	16.0 (36.6)	16.0 (36.6)	12.4 (33.0)	9.1 (28.7)
Withdrawal amount (% of comp.)	3.0 (21.9)	1.7 (12.6)	3.2 (22.8)	2.3 (14.8)	3.5 (25.0)	1.7 (14.0)	3.2 (19.2)	1.7 (11.9)	1.7 (11.8)	1.5 (11.8)
... Amount > \$1000	29.0 (62.1)	20.8 (39.6)	31.5 (65.2)	25.8 (42.7)	24.1 (61.2)	26.8 (49.4)	30.3 (51.3)	19.5 (35.9)	21.2 (35.7)	20.8 (38.5)
Net Contribution Rate (% of comp.)	2.4 (22.5)	2.7 (13.4)	2.8 (23.5)	1.7 (15.3)	0.1 (25.1)	5.9 (15.6)	0.9 (19.6)	1.6 (12.3)	2.5 (12.5)	4.5 (13.0)
Number of unique individuals	1,582,000	435,000	1,109,000	179,000	165,000	93,000	8,200	78,000	89,000	91,000

Notes: The table reports summary statistics for wage earnings data from the merged employee and employer data, which covers the 2008–2017 period. The “Race” columns correspond to the “Form 5500 sample”, while the “Parent Inc. Quintile” columns correspond to the “Parent–Form 5500 sample”. All the dollar values reported in the table are deflated to base year 2017 using the CPI, while the \$1,000 thresholds are in nominal terms. For more information about the different samples, please see Sections 3.3 and A.3.3.

Table 2: Wealth and wealth components, by population lifetime earnings bins

Value	(a) By race							(b) By parental income						
	Group	Q1	Q2	Q3	Q4	D9	D10	Group	Q1	Q2	Q3	Q4	D9	D10
Wealth from employee contributions (\$'000)	White	16.8	47.8	91.0	172.8	332.2	557.9	Bin 1	12.5	39.0	80.0	160.5	322.2	541.0
	Hispanic	11.3	35.7	75.0	148.3	310.2	524.1	Bin 3	15.1	44.0	86.6	168.2	331.7	556.3
	Black	8.6	25.7	50.8	108.1	244.4	443.8	Bin 5	17.3	49.4	95.3	185.9	363.5	614.7
	Asian	25.9	80.3	161.5	314.1	548.5	790.4							
Wealth from employer contributions (\$'000)	White	6.3	18.7	35.7	66.0	128.8	260.0	Bin 1	5.1	16.2	32.4	61.5	123.5	247.9
	Hispanic	5.1	16.1	32.6	60.3	124.7	246.8	Bin 3	6.0	18.0	34.7	64.7	128.5	257.4
	Black	4.3	12.5	23.9	47.2	103.5	207.2	Bin 5	6.6	19.2	37.1	70.2	138.1	280.3
	Asian	7.9	24.8	50.5	97.6	179.8	326.4							
Wealth from tax subsidies (\$'000)	White	3.8	13.4	28.0	54.8	98.5	238.5	Bin 1	3.0	11.7	26.5	52.6	96.0	227.5
	Hispanic	2.7	10.8	25.4	51.1	94.0	215.5	Bin 3	3.5	12.5	27.1	53.8	97.7	234.7
	Black	2.2	8.5	20.6	45.9	87.0	192.3	Bin 5	3.9	13.6	28.7	57.5	105.1	259.0
	Asian	5.5	19.5	40.6	73.5	133.0	310.3							
Total wealth from subsidies (\$'000)	White	10.1	32.1	63.7	120.8	227.3	498.5	Bin 1	8.1	27.8	58.9	114.0	219.5	475.4
	Hispanic	7.9	26.8	58.0	111.5	218.7	462.3	Bin 3	9.5	30.5	61.7	118.5	226.2	492.1
	Black	6.5	21.0	44.4	93.1	190.5	399.5	Bin 5	10.5	32.8	65.8	127.7	243.2	539.3
	Asian	13.4	44.3	91.1	171.1	312.8	636.7							
Baseline Total DC Wealth (\$'000)	White	26.8	79.9	154.6	293.6	559.5	1056.0	Bin 1	20.6	66.9	139.0	274.5	541.7	1016.0
	Hispanic	19.2	62.5	133.0	259.7	529.0	986.4	Bin 3	24.5	74.5	148.4	286.7	557.9	1048.0
	Black	15.0	46.7	95.2	201.2	434.9	843.3	Bin 5	27.8	82.2	161.1	313.7	606.8	1154.0
	Asian	39.3	124.6	252.6	485.2	861.3	1427.0							
Social Security Wealth (\$'000)	White	209.1	305.6	385.8	485.2	576.9	649.7	Bin 1	208.7	309.4	392.4	491.4	579.5	646.2
	Hispanic	212.3	302.3	382.6	481.7	574.7	644.0	Bin 3	209.1	303.3	384.8	486.1	579.3	649.0
	Black	204.2	301.3	381.9	481.7	573.3	633.3	Bin 5	209.3	300.2	378.5	479.0	573.9	653.7
	Asian	222.8	312.2	400.5	500.5	588.3	663.7							

Notes: This table presents average DC wealth (total and decomposed into its three components) and Social Security wealth by race (panel (a)) and parental income (panel (b)). The first sub-panel of each table shows average values for each component of DC wealth. The middle sub-panel gives total DC wealth. The third sub-panel is the average value of Social Security. Please note in panel Bins 1, 3, and 5 correspond to the bottom, middle, and top parental income quintiles. Columns show results by own lifetime earnings. There are six lifetime earnings bins—the bottom four quintiles and the top two deciles. Earnings bins are defined at the population level. All the dollar values reported in the table are deflated to base year 2017 using the CPI.

Table 3: Change in DC wealth at retirement under the counterfactual tax and employer contribution policy, population bins

(a) By race

Value	Group	Q1	Q2	Q3	Q4	D9	D10
Absolute change in DC Wealth (\$'000)	White	17.2 (0.2)	24.2 (0.3)	22.5 (0.3)	12.1 (0.6)	-7.3 (0.8)	-45.4 (1.8)
	Hispanic	17.9 (0.2)	26.3 (0.3)	25.2 (0.8)	17.2 (1.0)	-1.9 (1.8)	-48.9 (3.7)
	Black	19.2 (0.3)	32.6 (0.5)	37.0 (1.2)	31.1 (1.3)	14.4 (4.1)	-31.8 (10.9)
	Asian	16.0 (0.4)	16.6 (0.9)	8.2 (1.1)	-10.5 (1.0)	-52.3 (2.0)	-148.0 (7.1)
Counterfactual DC Wealth Gap	H-W Gap	15.8%	14.7%	10.7%	9.4%	4.5%	7.3%
	B-W Gap	22.3%	23.8%	25.4%	24.0%	18.6%	19.7%
	A-W Gap	-25.6%	-35.6%	-47.3%	-55.3%	-46.5%	-26.5%
Relative change in the racial DC wealth gap	H-W Gap	-44.3%	-32.3%	-23.6%	-18.4%	-16.6%	10.3%
	B-W Gap	-49.3%	-42.7%	-33.9%	-23.7%	-16.3%	-2.0%
	A-W Gap	-44.9%	-36.3%	-25.4%	-15.3%	-13.8%	-24.5%

(b) By parental income

Value	Group	Q1	Q2	Q3	Q4	D9	D10
Absolute change in DC Wealth (\$'000)	Bin 1	17.9 (0.2)	27.2 (0.3)	26.4 (0.6)	18.1 (1.0)	-0.1 (1.5)	-33.6 (6.5)
	Bin 3	17.6 (0.3)	25.3 (0.4)	24.0 (0.6)	14.2 (0.8)	-4.9 (1.3)	-42.4 (5.0)
	Bin 5	17.4 (0.2)	23.6 (0.4)	20.6 (0.6)	7.0 (0.6)	-19.3 (1.0)	-80.1 (3.1)
Counterfactual DC Wealth Gap	1-5 Gap	14.9%	11.1%	9.0%	8.8%	7.8%	8.5%
	3-5 Gap	6.5%	5.7%	5.2%	6.1%	5.9%	6.3%
Relative change in the parental income DC wealth gap	1-5 Gap	-42.7%	-40.7%	-34.6%	-29.9%	-27.2%	-29.0%
	3-5 Gap	-43.6%	-39.2%	-34.4%	-28.6%	-27.1%	-31.1%

Notes: This table presents the effect on wealth of our counterfactual exercise. Panel (a) gives results by race, and panel (b) gives results by parental income quintiles (bins) with Bins 1, 3, and 5 shown. Value row 1 shows the absolute change in DC wealth under the counterfactual. Value row 2 gives the counterfactual gap as a percentage of the White level (panel (a)) and the average level for those with the richest parents (panel (b)). Value row 3 gives the relative change in the percentage gaps obtained in moving from the baseline racial DC wealth gap to the counterfactual racial DC wealth gap. In both panels, each row is divided into six bins—the bottom four quintiles and the top two deciles. Earnings bins are defined at the population level. All the dollar values reported in the table are deflated to base year 2017 using the CPI.

Online Appendix for
“Who Benefits from Retirement Saving Incentives in the U.S.?
Evidence on Gaps in Retirement Wealth Accumulation
by Race and Parental Income”

Taha Choukhmane, Jorge Colmenares, Cormac O’Dea,
Jonathan Rothbaum, and Lawrence D.W. Schmidt

A Data

Appendix A.1 introduces our three data sources: the American Community Survey (ACS), the administrative tax data, and our codified Form 5500 filings. Appendix A.2 defines the variables used in our analysis. Appendix A.3 introduces our data construction and outlines which years of data we use, how we define our samples, and how we weight. Appendix A.4 discusses the representativeness of our data.

A.1 Data sources

A.1.1 American Community Survey (ACS)

Our individual-level build begins with all American Community Survey respondents aged 25-59.5 from the 2008-2017 waves (please see column (1) from Table A.1 for the total number of respondents). The ACS provides data on respondent age, education, gender, occupation, and county of residence. We supplement this with administrative tax records and Form 5500 regulatory filings, which we introduce in the next two subsections.

A.1.2 Administrative tax data

Our tax data comes from data from the following forms: W2s, 1099Rs, and 1040s.

We obtain data on earnings and deferred compensation from form W2. The W-2 extracts available at the Census Bureau have information from Box 1 on taxable wages, tips, and other compensation. These W-2 extracts also have an aggregate measure of deferred compensation from Box 12 that primarily consists of employee contributions to DC retirement plans. We cannot distinguish between contributions to different plans, but aggregate IRS data indicates that 93% of contributions are to 401(k), 403(b), or 457(b) plans, the dominant DC employer-sponsored plans offered by the employers in the private, non-profit, and public sectors, respectively.⁴³

⁴³Our aggregate measure of Box 12 aggregates elective deferrals to plans under Box 12 codes D: 401(k), E: 403(b), F: 408(k)(6), G: 457(b), and H: 501(c)(18)(D). The items in boxes E-F (403(b), 408(k), and

We obtain data on withdrawals from DC accounts using form 1099Rs. Form 1099-R filings (“Distributions From Pensions, Annuities, Retirement or Profit-Sharing Plans, IRAs, Insurance Contracts, etc.”) contain information on withdrawals from DC plans and payments from DB pensions. On the 1099-R extracts available to us, we observe the sum of withdrawals and distributions in two categories: 1) gross distributions from employer-sponsored plans and 2) IRA withdrawals.⁴⁴

We link individuals to their spouses and parents using form 1040. We link individuals to 1040 tax filings (from the years 1994, 1995, and 1998-2017), both contemporaneously (in the year we observe their earnings) and for a subset of younger workers (under age 42 in 2020), to the 1040 filings of their parents when they were claimed as dependents. We include non-filers who do not receive W-2s. From the contemporaneous 1040s of tax filers, we can observe marital status (from filing status) and link spouses through the PIK of the other filer on the tax return. We then link the spouses’ W-2s to observe their earnings as well. Section A.2.2 provides more details about how we make this link.

To construct intergenerational linkages and observe parental resources, we use the dependent information on 1040 tax returns, which is available for 1994, 1995, and from 1998 onwards. We begin with the 1978-1994 birth cohorts and, similar to Chetty et al. (2020), create a dependent claiming history that identifies any parent(s) that claimed each individual at all observed ages up to 18. In this paper, we define parent(s) as those who claim the child when they are closest to the age of 16. Therefore, we can link individuals with their parent(s), conditional on the parent(s) filing a 1040 in which they claim them as a dependent at some point during their childhood.

A.1.3 Retirement plan data

All retirement plans must submit an annual regulatory form (i.e., Form 5500) to the federal government. For plans with more than 100 participants, this form must include a narrative description of the retirement plan characteristics including details on the match schedules, vesting schedules, and auto features. These descriptions have been made publicly available by the Bureau of Labor, but in their original form (free-form text) they are not amenable

457(b) plans) are DC plans that primarily differ from 401(k)s in which employers can provide them (such as nonprofits and local, state, and federal governments). 501(c)(18)(D) contributions cover future payments under certain defined benefit (DB) plans. From 2008 to 2018, the average share of those dollars by Box 12 Code are D: 76 percent, E: 12 percent, F: 0.1 percent, G: 5.6 percent, and H: 0.02 percent. See IRS Statistics of Income Tax States for Individual Information Return Form W-2 Statistics, Table 7.A of Internal Revenue Service (2023), accessed 09/20/2023.

⁴⁴The IRS also excludes distributions, such as direct rollovers, Section 1035 exchanges, and Roth conversions from the 1099-R extract we use. For more information on the 1099-Rs, including separating DB and DC plans in the data, see Bee and Mitchell (2017).

to empirical analysis.⁴⁵ The data set that we use (described further in Arnoud et al. (2021), and Choukhmane et al. (2025)) was constructed from these files for the largest 4,800 defined contribution plans and a random sample of 1,000 smaller plans. For completeness, we reproduce several details about the Arnoud et al. (2021) data construction here. Details for each plan were codified in a consistent fashion. The plan-level data contain details on the full matching schedule, the vesting schedule, and any automatic features (auto-enrollment or auto-escalation). These very large firms cover a large number of employees—in 2017, 37 million employees were eligible to contribute to one of these large plans, collectively accounting for 55% of the population of workers enrolled in private and non-profit sector DC retirement plans.

We link these plan-year level variables to the Census firm infrastructure via a multi-stage procedure which incorporates information on numeric identifiers such as EIN and telephone number as well as fuzzy matching on name and address fields. We are able to match around 5,000 plans (Table A.1). We drop firms that have different match formulas for different employees, that change match formulas mid-year, or for which we cannot find match formulas. As a back-stop to our fuzzy linking, we further conduct internal consistency checks with our universe of W-2 filings, described further in section A.3.2 below, which leaves us with about 2,600 unique plans.

A.2 Variable definitions

All variables (unless otherwise specified) in dollar terms (including those used to compute rates) are deflated to base year 2017 using the Consumer Price Index provided by the Bureau of Labor Statistics.⁴⁶

A.2.1 Outcome variables

Employee contributions This is deferred compensation reported in Box 12 of the W-2 tax form. This dollar amount generally corresponds to contributions to an employer-sponsored contribution plan (such as a 401(k) plan). We define the employee contribution *rate* as the percentage of salary, using the ratio of the real employee contribution reported in Box 12 divided by the sum of the real taxable wage reported in Box 1 of the W-2 and the real employee contribution. The formula is $\frac{\text{employee deferred compensation}}{\text{employee deferred compensation} + \text{employee W-2 wages}}$. For the extensive margin, we define **Participation** as a dummy equal to one if the individual has strictly positive employee contribution. For the intensive margin, we look at the contribution rate conditional on participation equaling one. We refer to this variable as “Employee

⁴⁵See <https://www.dol.gov/agencies/ebsa/about-ebsa/our-activities/public-disclosure/foia/form-5500-datasets>.

⁴⁶See <https://www.bls.gov/cpi/research-series/r-cpi-u-rs-home.htm>.

contribution rate $| > 0$.”

Employer matching contributions This is the imputed match contribution implied by the employer matching formula collected from the employer’s Form 5500 filing. If an individual works more than one job, we match the employer matching formula to the highest-salary job. We apply the match formula to the three highest-earning jobs separately. We then aggregate the imputed contribution to generate the real employer match contribution. This is then added to the real employee contribution for the combined employee and employer matching contributions. We also refer to it as “Employer match,” “Employer contribution,” and “DC match contribution.” **Employee plus employer matching contribution** is the dollar amount of the sum of employee contributions and employer matching contributions. We also refer to it as “Total contributions.” The *rate* is defined as $\frac{\text{employee contributions} + \text{employer match}}{\text{employee contributions} + \text{employee W-2 wages}}$.

Early withdrawals We observe DC-plan withdrawals (and payments from pension plans) in Form 1099-R filings, which we treat as potential early withdrawals from DC plans. We take early withdrawals from the year after individuals appear in the ACS survey. We apply two main criteria: i) individuals must be younger than 55 *at the time of* the early withdrawal, and ii) individuals must withdraw more than \$1,000 to be classified as an early withdrawal. We apply these restrictions because federal law allows employers to automatically disburse individuals with under \$1,000 in deferred compensation upon separation and workers 55 years and older are allowed to take early withdrawals without incurring the tax penalty during separation from their employer. We define the **Withdrawal amount** as the real dollar amount of withdrawals in year $t + 1$ divided by real total worker compensation (i.e., wages plus employee contribution) in year t . Please note, those with positive early withdrawals but below \$1,000 are recorded as having 0% withdrawal amounts (just as those who take zero withdrawals). Similar to contributions, we measure withdrawals on the extensive margin (“Early withdrawal probability”) and the intensive margin (“Early withdrawal amount $|$ withdrawal $> \$1000$ ”). Also, since one needs to have funds in order to withdraw, we zoom in on these three withdrawal outcomes for “past DC savers,” which we define as people who have at least \$1,000 in cumulative contributions over the four years prior to taking the withdrawal.

Net contributions Since workers can contribute and withdraw early from their retirement accounts, we want to measure the *net* contributions in a given year. We define it as employee plus employer contributions in year t minus the withdrawal dollar amount in $t + 1$. The formula for the net contribution *rate* then is $\frac{\text{employee contributions}_t + \text{employer match}_t - \text{withdrawal dollar amount}_{t+1}}{\text{employee contributions}_t + \text{employee W-2 wages}_t}$. Please note that the rate is relative to the employee’s total compensation in year t .

DC or DB Offered We construct an indicator for working at an employer sponsoring a DC plan from the universe of W-2 filings. We define a firm as offering a DC plan in a given year if at least 5% of its employees have positive deferred compensation. As a robustness check for this threshold, we re-run the analysis with a higher threshold (25%) in Appendix Figure A.5. We define a firm as offering a DB plan if (i) it has filed a Form 5500 for a DB plan and (ii) the plan is open to new participants.

A.2.2 Additional Observable Mediating Variables

Race/Ethnicity Racial/ethnic groups are constructed using the answers to two questions asked in the ACS about race and Hispanic origin. The “Hispanic” group includes any individuals who answer yes to the question about Hispanic origin. In the paper, we let “AIAN”, “Asian”, “Black”, and “White” refer to non-Hispanic American Indian and Alaska Native, Asian, Black, and White individuals, respectively. Though we mainly reference the aforementioned groups in the paper, we also include in the categorical variable the remaining 1.8% of the population: non-Hispanic Native Hawaiian and Other Pacific Islander, non-Hispanic Some other Race, and non-Hispanic Two or more, which is any individual who identifies as more than one non-Hispanic group.

Parental income Parental income is defined as real adjusted gross income for parents that we can link to ACS respondents in 1040 filings. They are linked closest to when a person is claimed at age 16. We generate parent income bins (often ventiles in the paper) by year and child’s birth year (as a proxy for child age) from W-2s. Note that we do not incorporate ACS weights in our calculation of parent income bins.

Year The ACS provides the survey year.

Age bin We generate age from the ACS birth years and the ACS survey year. We bin people into ages 25-29, 30-34, 35-39, 40-44, 45-49, 50-54, and 55-59.5.

Income bin Income is defined as the sum of total real Box 1 wages and Box 12 deferred compensation on W-2 filings. We generate income deciles from the total compensation distribution per year and individual’s age bin, incorporating ACS weights.

Education We generate four educational categories from the ACS education variable, corresponding to whether a respondent has completed less than a high school degree, is a high school graduate, has a college degree, or has a graduate degree. Those who have some college but did not graduate college are included in the high school graduate category.

Gender The ACS provides gender (categorized as either male or female) for the 2001-2019 surveys. We generate a dummy for female.

Occupation The ACS provides several hundred occupational categories. The IPUMS 2010 crosswalk provides occupation codes that are consistent over time. We match the ACS occupation codes to the consistent IPUMS 2010 codes.

County The ACS provides the county of residence, so we construct categorical variables capturing the concatenation of a FIPS state code and the county code.

EIN W-2 forms are filed by employers who are required to report their Employer Identification Numbers (EINs). We take the EIN for the highest-earning job if an individual worked multiple jobs. We associate a worker with the retirement plan characteristics of the highest-earning EIN in the survey year.

Tenure Tenure is constructed by matching all ACS individuals with their employers from 2005-2020. Our W-2 filings report employers (by EIN) in order of most wages earned. We take the earliest known year for each individual-employer combination. We match the start year with the individual’s first EIN (employer from whom the individual earned the highest earnings) during the ACS survey year. Since our universe of W-2s begins in 2005 and our build begins in 2008, to avoid censoring issues we classify tenure at the main employer into four main categories: 1) working less than one year, 2) between 1-2 years, 3) between 2-3 years, and 4) at least 3 years.

Number of employees We bin workers from firms that have <10, <100, <1,000, <10,000, or 10,000+ employees. For brevity, we refer to this variable as “Num. Emp.”

Family structure We construct family structure from 1040 filings. The five main groups are single filer, no kids; single filer, with kids; dual filer, no kids; dual filer, with kids; and non-filers. Non-filers are those individuals who may receive W-2s but either forget or choose not to file 1040s.

Spousal income Spousal income is linked using 1040 filings from the ACS observation year. Spousal income is the sum of total real Box 1 wages and Box 12 deferred compensation from W-2 filings. Spousal income bins are classified by year and age using ACS weights into 12 main indicators: i) 0 percentile (spouses who report \$0 in earnings), ii) 10, 20, ..., 100 percentiles (spouses for whom we have nonzero earnings), and iii) missing (individuals who are either single, non-filers, or for whom we cannot match spousal income).

Auto-enrollment Auto-enrollment is taken from our universe of W-2 filings and matched firm data. Our Form 5500 filings report whether a 401(k) plan offers auto-enrollment in a given year. We classify Form 5500 filings that do not report an auto-enrollment start date after 2005 as not offering auto-enrollment. Individuals who start at their main firm after firms enact an auto-enrollment policy are classified as having auto-enrollment. Individuals

who start at their main firm before an auto-enrollment policy begins or work at firms without auto-enrollment policies are classified as not having auto-enrollment. Due to censoring issues, individuals who are observed starting at a firm in 2005 and work at firms where auto-enrollment begins either before or during 2005 are classified as unknown.

A.3 Data construction

A.3.1 Years

While some of our data cover a broader time range, we restrict our analysis to individuals observed in survey years from 2008-2017 due to censoring issues. Given that our W-2 filings begin in 2005, two key variables, job tenure and early withdrawals, depend on having a panel of at least four years. Job tenure is categorized into < 1 year, 1 year, 2 years, and 3+ years. Early withdrawals condition on more than \$1,000 in nominal deferred compensation over the four years prior to the early withdrawal. Both require W-2 Box 1 and EIN information from the three years prior to appearing in the ACS survey. Including pre-2008 individuals would select for higher income employees who can contribute more in a given year and would attenuate their tenure. We cap our observation years at 2017 due to our retirement plan-EIN crosswalk ending in 2017. Since the early withdrawal probability is computed in the year after the ACS year, our estimates for the probability of taking an early withdrawal are computed using data from 2009-2018.

A.3.2 Sample restrictions imposed throughout the paper

We apply three main restrictions to the ACS microdata from the survey years of 2008-2017. First, we restrict attention to survey respondents who are between ages 25 and 59.5 in the ACS survey year. Second, for compensation, we require the nominal sum of Box 1 wages and deferred compensation to be greater than \$8000 and require Box 1 wages to be strictly greater than \$0. This eliminates people who have zero wages but have deferred compensation (likely people with high wealth who are exploiting employer matches or instances in which box 1 wages were incorrectly parsed into the Census database). Third, we do not consider individuals who are associated with more than 11 W-2s in a given year and those who do not abide by the contemporaneous deferred compensation caps. This is referred to as the restricted “ACS sample.” For more information, please see the notes of Appendix Tables A.1 and A.2.

For analyses which use the linked information about retirement plans, we also impose a restriction which checks for internal consistency of the administratively reported level of employer contributions coming from Form 5500 with a comparable measure that we calculate internally by applying the matching formulas to the population of deferred compensation

levels coming from linked W2 forms. To reduce the potential impact of linking/measurement errors, we restrict analysis to plans for which difference in the calculated ratios of employee contributions to total contributions obtained from the two sources is smaller than 15 percentage points. Figure A.1 illustrates the strong concordance between two firm-level measures in our analysis sample by reporting a binned scatter plot comparing our imputations of employer contributions on the vertical axis with the actual reported Form 5500 measures on the horizontal axis. The two measures are very highly correlated, though there is a modest discrepancy for firms which have very low employee shares of contributions, likely due to the presence of additional, nonelective (non-matching) contributions which are excluded from our measurement. Please see Appendix Table A.1 for information on how the sample sizes change as we apply the sample restrictions.

A.3.3 Samples

We have two primary samples:

1. **Form 5500 sample** This is our main sample for our analysis of gaps in saving by race. It contains all individuals in the ACS for whom we match Form 5500 filings that meet our match formula and internal consistency restrictions. Analysis using this sample uses combined ACS and firm-level analytic weights, discussed below. The total number of unique individuals (after dropping missing individual variables required for our analyses) is approximately 1,582,000.
2. **Parent-Form 5500 sample** This is our main sample for our analysis of gaps in saving by parental income. It contains all individuals in the Form 5500 sample who are born after 1978 and to whom we can match a non-missing level of parental income. Analysis using this sample uses combined ACS and firm-level analytic weights, discussed below. The total number of unique individuals (after dropping those missing individual variables required for our analyses) is approximately 435,000.

To assess whether selection into the Form 5500 linked employer-employee sample matters for our findings, we compare summary statistics and key results from our baseline sample to two broader samples. These are our restricted “ACS sample” and “ACS sample with DC offered” which contains all respondents from the restricted ACS sample who are at firms where at least 5% of employees report deferred compensation (in which case we assume that the employer sponsors a DC plan). Before we run any regressions, we also require all potential mediating variables to be nonmissing to ensure consistency across all regressions. Please see the notes for Appendix Figure A.7 for more information.

In order to assess whether other plan features, such as vesting or the firm also offering a

DB plan, could be driving the main findings, we look at two different subsets of the Form 5500 sample: Fully-Vested and No DB. The Fully-Vested sample only includes workers that are 100% vested in the plans they are linked to in year t , while the No DB sample only contains workers at firms that also do not offer a DB plan. Appendix Figure A.8 compares the racial gap estimates across these different samples and finds them to be nearly identical.

A.3.4 Weighting

Individual weights. The ACS microdata include person-level analytic weights which enable researchers to produce estimates which are representative of the US population. For regressions and other summary statistics which do not use any linked retirement plan information, we use these person-level weights to construct estimates which are nationally representative.

Note that several of the individual-level income variables are converted into deciles or quintiles. To construct these categories, we first apply earnings and age restrictions (as explained in Appendix A.3.2), then compute weighted percentiles to use as breakpoints by year and age bin using the ACS sample weights. These decile assignments are therefore computed to be representative of all people in the ACS who match our sample requirements.

Retirement plan weights. Our employer data combines two different samples: a certainty sample of the 4,730 largest plans and a random sample of 1,471 plans from all remaining firms. For the random sample, we sampled firms with probabilities proportional to the number of participants. To ensure that our estimates are representative of the full population of firms filing the long version of form 5500, we calculate firm-level weights which are equal to the inverse of the probability of being selected into our sample.

Combining individual and plan weights. In our analyses which link the ACS and plan-level data (e.g., the Form 5500 sample) we combine person and firm-level weights to compute a combined measure to use at the individual level. Each individual’s probability of appearing in our matched build is the joint probability of being in a sampled firm and a sampled employee. Since the two samples are drawn independently of one another, the matched individual’s probability weight is the product of the ACS probability weight with the plan probability weight.

A.4 Data representativeness

This section presents some additional information which speak to the representativeness of our results, most of which are computed for a sample of ACS workers who are linked with firms whose retirement plans are included in our sample. Here, we characterize some differences between our analysis samples and broader populations of US workers.

Table A.2 provides information on averages for several main variables in the paper which are computed for various samples. We begin with the set of workers who are in the restricted ACS sample and, thus, satisfy the basic income and age restrictions we impose throughout (see Appendix A.3.2). As we move from left to right in the table, we see how sample means as we impose additional restrictions which are required to perform our analysis. Moving from column 1 to 2 shows the impact of a substantive restriction, namely that the worker receives income from an EIN for which at least 5% of its employees report positive levels of deferred compensation, which is our administrative proxy for working at an employer that offers a DC plan. Imposing this restriction is associated with higher income and higher average savings rates, driven by higher participation on the extensive margin.

Column 3 imposes that we can successfully link ACS respondents to a plan in our form 5500 sample. Relative to column 2, this sample excludes workers in small firms because employers with less than 100 employees are not required to submit a detailed description of their plan alongside their Form 5500 filing. Our sample weights are intended to make our estimates representative of the set of workers who are employed at the set of firms that offer DC plans and have more than 100 participants.

Overall, we see some modest changes in sample means between the sample of all ACS respondents whose employer sponsors a DC plan and our Form 5500 sample (for which we observe retirement plan details). In the latter sample of larger employers, workers earn about \$5,000 more in labor income and save at slightly higher rates (mostly driven by a higher probability of having positive contributions), and we also see that the propensity to take early withdrawals is slightly higher in this sample. The similarities between the two samples suggest that our estimates from the Form 5500 sample are quite representative of the broader population covered by the nationally representative ACS. Appendix Figure A.7 further shows that racial gaps in contribution and withdrawal behavior are also very similar across our full sample of ACS respondents whose employer sponsors a DC plan and our form 5500 sample.

Finally, column 4 additionally restricts to the subset of individuals with parental income available. Given that we can only match younger cohorts to their parents, this sample is unsurprisingly younger, has lower earnings, and saves at lower rates.

B Discussion of Mediating Variables in Regressions

B.1 Rationale for and impact of mediating variables in regression

In Section 4 of the main text, we report estimates of gaps in contributions by race and

parental income which include a number of observable characteristics, including dummies for age, year, deciles of labor income, gender, educational group, occupation, county, employer identification number (EIN), and tenure bin. In this section, we discuss potential economic rationales for why each of the individual-level characteristics that we include may impact DC savings rates, as well as the relationship between these variables and average savings rates in our data.

Year: Recent years have seen a substantial evolution in the DC landscape (e.g., the growth of auto-enrollment). To account for these, as well as savings differences over the business cycle, we include year fixed effects. However, we do not expect (and do not find) that the inclusion of year fixed effects affects our gaps as the composition of race and parental income groups is quite stable over our sample period.

Age: Age is an important driver of retirement saving; financing consumption in retirement is likely to be a central financial objective for older workers, whereas younger workers face a number of other competing savings objectives. We find, as expected, that savings rates are increasing in age (see, e.g., Gourinchas and Parker, 2002). Black and Hispanic workers are, on average, younger than White workers, and so understanding the extent to which age differences account for the gaps we observe is important.

Income: Income has been a traditional focus of the regulatory system and of the literature on the distributional analysis of the U.S. retirement system. It is well established that the rich save more (Dynan et al., 2004), and there are many reasons why this would be the case. Replacement rates from Social Security decline in income, the tax benefits are higher for those facing higher marginal tax rates (Congressional Budget Office, 2021), income risk tends to decline with income over most of the distribution outside of the top decile (Guvenen et al., 2014), and financial literacy is typically increasing in income (Lusardi and Mitchell, 2014; Lusardi et al., 2017). Furthermore, there are well-established differences in the distribution of income across races and by parental income, and so we find (as expected) that including income in the regression attenuates the gaps.

Education: Educational attainment could affect saving through channels beyond its correlation with income levels: life-cycle trajectories in expected income levels and income risk vary with education, and financial literacy increases in education. We consider the role of the highest degree attained, which we capture via four dummies for less than high school, a high school degree, a college degree, and a graduate degree. We find a strong relationship between educational attainment and savings.

Gender: Men and women may save different amounts for a variety of reasons such as differences in life-cycle earnings profiles (Goldin, 2021), risk preferences, life expectancy, and/or expected retirement benefits (Barber and Odean, 2001; Watson and McNaughton, 2007). Nevertheless, given that gender ratios are similar for workers across the racial and parental income groups we consider, gender has little impact on our estimated contribution gaps.

Occupation and County: Occupation may be relevant for savings as it can correlate with expected future earnings, income risk, and potential differences in risk or time preferences. Racial and parental income distributions differ across space, which may correlate with various factors such as the cost of living in retirement, so we additionally absorb county fixed effects.

Employer (EIN): Our data allow us to absorb EIN fixed effects, which enables us to identify racial contribution gaps among coworkers within the same employer. In addition to a number of economic characteristics that may differ across firms (for example, expected income trajectories and employment stability), a natural possibility is that workers sort into firms that differ in terms of the quality of the retirement benefits that they offer. For example, there is substantial heterogeneity across firms in the generosity of matching incentives, the nature of vesting schedules, and auto-enrollment and other default policies. Absorbing EIN fixed effects allows us to hold many of these features constant.

Tenure: The final economic characteristic that we consider is job tenure, which we split into bins for less than 1, 1, 2, and 3+ years. Tenure may relate to saving through its correlation with employment risk (e.g., Farber, 1994, shows that the probability of job separation decreases for workers with higher tenure), the probability that a worker’s contributions will vest, and workers’ awareness of plan benefits, among other channels.

C Comparison with Survey of Consumer Finances

We reproduce our baseline analysis using data from the Survey of Consumer Finances (SCF), the gold-standard source of survey information on wealth in the U.S. Given that the SCF does not contain information on parental background, we focus on differences in retirement contributions by race.

Sample. We use data from the 2010, 2013, and 2016 waves of the SCF, which cover a similar period to our administrative data. We impose the same restrictions as in our baseline analysis using administrative data: we focus on respondents aged 25 to 60 who make at least \$8,000 (in nominal terms) in wage income. For the regression analysis we further restrict the sample to those who report having access to a DC plan through their employer. This

restricted sample contains 4,097 respondents across the three SCF waves, of whom 512 are Black and 338 are Hispanic.

Descriptive statistics. Table A.3 compares summary statistics across the SCF sample and our sample of ACS respondents linked with administrative tax records. Demographics are broadly similar across the two samples, although the SCF sample has slightly higher labor earnings. Access to, and participation, in DC plans are significantly lower in the SCF (respectively 49.6% and 35.3%) relative to our ACS sample (78.2% and 45.2%). This is consistent with Dushi and Iams (2010) finding that survey responses underestimate access and participation in DC plans. They find that access to and participation in a DC plan are measured to be, respectively, 17 p.p. and 11 p.p. higher when complementing responses to the 2006 Survey of Income and Program Participation (SIPP) with respondents’ W2 records. The National Compensation Survey (NCS), which is based on responses from employers rather than employees, also reports higher levels of participation and access than the SCF (Topoleski (2018)). Among full-time civilian workers in the 2017 NCS—who are more comparable to our sample of workers earning more than \$8,000—68% have access to, and 48% participate in a DC plan.

Regression results. Figure A.11 compares the racial gaps in employee contributions estimated in the SCF to those estimated in the administrative data. Our specification using SCF data is the same as in the first set of bars in the baseline regression cascade (reported in Figure 1(a)), with standard errors adjusted for both imputation and sample variability errors.⁴⁷ Contribution gaps are qualitatively similar across the two datasets, although confidence intervals are much wider for the estimates using the SCF. In particular, it is hard to make precise statements about gaps conditional on age and income using the SCF: confidence intervals are large and not only overlap with zero (in the case of the Black-White gap) but also come close to and even overlap with our estimates using administrative data (in the case of the Hispanic-White gap). This suggests that survey data that has been typically used to study this question might be underpowered to detect (even sizeable) differences in retirement contribution rates by race.

D Micro-simulation Model

D.1 Overview

To understand the implications of differential saving and match patterns over the whole life cycle, we need to know the full life cycles of earnings, DC retirement plan features, and

⁴⁷Income bins in the SCF are constructed by year and 5-year age bins.

retirement contributions and withdrawals in the population. However, we have a maximum of 13 years of observations per individual. We use these partial life cycles and a simple hot deck imputation strategy to construct panels of synthetic life cycles, described in Section D.2.

With this data, we develop a micro-simulation model, described in Section D.4, which has three objectives. The first is to use the data on observed flows (earnings, contributions to DC accounts, and withdrawals from DC accounts) and a model of the economic and policy environment to generate simulated data for objects that we do not directly observe: the stock of resources for retirement, Social Security entitlements in retirement, and the trajectory of withdrawals from retirement accounts.

The second objective is to evaluate the counterfactual differences in wealth at retirement in a world where the individual saved in a taxable brokerage account rather than the tax-advantaged DC account. This allows us to build a measure of the value of tax expenditure at the individual level and its distributional incidence.

The third is to evaluate the distributional impact of changes to retirement savings institutions in the U.S. We consider three counterfactual policies. In the first, we break the link between saving and remuneration by calculating each firm’s counterfactual employer contribution that, if paid to every employee in proportion to their earnings, would cost the same to the employer as their current matching contributions. We evaluate the distributional impact of moving from the status quo to a system where all employees received that same proportional contribution. The second counterfactual setting breaks the link between government contributions to retirement accounts and savings choices by redistributing the tax expenditure so that it is proportional to lifetime income, once again regardless of the taxpayer’s retirement savings choices. The third counterfactual combines both reforms. In the interest of brevity, in the main paper, we focus on the combined counterfactual, but show selected results in the Appendix for the individual match and tax counterfactuals.

D.2 Modeled lifetime paths of earnings, retirement plans, and withdrawals

To estimate our micro-simulation model and evaluate the distribution of tax and wealth impacts of Defined Contribution (DC) retirement plans, we need to capture the distribution of paths of individual earnings, whether employers offer DC plans and matching contributions, and the amounts of DC plan contributions and withdrawals in the population. However, our data is limited in several respects. First, for many workers who are now close to retirement, DC plans were not widely used at the onset of their working careers. Furthermore, Form W-2s, our data source for individual wage and salary earnings and contributions to DC plans,

are only available starting in 2005. Our information on plan characteristics from the Form 5500 is only available through 2017. This leaves us with at most 13 years (2005 to 2017) to simultaneously observe earnings and DC contributions from W-2s, plan characteristics and matching from the Form 5500s, and retirement account withdrawals on Form 1099-Rs. Our aim is to convert these shorter windows of information into plausible lifetime trajectories spanning a working life cycle from age 25 to 65.

To construct the plausible lifetime trajectories, we use a simple hot deck imputation strategy. We partition ages starting at age 25 into overlapping bins of 4 years (25-28, 27-30, 29-32..., 63-66) For a given age bin b , we observe their ages at t , $t + 1$, $t + 2$, and $t + 3$. For individuals in bin $b + 1$, we observe their ages in $t + 2$, $t + 3$, $t + 4$, and $t + 5$. The individuals in each bin are observed in the ACS at the first two ages in that bin. So bin 1 has earnings at 25-28 for individuals we observe in the ACS at 25 and 26. We use the information from individuals in bin $b + 1$ to impute earnings, whether their employers offer a DC plan, contributions, characteristics, and withdrawals to individuals in bin b . We do so by matching individuals in bin b to similar individuals in $b + 1$ using the information observed at the overlapping ages ($t + 2$ and $t + 3$) and appending the information from the later non-overlapping age ($t + 4$ and $t + 5$) to bin b individuals.

As an example, suppose Person A had annual of \$25,000 at age 25, increasing \$1,000 each year to \$28,000 at age 28. Their employer did not offer a 401(k) plan and thus the person made no contributions to or withdrawals from a plan. Now suppose Person B earned \$26,500 at age 27, with annual increases of \$1,500 so that their salary was \$31,000 at age 30. Person B likewise had no access to a DC plan. Persons A and B had similar earnings and plan access during their observed overlapping ages, such that $y_{A,27} = \$27,000$ and $y_{A,28} = \$28,000$ compared to $y_{B,27} = \$26,500$ and $y_{B,28} = \$28,000$. As these workers had similar observable characteristics during their overlapping years, we impute to Person A the salary and contributions information from Person B for ages 29 and 30. This allows us to lengthen the number of years of “observed” earnings for Person A from four (covering ages 25-28) to six (covering ages 25-30). We can then repeat this process by imputing earnings for Person A at ages 31 and 32 using individuals in the next age bin covering ages 29 to 32. For a visual representation of how this works in practice, see Figure A.19. By repeating this process, we construct synthetic lifetime “histories” of earnings, access to DC plans, and employee and employer plan contributions.

For early retirement withdrawals by working-age individuals, we conduct an additional imputation step. Starting from the imputed lifetime earnings trajectories, we impute withdrawals in each year relative to contributions in the prior years to better align withdrawal

amounts to contributions. This helps reduce the number of cases in the model where withdrawals substantially exceed recent contributions. We also condition on earnings in the prior year, earnings in the year of imputed withdrawals (to better estimate withdrawals that result from earnings changes), and firm plan characteristics. However, because we do not observe returns or contributions in the distant past, there will be many cases in the data where withdrawals exceed recent contributions even with contributions observed over a longer time horizon than we use in the imputation.

D.2.1 Imputing DC plan access and matching rules for all firms

The hot deck model described in Section D.2 requires information on firm matching rules and DC plan availability for all firms. However, our data set of firm matching schedules from publicly available Form 5500 filings covers only a subset of firms, including the largest approximately 5,000 firms and a random sample of 1,000 smaller firms. We use this data to impute access to DC plans and plan matching rules for the remaining firms. Because we are interested in simulating lifetime trajectories for workers under the current system, we restrict to the plan characteristics in the most recent year for each firm linked to the Form 5500. For all firms, we summarize the distribution of deferred contributions across their workers. As an example, suppose that in a given firm 90% of workers have 0 deferred compensation and 10% contribute exactly 3 percent of their earnings to a DC plan. We summarize the share of workers in each firm that contribute between 0% and 10% of their earnings to DC plans with separate bins for 0 contribution and > 10 percent (i.e., bins of 0, (0-1) percent, [1-2) percent, [2-3) percent, and so on). We use k-means clustering to separate firms into 10 distinct groups based on the distribution of worker deferred contributions in these bins. Finally, we impute DC plan access and firm match schedules to those firms without available Form 5500 data using a hot deck matching on the worker DC contribution clusters, firm size, and average earnings for workers. This means that if two firms, A and B, have a mass of contributions at around 3 percent of earnings, they are likely to be in the same worker contribution cluster. Suppose Firm A has plan details available from Form 5500, with matching contributions of 100 percent up to 3 percent of earnings and 0 percent thereafter. Suppose further that Firm B, on the other hand, does not have available Form 5500-based plan information. Firm A would then be a likely “donor” of its match schedule to Firm B.

D.2.2 Accounting for Uncertainty from Imputed Lifetime Earnings and Firm Plan Characteristics

Imputation addresses missing data problems by drawing plausible values for the missing values. In our case, we are 1) imputing earnings for younger workers later in their career conditional on their earlier earnings (and other characteristics) and 2) match schedules for

firms conditional on the distribution of worker contributions to DC plans. However, there is necessarily uncertainty in the imputed values. A given distribution of DC plan contributions could be consistent with multiple match schedules and autoenrollment policies at a firm. Workers with the same characteristics in years t and $t + 1$ could go on to have very different earnings in years $t + 2$ and $t + 3$. We can take multiple independent draws to approximate this uncertainty. This is the intuition behind multiple imputation.

We repeat our entire imputation process with different initial seeds to get 5 independent “implicates”, or draws, from our imputation process.⁴⁸ We then use Rubin’s Formula (Rubin, 1987) to combine the results from the independent draws to estimate standard errors for statistics estimated on the imputed data. This formula combines the variance estimated for each statistic in the individual draws (the within imputation variance) with the variance in the estimate across the draws (the between imputation variance) to estimate the total variance for each estimate.⁴⁹

D.2.3 Assumptions and Limitations of the Imputation Model for Lifetime Earnings

The imputation model relies the assumption that we can accurately capture the trajectory of lifetime earnings by estimating earnings in years $t + 4$ and $t + 5$ by conditioning on earnings in years $t + 2$ and $t + 3$. This can affect the estimated variation in lifetime earnings depending on whether it is necessary to condition on more years of earnings ($t + 1$, t , $t - 1$, etc.) to model earnings in subsequent years. Additionally, because of the limited number of years of data available for each worker, we must assume that older cohorts are good proxies for the future earnings of younger cohorts. Formally, suppose we want to estimate the earnings in a given year (y_t) for an individual in cohort c given prior earnings and other observable information (X). We can impute that value by taking a draw from the distribution of $f_c(y_t|X)$. We are assuming that the distribution of earnings for the older cohort (O) and the younger cohort (Y) are the same, or $f_{c=O}(y_t|X) = f_{c=Y}(y_t|X)$. However, if that is not true and younger cohorts will have different long-term earnings trajectories than we would estimate from the partially overlapping shorter spells we observe from different cohorts in our data. This is a potential concern for racial subgroups, such as Asian workers, whose population has grown considerably in recent years (Krogstad and Im, 2025). For Asian workers, older cohorts

⁴⁸We chose 5 draws for practical reasons, as the full imputation process, tax modeling, and analysis takes several hours to complete and it would be impractical to scale the process up to 20 or 50 or 100 implicates. In terms of statistical efficiency, 5 implicates is often sufficient (see Schafer (1999)), however, more implicates would provide better statistical power (Dong and Peng, 2013). Five implicates has also become somewhat standard for large, complicated imputation processes - the SCF and Consumer Expenditure Survey both use Multiple Imputation with 5 implicates to account for uncertainty in their imputation models for survey nonresponse.

⁴⁹For reference, the formula can be found in equations 2.1-2.3 in Schafer (1999).

may differ from younger cohorts due to recent migration, and we might be concerned that $f_{c=O} \neq f_{c=Y}$.

Hot deck models are extremely flexible - by matching “recipients” (those getting values imputed) to “donors” (those with the information observed) that match on all model characteristics, we make very few parametric assumptions about the data. Suppose we match on 10 characteristics. We only impute values to recipients from donors that match on all 10 characteristics, and make no parametric assumptions about how we donors that match on fewer characteristics could provide information to improve our imputation. This is equivalent to estimating a fully saturated model with all possible 2-, 3-, 4-,..., 10-way interactions between all the model variables and then imputing a value by drawing from the error of the regression within sets of observations that match on all 10 variables. The cost of this flexibility is the curse of dimensionality. With 3 categories in each of the 10 characteristics, we would have $3^{10} = 59,049$ distinct cells. If a recipient is of a type (in a cell) with no donors, we cannot impute a value for that individual. To draw a value for an individual that does not have a donor that matches on all characteristics, we must drop a variable from the model or coarsen the match on another (such as matching on wider prior earnings bins).⁵⁰ With our baseline sample of 1.6 million individuals (1), we have approximately 44,000 individuals in each age cohort (1.6 million divided by the number of age cohorts in our sample, given the age range of 25-59.5 $\approx$ 44,000). That limits our ability to impute lifetime income trajectories for smaller groups, such as Asian or American Indian or Alaskan Native (AIAN) workers, as their small samples in our data would require matching on fewer additional variables (like education, prior earnings, gender, etc.) and result in less realistic lifetime income trajectories.

D.3 Summary and Output

The result of this procedure is a simulated data set for individuals i age $t \in \{25, \dots, 90\}$, where 90 is assumed to be the last age of life and in which mortality is deterministic.

Variables that we observe (with the associated notation given for objects that will feature in the treatment below) are:

- Demographic measures: age (t), race, and parental income
- Compensation measures: earnings (e) and contributions the employee elects to make to their employer-sponsored defined contribution account (dc^{ee}),
- Whether the individual works in a firm offering a DC plan and, if so, the match schedule

⁵⁰This is a major constraint in practice. The earnings imputation model in the CPS ASEC has 15 variables and 620 billion possible cells in the first match level, with a sample of under 200,000 observations. Each subsequent level will have fewer cells until all recipients have been matched to a donor.

$(dc_f(.))$, and

- Withdrawals from DC accounts before retirement (w).

D.4 Model Description

This section outlines the model. Table A.4 gives a summary guide to all the model parameters.

D.4.1 Savings Vehicles

Central to the exercise is to compare outcomes under the status quo (in which the deferred compensation is paid into a tax-deferred defined contribution account) with a counterfactual setting (in which tax-favored DC accounts are not available, and those same contributions are instead paid into a (taxable) brokerage account). We evaluate each individual's savings trajectory under two systems of taxation, indexed by $j \in \{DC, BK\}$. The superscript $j = DC$ indicates that the individual is saving in a tax-deferred 401(k) account, and $j = BK$ indicates that they are saving in a brokerage account. Savings in the tax-deferred (DC) account benefit from the fact that income tax is deferred until the funds are withdrawn and that investment returns accumulate free from income and capital gains taxes. Savings in the brokerage account come from taxed income, have returns that are subject to tax, and have income tax-free withdrawals.

Below we refer to the 'DC saver' and the 'brokerage saver' as shorthand for the saver in a setting where DC accounts are available and not, respectively.

D.4.2 Observable: Earnings, contributions, and withdrawals

Employees receive compensation that can be divided into earnings $e_{i,t}$ and deferred compensation $dc_{i,t}^{ee}$. Employees may also receive an employer match, which is a firm-varying function indexed by f : $dc_f^{er}(dc_{i,t}^{ee})$. For ease of notation, we suppress the dependence of the employer contribution on the employee contribution and denote the employer contribution made on behalf of individual i at age t as $dc_{i,t}^{er}$.

Withdrawals from retirement accounts are denoted $w_{i,t}^j$, with j indexing the nature of the account (DC or brokerage). We observe withdrawals made by our agents up to age 65. These observed withdrawals in the data are from the DC account and recorded before the deduction of income tax.

D.4.3 Wealth

Wealth balance at the beginning of the period is given by $B_{i,t}^j$ and is initialized to zero at

age 25. Net flows into the wealth vehicle are denoted by $f_{i,t}^j$:

$$f_{i,t}^j = dc_{i,t}^{ee} + dc_{i,t}^{er} - \tau_{i,t}^{c,j} - w_{i,t}^j, \quad (2)$$

where dc^{ee} and dc^{er} are, respectively, deferred compensation by the employee and the employer-match contributions. There are two deductions from these gross flows. The first ($\tau^{c,j}$) are taxes on these contributions. This object will be defined in detail below, but, in brief, note that dc^{ee} and dc^{er} are measured as gross-of-tax. For the DC saver, no income tax is owed on these flows and so $\tau_{i,t}^{c,DC} = 0$. For the brokerage saver, income tax must be paid before contributions are made. The second deduction, $w_{i,t}^j$, are withdrawals from the account. These are observed before the age of 65; in Section D.4.6, we propose a model of withdrawals which fills these in for after the age of 65.

The law of motion for wealth balance is given by:

$$B_{i,t+1}^j = (B_{i,t}^j + f_{i,t}^j)(1 + \rho_t) - \tau_t^{r,j}, \quad (3)$$

where ρ_t is a rate of return that depends on age (with time dependence due to the changing mix of assets in the portfolio), and $\tau_t^{r,j}$ represents the taxes paid on that return in that period. This will be zero for the DC saver, and we will describe it for the brokerage saver in the next subsection.

D.4.4 Investment returns

Two comments are needed on the investment returns. First, they vary with age. Each age t is associated with a portfolio composition between equities, bonds, and bills, with shares given by s_t^k , s_t^b , and s_t^m . During working years, these shares are interpolated from Fidelity target date funds.⁵¹ In retirement, we assume exclusive investment in bonds. The age profile of investment composition is shown in Figure A.20a, and the associated age profile of real rate of return is shown in Figure A.20b. Real rates of return for these asset types (ρ^k , ρ^b , and ρ^m , respectively) are taken from Jordà et al. (2019). The combination of these assumptions yields age-specific rates of return ρ_t :

$$\rho_t = \rho^k \cdot s_t^k + \rho^b \cdot s_t^b + \rho^m \cdot s_t^m. \quad (4)$$

⁵¹We use asset allocations of the Fidelity Freedom Funds ranging from retirement years 2005 to 2065 between equities, bonds, and short-term debt as of year-end 2022. Distance to retirement is thus the target date minus 2023. A one-dimensional Akima interpolator was used to calculate shares between observed age distances to retirement. Our shares may be compared to Fidelity's own description of their glide path (Fidelity, 2023).

The second comment on returns is the division of returns into unrealized capital gains, distributions taxed as long-term capital gains, and returns taxed as income (e.g., ordinary dividend income).⁵² Distinguishing between the nature of the return will be important in our treatment of the brokerage saver's taxable returns. The share of returns represented by each of these is given by χ^g , χ^k , and χ^i , respectively, which sum to 1. The dollar flows associated with each of these three types of return are given below:

$$r_{i,t}^{g,j} = (B_{i,t}^j + f_{i,t}^j) \cdot \chi^g \cdot \rho_t, \quad (5)$$

$$r_{i,t}^{k,j} = (B_{i,t}^j + f_{i,t}^j) \cdot \chi^k \cdot \rho_t, \quad (6)$$

$$r_{i,t}^{i,j} = (B_{i,t}^j + f_{i,t}^j) \cdot \chi^i \cdot \rho_t. \quad (7)$$

Accumulation and withdrawal of untaxed capital gains When individuals withdraw funds from their accounts, they realize some (previously unrealized) capital gains. This has tax implications for the brokerage saver, making it necessary for us to keep track of that part of the account balance formed of unrealized capital gains. We divide the account balance $B_{i,t}^j$ into principal $B_{i,t}^{p,j}$ and (thus far untaxed) capital gains $B_{i,t}^{g,j}$. We define the latter recursively as:

$$B_{i,t+1}^{g,j} = B_{i,t}^{g,j} + r_{i,t}^{g,j} - w_{i,t}^{k,j}, \quad (8)$$

where $B_{i,t}^{g,j}$ is the cash value of the stock of unrealized capital gains in the account balance, $r_{i,t}^{g,j}$ are additional untaxed gains attained in year t , and $w_{i,t}^{k,j}$ are gains actually realized when a withdrawal is made.

Whenever a withdrawal $w_{i,t}^j$ is made, we assume that the withdrawal comprises untaxed capital gains $w_{i,t}^{k,j}$ and principal $w_{i,t}^{p,j}$ in proportions that equal their share of the stock of wealth. That is, the share of any withdrawal by the brokerage saver that is subject to capital gains tax is equal to the share of unrealized capital gains in wealth:

$$\frac{w_{i,t}^{k,j}}{w_{i,t}^j} = \frac{B_{i,t}^{g,j}}{B_{i,t}^j}. \quad (9)$$

D.4.5 Social Security Income

We assume all individuals stop earning when they turn 66 and begin claiming Social Security benefits. Central to the determination of Social Security benefits is 'Average Indexed

⁵²The second component—distributions taxed as long-term capital gains—does not represent returns which are realized for a withdrawal. Rather, they are the gains realized as mutual fund managers trade assets and passed on to investors. See Fidelity's description of these distribution types at www.fidelity.com/learning-center/investment-products/mutual-funds/taxes.

Monthly Earnings' (*aime*), calculated as the average of the best 35 years of total compensation.⁵³ Consistent with Social Security rules, the measure of earnings that enters this calculation is capped at a value e^{max} :

$$aime_i = \frac{1}{35} \sum_{k \in \text{best } 35} \left\{ \frac{\min(e + dc_{i,t}^{ee}, e^{max})}{12} \right\}. \quad (10)$$

Monthly Social Security benefits are equal to 90% of *aime* up to the first 'bend point' (\$895 in 2018), 32% of any *aime* above the first bend point and below the second point (\$5,397 in 2018), and 15% of any *aime* above the second bend point.

D.4.6 Withdrawals

We distinguish between 'early withdrawals' and 'retirement withdrawals.' The former are those taken before age 65, which we observe in our data. The latter are taken after age 65. These are not observed and so must be modeled.

Early withdrawals We define early withdrawals as all withdrawals before age 65.⁵⁴ The measure that we observe in our data (denoted $w_{i,t}^{DC}$) is that before income taxation, which must be paid on all withdrawals from DC accounts. For the equivalent withdrawal applied to the brokerage saver (denoted by $w_{i,t}^{BK}$), we calculate the after-tax quantity retained by the DC saver.

One complication arises when the early withdrawal that we see would lead to the brokerage saver having a negative balance. This occurs in only a small share of cases (14%). In these cases, we adjust the measure we see in our data to be the largest number that avoids the brokerage saver going negative.

Retirement withdrawals Individuals retire at the beginning of age 66 with balance in their account of $B_{i,66}^j$. They employ a consumption rule each year to determine how much to withdraw each period t . We set this rule such that consumption for the DC saver is constant each period.

In particular, the withdrawal each period is equal to:

$$w_t^j = \frac{1 - \alpha}{1 - \alpha^{90-t+1}} B_{i,t}^j, \quad (11)$$

where $\alpha = \frac{1}{(1+\rho^b)}$ is defined using the return on bonds ρ^b .⁵⁵ This rule, which we illustrate

⁵³All variables are expressed in real terms, and we assume a stationary earnings distribution. As a result, there is no indexation of earnings in equation (10).

⁵⁴Not all of these will be subject to an early withdrawal penalty. We return to this when we discuss the taxation of withdrawals in Section D.6.1.

⁵⁵This consumption rule is that obtained from a cake-eating problem in which life-span is deterministic

in Figure A.21, keeps pre-tax withdrawals constant. We assume that individuals consume their withdrawal, net of taxes:

$$c_t^j = w_t^j - \tau_{i,t}^{w,j}, \quad (12)$$

where $\tau_{i,t}^{w,j}$ are taxes incurred by withdrawing money from account j and defined in the next section. Constant (pre-tax) withdrawals keep post-tax consumption constant for the DC saver (as income does not change in retirement) and close to constant for the brokerage saver (for whom small changes in average tax rates will occur as wealth is decumulated).

D.5 Summary

The data that we construct, together with the features outlined above, yield two parallel data sets: one representing the earnings, savings, account balance, and withdrawals of the DC saver, and one representing the same objects for the brokerage saver. We represent these by the following:

$$\left\{ \{e_{i,t}, dc_{i,t}^{ee}, dc_{i,t}^{er}, B_{i,t}^{DC}, w_{i,t}^{DC}\}_{t=25}^{90}; \{c_{i,t}^{DC}\}_{t=66}^{90} \right\} \quad \left\{ \{e_{i,t}, dc_{i,t}^{ee}, dc_{i,t}^{er}, B_{i,t}^{BK}, w_{i,t}^{BK}\}_{t=25}^{90}; \{c_{i,t}^{BK}\}_{t=66}^{90} \right\},$$

where the first three objects are common across the two tuples, but the balances, withdrawals, and consumption profiles differ due exclusively to the different forms of taxation faced by the two savers.

D.6 Taxation

The previous section concludes by noting our data and micro-simulation model yield, for each individual in our data, two trajectories of wealth accumulation and decumulation—one if they save in a DC account and one if they saved the same quantities in a taxable brokerage account. Due to their access to preferential taxation, the DC saver will have higher consumption in retirement. This section shows how we measure these differences in tax treatment across the life cycle.

At the most general level, we take the flow of income, saving, and returns and use TAXSIM to evaluate the taxes. This allows us to construct our summary measure of wealth at retirement: the present discounted value of consumption facilitated by accumulated wealth at retirement. This section provides the interested reader full details on how we measure that.

D.6.1 Decomposing the overall tax burden into components

We denote our modelled tax function, which distinguishes between the three forms of in-

and the discount rate is set equal to the interest rate.

come that agents in our model earn, as $T(N, K, S)$. N denotes inflows taxed according to the income tax schedule (e.g., wage income during working life and 401(k) distributions in retirement); K denotes income taxed as long-term capital gains; and S denotes Social Security benefits.⁵⁶

We decompose the total tax burden (denoted T) into shares that can be ascribed to earnings (τ^e), contributions to retirement accounts (τ^c), investment returns (τ^r), taxes owed on Social Security benefits (τ^s), and withdrawals from retirement accounts (τ^w). Earnings, contributions, returns and withdrawals, of course, interact in a non-linear (and quite complex) manner to generate overall tax liability. This means that there is no unique decomposition such that the total tax burden T can be written as the sum of these components. This section explains how we obtain one such decomposition.

We use rules for tax year 2018 according to NBER’s TAXSIM 32 tool to calculate federal income tax owed by each simulated individual.⁵⁷

Taxation of Earnings We first define taxes on earnings ($\tau_{i,t}^{e,j}$) as follows:

$$\tau_{i,t}^{e,j} = \begin{cases} T(e_{i,t}, 0, 0) & \text{if } t < 66 \text{ for } j = DC, BK; \\ 0 & \text{if } t \geq 66 \text{ for } j = DC, BK. \end{cases} \quad (13)$$

This does not differ by the type of saver, and the second equality follows from our assumption of no earnings from the age of 66.

Taxation of Social Security We define the tax on Social Security as the tax that would be paid if an agent had their Social Security income and no other income:

$$\tau_{i,t}^{ss,j} = T(0, 0, ss_{i,t}) \quad \text{if } t \geq 66 \quad \text{and } j = DC, BK, \quad (14)$$

which also does not differ by type of saver.⁵⁸

Taxation of Contributions Our definition of taxable earnings excludes that part of earnings which was saved for retirement: an employee’s choice of deferred compensation and any associated employer match $dc_{i,t}^{ee} + dc_{i,t}^{er}$. For the DC saver, income contributed to the

⁵⁶Note that effective tax rates in retirement are usually very low (Chen and Munnell, 2020) due in part to the favorable tax treatment of Social Security benefits, on which many households pay no tax at all (Joint Committee on Taxation, 2019).

⁵⁷The N , K , and S income sources are fed into the *pwages*, *ltcg*, and *gssi* fields in TAXSIM, respectively. We assume that all individuals take the standard deduction and do not claim any other credits or deductions. See Feenberg and Coutts (1993) for a description of the TAXSIM model.

⁵⁸As it happens, $\tau_{i,t}^{ss,j}$ will be zero for everyone in our sample—an individual with maximum Social Security income and no other income will not face any income tax. We retain the variable for completeness and because its exclusion may obscure some features of the exposition.

account is untaxed, so $\tau_{i,t}^{c,DC} = 0$. For the brokerage saver, the tax we ascribe to contributions is equal to the additional income tax the saver would have paid by taking compensation as earnings. This is given by the second line in:

$$\tau_{i,t}^{c,j} = \begin{cases} 0 & \text{for } j = DC, \\ T(e_{i,t} + dc_{i,t}^{ee} + dc_{i,t}^{er}, 0, 0) - \tau_{i,t}^{e,BK} & \text{for } j = BK, \end{cases} \quad (15)$$

where the positive term in the second line gives the income tax owed from earnings that include deferred compensation, and the negative term nets off that tax already ascribed to earnings, defined in equation (13).

As we assume that there are neither earnings nor contributions after retirement, for both savers we obtain $\tau_{i,t}^{c,j} = 0$ for all $t \geq 66$.

Taxation of withdrawals The taxation of withdrawals depends on whether they are ‘early withdrawals’ (those made up to the age of 65) or ‘retirement withdrawals’ (from the age of 65). In the former case, the DC saver must pay income tax and may face a tax penalty. This penalty is incurred at a rate p_t , which is equal to 10% for non-exempt withdrawals before age 59.5 and 0 for withdrawals after age 65. The first line of equation (16) gives this quantity. The positive terms are the regular income tax on earnings and DC withdrawals and the tax penalty; the negative term subtracts taxes already ascribed to earnings.

The brokerage saver need not pay income tax on withdrawals but must pay capital gains taxes on gains realized to withdraw their funds ($w_{i,t}^{k,BK}$). This quantity is defined in the second line in equation (16), where the first term gives the tax liability from earnings, contributions, and capital gains and the negative term subtracts taxes already ascribed to earnings and contributions:

$$\tau_{i,t}^{w,j} = \begin{cases} T(e_{i,t} + w_{i,t}^{DC}, 0, 0) + p_t w_{i,t}^{DC} \mathbf{1}(t < 60) - \tau_{i,t}^{e,DC} & \text{if } j = DC \text{ and } t < 66, \\ T(e_{i,t} + dc_{i,t}^{ee} + dc_{i,t}^{er}, w_{i,t}^{k,BK}, 0) - (\tau_{i,t}^{e,BK} + \tau_{i,t}^{c,BK}) & \text{if } j = BK \text{ and } t < 66. \end{cases} \quad (16)$$

In retirement, the DC saver pays regular income taxes on withdrawals (see the first line of equation (17)), while the brokerage saver pays capital gains taxes on the share of withdrawals that represent previously unrealized gains ($w_{i,t}^{k,BK}$). Both savers are also claiming their Social Security payments, which enter as the third argument of the tax function:

$$\tau_{i,t}^{w,j} = \begin{cases} T(w_{i,t}^{DC}, 0, ss_{i,t}) - \tau_{i,t}^{ss,DC} & \text{if } t \geq 66 \text{ and } j = DC, \\ T(0, w_{i,t}^{k,BK}, ss_{i,t}) - \tau_{i,t}^{ss,BK} & \text{if } t \geq 66 \text{ and } j = BK. \end{cases} \quad (17)$$

Taxes on investment returns All returns on funds in DC accounts are untaxed. That is, there is no taxation of unrealized gains ($r_{i,t}^{g,j}$), no income tax on dividend income ($r_{i,t}^{i,j}$), and no capital gains tax for distributions ($r_{i,t}^{k,j}$). The taxes paid by the *DC* saver on returns are therefore zero.

For the brokerage saver, while the unrealized capital gains ($r_{i,t}^{g,j}$) incur no immediate tax liability, income tax is paid on dividend income ($r_{i,t}^{i,j}$), and capital gains tax is paid on realized gains. As described in Section D.4.4, the latter comes in two parts—that part of the return which is distributed even in the absence of a withdrawal ($r_{i,t}^{k,j}$) and that part of the return which is realized when a withdrawal is made ($w_{i,t}^{k,j}$).

Taxes on portfolio returns for the brokerage saver are given in (18). In both lines (representing taxes before and after retirement, respectively), the first term gives all taxes due in a particular period (on earnings, contributions, withdrawals, and returns), and the second term nets off those taxes already ascribed to earnings, contributions, and withdrawals:

$$\tau_{i,t}^{r,BK} = \begin{cases} T \left(e_{i,t} + dc_{i,t}^{ee} + dc_{i,t}^{er} + r_{i,t}^{i,BK}, r_{i,t}^{k,BK} + w_{i,t}^{k,BK}, 0 \right) \\ \quad - \left(\tau_{i,t}^{e,BK} + \tau_{i,t}^{c,BK} + \tau_{i,t}^{w,BK} \right) & \text{if } t < 66, \\ T \left(r_{i,t}^{i,BK}, r_{i,t}^{k,BK} + w_{i,t}^{k,BK}, ss_{i,t} \right) - \left(\tau_{i,t}^{ss,BK} + \tau_{i,t}^{w,BK} \right) & \text{if } t \geq 66. \end{cases} \quad (18)$$

D.7 Lifetime Measures

D.7.1 Implied post-tax interest rate

Our model contains multiple interest rates that could be used to evaluate the present value of future flows. To do this, we define an interest rate $\hat{r}_{i,t}$ as the post-tax rate of return that the brokerage saver would pay if their deferred gains each period were realized as long-term capital gains.⁵⁹ We first define the hypothetical taxes on portfolio returns in this case as:

$$\widehat{\tau}_{i,t}^{r,BK} = \begin{cases} T \left(e_{i,t} + dc_{i,t}^{ee} + dc_{i,t}^{er} + r_{i,t}^{i,BK}, r_{i,t}^{g,BK} + r_{i,t}^{k,BK} + w_{i,t}^{k,BK}, 0 \right) \\ \quad - \left(\tau_{i,t}^{e,BK} + \tau_{i,t}^{c,BK} + \tau_{i,t}^{w,BK} \right) & \text{if } t < 66, \\ T \left(r_{i,t}^{i,BK}, r_{i,t}^{g,BK} + r_{i,t}^{k,BK} + w_{i,t}^{k,BK}, ss_{i,t} \right) - \left(\tau_{i,t}^{ss,BK} + \tau_{i,t}^{w,BK} \right) & \text{if } t \geq 66, \end{cases} \quad (19)$$

where this expression is the same as that in equation (18) except for the inclusion of $r_{i,t}^{g,BK}$

⁵⁹This assumption ensures that interest rate we choose for discounting does not depend on patterns of withdrawals that we observe in our data.

each period in the second argument. The implied, post-tax interest rate is then

$$\hat{r}_{i,t} = r_t - \frac{\widehat{\tau_{i,t}^{r,BK}}}{B_{i,t}^{BK} + f_{i,t}^{BK}}. \quad (20)$$

This rate is used across all counterfactuals.

D.7.2 Wealth

We have two measures of resources in retirement: a) ‘DC wealth’ (the value in DC accounts) and b) ‘Broad Retirement Wealth,’ which also includes Social Security wealth.

Wealth Our measure of wealth is the present discounted value of after-tax withdrawals facilitated by the account balance. We express this as recursively, backwards from age 90 with $A_{i,90}^j = 0$ for both savers $j = BK, DC$:

$$A_{i,t}^j = \frac{A_{i,t+1}^j}{1 + \hat{r}_{i,t+1}} + \left(w_{i,t+1}^j - \tau_{i,t+1}^{w,j} \right) - \left(dc_{i,t+1}^{ee} + dc_{i,t+1}^{er} - \tau_{i,t+1}^{c,j} \right) \quad (21)$$

as the present value of future post-tax withdrawals less future post-tax contributions.

This is private retirement wealth and does not include wealth held in the form of Social Security benefits. We define Social Security wealth as:

$$SS_{i,t} = \frac{SS_{i,t+1}}{1 + \hat{r}_{i,t+1}} + \left(ss_{i,t+1} - \tau_{i,t+1}^{ss} \right). \quad (22)$$

Our broad measure of wealth takes into account both wealth in private accounts and Social Security wealth:

$$A_{i,t}^{BR} = A_{i,t}^{DC} + SS_{i,t}. \quad (23)$$

D.8 Decomposing retirement wealth

In this subsection we define how we decompose retirement wealth into three components: wealth that flows from employee contributions, wealth that can be ascribed to employer contributions, and wealth due to the favorable tax treatment of DC accounts.

D.8.1 Value of DC tax treatment

The total tax benefit to an individual i is defined as the difference between the retirement wealth of the DC saver and that of the brokerage saver:

$$A_i^T = A_{i,65}^{DC} - A_{i,65}^{BK}. \quad (24)$$

To find the retirement wealth concept attributable to the employee alone, we need to find the proportion of contributions that are from the employee for each individual in our data. The value at retirement of the respective contributions made by the employee and the employer are:

$$DC^{ee} = \sum_{t=25}^{65} dc_{i,t}^{ee} \left(\prod_{\tau=t}^{65} (1 + \hat{r}_{\tau}) \right) \quad DC^{er} = \sum_{t=25}^{65} dc_{i,t}^{er} \left(\prod_{\tau=t}^{65} (1 + \hat{r}_{\tau}) \right). \quad (25)$$

These can then be used to calculate the respective proportions of retirement wealth for the brokerage saver (i.e., after tax benefits have been removed) coming from employee and employer contributions:

$$A_i^{EE} = \frac{DC^{ee}}{(DC^{ee} + DC^{er})} \cdot A_{i,65}^{BK} \quad A_i^{ER} = \frac{DC^{er}}{(DC^{ee} + DC^{er})} \cdot A_{i,65}^{BK}. \quad (26)$$

D.8.2 Benchmarking our estimated value of the tax expenditure

The Treasury Department provides estimates of the present value of tax benefit given to DC savers each year (U.S. Department of the Treasury (2025)). Between 2008 (the first year of our sample) and 2022 (the most recently available estimates at time of writing) this estimated cost has ranged from approximately \$68bn to \$160bn. As a rough check on our model, we compare our estimate of the tax benefit to this official estimate. Using an annuitization factor based on our model interest rate, we transform the mean lifetime tax benefit $A_i^T = \$52,857$ to an annual measure by dividing it by a factor of 57.6. This results in a mean *annual* tax benefit of about \$918 for the population represented by our simulated data, where population DC coverage is estimated to be of those currently in their 20s.⁶⁰ To convert our number to one which can be considered reflective of the current US population (who are the basis for the Treasury's numbers), we divide our average annual tax benefit by the ratio of the population DC savings rate to the hot deck sample DC savings rate (approximately 1.6). This yields a comparable mean annual tax benefit of \$576 per worker. Finally, we multiply this by an estimate of the civilian population engaged in work at any time in 2018 from the public CPS ASEC, around 166 million people. Our model estimate of aggregate annual tax benefit to DC savers is then \$96 billion: in the range of official estimates

D.9 Model Limitations

While our model aims to capture key features of the federal income tax expenditure, it

⁶⁰Our hot deck imputation model matches younger people to older people based in part on access to a DC plan. The fact that younger people are more likely to work at employers sponsoring DC plans makes access to DC plans more prevalent in our sample than in the population.

is an approximation that abstracts from several potentially relevant factors. For example, we assume no inflation, which is relevant given that the income tax code taxes nominal rather than real gains. We also do not model state income taxes, whereas in reality, many states have tax codes that favor DC saving. Both factors will lead us to understate the tax benefit. Other assumptions, such as the absence of heterogeneity in asset allocation, have an ambiguous effect on the size of the tax expenditure.

Our model of Social Security assumes all agents have similar life expectancy and claim at the same age: whereas different life expectancies and claiming ages can affect how an income stream is mapped into a wealth level (see, for example Brown (2002)). We also assume that individuals claim Social Security on their own (rather than their spouse’s) records, meaning we do not account for spousal claiming strategies or survivors benefits. This omission may differentially affect our estimates across racial and ethnic groups, as marriage rates vary considerably: Asian individuals have the highest marriage rates, followed by White, Hispanic, and Black individuals. Since married individuals have access to additional benefit options, our model likely understates Social Security benefits more for groups with higher marriage rates. Finally, our measure of Social Security is a wealth measure: we do not account for the longevity insurance that it provides: this can vary by race and socio-economic characteristics (Arapakis et al. (2023)).

D.10 Tax Counterfactual

The tax counterfactual considers how distributing the aggregate tax expenditure proportionally to lifetime earnings would affect wealth. This would break the link between saving decisions and a worker’s share of this tax expenditure but would not otherwise increase redistribution across lifetime income groups. Every individual would receive a government contribution to her DC account proportional to her lifetime earnings. This uniform proportion is chosen such that the total cost of these government contributions matches the total cost incurred under the existing tax-favored system.

Let the value of lifetime total earnings be:

$$LE_i = \sum_{t=25}^{65} (comp_{i,t}) \left(\prod_{\tau=t}^{65} (1 + \hat{r}_{\tau}) \right), \quad (27)$$

where $comp_{i,t} = e_{i,t} + dc_{i,t}^{ee}$ is the sum of earnings and deferred compensation. We define a redistributed tax advantage that allocates the total tax benefit in the economy proportionally

to lifetime income:

$$A_i'^T = \frac{LE_i}{\sum_n LE_n} \cdot \sum_n A_n^T \quad (28)$$

where the first term is an individual's share of aggregate lifetime earnings and the second term is the aggregate tax expenditure. We assume that these redistributed tax benefits are fully illiquid before retirement and cannot be withdrawn early. In our baseline, we assume no behavioral response to the change in the tax treatment of retirement contributions so that employee and employer contributions are unchanged. We indicate aggregates under this counterfactual with a $'$ superscript. DC wealth and broad retirement wealth in this counterfactual are therefore:

$$A_i'^{DC} = A_i'^{EE} + A_i'^{ER} + A_i'^T \quad A_i'^{BR} = SS_i + A_i'^{DC}.$$

D.11 Match counterfactual

In the presence of an employer match for retirement contributions, those who save more receive higher total compensation from their employer. Our employer match counterfactual breaks this link and considers the effect of a noncontingent employer contribution that is proportional to employee earnings. Every worker receives an employer contribution to her DC account proportional to her current earnings, regardless of her contributions. This percentage would be the same for all workers at the same employer but varies across employers.

For our employer-match counterfactual, we calculate the proportional contribution that, if given to all employees in the firm, would cost the same as their actual matching contributions. That is, for each time period t we calculate the ratio of total matching contributions to total income for each firm and multiply that by individual income. Denoting an employee i working in firm f with an employer match of $dc_{i,t}^{er}$,⁶¹ instead of receiving $dc_{i,t}^{er}$ in period t , the employee receives:

$$dc_{i,t}^{*er} = \frac{comp_{i,t}}{\sum_{i \in f} comp_{i,t}} \cdot \sum_{n \in f} dc_{n,t}^{er} \quad (29)$$

where the first term is individual i 's share of compensation in their firm in period t and the second is the aggregate matching contribution made by their employer in period t . We then calculate all modeled objects as described above assuming that employees receive the

⁶¹This will be linked to the employee's contribution ($dc_{i,t}^{ee}$) by a function that gives the employer match: $dc_{i,t}^{er} = m_f(dc_{i,t}^{ee})$.

counterfactual match $dc_{i,t}^{*,er}$. Accounting for this and for the fact that taxation trajectories will be different will yield different levels of wealth at retirement. All stocks in this model are denoted as in the baseline model but with the addition of a $*$ superscript. We denote the employers' counterfactual contributions and due to the tax expenditure as $(A_i^{*,ER}$ and $A_i^{*,T})$, respectively, so that the new levels of wealth and broad wealth in retirement equal:

$$A_i^{*DC} = A_i^{EE} + A_i^{*,ER} + A_i^{*,T} \quad A_i^{*BR} = SS_i + A_i^{*DC}. \quad (30)$$

D.11.1 Alternative counterfactual

We also consider an alternative scenario that makes a different counterfactual distribution of employer contributions. In this second counterfactual (reported in Figure 8 (b) and (d)), instead of redistributing the match so that in most firms each worker receives the same employer contribution rate, we redistribute so that each worker in each firm and earnings decile receives the same employer contribution rate. Letting d be firm decile, the analogue to equation 29 in this case becomes

$$dc_{i,t}^{*,er} = \frac{comp_{i,t}}{\sum_{i \in f,d} comp_{i,t}} \cdot \sum_{n \in f,d} dc_{n,t}^{er}. \quad (31)$$

The reference to ‘most’ firms in the above reflects the facts for small firms, accurately estimating an income gradient is difficult given the small sample sizes. For firms with fewer than 200 workers we thus retain our previous approach of redistributing and equalizing employer contribution rates (as a proportion of earnings) at the firm level.⁶²

D.12 Combined Counterfactual

Our combined counterfactual equalizes the employer match contribution and the tax subsidy. To do this, we first obtain the brokerage saver's wealth under the employer match counterfactual $B_{i,t}^{\dagger BK}$. We add the redistributive tax subsidy calculated in tax counterfactual ($A_i^{\dagger T}$). Denoting all aggregates under the combined counterfactual with an $\dagger$ superscript (though note that $A_i^{\dagger ER} = A_i^{*,ER}$), we obtain:

$$A_i^{\dagger DC} = A_i^{EE} + A_i^{\dagger ER} + A_i^{\dagger T} \quad A_i^{\dagger BR} = SS_i + A_i^{\dagger DC}. \quad (32)$$

⁶²In Form 5500 data, plans with more than 200 eligible participants account for more than 93% of aggregate employer dollars in 2017.

D.13 Additional Discussion and Analysis Related to Behavioral Responses

D.13.1 Related Evidence from the Literature

There is no consensus in the literature on how much private saving responds to employer matching and tax incentives. Engen et al. (1996) and Poterba et al. (1996) discuss the implications of the early literature on saving incentives. In a more recent contribution, Choi (2015) reviews the literature on matching and finds that it is associated with a small positive effect on participation and an ambiguous effect on average contribution rates. Engelhardt and Kumar (2007) use cross-sectional data to estimate that an increase in the match rate of 25 cents per dollar increases 401(k) participation rates by 5 p.p., while Duflo et al. (2006), in a randomized controlled trial with a one-time saving subsidy, find that increasing the match rate from 0% to 50% increases take-up by 11 p.p. However, the positive effect of matching on take-up and employee contributions may not translate into higher wealth accumulation if employees reduce their nonretirement saving or increase borrowing in response. Choukhmane and Palmer (2025) estimate that approximately two-thirds of increased employee pension contributions in the UK are financed through reduced nonretirement saving and increased credit card borrowing.

Regarding tax incentives, a review by Friedman (2015) notes that “tax subsidies appear to primarily affect the allocation of savings across accounts, rather than the total amount of savings.” Ramnath (2013) finds no statistically significant effect of the U.S. saver’s tax credit on the level of retirement contributions. Similarly, Chetty et al. (2014), using administrative data from Denmark, estimate an increase in net saving of approximately 1 cent per Danish kroner (DKr) of tax expenditure on subsidies for retirement saving.

D.13.2 Counterfactuals with Alternative Behavioral Response Assumptions

Given the lack of consensus and overall small effects found in the empirical literature, our baseline assumption of no behavioral response in private saving will likely be a reasonable approximation. In an extension, we recalculate the results assuming that each dollar of employer matching or tax subsidies generates either 10 cents (which corresponds to the upper bound of the 95% confidence interval in Chetty et al. (2014)) or 30 cents of additional employee savings. The results are shown in Figure A.15, and indicate that, even with an elasticity of employee saving to financial incentives higher than empirical estimates, our counterfactual policy raises retirement wealth accumulation in the bottom half of the lifetime earnings distribution (and especially for those who are Black, Hispanic or have lower income parents). Furthermore, Figure A.16 shows that the relative change in DC wealth gaps by race

and parental income is quantitatively very similar when assuming a 10% or 30% elasticity, reflecting the fact that the reduction in employees' saving is larger for groups that benefit more from saving incentives in the baseline.

D.14 Parameterization

D.14.1 Rates of return

Total investment return is given by an age-varying interest rate r_t . Each age t is associated with a portfolio composition between equities, bonds, and bills, parameterized by σ_t^k , σ_t^b , and σ_t^m . During working years, these shares are interpolated from Fidelity target date funds (see, for example, the 2040 Target Date Fund in Fidelity (2023)). In retirement, we assume exclusive investment in bonds. The age-profile of investment composition is shown in Figure A.20a, and the associated age-profile of the real rate of return is shown in Figure A.20b. Real rates of return for these asset types (ρ^k , ρ^b , and ρ^m) come from Jordà et al. (2019). The combination of these assumptions yields age-specific rates of return r_t according to:


$$r_t = \rho^k \cdot \sigma_t^k + \rho^b \cdot \sigma_t^b + \rho^m \cdot \sigma_t^m. \quad (33)$$

Note that in retirement $r_t = \rho^b$. We derive the decomposition of returns into these shares by studying the historical price trends and distributions of the Fidelity Freedom Funds Fidelity (2023).⁶³

⁶³Our breakdown of 50% price change, 40% distribution taxed as long-term capital gains, and 10% taxed as income is very similar to the 48/43/9 breakdown found by Sialm and Zhang (2020) under the assumption that 95% of dividends are non-qualified.

E Supplemental Figures and Tables


Figure A.1: Bin scatter of W2-imputed vs. Form 5500-reported employee contribution share




This graph presents a binned scatter plot comparing the W2-imputed firm-level employee share of contributions ($\frac{\text{total employee deferred compensation}}{\text{total employee deferred compensation} + \text{total employer match}}$) against the publicly-filed Form 5500 average employee share of contributions, which is used as the running variable to compute 20 ventiles on the horizontal axis. The dashed line is a line of best fit through the individual points on the scatter plot.

Figure A.2: Savings and Withdrawal Probability Gaps, by all racial/ethnic groups

(a) Employee+Match DC Contribution Rate




(b) Withdrawal amount | past DC savers




Notes: We show for the headline outcomes the gaps for all the racial/ethnic groups included in the analysis. Note, “Two or more” refers to individuals who identify as more than one racial group (i.e., not including Hispanic individuals), and “NHPI” refers to Native Hawaiian and Other Pacific Islander. For more information, please see the Figure Notes for Figures 1 and 4.

Figure A.3: Savings and withdrawal amount gaps showing all mediating factors




Notes: Please see Figure Notes for Figures 1 and 4. Here, we show each layer of the cascade with all the individual specifications (education, gender, occupation, county, EIN, and tenure).

Figure A.4: Savings gaps and racial composition across the parental income distribution




Notes: In panel (a), we plot the coefficients for each parent income decile dummy in the fully saturated model, which includes individual and household characteristics (i.e., the basis for the final set of columns in Figure 2). In panel (b), for each parental income decile, we report the ratio of the group's share in the decile to its share in the overall population.

Figure A.5: DC or DB Access with higher threshold




Notes: Please see Figure Notes for Figure 3. Here, DC plan access is measured for those at firms where at least 25% of the employees have strictly positive deferred compensation.

Figure A.6: DC Access and Firm Size



Notes: In panels (a) and (b), we document DC plan access across the income and firm size distributions, respectively. In panel (c), we report the racial shares across the firm size distribution.


Figure A.7: Racial gap estimates in Form 5500 vs. ACS sample with DC offered




Notes: This figure presents robustness checks for our main results (Figure 1). We use the progressive specification (Equation 1) for the Form 5500 sample and the restricted ACS sample of employees (conditional on non-missing control variables so the sample is consistent across columns) and the subset of those with DC Offered. In panels (a) and (c), we report the same figures as in Figures 1(a) (note, here with just the employee DC contribution rate) and 4(c). In panel (b), we show the same estimates as in panel (a) but on a larger sample, which is the ACS sample with DC Offered. In panel (d), we similarly show the same estimates as in panel (c) but on the entire ACS sample. We report the ACS sample in panel (d) because our sample inclusion restriction for analysis of early withdrawals conditions on having recently made contributions and therefore implies having had access to a DC plan (please see Appendix A.2.1 for more information). As in previous figures, the numbers in bold represent the percentage difference relative to the average level for the omitted category (i.e., White workers). For more information on this methodology, please refer to the notes of Figures 1 and 4, and for more information on the differences between the samples, please see Appendix A.3.3.

Figure A.8: Racial gap estimates in Form 5500 vs. Fully-Vested vs. No DB samples

- (a) Employee + Match DC Contribution Rate, Form 5500 sample (b) Employee + Match DC Contribution Rate, Fully-vested sample




- (c) Employee + Match DC Contribution Rate, No DB sample




Notes: This figure is analogous to Figure A.7. We now compare the Form 5500 sample results with the sub-sample that are fully-vested and do not offer a concurrent DB plan, respectively. For more information on the different samples, please see Appendix A.3.3.

Figure A.9: Contribution rates across different specifications




Notes: This figure presents the average racial contribution rates (employee plus employer contribution rates) across different age bins (panels (a) and (b)) and own income deciles (panel (c) and (d)). The estimates come from the raw specification without other mediating channels. The model is $y_{it} = \alpha + \beta_1 group.i + \zeta(group.i \cdot race_i) + \epsilon_{it}$. We add the intercept to put these numbers into levels rather than differences. 95% confidence intervals are included; standard errors are clustered by EIN. Please note that the confidence intervals reflect differences relative to the base category, which varies depending on the specification. For the race (parental income) by age results, it is White (parental income Q5) workers of ages 25-29. For own income, it is White (parental income Q5) employees in the 5th income decile. For the age specification, we nonparametrically reweight to balance income and calendar time distributions across groups, while for the income specification, we reweight to balance age and calendar time distributions.

Figure A.10: Withdrawal probabilities across different specifications




Notes: This figure is analogous to Figure A.9 except the outcome here is the probability of taking an early withdrawal among past DC savers. Please refer to the Figure A.9 Notes for details on the specification and Figure 4 Notes for details on the outcome variable.

Figure A.11: Comparison of our results with the Survey of Consumer Finances




Notes: This figure compares our baseline results from the Form 5500 sample with those from the SCF. Panel (a) shows the estimates of the Black-White gap, while panel (b) is for the Hispanic-White gap. Standard errors are calculated by bootstrapping with 500 replications, using the `scfcombo` command. For more information, please see the figure notes for Figure 1. The only departure from the structure of Figure 1 is that here we do not reweight to balance the distributions and all controls are linearly added.

Figure A.12: Microsimulation model: Key outputs by Race




Notes: This figure illustrates the features of main outputs from our micro-simulation model. Panel (a) shows mean values by race and age bins 25-29, 30-34, ..., and 60-65. Note that the last age bin contains six ages. In panel (a), earnings are the sum of wage income and deferred compensation. Panel (b) shows DC wealth at retirement divided by the simple average of earnings during working years for each race and population earnings bin group. Lifetime earnings groups are divided into six bins—the bottom four quintiles and the top two deciles. Panel (c) illustrates the heterogeneity in DC wealth at retirement within each race and lifetime earnings group. Percentiles shown are p10, p25, p50, p75, and p90. The measure of wealth shown is the same as in panel (b), DC wealth at retirement divided by average lifetime earnings. Panel (d) shows the same percentiles for the present value of all Social Security distributions over average lifetime earnings.

Figure A.13: Microsimulation model: Key outputs by Parental Income




Notes: This figure illustrates the features of main outputs from our micro-simulation model. Panel (a) shows mean values by parental income bin and age bins 25-29, 30-34, ..., and 60-65. Note that the last age bin contains six ages. In panel (a), earnings are the sum of wage income and deferred compensation. Panel (b) shows DC wealth at retirement divided by the simple average of earnings during working years for each parental income bin and population earnings bin group. Lifetime earnings groups are divided into six bins—the bottom four quintiles and the top two deciles. Panel (c) illustrates the heterogeneity in DC wealth at retirement within each parental income bin and lifetime earnings group. Percentiles shown are p10, p25, p50, p75, and p90. The measure of wealth shown is the same as in panel (b), DC wealth at retirement divided by average lifetime earnings. Panel (d) shows the same percentiles for the present value of all Social Security distributions over average lifetime earnings.

Figure A.14: Contributions of employer and tax subsidies to retirement wealth




Notes: This figure shows lifetime employer and tax subsidies as a percentage of average annual lifetime earnings, by own earnings level and by either race or parental income. Panel (a) shows these subsidies by race, and panel (b) shows them by quintiles (“bins”) of parental income. In both panels, the darker bars show average employer matching subsidies, and the lighter bars show average tax subsidies to retirement savings. Lifetime earnings groups are divided into six bins—the bottom four quintiles and the top two deciles. Earnings bins are calculated at the population level.

Figure A.15: Change in retirement wealth measures under alternative assumptions about the elasticity of employees' savings to incentives




Notes: This figure shows alternate versions of Figure 8, panels (a) and (c), under different assumptions about behavioral responses to a reduction in savings incentives. We assume that each additional dollar of tax or employer matching incentives creates an additional \$0.1 (left panels) or \$0.3 (right panels) of additional employee savings. Moving from subsidies to flat contributions, therefore, crowds out some employee contributions (i.e., equal to either 10% or 30% of the baseline level of tax and matching subsidies). The change in retirement wealth under the counterfactual policy is expressed as a percentage of average annual lifetime earnings by population lifetime earnings quintiles, with the top quintile split into two deciles. In all panels, the height of bars in solid colors represents the change in wealth, taking into account the behavioral responses, while the transparent portion represents the wealth crowded out due to these behavioral responses.

Figure A.16: Relative change in DC wealth gaps under alternative assumptions about the elasticity of employees' savings to incentives, by group and parental income bins.



Notes: This figure shows the relative change in the DC wealth gaps under different assumptions about the behavioral response to a reduction in savings incentives. We assume that each additional dollar of tax or employer matching incentives creates an additional \$0, \$0.10 or \$0.30 of additional employee savings. The first of these returns the mechanical effect of removing and redistributing the subsidies, and the change in the gaps under this assumption are those in our baseline (see the last rows in both panels of Table 3). The second two assumptions imply that moving from subsidies to flat contributions crowd out some employee contributions (i.e., equal to either 10% or 30% of the baseline level of tax and matching subsidies). The change in DC wealth gaps under the counterfactual policy is expressed as a percentage of the racial and parental income gap for each lifetime earnings quintiles, with the top quintile split into two deciles. The lifetime earnings quintiles and deciles are defined at the population level.


Figure A.17: Counterfactual and alternative counterfactual changes in wealth.




Notes: This figure considers the counterfactual in panels (a) and (c), and the alternative counterfactual in panels (b) and (d). For more details, please see Figure 8 notes. Changes in wealth are measured as a percentage of average annual lifetime earnings, by own earnings level and by either race or parental income. The top row shows results by race/ethnicity, and the bottom row shows it by parental income quintiles. Lifetime earnings groups are divided into six bins—the bottom four quintiles and the top two deciles. Earnings bins are calculated at the population level. Error bars are at the 95% confidence interval.

Figure A.18: Change in wealth from tax and match counterfactuals separately

(a) Δ broad ret. wealth under tax counterfactual, by race (b) Δ broad ret. wealth under match counterfactual, by race




(c) Δ broad ret. wealth under tax counterfactual, by parental inc. (d) Δ broad ret. wealth under match counterfactual, by parental inc.




Notes: This figure illustrates the impact of two supplementary counterfactual exercises on measures of retirement wealth. The ‘tax’ counterfactual (panels a and c) distributes the aggregate federal tax expenditure so that all workers receive a contribution that is in proportion to their lifetime earnings. The ‘match’ counterfactual (panels b and d) distributes the aggregate employer matches in each firm so that all workers in that firm receive the same proportion of their earnings. Changes are measured as a percentage of broad retirement wealth, by own earnings level and by race (panels a and b) or parental income (panels c and d). Lifetime earnings groups are divided into six bins—the bottom four quintiles and the top two deciles. Earnings bins are calculated at the population level. Error bars are at the 95% confidence interval.

Figure A.19: Simulating lifetime trajectories from shorter panels




Notes: This figure shows a schematic of the imputation model used to simulate lifetime trajectories of earnings, deferred compensation, and DC plan withdrawals for workers aged 25 to 65 from the shorter panels available to us for individual workers. We construct full lifetime trajectories by repeatedly matching individuals across overlapping age bins. For example, in 1. Persons A and B have similar earnings and job characteristics in the overlapping ages (27 and 28), so we append Person B's information at 29 and 30 to Person A and therefore add two additional years of data to the trajectory of Person A. We repeat this process at increasing ages (31-32, 33-34, ..., 65-66) to create a full lifetime path of earnings, employee and employer contributions to DC plans, and early withdrawals from ages 25 to 65.

Figure A.20: Portfolio shares and rate of return



Notes: This figure shows underlying parameterizations for portfolio composition and returns in the micro-simulation model described in Appendix D. The points of panel (a) show actual portfolio shares for Fidelity Freedom Funds at each age. We interpolate shares between these observations for each integer age, given by the lines in panel (a), although we assume exclusive investment in bonds at retirement. Panel (b) shows the real rate of return by age, which is determined by the portfolio composition and the associated returns of each component.

Figure A.21: Withdrawal Path



Notes: This figure shows the process for estimating withdrawals in retirement in the micro-simulation model described in Appendix D. For the purposes of illustration, we suppose an individual retires with wealth balance of \$100,000, which they draw down until their last year of life at age 90. The left axis corresponds to the green line, showing the wealth balance at the start of each period. The right axis corresponds to the blue line, showing the proportion of remaining wealth balance that is withdrawn each period. This process ensures constant withdrawals each period and a smooth draw-down of wealth in retirement.

Table A.1: How the sample size changes sequentially with the sample selection criteria for the Form 5500 sample

	Full ACS	Positive W2	Restricted sample	Link to Form 5500	Usable plans	Consistent plans	Final sample
Number of PIKs	18,470,000	14,010,000	12,490,000	3,326,000	2,118,000	1,606,000	1,582,000
Number of EINs	2,311,000	2,308,000	1,928,000	30,000	20,000	15,500	15,500
Number of Plans	—	—	—	5,000	3,500	2,600	2,600

Notes: This table presents how the number of PIKs, EINs, and plans (rounded for disclosure reasons) changes as we impose the sample selection criteria for the main sample we use in the paper: workers in the ACS that we can link to Form 5500 data. To arrive at each of these samples, we start with the “Full ACS”, which includes the most recent survey of all participants between the ages of 25-59.5 in the 2008-2017 waves of the ACS. For “Positive W2”, we then require that they have a strictly positive wages from their W2 information. For the “Restricted Sample”, we ensure W2 wages plus deferred compensation is at least \$8,000 (in nominal terms), no one filed more than 11 W2s that year, and their recorded deferred compensation complies with the caps for that year. Next, we merge on the Form 5500 plans, “Link to Form 5500” gives the plans (of our 6,201 coded plans), which we can link and which have an ACS respondent employed in them. “Usable plans” refers to workers at plans with straightforward policies that apply to all employees and have reasonable match formulas, e.g., match caps below 20% and match rates below 300%. To make sure it is a “Consistent Plan,” we restrict to plans for which the difference in the calculated ratios of employee contributions to total contributions obtained from the two sources is smaller than 15 p.p. Lastly, for the “Final Sample,” we require all potential mediating variables to be non-missing to ensure consistency across all regressions. For more details on the selection criterion, please see Section 3.3 and Appendix A.3.

Table A.2: Averages across samples

Variable	ACS sample		Form 5500 sample	
	Restricted	DC Offered	All	Parent
Age	41.6	41.8	41.2	30.0
W-2 total comp.	\$61,410	\$67,540	\$72,070	\$49,410
Spousal W-2 total comp.	\$9,902	\$9,988	\$9,912	\$9,753
Participation dummy	45.2%	57.1%	65.1%	59.2%
Employee contrib.	\$2,221	\$2,838	\$3,315	\$1,858
Employee contrib. rate (% of comp.)	2.7%	3.4%	3.8%	2.8%
Employee contrib. rate > 0	5.9%	5.9%	5.8%	4.7%
Withdrawal _{t+1} > \$1000 dummy	8.0%	9.0%	10.3%	8.1%
... past DC saver	11.8%	11.8%	13.6%	12.1%
Withdrawal amount (% of prev. comp.)	2.8%	3.0%	3.0%	1.7%
Withdrawal amount > \$1000	35.1%	33.0%	29.0%	20.8%
Number of unique individuals	12,490,000	9,708,000	1,582,000	435,000

Notes: This table compares the averages for the main variables in the paper across the different samples. From left to right, this table presents our (1) the restricted sample of ACS respondents linked to administrative tax records, (2) the subsample of ACS restricted that likely have access to a DC plan, (3) the Form 5500 sample, and (4) the Parent-Form 5500 sample. For more information about the different samples, please see Section 3.3, Appendix Section A.3.3, and the table notes from Table A.1. Spousal W-2 total compensation includes spouses claimed on Form 1040 who made \$0 in earnings (41% of our Form 5500 sample). As a reminder, the outflow measures are limited to employees below the age of 55 but the in-flow contribution measures include those from the ages of 55-59.5. All the dollar values reported in the table are deflated to base year 2017 using the CPI.

Table A.3: Administrative Data and SCF Comparison

Outcome	Statistic	Worker Samples		DC Offered Samples	
		ACS, restricted	SCF Worker	Form 5500 sample	SCF DC Offered
Age	Mean	41.6	42.6	41.2	42.5
Wage Compensation	Mean	\$61,410	\$70,797	\$72,070	\$86,814
DC Offered	Mean	78.2%	49.6%	100.0%	100.0%
DC Participation	Mean	45.2%	35.3%	65.1%	71.3%
Employee contribution rate	Mean	2.7%	2.3%	3.8%	4.8%
Count		12,490,000	8,517	1,582,000	4,097

Notes: This table presents means and observation counts for our administrative and SCF samples. The worker samples compare ACS and SCF workers, independent of whether or not their employer offers a DC plan. The DC Offered samples compare our Form 5500 sample with SCF workers who have access to a DC plan. All the dollar values reported in the table are deflated to base year 2017 using the CPI.

Table A.4: Parameter and variable definitions

Earnings, Wealth, Social Security		State Variables	
$e_{i,t}$	Earnings	i	Individual
α	Discount rate in retirement	$t \in \{25, \dots, 90\}$	Age
c_t^j	Consumption in retirement	$j \in \{DC, BK\}$	Type of savings vehicle
$aim e_i$	Average indexed monthly earnings		
e^{max}	Social Security taxable maximum		
δ_1	First PIA bend point	$T(\cdot, \cdot, \cdot)$	Federal income tax function
δ_2	Second PIA bend point	$\tau_{i,t}^{e,j}$	Taxes owed on earnings
$ss_{i,t}$	Annual Social Security benefits	$\tau_{i,t}^{ss,j}$	Taxes owed on Social Security Benefits
		$\tau_{i,t}^{c,j}$	Taxes owed on savings
		$\tau_{i,t}^{r,j}$	Taxes owed on returns
		$\tau_{i,t}^{w,j}$	Taxes owed on withdrawals
		$\hat{\tau}_{i,t}^{r,j}$	Hypothetical taxes owed on returns
	Wealth Flows		Lifetime Measures
$dc_{i,t}^{ee}$	Employee savings	$A_{i,t}^{DC}$	DC Wealth
$dc_{i,t}^{er}$	Employer savings	$SS_{i,t}$	Social Security Wealth
$w_{i,t}^j$	Savings account withdrawals	$A_{i,t}^{BR}$	Broad Retirement Wealth
$f_{i,t}^j$	Flow into retirement account	$A_{i,t}^{BR,BK}$	Broad Retirement Wealth brokerage WC
$B_{i,t}^j$	Wealth balance	$A_{i,t}^T$	DC tax subsidy
$B_{i,t}^{p,j}$	Principal part of wealth balance	$DC_{i,t}^{EE}$	Value of employee contributions
$B_{i,t}^{g,j}$	LTCG part of wealth balance	$DC_{i,t}^{ER}$	Value of employer contributions
$w_{i,t}^{k,j}$	LTCG portion of withdrawal	$A_{i,t}^{EE}$	Wealth attributable to employee
		$A_{i,t}^{ER}$	Employer subsidy
		LE_i	Value of lifetime income
	Rate of Return		Counterfactuals
ρ_t	Rate of return at age t	$A_{i,t}'^T$	Tax subsidy under tax CF
$r_{i,t}^{g,j}$	Return from unrealized capital gain	$A_{i,t}'^{DC}$	DC Wealth under tax CF
$r_{i,t}^{k,j}$	Return from LTCG distributions	$A_{i,t}'^{BR}$	Broad Retirement Wealth under tax CF
$r_{i,t}^{i,j}$	Return from interest distributions	dc^{*er}	Counterfactual employer match
s_t^k	Portion of assets invested in stocks	$A_{i,t}^{*DC}$	DC wealth under ER CF
s_t^b	Portion of assets invested in bonds	$A_{i,t}^{*BR}$	Broad Retirement Wealth under ER CF
s_t^m	Portion of assets invested in money	$A_{i,t}'^T$	Tax subsidy under combined (CB) CF
ρ^k	Real rate of return on stocks	$A_{i,t}^{\dagger DC}$	DC Wealth under combined (CB) CF
ρ^b	Real rate of return on bonds	$A_{i,t}^{\dagger BR}$	Broad Retirement Wealth under CB CF
ρ^m	Real rate of return on money		
χ^g	Share from unrealized capital gain		
χ^k	Share from LTCG distributions		
χ^i	Share from interest distributions		
$\hat{r}_{i,t}$	Implied post-tax rate of return		

Notes: This table gives selected variables that enter the micro-simulation model, outlined in Appendix D.

Table A.5: Parameter values and sources

Parameter	Value	Source
e^{max}	\$128,400	Social Security Administration (2023b)
δ_1	\$895	Social Security Administration (2023a)
δ_2	\$5,397	Social Security Administration (2023a)
ρ^k	0.0688	Jordà et al. (2019)
ρ^b	0.0253	Jordà et al. (2019)
ρ^m	0.0103	Jordà et al. (2019)
$\sigma_t^k, \sigma_t^b, \sigma_t^m$	Figure A.20a	Fidelity (2023)
χ^g	0.5	Yahoo Finance, Sialm and Zhang (2020)
χ^k	0.4	Yahoo Finance, Sialm and Zhang (2020)
χ^i	0.1	Yahoo Finance, Sialm and Zhang (2020)

Notes: This table gives the values for parameters used in the micro-simulation model, outlined in Appendix D.