

The Effect of Increasing Retirement Saving on Consumption, Balance Sheets, and Welfare*

Taha Choukhmane[†] Christopher Palmer[‡]

September 2025

Abstract

Does raising retirement contributions increase *net* wealth accumulation and improve lifetime welfare? The answer depends on how individuals finance the additional contributions: by cutting spending, reducing non-retirement savings, or increasing debt. We use newly linked deposit, credit, and pension data from a large UK financial institution to study a gradual increase in minimum retirement contributions from 2% to 8% of salary. On average, only one-third of the increase is funded through reduced spending (especially in discretionary categories like restaurants and leisure), with the rest of the decrease in take-home pay financed through higher borrowing and lower non-retirement saving. Low-liquidity individuals primarily cut spending, while those with high-liquidity shift existing savings across accounts with minimal impact on spending over two years. We show that these behavioral responses are crucial for welfare analysis in a sufficient statistics framework: well-targeted policies raise contributions among those who (i) exhibit less crowd out and (ii) are more prone to undersaving (e.g., due to present bias). We identify these groups in a lifecycle model by targeting the observed spending responses (to identify low crowd-out individuals) and the persistence of simultaneous retirement saving and credit card borrowing (to identify present bias). Using the estimated model, we find that standard financial incentives for retirement saving are poorly targeted, as they are disproportionately taken up by high-liquidity, time-consistent individuals who exhibit high crowd-out. Conversely, asset tests on tax incentives help screen out such individuals, and retirement income floors and annuities provide commitment for present-biased households.

*This paper was previously circulated under the title “How do consumers finance increased retirement savings?”. For helpful comments, we thank discussants Stefano DellaVigna, Eduard Boehm, Soojin Kim, Paolina Medina, Michaela Pagel, and Mingli Zhong; James Choi, Mariacristina De Nardi, Eric French, Erzo Luttmer, Brigitte Madrian, Olivia Mitchell, Gianluca Violante, Redis Zaliauskas, Jonathan Zinman; and seminar and workshop participants at Bocconi, Cowles, Dartmouth, EIEF, ETH Zurich, GWU, HBS, HEC Montreal, IFS, Michigan, MIT, NBER (Public Economics, Aging, and Innovative Data in HF), Nest, NTA, PUC Chile, Sciences Po, Stanford SITE, UGA, UT Austin, and Wharton. Xander de la Bruere, Raghav Saraogi, William Liu, and Brice Green provided helpful research assistance. This project received exempt status from MIT’s IRB. Palmer acknowledges funding from NSF CAREER Grant 1944138.

[†]MIT Sloan School of Management and NBER; tahac@mit.edu

[‡]MIT Sloan School of Management and NBER; cjpalm@mit.edu

1 Introduction

Governments around the world are heavily invested in promoting retirement savings, through mandatory contributions, financial incentives and, increasingly, behavioral nudges. Every OECD government provides financial incentives for private retirement savings, such as tax relief and matching subsidies, while automatic enrollment is now mandated in an increasing number of countries including the UK (for all private sector employees) and the US (for new 401(k) plans). These large-scale policies rest on a critical assumption: that increasing contributions to retirement accounts will meaningfully raise the resources available to retirees and improve households' lifetime welfare. Yet retirement accounts are just one component of households' balance sheets; whether retirement contributions raise *net* savings, and ultimately lifetime welfare, depends on how the other components of the balance sheet react. If consumers respond by cutting their spending, increasing retirement contributions can significantly raise total household savings, aligning with policymakers' oft-stated objective of increasing resources available in retirement. However, if consumers cope with lower take-home pay by increasing borrowing or reducing other forms of savings, higher retirement contributions may not increase overall savings and even leave some consumers more financially vulnerable with a higher debt burden.

In this paper, we ask: how do policies that increase retirement contributions affect saving, borrowing, consumption, and most importantly, lifetime welfare? To address this question, we use a new dataset linking (1) the pension contributions of UK workers affected by a gradual increase in minimum retirement contributions from 2% to 8% in 2018 and 2019, (2) their comprehensive spending patterns derived from transaction-level bank data, and (3) their deposits, savings, and credit account balances. This unique combination allows us to observe spending and balance sheet responses to retirement policy changes, something previous studies have been unable to capture in such detail using administrative data. We use this dataset to trace each pound moving into retirement accounts back to its source: whether from foregone consumption, lower non-retirement savings, or newly accumulated debt. We then use these measured elasticities and the comprehensive personal finance data to evaluate the distributional and welfare effects of alternative policies designed to promote retirement saving.

We divide our analysis into three parts. First, we build on the behavioral public finance framework of Allcott et al., 2025 initially applied to soda taxes to study the welfare effects of perhaps the most consequential behaviorally motivated policies: paternalistic interventions for retirement saving. We show that when policymakers are concerned about undersaving for retirement, the consumption response to a policy serves as a sufficient statistic for its welfare

impact. Intuitively, if an individual hasn't reduced her consumption, then the policy has failed to address her undersaving bias; it has merely reshuffled her balance sheet. Moreover, we show that the welfare impact of increasing retirement contributions depends on two key covariances: whether those induced to save more (i) actually reduce consumption more, and (ii) suffer from stronger undersaving biases, such as present bias. A budget-neutral policy that successfully raises retirement contributions can even reduce welfare if it primarily benefits patient individuals who simply shift their existing savings to capture more government subsidies. This illustrates why the average treatment effect on retirement account balances, which is often the metric of an intervention's success in the empirical literature, is a poor indicator of *welfare* impact. These theoretical results motivate our empirical focus on two key dimensions: documenting the consumption responses that determine welfare effects and identifying the heterogeneity in these responses that can help identify which policies are well targeted.

In the second part of the paper, we exploit two policy changes implemented as part of the UK national auto-enrollment policy, which raised the minimum combined employee and employer default contribution rate from 2% to 5% of salary in April 2018 and from 5% to 8% in April 2019. This policy changed both the default autoenrollment contribution rate for employees and the financial incentives for contributing: each step-up introduces a notch in the financial incentives of retirement contributions. Employees cannot participate in the retirement scheme at a lower contribution rate than the minimum default rate such that opting-out of participation leads to losing an increasingly large employer contribution, up from a minimum of 1% to 2% of salary in April 2018 and to 3% of salary after April 2019.

The increases in contribution-rate minimums were binding for some but not all employees and employers. This allows us to compare the behavior of consumers affected and not affected by the increase in the minimum retirement contribution. We find that for every £1 decrease in monthly take-home pay induced by the policy change, consumers respond by cutting their spending by £0.34 and financing the remainder with lower deposit account balances and higher credit card balances. Overall, relatively discretionary non-durable spending, such as restaurants and leisure, are the most elastic to the decrease in income net of pension contributions. We also find evidence of substantial treatment-effect heterogeneity. The most liquidity-constrained customers (i.e., those with lower deposit balances) cut their consumption the most. In contrast, those with significant liquid savings finance the increased pension contributions by running down their deposit balances.

In the third part of the paper, we move beyond these short-run responses to examine long-run welfare implications. We estimate a structural life-cycle model with heterogeneously present-biased agents. This allows us to simulate the long-run dynamics of our treatment ef-

fects, allow for alternative assumptions about the incidence of higher employer contributions, and estimate and contrast the welfare effects of various retirement policies. A key innovation of the paper is our structural identification of naïve present bias, which we estimate to affect roughly half of our sample. We target the share of individuals who consistently contribute to retirement accounts while revolving credit card debt, a pattern our model can match without time-inconsistent preferences only in the short run but not over a six-year horizon. Doing so, we show how extending the identification approach of Laibson et al. (2024) to panel data sharpens the identification of quasi-hyperbolic preferences in lifecycle consumption-saving models.

Using our estimated life-cycle model, we evaluate the welfare effects of different budget-neutral approaches to raising retirement savings. We derive three main insights for retirement policy design. First, widespread financial incentives often fail to significantly improve lifetime welfare: doubling employer match rates produces virtually no aggregate welfare gains due to poor targeting. Liquidity-constrained and present-biased workers who would benefit most rarely take up these financial incentives, while time-consistent workers who already save sufficiently capture most subsidies and exhibit maximum crowd-out. Second, behavioral nudges, such as auto-enrollment alone, similarly yield minimal impact. By contrast, the UK’s actual reform combining auto-enrollment with increased minimum contributions performs better, generating aggregate welfare gains equivalent to 0.5% higher retirement consumption (1.4% for present-biased individuals). This combined approach works better precisely because minimum contributions bind on present-biased low-savers while having little direct effect on high savers who already contribute above the minimum. To reduce the policy’s sizable fiscal and employer costs, we find that adding asset limits that restrict tax advantages for high-balance individuals not only reduces the policy cost but further improves targeting, as these limits screen out patient time-consistent individuals. Third, we find that the largest welfare gains come from policies addressing both the accumulation and decumulation phases of the life cycle. Raising government-provided income floors by £650 annually (a form of forced annuitization) generates aggregate welfare gains equivalent to 1.7% higher retirement consumption. Such policy serves dual roles as both insurance against longevity risk that benefit even time consistent workers and as a commitment device preventing present-biased retirees from rapidly depleting retirement assets in the first few years of retirement.

Our findings contribute to a long-standing literature on how retirement saving policy affects total wealth accumulation, where empirical evidence remains scarce.¹ A notable

¹Canonical papers documenting the effect of autoenrollment on within-retirement-plan savings include Madrian and Shea (2001) and Choi et al. (2004). See Choukhmane (2025) and Choi et al. (2024) for

exception is Chetty et al. (2014), who use comprehensive data from Denmark on wealth, savings, and income to study how households react to increased *employer* pension contributions and changes in mandated public-pension savings. However, the discontinuity in the mandated savings studied by Chetty et al. (2014), which, unlike employer contributions, did directly impact take-home pay, was relatively modest, around 50 USD annually. A limited literature has focused on whether *forced* savings—for example, from higher social security contributions—crowds out private savings (Feldstein, 1974; Attanasio and Brugiavini, 2003; Attanasio and Rohwedder, 2003; Jensen et al., 2025).² Recent research has studied the effects of autoenrollment on unsecured borrowing with mixed results.³ Beshears et al. (2022) find limited effects of autoenrollment on unsecured debt for civilian employees of the US military, and Beshears et al. (2024) finding an increase in unsecured debt and mortgages from autoenrollment in the UK.⁴ There is no direct evidence on how spending, borrowing, and deposits jointly change in response to retirement savings. We add to this existing empirical evidence in multiple ways: (i) we can directly observe spending and spending categories in our data, (ii) we study a policy that caused a moderate decrease in take-home pay for workers, and (iii) our life-cycle model allows us to quantify welfare and evaluate the targeting properties of common policy approaches to increasing retirement savings.

Answering these questions is usually complicated by two significant obstacles. First, it is rare for analysts to simultaneously observe data on a worker’s income, pension, spending, and liquid deposits to trace out the effects of increased pension contributions on a range of financial behaviors. Without such data allowing for a holistic view of consumer behavior, policymakers risk being unaware of significant side effects of pension regulations and unable to assess the targeting efficiency of retirement savings policy. As we discuss below, new data from a large UK bank uniquely permits such joint analysis. The second challenge is an empirical research strategy that permits characterizing the causal effect of increased contributions on other financial outcomes. We exploit changes in a national policy that changed the minimum default contributions but was binding only for employers and employees who originally contributed less than the new minimum default option.

The remainder of the paper is organized as follows. We present a sufficient statistics

reevaluations.

²We note that a priori, even full crowd out has advantages and disadvantages. On the one hand, saving in a tax-advantaged illiquid account has financial and behavioral benefits. On the other, the reduced liquidity from saving in a retirement account relative to a deposit account could make vulnerable households less resilient in the face of negative financial shocks.

³Outside of the retirement saving setting, Medina and Pagel (2025) also fail to reject that increased deposits from text message nudges are financed with unsecured debt.

⁴Workers who have selected into low-risk occupations may be more risk averse and potentially have a more conservative response to increased retirement savings among (Fuchs-Schündeln and Schündeln, 2005).

approach to analyzing welfare in our context in Section 2, highlighting the important covariances previously unknown to researchers that are essential for assessing the targeting properties and welfare effects of retirement policies. In section 3, we describe the administrative bank data that facilitate our analysis. Section 4 explains the institutional setting and introduces our analytical methodology. Section 5 presents our analysis of the effects of increased default contribution rates on take-home pay and behavioral responses, respectively. Section 6 presents the lifecycle model, Section 7 details the estimation procedure and results, and Section 8 applies the estimated model to analyze welfare. Section 9 concludes with a summary of our findings and policy implications.

2 Conceptual Framework

In this section, we outline a conceptual framework to motivate our approach and highlight the necessary ingredients for welfare analysis missing from the literature and motivate our empirical analysis. We adopt a behavioral public finance approach (Bernheim and Taubinsky, 2018), and in particular build on recent work by Allcott et al. (2025). We consider a two-period utility maximization problem with consumer i maximizes utility by choosing consumption c_i , retirement contributions ret_i and liquid savings or borrowing liq_i with second-period indirect utility $V(\cdot)$ discounted by an individual-specific discount factor β_i . Each asset types can have different properties (e.g., liquidity and taxes) and enter separately into an individual's second-period utility. Decisions are made taking as given the generosity γ of government retirement savings subsidies $s(ret_i, \pi_i, \gamma)$ and taxes $\tau_i(liq_i, \pi_i, \gamma)$:

$$\max_{c_i, ret_i, liq_i} u(c_i) + \beta_i V_i(ret_i, liq_i, \pi_i), \quad (1)$$

where $V(\cdot, \cdot, \cdot)$ is the indirect utility in period two that depends on how much an individual has saved in a retirement account and a liquid account and on π_i , which is a state variable capturing all individual characteristics relevant for determining their utility and choices (following Kolsrud et al. (2021)). Consumer optimization is subject to a budget constraint depending on the income y_i of each agent net the amount set aside in liquid and illiquid accounts, savings incentives received, and taxes paid

$$c_i = y_i - liq_i - ret_i + s(ret_i, \pi_i, \gamma) - \tau_i(liq_i, \pi_i, \gamma). \quad (2)$$

Normative preferences can differ from decision utility, and the social planner can be more patient than individuals. This could be due to a desire to correct individuals present bias

or a response to the externalities under-saving can create for social-safety programs (Moser and Olea de Souza e Silva (2019)). The degree of paternalism p_i determines the difference between an individual discount factor β_i and the social discount factor $\beta_i(1 + p_i)$. The social planner's objective is to set γ to maximize welfare, defined as:

$$W(\gamma) = \int_i \omega_i [u(c_i(\gamma)) + \beta_i(1 + p_i)V_i(\text{ret}_i(\gamma), \text{liq}_i(\gamma))] di + \mu \int_i (\tau(\text{liq}_i(\gamma), \gamma) - s(\text{ret}_i(\gamma), \gamma)) di \quad (3)$$

where ω_i is the welfare weight of individual i , $p_i \geq 0$ captures the degree to which agent i is too impatient relative to the social planner, and μ is the marginal value of government revenue. Note that in the planner's problem, each of the three objects chosen by agents—consumption, liquid savings, and retirement savings—explicitly depend on γ , which the agent takes as given. However, because agent decisions depend on preferences that are different from the social planner's, envelope conditions do not hold and second-order effects can be welfare relevant.

To consider the welfare effect of a small reform that increases the generosity of retirement savings incentives, we examine

$$\begin{aligned} \frac{dW(\gamma)}{d\gamma} = \int_i \omega_i \left\{ \underbrace{\frac{dc_i}{d\gamma} u'(c_i)}_{\text{cons. response}} + \beta_i(1 + p_i) \left[\underbrace{\frac{d\text{ret}_i}{d\gamma} V'_1}_{\text{ret. savings response}} + \underbrace{\frac{d\text{liq}_i}{d\gamma} V'_2}_{\text{crowd-out liq. savings}} \right] \right\} di \\ + \mu \int_i \left\{ \underbrace{\frac{d\tau_i(\gamma)}{d\gamma} - \frac{ds_i(\gamma)}{d\gamma}}_{\text{fiscal effect}} \right\} di \end{aligned}$$

where V'_1 and V'_2 are the derivatives of the indirect utility function $V(\cdot, \cdot, \cdot)$ with respect to its first and second arguments, respectively. The welfare effect of changing the incentives for retirement saving γ can be decomposed into four effects: the consumption response, the retirement savings response, the liquid savings response, and the effect on government revenue. This decomposition highlights the importance of characterizing the degree to which increased retirement savings are offset by decreases in non-retirement savings.

When individuals first-order conditions hold, the consumption response to a small change in retirement incentives is a sufficient statistic for welfare:

$$\frac{dW(\gamma)/d\gamma}{\mu} = \int_i \left\{ g_i p_i \left[\underbrace{\left(-\frac{dc_i}{d\gamma} \right)}_{\text{cons. response}} + \underbrace{\frac{ds_i}{d\gamma} - \frac{d\tau_i}{d\gamma}}_{\text{mechanical effect}} \right] \right\} di + \int_i (g_i - 1) \underbrace{\left[\frac{ds_i}{d\gamma} - \frac{d\tau_i}{d\gamma} \right]}_{\text{redistribution effect}} di$$

where $g_i = \omega_i u'(c_i)/\mu$ denotes the marginal social welfare weight on agent i . Abstracting away from any motives for redistribution and setting $g_i = 1$, the expression becomes:

$$\frac{dW(\gamma)/d\gamma}{\mu} = \int_i \left\{ p_i \left[\underbrace{\frac{dret_i}{d\gamma} \left(-\frac{dc_i}{dret_i} \right)}_{\text{change in behavior}} + \underbrace{\frac{ds_i}{d\gamma} - \frac{d\tau_i}{d\gamma}}_{\text{mechanical effect}} \right] \right\} di \quad (4)$$

where we replace the reduced-form effect of policy on consumption $dc_i/d\gamma$ with the product of the effect of retirement policy generosity on retirement savings $dret_i/d\gamma$ and the effect of increased retirement savings on consumption $dc_i/dret_i$. Readily apparent from (4) is that when the planner is not paternalistic (i.e., $p_i = 0$ for all i) then the envelope theorem holds and there is no effect on welfare of retirement savings policies.

What does this framework tell us about the mapping from the empirical objects traditionally estimated in the literature to welfare? Consistent with the theoretical framework of Allcott et al. (2025), the average treatment effect on retirement savings $E(dret_i/d\gamma)$ is only a partial guide for welfare. Instead, what determines the welfare impact of an intervention are the covariances between retirement savings responses ($\frac{dret_i}{d\gamma}$), the elasticity of consumption to increased retirement savings ($\frac{dc_i}{dret_i}$), and the degree of undersaving for retirement (p_i) and consumption responses to any increased savings (Allcott et al. (2025)).

When the level of bias is assumed to be homogeneous in the population ($p_i = p \forall i$), a well-targeted policy is one that that increases the retirement contributions of those with the larger consumption responses—such that $Cov\left(\frac{dret_i}{d\gamma}, -\frac{dc_i}{dret_i}\right) > 0$.

When the level of bias is heterogenous in the population, welfare is also determined by the covariances with both the level of bias and the mechanical effect of the policy. Even when there is no crowd-out of liquid saving by retirement saving such that all retirement savings are financed with consumption decreases ($dc_i/dret_i = -1$), a budget-neutral increase in the generosity of retirement savings incentives could decrease welfare if net subsidies accrue predominantly to the least biased individuals, such that $Cov(p_i, ds_i/d\gamma - d\tau_i/d\gamma) < 0$. This would be the case, for instance, if tax incentives for retirement savings are mostly taken up by the most patient agents, for whom the discount factor wedge p_i is smallest. Conversely, even when the consumption response to retirement savings is zero, i.e., $dc_i/dret_i = 0$ —as would be the case if people do not take-up retirement savings subsidies or finance it entirely by decreasing their liquid savings—a policy can be welfare improving when $Cov\left(p_i, \frac{ds_i}{d\gamma} - \frac{d\tau_i}{d\gamma}\right) > 0$.

Overall, this framework highlights that understanding the degree to which a given retirement savings policy increases welfare depends crucially on understanding both the con-

sumption response and how it covaries in the cross-section with individual characteristics. A well-targeted intervention increases contributions among those that are undersaving for retirement ($Cov(p_i, dret_i/d\gamma > 0)$) without crowding out other assets ($dc_i/dret_i \approx -1$). Conversely, a poorly targeted intervention increases contributions and transfers resources to those that aren't undersaving ($Cov(p_i, dret_i/d\gamma > 0)$) or merely shifts savings from a non-retirement account to a retirement account ($dc_i/dret_i \approx 0$). In this paper, we overcome the two main obstacles to estimating the necessary objects by developing new data that jointly observes c_i and ret_i and embeds policy variation that changes retirement savings incentives ($d\gamma$).

3 Data

We use proprietary data on bank customers who have pensions with the bank subsidiary, which provides us broad picture of the personal finances of these customers as well their demographics. For each individual, we have four types of data through the bank. First, we observe monthly aggregate pension contributions at the pension account level, which we can match to anonymized individuals using unique customer identifiers. Second, the data provides monthly aggregates of various categories of cash flows to and from current accounts and credit cards, including income, several categories of spending, debt payments, and bank transfers. Third, we observe the month-end balances of current accounts (the UK equivalent to checking accounts), savings accounts, and credit cards.⁵ Unlike the flow data, month-end balances are rounded to the nearest £100. For customer demographics, we observe only age and gender. The raw data is available through mid-December 2019, with a few consumers having data available as early as 2011.

The original raw datasets from the bank provide information on 614,000 unique individuals. We initially clean this data by excluding observations where net wage income is missing or zero, pension contributions are missing, or where the size of the contribution as a fraction of net wage income is not between 1.5% and 15%. We also impose a sample window of January 2016 to November 2019 to have a roughly similar number of observations per person. To construct our final analysis sample, we further filter the data to a smaller sample of bank customers to limit our analysis to individuals that use their accounts with our partner bank frequently and whom the data covers relatively continuously (i.e., without large gaps or missing information). The steps to filter to this analysis sample are detailed in the Appendix.

⁵These categories also include product variants within these account types, such as “added value accounts,” which are current accounts packaged together with other services such as insurance.

To visualize the granularity available using these data, Figure 1 plots the cash flows in and out of the average checking account of someone in the medium tercile of monthly net income. The green bar on the left is the average income for a middle tercile worker of £1,658. The first outflow bar represents net transfers. At -£126, the data indicate that the average middle tercile worker transfers £126 out of her account each month that cannot be otherwise categorized. These transfers could be transfers to accounts at other banks, checks or electronic bank transfers to landlords, friends, contractors, etc. For the purposes of simplifying this figure, we have combined several spending categories. The average middle-tercile worker spends £411 per month shopping at stores and supermarkets and £219 per month on the travel, leisure, and restaurant category. Transportation and utilities expenses average £315 per month and other categorized spending from several smaller categories totals an average of £174 per month. Cash withdrawals are small on average for this group—£9 per month. The average middle tercile worker makes credit card and other loan payments of £179 and £220, respectively. All told, these cash outflows account for almost the entirety of the cash inflow from the average worker’s paycheck. On average, checking account balances decrease by £1 per month and credit card balances increase by £4 per month.

One critical question that must be answered is whether we are comprehensively capturing the earning, spending, and saving behavior of the bank customers in our data. Specifically, since only accounts held with our partner bank are observed, we might be concerned about potential leakage from accounts held by households in our data at other financial institutions. Our choice of sample selection criteria helps address the income side of this question: by conditioning on (at least some) wages being paid into customers’ current accounts with our partner bank, we expect the data to almost entirely capture the incomes of the individuals in our sample. We would only get an incomplete and potentially biased picture of incomes for people who are partly paid by direct credit into accounts with our partner bank and also partly paid in cash or whose wages are also partially paid into accounts with other providers. Similarly, the sample selection criteria also help ensure that we are adequately capturing the spending behavior of our sample. We ensure that our results are robust to restricting the data sample to only individuals with a credit card with our partner bank, who use their observed accounts for food expenses and transact frequently using observed accounts. Such individuals are likely to mainly use their accounts with our partner bank when spending, especially because conditioning on wages being paid into their observed current accounts suggests their primary current accounts are at our partner bank. Finally, since total savings can be indirectly imputed from income and spending data, capturing the behavior of the latter two means we can adequately capture savings behavior even when we only directly observe spending from accounts in our data.

Another important question to contextualize conclusions derived from this data is the degree to which our partner bank’s customers are representative of the broader UK population. To assess whether spending by workers in our sample is representative of the average UK consumer, we calculate monthly budget expenditure shares using our spending data and compare these spending shares with nationally representative data from the UK Living Costs and Food Survey in Appendix Figure A1. Overall, we find that our sample closely matches the consumption profile of the average UK resident. Each point represents an expenditure share for the indicated category in a year between 2016-2018. For most categories, the plotted points are close to the 45-degree diagonal line, suggesting that budget shares are quite similar in the two datasets. The only outlier is the Other category, which is significantly below the 45-degree line, suggesting that this category is underrepresented in our data. We interpret this as evidence that, while our data captures many spending categories, there is some leakage from our partner bank’s customers paying for some things with money not initially held by our partner bank. Appendix Figure A2 benchmarks our data coverage and representativeness with nationally representative data from the Office for National Statistics on the share of consumers that have financial debt, property debt, and any debt. Overall, consumers in our data are 12 and 15 percentage points less likely to have financial and property debt, respectively, consistent with our partner bank’s clientele being higher-income on average.

Table 1 reports summary statistics for three categories of variables: income and spending, debt and account balances, and individual characteristics in panels I-III, respectively. The average worker in our data has a monthly net wage income of £2,300, implying an annual average wage income of £27,600, although there is substantial heterogeneity and skewness in the data. On average, workers in our data spend £1,400 per month out of their current accounts and £130 on their credit card with our partner bank. Looking across spending categories, the other spending category is the largest and most variable, with £490 per month on average. We also categorize observed spending into categories for consumer retail spending, utilities, supermarket purchases, restaurants, and leisure.

Panel II of Table 1 reports that the average current account balance is £4,200. However, even more so than any other variable, this average is significantly driven by high-balance outliers: the median current account balance is £1,600. From these accounts, workers pay an average of £360 and £140 per month of credit card payments and loan payments. The average credit card balance is £650, although more than half of consumers carry no balance in a typical month. In a given month, an average of 35% of individuals do not have a savings account with our provider. Finally, panel III shows that customers in the data have an average estimated age of 41 and 39% are estimated to be female.

Appendix Figure A3 documents the extent to which workers in our data use their credit card or current account for each spending category. Overall, workers in our data (and UK consumers more broadly) are much more likely to use their current account for their spending. This feature of UK payment modes (in contrast, e.g., to the average US consumer) supports our ability to make inference about the responsiveness of various spending categories to pension contributions. Consumer retail, leisure, and other spending have the highest share of spending observed on credit cards, each with around 10% of expended pounds being charged to a credit card on average.

4 Institutional Setting and Empirical Strategy

To identify how workers finance increased retirement savings, we leverage a natural experiment created by a legislative mandate of increased retirement contributions in the UK. After explaining the changes induced by the policy, we explain how we contrast workers affected and unaffected by these changes to identify the impact of increased savings on spending, borrowing, and saving.

The UK Pension Act of 2008 went into effect in 2012 and requires employers to automatically enroll their employees into a workplace pension scheme. Data from the Annual Survey of Hours and Earnings indicates that as of 2019, 77% of UK workers were participating in a workplace pension scheme. The Act initially set the minimum employee default contribution rate at 1% of qualifying earnings and the minimum employer contribution at 1%, although the minimum employee contributions include the tax relief from pension contributions being pre-tax.⁶ Each employer was assigned a staging date based on its number of employees, by which time employers were required to enroll all employees working in the UK aged between 22 and the state pension age in a workplace pension plan.⁷ As of 2017, the regulations applied to any worker earning over £10,000 a year. While employees can opt-out of their employer’s pension scheme at any time, the law requires employers to automatically reenroll all eligible opted-out employees every three years.

Subsequent revisions to the Pension Act of 2008 increased the default contribution levels. On 6 April 2018, the minimum default employer and employee contribution rates increased to 2% and 3%, respectively, such that the minimum total contribution rate increased from

⁶Qualifying earnings are defined as earnings over a lower threshold, subject to a maximum amount. In the 2021/22 tax year, the minimum contribution applies to earnings over £6,240 up to a cap of £50,270. Employers can optionally use total income instead of qualifying earnings to calculate contributions, in which case weakly lower default minimum contribution rates apply.

⁷Unlike in the US, the autoenrollment requirement applies to both new hires and any non-participating seasoned employees.

2% to 5%.⁸ On 6 April 2019, the default employer and employee contribution levels were increased to 3% and 5%, respectively, such that the minimum total contribution rate increased from 5% to 8%.⁹ Employees can choose to opt-out (but cannot contribute less than the combined employee and employer contribution minimum), and their employers are required to reenroll them at the minimum contribution limits every three years. Therefore this policy has stronger bite than a typical auto-enrollment default contribution nudge: in addition to changing the default option for contributions, the policy restricts the contribution choice set (i.e. employees cannot contribute below the new minimum) and changes the financial incentives for contributing (by raising employer contribution levels). Cribb and Emmerson (2016) find that UK autoenrollment substantially increased pension plan participation and contribution rates.

To develop a laboratory that facilitates learning about the causal effects of increased pension contributions on financial behavior, we divide pension plan participants in our data into four groups based on their contribution rates in March 2018. As the first step-up in default contribution rate was in April 2018, this allows us to compare workers who, before the change, were slated to be directly affected by the law change and those for whom the increase in default minimums should have no effect because their contribution rate was already quite high. Although workers with high and low pre-period total contribution rates undoubtedly differ on many dimensions, we can use pre-period data to characterize the differences between these groups at baseline and then study how the difference in these groups changes from this baseline.

We construct four groups, each based on a worker’s total contribution rate in March 2018, defined as the sum of employer and employee contributions each month divided by qualifying earnings for that worker in that month. The lowest and highest contribution rates are 1.5% and 15%.¹⁰ We set the lower bound of the lowest contribution rate to be at 1.5% to focus

⁸Optionally, employers are able to satisfy the legislation’s increased default contribution requirements by contributing some or all of the employee contribution, so long as the total minimum contribution requirements are set with the employer share being at least the indicated minimum.

⁹Technically, the Act mandates different minimum contributions depending on the definition of earnings used. Measured using either basic pay or qualifying earnings, the minimum employee contribution rate was 0.8% before April 2018, when it increased to 2.4% and then to 4% in April 2019. Measured against gross earnings, the minimum default employee contributions were also 0.8% and 2.4% in fiscal years 2017 and 2018 but increased to only 3.2% in fiscal year 2019. Employer minimum contributions increased from 2% to 3% in April 2018 and to 4% in April 2019 using basic pay. Measured using qualifying earnings or gross earnings, employer required minimum contributions increased from 1% to 2% in April 2018 and to 3% in April 2019. The tax relief at the source offered also depends on the income measure being used. For all income measures, the tax relief increased from 0.2% to 0.6% in April 2018. In April 2019, the tax relief increased to 1% for basic pay and qualifying earnings and to 0.8% for gross earnings. In our data, qualifying earnings seems to be the most common income definition used, based on the prevalence of employees contributing 4% in 2019 and employers contributing 3% in 2019.

¹⁰Increasing the upper bound to a 20% contribution rate would add an additional 10,000 individuals to

on employees who are participating in their employer’s pension scheme (non-participating employees are presumably less comparable to participating ones). Because UK law prescribes 1% as minimum employer contribution for most firms and schemes and 1% is generally the minimum employee contribution amount conditional on contributing anything, workers who appear to have less than 1.5% of income contributed to their pension likely have mismeasured income or aren’t currently participating in their pension. This 1.5%-15% interval is divided into four intervals, one for each group. The boundaries between the successive intervals are: 2.5%, 4.5%, and 7.5% because, at these boundaries, the numbers round to 3, 5, and 8. Similarly, 1.5% is chosen as the lower boundary since, at this boundary, the numbers round to 2. The four groups, 2, 3, 5, and 8, thus refer to the lowest integers in each interval and correspond to the various values that the minimum total, employee, and employer default contributions rates take on in recent UK history.

Figure 2 plots the distribution of contribution rates, highlighting the segment of contribution rates corresponding to each group. There is a large mass of workers with a total contribution rate at 2%, which is intuitive, given that before April 2018, the default minimum total contribution rate specified by law was 2%. The share of consumers with each contribution rate declines steadily from 2% onward. However, as Table 2 indicates, group 8 has the most individuals given the wide range of contribution rates represented. Our use of total contributions for group assignment is because of data limitations; our data is generally unable to differentiate employer and employee contributions. However, using data from the 20% of workers for whom the data does differentiate employer and employee contributions, Figure 3 shows that the vast majority of variation across groups defined by total contribution rates is driven by the employer contribution. This suggests a significant share of the variation in March 2018 group assignment is determined by the decision of employer to contribute above the minimum rather than wait for the policy to become binding. However, we caveat that employer contributions are not always predetermined by the employer and some plans, known as salary sacrifice schemes, allow employees to reduce their salary in exchange of a larger employer contribution.

4.1 Comparability of Contribution-Rate Groups

To further examine the comparability of these groups, Table 2 reports summary statistics for each contribution rate group. Using the group definitions described above, there are roughly 28,000 workers in the 2% group, 21,000 in each of the 3-4% and the 5-7% groups and 36,000 in the 8% group. The 2% contribution rate group has significantly lower March 2018 net wage

the sample—only a 1.6% increase in sample size at the potential expense of reduced comparability across contribution-rate groups.

income and spending, with roughly £400 per month less in income and £400 per month less in spending than the other groups. Combining lower average income with lower contribution rates, the 2% contribution rate group also has the lowest pension contributions at £42 per month compared to £85, £153, and £271 for the 3-4%, 5-7% and 8% groups, respectively.

Figure 4 summarizes the logic behind studying these groups by plotting the median monthly contribution amount for each of the four groups. The vertical lines represent the two policy-mandated increases in default contribution limits. Our methodology is to characterize the difference between the groups before the first policy change and then use the dynamics of how this differentially changed as a result of the policy to understand the causal effects of the policy. We then extend this same strategy to other financial outcomes.

As evident in Table 2, there are significant baseline differences prior to April 2018 in the contribution amount for each group, with contribution amounts tending to increase for successively higher rate groups. We also note that the highest contribution rate group has a slight upward time trend to its median contributions compared with the relatively flat time trends for the other groups. As expected, the lower contribution rate groups have their median pension contributions increased significantly by the increase in default minimum contribution rates in both periods. For example, the 2% contribution rate group’s median contribution amount more than doubles after the April 2018 increase from £42 per month to over £100 per month. The 3-4% group was also affected by the both policy changes, as expected given that after the first increase in contribution rates, their total contribution rate was less than the 5% mandated to be the new default minimum rate.

In contrast to the lower two contribution rate groups, the 5-7% contribution rate group was much less affected by the April 2018 increase because their total contribution rate was already over the minimum. However, 5-7% was below the required minimum 8% total contribution rate after the April 2019 increase, and we see a noticeable increase in contribution amounts for this group after the April 2019 policy change. The highest contribution rate group was entirely above the required minimum after the 2018 increase and mostly above the required minimum after the April 2019 increase, such that we see little deviation from trend for the highest contribution rate group.

Using the Sun and Abraham (2021) event-study methodology described below in section 4.2, we plot estimates of a retirement contributions event study in the left-hand panel of Figure 5. By normalizing March 2018 to zero, the event-study coefficients estimate the change in retirement contributions for workers in treated contribution-rate groups relative to untreated workers in March 2018 who were already contributing above the minimum. The graph shows that—with no discernible pre-trend or anticipation—immediately after the scheduled increase in required default minimum contributions, the average affected worker’s

retirement savings went up by £30. The subsequent trend shows that contributions for affected workers stayed about £30 higher than March 2018 levels until April 2019, when default minimum contributions again increased as legislated and contributions increased by another £30-40 per month. The right-hand panel of Figure 5 reports effects on cumulative contributions. By the end of our sample period, treated workers have increased their contributions by approximately £1,200 more than control-group workers. Lastly, we use the life-cycle model in section 6 to show that in a world where there is no selection into contribution-rate groups, we would anticipate finding very similar effects.

4.2 Estimation Strategy

To characterize how workers finance increased retirement savings, we adopt two complementary estimation strategies. First, we estimate the reduced-form effects of the change in retirement savings policy—analogue to $dc_i/d\gamma$ —using the Dynamic Event Study difference-in-differences approach of Sun and Abraham (2021). For a given outcome, we estimate

$$Outcome_{it} = \beta \sum_{\ell} \mu_{\ell} 1(t - PolicyDate_{k(i)} = \ell) + \alpha_i + \delta_t + \varepsilon_{it} \quad (5)$$

where i denotes each worker, t denotes calendar months between 2016 and 2019, $1(\cdot)$ is an indicator function, and α_i and δ_t are individual and month fixed effects, respectively. The event study coefficients μ_{ℓ} capture how the average outcome evolved relative to the month $PolicyDate_{k(i)}$ that contribution group k had its minimum pension contributions increased by the Pension Act of 2008. For example, the indicator function in (5) is turned on to estimate μ_2 for groups 2 and 3 in June 2018 (two months after their first increase in minimum default contribution rates) and for group 5 in June 2019 (two months after the second increase in minimum default contribution rate, the first one that would have affected group 5). For group 8, which already had sufficiently high contributions, the indicator function in (5) is always off and these control-group observations help to identify δ_t .

Second, to use this variation in estimation of the effect of an increase in retirement contributions on a given outcome, we estimate

$$Outcome_{it} = \beta \cdot PensionContributions_{it} + \alpha_i + \delta_t + \varepsilon_{it}$$

$$PensionContributions_{it} = \sum_{s \in \{1,2\}} \sum_{k \in \{2,3,5\}} \pi_{ks} Group_i^k \times Post_t^s + \psi_i + \phi_t + v_{it}$$

by two-stage least squares. Using contribution rates as of March 2018 defined as above, we define $Group_i^k$ is an indicator for whether borrower i was in contribution rate k in March

2018. There are two post periods, and $Post_t^1$ and $Post_t^2$ are indicators for whether month t fell on or after the two policy change months in April of 2018 and 2019. The outcome in the second stage could be a cash flow such as spending or a deposit or debt balance. When the outcome in the second stage is a worker’s take-home pay, $\hat{\beta}$ estimates the fraction of each £1 of additional pension contributions induced by the UK Pensions Act escalations that was a contribution of an employee’s wages. For spending category outcomes, we will often interpret effects by dividing by $\hat{\beta}_{income}$ to characterize the share of each £1 of an employee’s take-home pay difference that was financed with a change in the indicated category of spending.

5 Effects of Increased Pension Contributions

The first direct effect of increasing pension contributions is to reduce employees’ take-home pay. However measuring this effect is not straightforward; the increase in pension contributions in April 2018 and 2019 (shown in Figure 4) reflects both an increase in employer and employee contributions, but only the increase in employee contributions reduces take-home pay initially. We implement a regression analysis to measure the extent to which the mandated increase in pension contributions reduced take-home-pay. Our analysis compares the change in take-home pay along two dimensions: (i) comparing individuals’ outcomes to their behavior prior to the policy change allows us to control for all time-invariant individuals characteristics (such as education, gender, attitudes toward saving, etc.) and (ii) comparing groups with low initial contribution rates (who are affected by the contribution rate step-up) and groups with higher initial contribution rates (who are unaffected) allows us to control for trends that affect all individuals at the same time (such as overall growth in income or seasonal variation in compensation).

How should we expect the increases in default contribution rates for employers and employees mandated by the UK Pension Act to affect workers’ take-home pay? Using the qualifying earnings definition of income, total contributions increased by 3 percentage points in April 2018. This translates into a 1.6% increase in employee contributions when the tax-relief is done at source (and added directly to the pension pot) and a 2% increase in employee contributions otherwise. Therefore, we might expect between take-home pay to decrease by £0.53 or £0.67 (depending on the tax-relief method) for every £1 increase in contributions to an employee’s pension. The remainder being financed by employer contributions and tax relief. We should expect the reduction in take-home pay to be larger if the incidence of increased employer contributions falls more on the affected employees relative to our control group. Table 3 displays the results of the regression analysis on the effect of increased pension contributions on take-home pay (panel I) and spending (panel II). The results are expressed

in terms of the effect of increasing pension contributions by £1. Consistent with our expectations above, we indeed find in the data that increasing pension contributions by £1 reduced take-home-pay by around £0.67 (panel I). The -£0.23 coefficient on total spending in the first row of panel II implies that a third of the drop in take-home pay is financed by a reduction in total spending. Figure 7 plots the category-specific spending effects. The reduction in spending is particularly significant for categories that capture discretionary spending, such as restaurant, leisure, and retail spending. Leisure spending includes spending on sports, hobbies, gambling, and entertainment. This suggests that it is easier for individuals to adjust their discretionary spending in response to a reduction in take-home pay as opposed other categories such as housing and utilities, for which we do not observe significant effects.

In addition to reducing their spending, consumers also respond to the drop in take-home pay by reducing their monthly credit card payments by £0.22 (Panel III) and their net checking account deposits by £0.34 (although this estimate is very imprecise). These changes in monthly flows translate into significant reduction in current account balances and an increase in credit-card balances. Finally, to characterize the timing of these balance changes, Figure 6 plots event study estimates of the cumulative change in checking account balances (panel I), credit-card balances (panel II), and non-mortgage, non-credit-card debt balances (panel III). For the average worker, we see a £100 relative decline in checking-account balances soon after the first increase in contribution rates, although our ability to make strong statements about effect timing is limited somewhat by the precision of our estimates. In contrast to the more immediate effects of the policy on retirement savings seen in Figure 5, the effects on credit-card debt and non-mortgage debt are more gradual, suggesting that workers draw on these sources of financing to cope with decreased take-home pay gradually over time. By the end of our sample 19 months after the first nationwide increase in minimum contribution limits, treated workers have approximately £150 less in their checking accounts, £100 more credit-card debt, and £150 more non-mortgage debt than would be predicted using the control group and baseline differences between treatment and control.

Figure 8 illustrates a combination of all of these effects to decompose how the average worker in our sample financed increased pension contributions. Total contributions to pensions increased by £1,247 for the average worker, with £816 (65%) of that increase coming from employee contributions and a corresponding decrease in take-home pay. Approximately 40% of that decrease in take-home pay was financed through lower spending, 19% through lower credit card balances, and 11% and 21% through higher credit card and loan balances, respectively. This leaves £79 (10%) of the cumulative decrease in take-home pay unaccounted for, potentially resulting from transfers from financing sources outside our data provider.

5.1 Heterogeneous Effects by Liquidity Status

Our results suggest that, on average, only 40% of the increase in retirement contributions was financed by reducing spending. There is, however, substantial heterogeneity: those with limited liquid savings primarily reduce spending, while those with substantial liquid savings shift existing savings from outside to inside retirement accounts, with minimal impact on spending. Figure 9 reports total spending effects by deposit tercile, showing that the effects are quite large for workers in the bottom third of deposit balances and statistically insignificant for workers in the top tercile of deposit balances. This reflects the fact that customers with large deposit balances can offset a reduction in take-home-pay by reducing their deposits, whereas individuals with low or no deposits are more constrained and must either reduce their consumption or resort to expensive borrowing. But even individuals with larger initial liquid balances cannot run down checking account balances indefinitely, therefore we use a quantitative model—calibrated to match the observed short-term reductions in account balances—to estimate the long run spending responses and draw implications for targeting retirement policies.

6 Life-Cycle Model

This section develops a life-cycle consumption-saving model to examine policy incidence, simulate long-run spending responses, and assess the welfare impact of alternative policies. Agents choose consumption, retirement account contributions, and liquid savings while facing age-varying income, employment, and mortality risks. The model, calibrated to the UK tax and pension system, extends Choukhmane (2025) by incorporating heterogeneity in both present bias and employer characteristics (with firm-specific earnings and matching formulas), a different specification of contribution adjustment costs, and captures the evolution of UK retirement policies from 2012 to 2019.

6.1 Life cycle model

Demographics. Time is discrete, and each period t corresponds to one calendar quarter. Individuals start their working life at age a_0 and transition into retirement at age a_r . They face age-dependent mortality risk m_a and die with certainty at age a_T .

Preferences. Individuals derive flow utility from consumption c , with a constant elasticity of intertemporal substitution σ such that $u_a(c) = n_a \frac{\left(\frac{c}{n_a}\right)^{1-\frac{1}{\sigma}} - 1}{1-\frac{1}{\sigma}}$, where n_a serves as

an equivalence scale that adjusts for changes in average household size over the lifecycle. Individuals can exhibit either exponential or quasi-hyperbolic time preferences, characterized by a sequence of intertemporal discount factors $\{1, \beta_i \delta, \beta_i \delta^2, \dots\}$. Exponential discounting corresponds to $\beta_i = 1$ and present bias to $\beta_i \in (0, 1)$.

Labor market. During their working lives, individuals face earnings and employment risk. The employment earnings depends on a deterministic component that depends on age $f(a)$, a firm specific component $\bar{y}_t^e$, and of a stochastic component θ , such that $y_t^e = \bar{y}_t^e \ln(f(a_t) + \theta_t)$. After age a_0 , the stochastic component evolves according to an AR(1) process $\theta_t = \rho \theta_{t-1} + \epsilon_t$, where $\epsilon_t \sim \mathcal{N}(0, \sigma_\epsilon^2)$ corresponds to persistent earnings innovations. Employer changes are modeled separately from ongoing employment because pension benefits typically reset when workers switch jobs. Individuals transition to a new job either directly or after an unemployment spell. The unemployment and job-to-job transition probabilities, respectively π^{EU} and π^{JJ} , depend on the worker's age a , productivity θ , and tenure ten . A job-to-job transition triggers a new productivity draw governed by $\theta_t = \rho \theta_{t-1} + \zeta_t^{JJ}$, where $\zeta_t^{JJ} \sim \mathcal{N}(\mu^{JJ}, \sigma_\xi^2)$ is a shock with a positive mean μ^{JJ} . This captures the earnings boost typically observed when workers transition from job to job. Unemployment spells lead to lower average earnings when an unemployed individual returns to employment, with an age-dependent probability π^{UE} . Productivity at the new job, following an unemployment spell, is governed by $\theta_t = \rho \theta_{t-1} + \zeta_t^{UE}$, where $\zeta_t^{UE} \sim \mathcal{N}(\mu^{UE}, \sigma_\xi^2)$ is a shock with a negative mean μ^{UE} .

Liquid assets. Individuals begin with zero wealth. In each subsequent period, they decide how much to hold in liquid assets, denoted l_{t+1} . When liquid asset holdings are positive, they accrue a risk-free after-tax rate of return $R^{liq} = 1 + (1 - \tau^k) r$, where τ^k is the tax rate on asset returns. Agents are allowed to access unsecured credit by holding negative liquid wealth, up to a lower bound $\underline{l}_a$. Borrowed amounts incur a higher interest rate, $R^{liq} = 1 + r^{CC}$.

Retirement accounts. Employees can contribute a share $s_t \geq 0$ of their wage to an employer-sponsored retirement account when employed, and withdraw a share $s_t \leq 0$ of their balance of retirement wealth DC_t when retired. Employers, indexed by e , make an employer contribution ec_t^e that depends on the employee contribution s_t . Employers also vary in terms of the tax treatment of retirement contributions, with the variable $relief_e$ an indicator equal to 1 if the employers allows for tax relief at source. Assets inside the retirement account grow tax-free with a rate of return $R^{dc} = 1 + r$.

Adjustment cost. In the model, individuals behave as though changing their contribution rate away from the default contribution rate entails a utility cost $\tilde{k}$. This cost can capture the effect of various decision-making frictions, such as the cognitive burden of

paying attention, and the effort required to identify the optimal choice. The default contribution rate is chosen by the employer in the first period of employment, and equal to one's last contribution rate later in tenure $d_t = s_{t-1}$. We also define default take home pay y_t^d as the agent's after-tax earnings when making the default retirement contribution $y_t^d = (1 - d_t) y_t - tax_t(d_t, y_t)$. To express the utility cost in terms of unit of consumption, we set $\tilde{k} = u_a(y_t^d + k) - u_a(y_t^d)$. This implies that the adjustment cost is equal to the utility cost of given up k units of consumption when consuming the default take-home pay.

Aggregate policy environment. We divide time into four periods, captured by an aggregate policy state $Pol_t = \{1, 2, 3, 4\}$. The first state, corresponds to the policy environment in the UK prior to April 2012 during which there was no minimum default contribution rate. The second (third) state, corresponds to the period between the second quarter of 2012 (2018) and the first quarter of 2018 (2019) when the minimum default contribution rate was 2% (5%) of salary. Finally, the fourth state corresponds to the policy environment after the second quarter of 2019 and is characterized by a minimum default contribution rate of 8%. Agents in the model do not anticipate the transition across aggregate policy states.

Taxes and benefits. Agents face a nonlinear income tax schedule $tax_t(s_t, y_t)$ on their labor earnings. Contributions to DC accounts reduce taxable income in working life, returns are not subject to taxation, but a portion of DC withdrawals in retirement is subject to income taxation. This captures the fact that, in the UK, can withdraw 25% of their retirement balance tax-free. During periods of unemployment, agents receive unemployment benefits ui_t . After retirement, they receive a constant government pension benefit $pens_t$.

6.2 Dynamic optimization problem

We present the recursive formulation of the household's dynamic optimization problem. The state space is represented by the vector X_t , which consists of 8 state variables: age (a), employment status (emp), job tenure (ten_t), employer DC plan type (e), labor productivity (θ), liquid assets (l), DC wealth (dc), and the default contribution rate (d). Households face uncertainty about survival (m_a), earnings shocks (θ), employment transitions (emp), and future employer plan types (e) conditional on job changes. The joint distribution over earnings, employment, and employer types is denoted $F(\theta_t, emp_t, e_t)$. Each period, we compute two versions of the value function: $V_t^{PB}(X_t)$, which incorporates present-biased preferences, and $V_t(X_t)$, which corresponds to time-consistent preferences that individuals (naïvely) assume govern their future choices. Because no closed-form solution exists, the problem must be solved numerically.

Individual problem in retirement (after age 65). Each period, retirees choose how

much to keep in liquid assets and how much to draw from their defined contribution (DC) account, subject to an uncertain mortality risk .

$$\begin{aligned}
V_t^{PB}(X_t) &= \max_{s_t \leq 0, l_{t+1} \geq l_a} u_a(c_t) + \beta \cdot \delta (1 - m_a) V_{t+1}(X_{t+1}) \\
s.t. \ l_{t+1} &= R^{liq}(l_{t+1}) [pens_t + (-s_t) \times dc_t + l_t - c_t - tax_t(s_t, pens_t)] \\
dc_{t+1} &= R^{dc}((1 - (-s_t)) dc_t)
\end{aligned}$$

Individual problem when employed. Workers may be employed after a job-to-job transition, after moving from unemployment into work, or by remaining continuously employed. They face both income and employment risk, and each period they must determine how much to allocate to liquid savings or unsecured borrowing, as well as how much to contribute to the DC account.

The decision process unfolds in two stages. In the first stage, for each feasible DC contribution choice $s \in S$, the worker solves for the optimal level of liquid asset holdings (or unsecured debt) consistent with the intertemporal budget constraint 6. The value associated with selecting contribution rate s is therefore:

$$\begin{aligned}
V_t^s(X_t) &= \max_{l_{t+1} \geq l_a} u_a(c_t) - \mathbb{1}_{(s_t \neq d_t)} \tilde{k} + \beta \cdot \delta (1 - m_a) \int V_{t+1}(X_{t+1}) dF(\theta_t, emp_t, e_t) \\
s.t. \ l_{t+1} &= R^{liq}(l_{t+1}) [(1 - s_t) y_t^e + l_t - c_t - tax_t(s_t, y_t^e, relieve)]
\end{aligned} \tag{6}$$

In the second stage, the worker chooses a contribution rate s to the DC account:

$$\begin{aligned}
V_t^{PB}(X_t) &= \max_{s_t \in S} \{V_t^s(X_{t+1})\} \\
s.t. \ dc_{t+1} &= R^{dc}(dc_t + s_t y_t^e + ec_t^e)
\end{aligned}$$

Individual problem when unemployed. Following an unemployment shock, individuals receive a constant unemployment benefit and face uncertainty regarding re-employment prospects, future wages, and the type of employer they will next join. They choose liquid asset holdings or unsecured borrowing, but withdrawals from retirement accounts are not permitted:

$$\begin{aligned}
V_t^{PB}(X_t) &= \max_{l_{t+1} \geq l_a} u_a(c_t) + \beta \cdot \delta (1 - m_a) \int V_{t+1}(X_{t+1}) dF(\theta_t, emp_t, e_t) \\
s.t. \ l_{t+1} &= R^{liq}(l_{t+1}) [wi_t + l_t - c_t - tax_t(wi_t)]
\end{aligned}$$

7 Estimation and Model Fit

7.1 Method of simulated moments

The estimation proceeds in two steps. First, we fix a subset of parameters outside the model, drawing their values either directly from the data or from established results in the literature. In the second step, we apply the method of simulated moments to recover five preference parameters of the model: the intertemporal discount factor (δ), the elasticity of intertemporal substitution (σ), the contribution adjustment cost (k), the share of present biased individuals (π^β), as well as the mean and shape parameters α_1 and α_2 that characterize the distribution of present bias, which is modeled as a Beta distribution ($\beta \sim \text{Beta}(\alpha_1, \alpha_2)$).

7.2 Identification of preference parameters

We now discuss the sources of variation in the data that help identify each preference parameter in the second stage.

Naïve present bias. We identify the three parameters governing the distribution of quasi-hyperbolic discounting from the persistence of simultaneous retirement saving and credit card borrowing. Although this behavior can be rationalized in the short run with time-consistent preferences, its continuation over a six-year horizon can help identify present bias in the model. Intuitively, a worker may find it optimal to revolve high-interest rate credit card debt while contributing enough to the retirement account to capture a large one-time employer contribution. However, if the debt is expected to be rolled over for several years, the compounding interest cost quickly outweighs the one-time employer match, making credit card repayment financially preferable. Naïve present-biased individuals mistakenly believe they will soon pay off their balances but instead continue revolving debt for extended periods. This approach extends the Laibson et al. (2024) strategy of identifying present bias to panel data, taking advantage of the longitudinal persistence of this pattern to achieve sharper identification. Specifically, we target the share of individuals who simultaneously contribute to retirement accounts while revolving credit card debt and the probability of maintaining this behavior throughout the three years before and after.

Time preferences. We estimate two preference parameters governing saving behavior: the intertemporal elasticity of substitution (σ) and the discount factor (δ). Since both affect saving decisions, identification is difficult using lifecycle profiles alone. We exploit the large changes in saving incentives introduced in April 2018 and 2019 to identify σ , which governs the responsiveness of saving to interest rate changes, separately from δ , which determines the overall saving level. Specifically, we match the distribution of contribution rates in the

last quarters of 2017, 2018, and 2019, when the minimum employer contribution rose from 1% to 2% to 3%, providing sharp variation in financial incentives to contribute to pensions.

Adjustment cost. We identify the size of the contribution adjustment cost (k) by targeting the bunching of contributions at the auto-enrollment default. As the adjustment cost value grows larger, we should expect more and more workers to remain at the auto-enrollment default contribution rate. In practice, we target the differences in participation rate and in bunching at the auto-enrollment default by age and by income, both before and after the initial roll-out of the national auto-enrollment policy.

7.3 First-stage parameters

We use data from the Annual Survey of Hours and Earnings (ASHE) to estimate the parameters of the labor earnings process and the heterogeneity in employer retirement contribution formulas. The remaining first-stage parameters (including demographic variables, the parameters of the tax and benefits system and the rate of return on assets) are estimated directly from the data or calibrated with reference to other papers. The estimation of the earnings process, treatment of the data, and calibration of other first-stage parameters are discussed in Appendix B.1.

7.4 Second-stage parameter estimation

The six second-stage parameters ($\delta, \sigma, k, \pi^\beta, \alpha_1, \alpha_2$) are jointly estimated with the method of simulated moments. The parameters are chosen to minimize the distance between the model-simulated and empirical saving and borrowing patterns under different policy regimes.

Estimation samples. We use two sources of data for estimating the model. Whenever possible, we rely on data from the nationally representative ASHE panel, restricting our attention to employees and firms subject to the policy rollout between 2012 and 2016: private sector workers earning at least £10,000 in annual gross pay, whose employers offers a DC plan and had at least 30 employees and were staged into auto-enrollment during this period. To capture simultaneous retirement saving and credit card borrowing, we use our baseline bank sample described previously.

Estimation moments. We use 4 sets of moment conditions (140 in total). First, we target the frequency distribution of contribution rates (from 0% to 15%) in April 2017, 2018, and 2019 from ASHE (48 moments in total). Second, we target both the participation rate and the fraction of workers contributing at the 2% auto-enrollment default under opt-in and auto-enrollment regimes, across 8 age groups and 8 income groups in ASHE (64 moment conditions). Third, we target the share of retirement savers (those with positive

contributions) who carry credit card debt across 8 age bins in the bank data (8 moments). Finally, among those observed both revolving debt and contributing to their retirement account in May 2015, we target the probability of maintaining this behavior in each quarter during the 36 months before and after that date (18 moments).

Weighting matrix. Since all the target moments have the same unit (percentages of participants), we use the identity matrix as our baseline weighting matrix in the estimation. We avoid the optimal weighting matrix due to its poor small-sample properties (Altonji and Segal, 1996).

7.5 Results and model fit

Estimation results. We report preference parameter estimates in Table 4. We estimate that 45% of our sample are exponential discounters (with $\beta = 1$), while the remaining 55% have present-biased preferences following a beta distribution with mean 0.547 and standard deviation 0.042. Our estimate of the average β aligns with existing estimates of quasi-hyperbolic discounting in lifecycle models (Laibson et al. (2024)). We estimate an elasticity of intertemporal substitution of approximately 1 and a quarterly long-term discount factor $\delta = 0.998$ (or $\delta^4 = 0.992$ annually). We estimate an adjustment cost of £35 (standard error: £31) for changing contribution rates, reflecting that in models with present bias, even small adjustment costs can generate significant inertia.

Model fit. We report the fit between our model simulations and empirical moments in Figures 10, 11, and 12. Figure 10 illustrates how the simultaneous occurrence of retirement saving and credit card borrowing. In a calibration without present bias, the model generates such behavior only among younger individuals (left panel), and this behavior is short-lived, largely driven by transitory liquidity needs (right panel). Conversely, a model with only present biased individuals generate too much persistence in this behavior, especially for lower values of β (left panel). Naïve present-biased individuals mistakenly believe they will soon pay off their credit card balance thus making capturing the employer retirement contribution more financially advantageous than immediately repaying high interest rate debt. In contrast, to these two calibrations, the estimated model with 55% of present biased individuals fits the data substantially better. At the estimated values reported in Table 4, the model fits both the lifecycle profile and persistence of simultaneous credit card borrowing and retirement saving (left and middle panels).

Figure 11, shows that the model broadly reproduces the distribution of contribution rates before and after the step-up in minimum contributions from 2017 to 2019. However, the model tends to overestimate opt-out rates in 2018 and 2019 after minimum rates were in-

creased. The level and shift in the distribution following large changes in financial incentives for saving helps identify the model time preference parameters.

Finally, Figure 12 validates the model’s ability to capture heterogeneous responses to auto-enrollment across age and income groups. Here we take advantage of the phased roll-out of the UK national auto-enrollment policy by employer size between 2012 and 2016. Panels (a) and (b), show that the model matches the dramatic increase in participation rates following after a employer’s auto-enrollment staging date for larger and lower income workers, while reproducing the more modest effects for older, higher-income workers. Panel (c) shows the model generates patterns of bunching at the 2% auto-enrollment default declining with both age and income, consistent with our assumption of modest fixed adjustment costs.

8 Welfare Analysis

8.1 Welfare framework

8.1.1 Social preferences

We allow social preferences to differ from individual preferences along three dimensions. First, the social planner adopts a long-term criterion and does not treat present biased preferences as welfare relevant $\beta^{SP} = 1$. Second, we capture social preferences for redistribution with exogenous Pareto weights φ following Saez (2002). If φ is equal across individuals, the social planner has no redistributive motive beyond declining marginal utility of consumption. Finally, following Goldin and Reck (2022), we assume that the policymaker considers that only a fraction π of the contribution adjustment cost is welfare-relevant. $\pi = 1$ corresponds to the case where the adjustment cost cost is fully normative (for example, if making a decision involves real costs, such as the fee of a financial adviser) and $\pi = 0$ if the adjustment cost cost is not welfare-relevant (for instance, if inertia reflects employees’ mistakes from the policymaker’s perspective).¹¹

The policymaker evaluates the welfare impact of different policies d^{SP} by comparing the respective values $V^{SP}(d^{SP})$, the *normative* expected lifetime utility from the perspective of period 0 (i.e. before any uncertainty is realized):¹²

$$V^{SP}(d^{SP}) = \mathbb{E}_0 \left[\sum_{t=1}^A (\delta^{SP})^t \left(\prod_{k=1}^{t-1} (1 - m_k) \right) \varphi_t u_t(c_t) - \mathbb{1}_{(s_t \neq d_t)} \pi \tilde{k} \right]$$

¹¹Goldin and Reck (2022) discusses the empirical plausibility of different values of π .

¹²Note that because preferences are homogeneous, lifetime utility is ex-ante identical across individuals.

subject to individuals' optimization defined in Section 6.2.

$$\{c_t, s_t, l_t\} = \underset{s_t, l_t}{argmax} V_t^{PB}(\vec{X}_t)$$

and subject to an aggregate government budget constraints (i.e., in expectations, the present value of tax revenues minus public pension benefits is constant):¹³

$$\mathbb{E}_0 \left[\sum_{t=1}^A \frac{(\tau_i(d^{SP}) - pens_t(d^{SP}))}{(1+r)^t} \right] = \bar{T} \quad (7)$$

and an employer budget constraint for each employer type e :¹⁴

$$\mathbb{E}_0 [Pf(d^{SP}) + W(d^{SP}) + Mtc(d^{SP})] = \bar{Y} \quad (8)$$

where $Pf(d^{SP})$ measures aggregate profits, $W(d^{SP})$ aggregate wages and $Mtc(d^{SP})$ aggregate employer matching contributions.

Welfare metric. We express changes in welfare relative to the status-quo 2017 regime in terms of lifetime consumption-equivalent $\gamma(d)$:

$$1 + \gamma(d) = \left(\frac{V^{SP}(d^{SP} = d)}{V^{SP}(d^{SP} = 0)} \right)^{\frac{\sigma}{\sigma-1}}$$

where σ is the elasticity of intertemporal substitution. The policymaker is indifferent between: (i) a policy reform d and (ii) continuation of the status quo with consumption multiplied by a factor $(1 + \gamma(d))$ in every period of life and in every state of the world.

8.1.2 Incidence and means of balancing the budget

To maintain the government and employers' aggregate budget constraint (7,8), we assume that both the government and employers respond to the policy in the following ways:

Tax adjustment. Marginal tax rates adjust proportionally to offset changes in government revenue from tax-deferred savings: that is, all tax brackets increase or decrease by the same percentage point amount to maintain budget balance. Since these adjustments are relatively small across all the policies we consider, we abstract from labor supply responses.

Incidence on employees. In our baseline specification, workers bear the full cost

¹³Note that since employers are ex-ante identical in this set-up, the aggregate employers' budget constraint is equal, in expectations, to each individual employer's budget constraint.

¹⁴Note that since employers are ex-ante identical in this set-up, the aggregate employers' budget constraint is equal, in expectations, to each individual employer's budget constraint.

of higher employer contributions through lower wages. For each firm type e , the employer-specific earnings component $\bar{y}^e$ adjusts to keep total compensation constant: $\sum_{i=1}^N (y_i^e + ec_i^e) = \bar{Y}^e \forall e$. An advantage of this approach is that it isolates the retirement policy’s direct welfare effects from any changes in total compensation.

Incidence on employers. As robustness, we also simulate a second scenario that assumes the policy’s incidence falls completely on employers who reduce their profits to cover higher employer contribution costs. This reduction in profits is assumed to have no direct effect on social welfare and no effect on the interest rate (which is calibrated to match the long-term yield on government bonds).

8.2 Policy Evaluation: Targeting, Commitment, and Welfare

In this section, we use our estimated model to evaluate the welfare impact of alternative policies aimed at raising retirement savings. We adopt a long-term welfare criterion, where the social planner evaluates outcomes using individuals’ long-run time preferences rather than their present-biased preferences. All policies are compared to a counterfactual benchmark in which the UK’s pre-2012 policy regime (with no auto-enrollment and no increase in minimum contribution rates) would have continued indefinitely.

We evaluate five policy approaches: expanding financial incentives, behavioral nudges, the actual UK reform, asset-based contribution limits, and retirement income floors. Figure 14 summarizes the welfare results for the average population and separately for present-biased individuals and time-consistent individuals.

Financial incentives. The first policy we evaluate doubles the average employer match rate from 50% to 100% (first set of bars in Figure 14). To maintain budget neutrality, the policy requires two adjustments: wages fall by an average of 1.4% to offset higher employer costs (varying across employers based on their matching formulas), and marginal tax rates increase by 1% (e.g., from a 20% to a 20.2% marginal tax rate) to compensate for lost tax revenue from increased tax-deferred savings. Despite these generous incentives, the policy generates virtually no average welfare gains. Present-biased individuals experience only modest benefits, equivalent to a 0.5% increase in retirement consumption, while time-consistent workers suffer welfare losses from reduced wages, higher taxes, and distorted intertemporal consumption paths.

This near-zero net effect, even under a paternalistic social welfare criterion, reflects the poor targeting properties of many widely adopted financial incentive programs. To illustrate this, Figure 15 presents a set of model experiments offering workers a large one-time financial incentive to increase retirement contributions by 1 percentage point of their salary. The

left panel shows that liquidity-constrained workers rarely take up even extremely generous incentives. (e.g., only a third of the lowest liquidity workers take up a one-time offer of a 500% match to raise contributions by 1 pp). Yet these same liquidity-constrained workers exhibit the largest spending reductions when they do increase contributions (middle panel) and are the most likely to be present biased (right panel), making them the ideal targets for paternalistic retirement savings policies. Conversely, high-liquidity workers are most likely to take up financial incentives but exhibit minimal spending adjustments and low levels of present bias.

From the sufficient statistics framework perspective (Section 2), this pattern creates negative covariances: contribution increases induced by matching incentives are negatively correlated with both spending responses and present bias, precisely the opposite of efficient targeting. The workers who would benefit most from paternalistic intervention (low-liquidity, present-biased individuals) are least responsive to incentives, while those who need no help saving (high-liquidity, time-consistent workers) capture the bulk of subsidies.

Behavioral nudge. Given the poor targeting of untargeted financial incentives, we next consider whether behavioral interventions perform better (second set of bars in Figure 14). We simulate a policy with auto-enrollment for new hires, with default contribution rates at the employer match threshold and with no change in financial incentives. Analyses of optimal default contribution rates often find this to be the optimal default contribution rate (Choukhmane, 2024, Bernheim and Mueller-Gastell, 2024). Again, this policy has very limited impact on aggregate welfare, equivalent to a 0.1% increase in consumption in retirement. This is consistent with the evidence in Choukhmane (2024) that in a dynamic lifecycle setting with modest adjustment costs, auto-enrollment nudges only create limited welfare gains.

UK policy. The limited impact of both financial incentives and nudges in isolation raises the question of whether combining multiple interventions improves outcomes. Here we evaluate the UK’s actual policy that combines both elements: auto-enrollment and increased financial incentives (third set of bars in Figure 14). We compare welfare in the steady state of the UK’s actual retirement savings policy post-2019 (with an 8% minimum total contribution) relative to maintaining the pre-2012 status quo (no auto-enrollment, no minimum employee or employer contributions). This comprehensive approach generates larger aggregate welfare gains than financial incentives or nudges alone, equivalent to 0.5% higher retirement consumption. Present-biased individuals benefit substantially more, experiencing a 1.4% increase in consumption-equivalent welfare relative to the status quo. However, these gains come at significant cost: wages fall by 1.5% on average as employers offset higher employer contribution costs, and income tax rates rise by 1.5% (e.g., from a 20% to a 20.3%

marginal tax rate) to compensate for lost revenue from expanded tax-deferred savings. These offsetting effects contribute to reducing welfare for time-consistent individuals, who experience the equivalent of a 0.7% drop in retirement consumption due to lower wages, higher taxes, and distorted savings decisions. The policy’s relative success stems from improved targeting through minimum contributions. By raising the floor from 2% to 8%, the policy primarily binds for low savers, who tend to be present-biased, while leaving high-saving, time-consistent workers largely unaffected, effectively directing paternalistic interventions toward those who would benefit most.

Asset limits. While the UK reform shows promise, can its targeting be further improved? Next, we explore how asset limits could improve the targeting (fourth set of bars in Figure 14). We simulate the actual UK policy rollout but this time limiting the ability to make new pension contributions to individuals whose retirement balances fall below the 95th percentile.¹⁵ This restriction serves as a revelation mechanism: workers with very large retirement balances are almost always time-consistent savers with ample liquidity. Present-biased individuals, by contrast, rarely accumulate such high balances. By limiting tax and employer match incentives earned by these workers, who have revealed their time-consistent type and save adequately without intervention, the policy concentrates resources on those who are more prone to undersaving and most likely to cut spending. We find that adding asset limits increases the UK policy’s average welfare gains by 20%, despite reducing the welfare of time-consistent workers by an additional 0.1%, while eliminating the need for offsetting tax increases. Indeed, the fiscal savings from increased taxes collected from high-balance workers are large enough that the government could reduce income tax rates while maintaining budget neutrality, transforming a fiscally costly intervention into one that better serves its paternalistic objectives at a negative fiscal cost.

Income floors and forced annuities. While the policies studied thus far can increase contributions by present-biased individuals during their working lives, their welfare benefits remain limited because present-biased individuals tend to rapidly deplete their accumulated assets early in retirement. Annuities offer a valuable commitment device by forcing gradual decumulation of retirement savings. We estimate that raising the government-provided income floor (i.e., the state pension) by £650 annually increases aggregate welfare equivalent to 1.7% higher retirement consumption (fifth set of bars in Figure 14). Remarkably, even time-consistent individuals benefit despite higher income taxes to finance the policy, as income floors provide insurance against longevity and income risk (O’Dea, 2018). While previous research has demonstrated substantial welfare gains from retirement income floors

¹⁵This policy resembles the lifetime contribution limits that existed in the UK between 2006 and were abolished in 2024.

purely through their insurance value, our results reveal an additional benefit: they serve as commitment devices that protect present-biased individuals from their own tendency to overspend upon gaining access to their accumulated DC assets at retirement. This dual role, providing both insurance and commitment, makes income floors particularly effective policy tools for improving retirement security across all types of savers.

9 Conclusion

This paper shows that when UK employees were induced to raise their retirement contributions, only one-third of the increase came from reduced consumption. Instead, workers primarily financed higher contributions through reduced non-retirement savings and increased borrowing. This financing pattern varies sharply by liquidity: workers with limited liquid savings cut spending substantially, while those with ample liquid balances simply reshuffled existing assets across accounts with minimal spending adjustment over a two-year period.

These empirical findings have important implications for evaluating retirement savings policies. Using a lifecycle model that reproduces the observed heterogeneity in spending responses and accounts for the prevalence of present bias, we find that many popular policies generate limited aggregate welfare changes. Financial incentives that match employee contributions produce virtually no aggregate welfare gains despite their high fiscal cost because they are taken up primarily by high-liquidity, time-consistent savers who merely shift assets to capture matching benefits.

Our analysis yields three key lessons for retirement savings policy design. First, untar-geted financial incentives often fail to raise welfare because they attract savers who merely shift the location of their assets rather than increase total savings. Second, effective policies benefit from incorporating targeting mechanisms that act differently on different behavioral types: minimum contributions bind on present-biased low savers, while asset limits screen out time-consistent high savers. Third, the largest welfare gains come from policies that address both the accumulation and decumulation phases of the life cycle. Policies like retirement income floors and annuitization provide commitment by preventing present-biased retirees from rapidly depleting their savings.

Tax incentives for retirement savings are now the the largest non-structural tax relief in the UK and the largest federal tax expenditure in the US. Given their importance for fiscal sustainability and households' financial security, our results highlight the value to policymakers and researchers of looking beyond average effects to better understand the targeting and welfare implications of these consequential policies.


References

- Allcott, Hunt, Daniel Cohen, William Morrison, and Dmitry Taubinsky, “When Do Nudges Increase Welfare?,” Technical Report 5 2025.
- Altonji, Joseph G and Lewis M Segal, “Small-sample bias in GMM estimation of covariance structures,” *Journal of Business & Economic Statistics*, 1996, 14 (3), 353–366.
- Attanasio, Orazio P and Agar Brugiavini, “Social security and households’ saving,” *the Quarterly Journal of economics*, 2003, 118 (3), 1075–1119.
- and Susann Rohwedder, “Pension wealth and household saving: Evidence from pension reforms in the United Kingdom,” *American Economic Review*, 2003, 93 (5), 1499–1521.
- Bernheim, B Douglas and Dmitry Taubinsky, “Behavioral public economics,” *Handbook of behavioral economics: Applications and Foundations 1*, 2018, 1, 381–516.
- and Jonas Mueller-Gastell, “Optimal default options,” *Journal of Public Economics*, 2024, 237, 105203.
- Beshears, John, James J Choi, David Laibson, and Brigitte C Madrian, “The Importance of Default Options for Retirement Saving Outcomes: Evidence from the United States,” in “Social Security Policy in a Changing Environment,” University of Chicago Press, June 2009, pp. 167–195.
- , – , – , – , and William L Skimmyhorn, “Borrowing to save? The impact of automatic enrollment on debt,” *The Journal of Finance*, 2022, 77 (1), 403–447.
- , Matthew Blakstad, James J Choi, Christopher Firth, John Gathergood, David Laibson, Richard Notley, Jesal D Sheth, Will Sandbrook, and Neil Stewart, “Does Pension Automatic Enrollment Increase Debt? Evidence from a Large-Scale Natural Experiment,” February 2024. NBER Working Paper No. 32100.
- Blumenstock, Joshua, Michael Callen, and Tarek Ghani, “Why Do Defaults Affect Behavior? Experimental Evidence from Afghanistan,” *American Economic Review*, October 2018, 108 (10), 2868–2901.
- Brière, Marie, James Poterba, and Ariane Szafarz, “Household Portfolio Choice: A Review of the Literature,” *Annual Review of Financial Economics*, 2022, 14, 365–390.
- Card, David and Alan Krueger, “Minimum Wages and Employment: A Case Study of the Fast-Food Industry in New Jersey and Pennsylvania,” *American Economic Review*, 1994, 84 (4), 772–93.
- Chetty, Raj, John N Friedman, Søren Leth-Petersen, Torben Heien Nielsen, and Tore Olsen, “Active vs. passive decisions and crowd-out in retirement savings accounts: Evidence from Denmark,” *The Quarterly Journal of Economics*, 2014, 129 (3), 1141–1219.
- Choi, James J, David Laibson, and Brigitte C. Madrian, “Plan Design and 401(k) Savings Outcomes,” *National Tax Journal*, June 2004, 57 (2), 275–298.

- , – , and – , “Reducing the Complexity Costs of 401(k) Participation Through Quick Enrollment,” in “Developments in the Economics of Aging,” University of Chicago Press, March 2009, pp. 57–82.
- , – , **Jordan Cammarota, Richard Lombardo, and John Beshears**, “Smaller than We Thought? The Effect of Automatic Savings Policies,” August 2024. NBER Working Paper No. 32828.
- Choukhmane, Taha**, “Default Options and Retirement Saving Dynamics,” 2024. Extended Working Paper Version.
- , “Default Options and Retirement Saving Dynamics,” 2025. Forthcoming, American Economic Review.
- , **Jorge Colmenares, Cormac O’Dea, Jonathan Rothbaum, and Lawrence Schmidt**, “Who Benefits from Retirement Saving Incentives in the U.S.? Evidence on Racial Gaps in Retirement Wealth Accumulation,” 2023. Unpublished working paper.
- Clark, Robert L., Emma Hanson, and Olivia S. Mitchell**, “Lessons for Public Pensions from Utah’s Move to Pension Choice,” July 2015. NBER Working Paper 21385.
- Cribb, Jonathan and Carl Emmerson**, “What Happens When Employers Are Obligated to Nudge? Automatic Enrolment and Pension Saving in the UK,” 2016. IFS Working Paper W16/19.
- Feldstein, Martin**, “Social Security, Induced Retirement, and Aggregate Capital Accumulation,” *Journal of Political Economy*, 1974, 82 (5), 905–926.
- Fuchs-Schündeln, Nicola and Matthias Schündeln**, “Precautionary Savings and Self-Selection: Evidence from the German Reunification “Experiment”,” *The Quarterly Journal of Economics*, 08 2005, 120 (3), 1085–1120.
- Goldin, Jacob and Daniel Reck**, “Optimal Defaults with Normative Ambiguity,” *Review of Economics and Statistics*, 2022, 104 (1), 17–33.
- Jensen, Svend E. Hougaard, Sigurdur P. Olafsson, Arnaldur Stefansson, Thorsteinn S. Sveinsson, and Gylfi Zoega**, “Does Mandatory Saving Reduce Voluntary Saving? Evidence from a Pension Reform,” *The Review of Economics and Statistics*, 05 2025, pp. 1–48.
- Kolsrud, Jonas, Camille Landais, Daniel Reck, and Johannes Spinnewijn**, “Retirement Consumption and Pension Design,” 2021.
- Laibson, David, Sean Chanwook Lee, Peter Maxted, Andrea Repetto, and Jeremy Tobacman**, “Estimating discount functions with consumption choices over the lifecycle,” *The Review of Financial Studies*, 2024, p. hhae035.
- Madrian, Brigitte C. and Dennis F. Shea**, “The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior,” *The Quarterly Journal of Economics*, 2001, 116 (4), 1149–1187.


- Medina, Paolina C and Michaela Pagel**, “Does Saving Cause Borrowing? Implications for the Coholding Puzzle,” *The Journal of Finance*, 2025.
- Mitchell, Olivia S., Stephen P. Utkus, and Tongxuan (Stella) Yang**, “Turning Workers into Savers? Incentives, Liquidity, and Choice in 401(k) Plan Design,” NBER Working Paper 13369, National Bureau of Economic Research 2007.
- Moser, Christian and Pedro Olea de Souza e Silva**, “Optimal paternalistic savings policies,” *Columbia Business School Research Paper*, 2019, (17-51).
- O’Dea, Cormac**, “Insurance, Efficiency and the Design of Public Pensions,” 2018.
- Saez, Emmanuel**, “Optimal income transfer programs: intensive versus extensive labor supply responses,” *The Quarterly Journal of Economics*, 2002.
- Sun, Liyang and Sarah Abraham**, “Estimating dynamic treatment effects in event studies with heterogeneous treatment effects,” *Journal of Econometrics*, 2021, 225 (2), 175–199.

Figure 1: Average Cash Flows by Category for Medium Tercile Income




Notes: The figure plots the average cash flows in and out of checking accounts for workers in the middle tercile of income for the unrestricted sample in the year 2016. Bar heights report the net cash flow for the indicated category.

Figure 2: Distribution of March 2018 Contribution Rates by Group




Notes: The figure plots the probability density of contribution rates in March 2018, with each contribution rate group plotted in a different color. See section 4 for an explanation of the contribution rate groups.

Figure 3: Pension Contribution Rates by Source and Total Contribution Rate Group




Notes: Figure plots the average employee (red) and employer (green) contribution rate for each contribution rate group for the subsample of workers with data differentiating employer and employee contributions. Contribution rate groups are defined based on each employee's March 2018 total contribution rate. See section 4 for an explanation of the contribution rate groups.

Figure 4: Average Monthly Pension Contributions by Contribution Rate Group




Notes: Figure plots the median monthly contributions by contribution rate group. See section 4 for an explanation of the contribution rate groups. The vertical lines indicate the increases in the default pension contributions on 6 April 2018 and 6 April 2019.

Figure 5: Event Study of Change in Monthly Contributions




Notes: Figure plots the estimated average change in monthly retirement contributions (panel I) and the average change in cumulative contributions (panel II) for treated contribution-rate group workers relative to control-group workers using the Sun and Abraham (2021) estimator, normalizing March 2018 to zero. In panel I, coefficients are normalized by subtracting the pre-period average. Dashed lines plot 95% confidence intervals clustered at the pension scheme level. Individual and calendar-month fixed effects included in all models.

Figure 6: Event Study of Change in Checking Account and Debt Balances




Notes: Plotted coefficients are the cumulative change in checking account balances (panel I), credit card balances (panel II), and non-mortgage loans (panel III) for treated contribution-rate group workers relative to control-group workers using the Sun and Abraham (2021) estimator. Coefficients are normalized by subtracting the pre-period average. Dashed lines plot 95% confidence intervals clustered at the pension scheme level. Individual and calendar-month fixed effects included in all models.

Figure 7: Effects on Spending by Expenditure Category



Notes: Coefficients are the instrumental variables coefficients from a regression of pension contributions on spending in the indicated expenditure category on contribution amounts along with error bars indicating 95% confidence intervals.

Figure 8: Event Study of Change in Checking Account and Debt Balances




Notes: Figure plots the cumulative change in pension balances by the end of our sample for the average worker, along with a decomposition of how that increase in pension balance was financed. Plotted coefficients are the cumulative change for treated contribution-rate group workers relative to control-group workers using the Sun and Abraham (2021) estimator, normalizing March 2018 to zero. Individual and calendar-month fixed effects included in all models.

Figure 9: Effects on Total Spending by Deposit Tercile




Notes: Plotted coefficients are the instrumental variables coefficients from a regression of total spending on pension contribution amounts for each tercile of March 2018 deposits level along with error bars indicating 95% confidence intervals.

Figure 10: Simultaneous Retirement Saving and Credit Card Borrowing: Data vs. Model Simulations



Notes: Each panel shows the share of workers who are observed simultaneously (i) making positive contributions to their employer-sponsored retirement plan and (ii) revolving credit card debt. The "Data" series corresponds to workers in our linked bank-pension account data, "Model" corresponds to model simulations under the baseline estimates from Table 4, and "Model w/ $\beta = 1$ " keeps all other parameters constant but assumes no present bias. Finally, the series $\beta < 0.5$ and $0.5 < \beta < 1$ correspond to simulations with heterogeneous present bias.


Figure 11: Distributions of contribution rates over time: Data vs. Model Simulations



Notes: Each panels panel shows total (employee and employer) contribution rates of workers observed in the second quarter of 2017, 2018 and 2019 (before and after the step-up in minimum contributions). The model series corresponds to simulations before and after changes in the aggregate policy states, while the empirical series is produced using ASHE data.

Figure 12: Contribution behavior before and after auto-enrollment: Data vs. Model Simulations


(a) Participation rate by age



(b) Participation rate by income




(c) Bunching at the auto-enrollment default by age and income




Notes: The data series corresponds to evidence from ASHE before and after an employer staging date (between 2012 and 2016). The model series correspond to the same simulations before and after the first change in the aggregate policy state.

Figure 13: Spending Effects by Deposit Tercile: Data vs. Model Simulations




Notes: Figure graphs effects on total spending by initial deposit tercile from an increase in employer and employee contribution rates. Red bars with 95% confidence intervals repeat quasi-experimental estimates from Figure 9. Green bars report average effects on total spending in the simulated model after two years when default employee contributions increase to 3% and then 5% a year later and default employer contributions increase to 2% and then 3% a year later. Blue bars report simulated effects after 20 years.

Figure 14: Welfare Impact of Alternative Retirement Saving Policies



Notes: This figure shows the welfare impact of alternative retirement savings policies measured as equivalent changes in retirement consumption relative to the pre-2012 UK policy baseline (no auto-enrollment, no minimum contributions). Each policy is evaluated under a paternalistic welfare criterion using individuals' long-run time preferences. The bars show effects for the average population (dark blue), present-biased individuals (orange), and time-consistent individuals (green). Financial parameters below each policy indicate changes in average wages and marginal tax rates required for budget neutrality.

Figure 15: Simulated Effects of a One-time Subsidy to Increase Contribution by 1% of Salary



Notes: The left panel plots the take-up of a one-time match subsidy to increase retirement contributions by 1% of salary by deciles of liquid wealth. On the left of the red vertical line are deciles with net negative liquid wealth (i.e., unsecured debt holders). The middle (right) panel corresponds to the average spending reduction per additional pound of pension contribution (average level of present bias), conditional on increasing retirement contribution by 1% of salary..

Table 1: Summary Statistics

	Mean	Median	Std. Dev.
<i><u>I. Income and Spending</u></i>			
Net Wage Income	2292.5	1869.7	2888.0
Total Spending	1492.5	1166.3	2196.3
Spending via Current Accounts	1360.6	1062.2	2138.3
Spending via Credit Cards	131.9	0.0	464.7
Housing and Utilities Spending	268.6	220.45	702.1
Restaurant Spending	88.9	45.0	141.4
Consumer Retail Spending	323.3	191.9	517.4
Supermarket Spending	234.0	167.0	238.0
Leisure Spending	88.0	20.4	510.1
Other Spending	489.7	272.3	1757.5
<i><u>II. Debt and Balances</u></i>			
Credit Card Payments	364.6	30.0	1136.4
Loan Payments	138.5	0.0	711.2
Current Account Balance	4212.9	1600.0	13286.4
Credit Card Balance	654.3	0.0	1850.6
Has Savings Account	0.35	0.0	0.48
<i><u>III. Individual Attributes</u></i>			
Female	0.40	0.0	0.49
Age	40.6	38.0	11.0
Number of Observations		3,887,397	
Number of Individuals		106,345	

Notes: Table reports summary statistics covering individuals from January 2016 through November 2019. All amounts are in nominal British pounds. See the Appendix for details on the data cleaning and sample selection procedures.

Table 2: Summary Statistics in March 2018 by Contribution Rate Groups

Contribution Rate Group	2	3	5	8
Contribution Rate	2.0 (0.28)	3.4 (0.57)	6.0 (0.88)	11.0 (2.24)
Net Wage Income	2101.1 (2322.2)	2478.8 (3089.0)	2567.8 (3000.6)	2471.6 (1990.3)
Pension Contribution Amount	41.5 (46.2)	84.9 (110.4)	153.3 (181.1)	270.7 (218.8)
Total Spending	1248.8 (1831.0)	1387.6 (1767.4)	1389.2 (2083.7)	1447.4 (2215.6)
Number of Individuals	27,533	21,473	20,889	36,450

Notes: Table reports means with corresponding standard deviations in parentheses by contribution rate group using data in March 2018. See section 4 for an explanation of the contribution rate groups.

Table 3: Effect of Pension Contributions on Income, Spending, Borrowing, and Deposits

Dependent Variable	Effect per £1 Increase in Contributions
<i>I. Income</i>	
Net Wage Income	-0.67*** (0.11)
<i>II. Spending</i>	
Total Spending	-0.23** (0.10)
Leisure Spending	-0.09*** (0.03)
Restaurants Spending	-0.04*** (0.01)
Utilities/Subscription Spending	-0.05* (0.03)
Consumer Retail Spending	-0.04 (0.03)
Supermarket Spending	0.00 (0.01)
Rent Spending	-0.02 (0.02)
Other Spending	-0.01 (0.01)

Notes: Table reports instrumental variables regression estimates of the effect of pension contribution increases on income (panel I) and spending (panel II). Each row reports the results of a separate regression with the indicated outcome variable. The estimation sample is described in Table 1 and contains 1,534,654 monthly observations from 44,122 individuals. Total Spending equals Spending via Current Accounts plus Spending via Credit Cards. Robust standard errors in parentheses. Individual and calendar-month fixed effects included in all models. * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$

Table 4: Preference parameter estimates

(1)	(2)	(3)	(4)	(5)
Annual discount factor	Elasticity of Intertemp. Subsit.	Share of present biased	Distribution of present bias	Contribution Adj. cost
δ^4	σ	π^β	$\beta \sim \text{Beta dist.}$	k
0.99	0.93	55%	Avg. β : 0.547 (s.d.: 0.04)	£35

A Data Appendix

In the raw data, all observations are at the account-month level, with the exception of demographic data which are captured at the account-year level. Some individuals have multiple observations on file in a given month. For individuals with multiple observations in a given month, we consolidate as follows: For pension accounts, we keep only the account having the lowest recorded pension contribution amount for a given month. When an individual has multiple accounts that share this lowest contribution amount, we keep only the account that has the most recent open date. For accounts with multiple demographic observations (usually due to changes in characteristics such as ownership, employment status, or marital status in a given year), observations with employment, observations marked as active, and observations marked as registered are prioritized in that order. Finally, for individuals with multiple observations corresponding to account flows for a given month, in all cases only one observation has a non-missing employer ID. We keep this observation only.

We then merge together and aggregate data on pension contributions, net flows, deposit account balances, and demographics from the different accounts of each individual. After these data sets are merged, we have a panel of workers at the individual-month level.

We then restrict the sample in the following ways to ensure that the individuals captured give a complete and reliable picture of the finances of individuals who would likely be affected by a change in pension contributions. First, we drop observations corresponding to an individual in a given year if that person never contributed to a pension account that year. Second, we drop individuals who have a nonpositive sum of current account balances over the sample period. Third, we drop individuals who never had recorded wages from employment during the sample period. Fourth, we drop observation corresponding to an individual in a given year if that person spent less than 30% of earned wages during the year. Fifth, we drop observations where pension contribution amounts are missing. Sixth, we drop observations that are prior to the first recorded employment month or later than the last recorded employment month for a given individual (determined by whether the observation has a non-missing employer ID). Seventh, we drop observations that have multiple employer IDs for a given individual and month. And finally, we drop observations for which information to calculate a pension contribution rate is missing, or which were not assigned to any contribution group (because their contribution rate rounds to below 2% or above 15%).

We then restrict the time period of the estimation sample to the months between January 2016 and November 2019. We drop the month of December 2019 from the estimation sample because observations corresponding to December 2019 from the raw data only contain information from part of the month. And while the panel is available from August 2011, January 2016 is chosen as the lower cutoff to limit the degree to which the panel is unbalanced because there are a very limited number of observations before 2016.

For most of our analysis, we also make two final restrictions. First, we drop individuals

that never used a credit card (as evidenced by having no non-missing credit card balances) during the sample period. Second, we drop observations correspond to an individual in a given year if the median monthly amount spent on food (restaurant and supermarket purchases) was less than 50 pounds. The estimation sample with these two final restrictions is referred to as the restricted sample. As a robustness check, we also provide estimates of some of our main results using the estimation sample without these restrictions. We refer to this larger estimation sample as the unrestricted sample.

Finally, we then clean the data as follows. First, we set all credit card variables to zero for individuals who do not have a credit card. We then set all other existing but missing cash flow and account balance observations to zero. We then aggregate the totals for each category by individual and month. In addition, since age is provided in segments (i.e., 21-25, 26-30, ..., 61-65, 66+), we calculate an age midpoint variable which is set to be the midpoint of these ranges (or 70 in the case of the 66+ group); this age midpoint variable is used as a proxy for age in our analysis. Finally, we winsorize pension contributions at the 0.1 and 99.9 percentiles so as to lessen the effects of extreme outlier values for pension contribution in the sample.

B Model Appendix

B.1 First-stage parameters calibration

In what follows, we describe the calibration of the model’s first stage parameter.

Employer matching formulas. We classify employers into eight types according to their contribution formulas. For all employers, we set the match rate at 50% and specify eight distinct matching schedules, calibrated to replicate the distribution of employer contributions observed in ASHE prior to 2012 (i.e., before the introduction of auto-enrollment and minimum contribution requirements). The resulting probability distribution of employer types is presented in Table A1.

UK income profile. We estimate the age-income profile using earnings data for UK private sector workers aged 22 to 65 using ASHE data from between 1997 and 2016. The estimation results are reported in Table A2.

Table A1: Probability distribution of employer types

Threshold on matching (cap_e)								
	No match	2%	3%	4%	5%	6%	7%	8%
Prob.	0.40	0.063	0.096	0.075	0.128	0.156	0.044	0.037

Notes: This table shows the probability distribution of match thresholds in the UK calibration of the model. This distribution is inferred from the empirical distribution of modal contribution rates in the three years prior to the auto-enrollment staging date. Data source: ASHE waves 2006 to 2016.

Table A2: UK age-earnings profile

Age component				Stochastic component			
δ_0	δ_1	δ_2	δ_3	ρ	$\sigma_{\xi_1}^2$	σ_{ξ}^2	σ_m^2
5.306	0.255	-0.0043	0.0000219	0.974	0.184	0.0125	0.10


Notes: Age-earnings profile estimated on a panel of workers continuously employed in the same job. Data source: ASHE waves 1997 to 2016.

Income taxation. Individuals' income tax liability is calculated according to the UK income tax schedule of 2012 for an individual claiming the personal allowance for singles. The quarterly income tax liability is equal to:

$$tax_t^i = \begin{cases} 0 & \text{if } 4y^{tax} \leq \pounds 8,105 \\ 0.22 \left(y^{tax} - \frac{8,105}{4} \right) & \text{if } \pounds 34,370 \geq 4y^{tax} > \pounds 8,105 \\ 0.22 \left(\frac{34,370 - 8,105}{4} \right) + 0.40 \left(y^{tax} - \frac{34,370}{4} \right) & \text{if } 4y^{tax} > \pounds 34,370 \end{cases}$$


Gouvernement benefit. We calibrate the level of the basis state pension ($pens_i$) and jobseeker allowance (ui_t) to their values in 2012.

Figure A1: Benchmarking Sample Consumption Data with National Survey Data




Notes: UK survey data comes from the Living Costs and Food Survey. Each point represents an annual average between the years 2016 and 2018, and the lines chronologically connect the points within each category.

Figure A2: Benchmarking Debt Data with National Survey Data



Notes: This figure compares our data with UK survey data from the Office for National Statistics (<https://www.ons.gov.uk/peoplepopulation-andcommunity/personalandhouseholdfinances/incomeandwealth/datasets/householddebt-wealthingreatbritain>, section 7.1) for the April 2016–March 2018 period. For our data, property debt is defined as mortgage balances; all other debt is categorized as financial debt.

Figure A3: Credit Card and Current Account Spending by Usage Category



Notes: Each category of spending represents the average of the share of total spending in that category through credit cards vs. current accounts. Cash spending is included in current account spending. The data is from January 2016 through November 2019.