How Do Consumers Finance Increased Retirement Savings?*

Taha Choukhmane† Christopher Palmer‡

March 2024

Abstract

Higher retirement savings might not translate into net wealth accumulation if, rather than cutting spending, individuals reduce their non-retirement savings and take on more debt. We use newly merged deposit-, credit-, and pension-account data from a large UK financial institution to examine a policy that gradually increased retirement contributions from 2% to 8% of salary between March 2018 and April 2019. For every £1 reduction in take-home pay due to higher employee contributions, employees cut their spending by only £0.34 (especially in the restaurant and leisure categories) and financed the remaining with lower deposit balances and higher credit card debt. Those with lower initial deposit balances cut their spending the most, while those with significant liquid savings draw down their deposits. Since employees cannot draw down their balances indefinitely, we use a lifecycle model–calibrated to match the observed short-term responses–to predict that long-run spending responses are larger but feature similar heterogeneity. A social planner concerned about undersaving for retirement due to heterogeneous present bias should, therefore, avoid targeting retirement interventions at high-liquidity individuals who are both less likely to cut their spending and less likely to be present-biased.

*For helpful comments, we thank James Choi, Stefano DellaVigna (discussant), Mariacristina De Nardi, Eric French, Olivia S. Mitchell, Redis Zaliauskas, and Mingli Zhong (discussant) and seminar and workshop participants at the ETH Zurich Microeconomics of Aging conference, the IFS, Michigan Research Data Center, MIT, Nest, NBER Summer Institute Aging, NTA 2021, PUC Chile, Sciences Po, Wharton, George Washington University, and Stanford SITE. 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. All errors and omissions are the authors' own.

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

‡MIT Sloan School of Management and NBER; cjpalmer@mit.edu

1 Introduction

Governments around the world are heavily invested in promoting contributions to retirement plans, employing methods from mandatory savings such as social security programs, financial incentives such as tax advantages and matching schemes, and behavioral interventions such as automatic enrollment. While a rich literature has studied the effects of such policies on savings inside retirement accounts, much less known is whether and how higher contributions to retirement accounts affect net wealth accumulation and welfare. Answering this question requires knowing how consumers react to higher retirement contributions and the accompanying reduction in their disposable income. If consumers respond by cutting their spending, increasing retirement contributions can significantly boost total savings, aligning with policymakers' objective of increasing total 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 with a higher debt burden and more financially vulnerable.

In this paper, we first characterize in a theoretical model the welfare consequences of policies aimed at promoting retirement savings. We show that when the social planner is concerned about undersaving for retirement—for instance due to present bias—the consumption response to the policy is a sufficient statistic for welfare. Consistent with models in behavioral public finance (Allcott et al., 2022), we demonstrate that a policy's average treatment effect on retirement savings, which is often used as a measure of an intervention's success, is a poor indicator for assessing welfare. Instead, judging a policy's economic efficiency depends on the covariances between retirement savings responses, consumption responses, and the level of individuals' undersaving bias. In particular, if the most patient (least biased) individuals are also the ones whose contributions are most responsive to a given incentive scheme, a policy that succeeds in increasing retirement contributions may nonetheless reduce welfare.

Measuring these welfare-relevant elasticities and covariances requires comprehensive data on consumption and savings along with retirement savings policy shocks. We use a new panel dataset from a large UK financial institution created by merging retail deposit and credit accounts with pension account data. We take advantage of two policy changes implemented as part of the UK national auto-enrollment policy, which raised the minimum default combined employee and employer contribution rate from 2% to 5% of salary in April 2018 and from 5% to 8%. This policy changed both the default contribution rate for employees as well as the financial incentives for contributing: each step-up introduces a notch in the financial incentives of retirement contributions. Employees cannot reduce their contribution below

the minimum default and opting-out of participation leads to loosing an increasingly large employer contribution, increased from a minimum of 1% to 2% of salary in April 2018 and from 2% to 3% of salary after April 2019. Therefore this policy variation speaks to both the effect of financial incentives and behavioral nudges.

This policy change was binding for some but not all employees and employers. This allows us to compare the behavior of consumers affected and not-affected by this significant increase in the minimum retirement contribution required to stay in the plan and benefit from the employer 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 £0.35 and financing the remaining with lower deposit account balances and with 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 find evidence of substantial treatment-effect heterogeneity. The most liquidity-constrained customers (i.e., those with lower deposit balances and high credit-card debt) cut their consumption the most. In contrast, those with significant liquid savings finance the increased pension contributions by running down their deposit balances. Finally, we simulate the long-run dynamics of our treatment effects using a structural life-cycle model. From a policy perspective, our results suggest that retirement interventions should target low-liquidity individuals (whose spending is more elastic to increased retirement savings and who are likely to have higher levels of undersaving or present bias). In contrast, interventions that increase the retirement contributions of high-liquidity individuals are both less efficient (due to large crowd-out) and more likely to be regressive.

How retirement saving policy affects total wealth accumulation has long interested researchers, yet the empirical evidence is scarce. A limited literature has studied whether forced savings—for example, from higher social security contributions—crowds out private savings (Feldstein, 1974; Attanasio and Brugiavinni, 2003; Attanasio and Rohwedder, 2003).[1] A notable 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 government savings. However, the discontinuity in these mandated savings, which - unlike employer contributions - directly impact take-home pay, is relatively modest, around 50 USD annually. In contrast to these papers, we can directly observe spending and spending categories in our data and we study a policy that materially decreased take-home pay for workers (rather than increase government payroll

[1] We 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.

taxes or employer contributions). While recent research has found limited effects of autoenrollment on unsecured borrowing (Beshears et al., 2021), there is no current evidence on how spending and deposits change in response to retirement savings.[2]

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. As we discuss below, a new database created by 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. Here we exploit changes in a national policy that changed the minimum default contributions and was binding only for employers and employees who contributed less than the new minimum default option.

The remainder of the paper is organized as follows. In section 3, we describe the 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. In section 6, we document important differences across consumer segments in their response to increased contribution rates. Section 7 sketches a life-cycle model to confirm that our quasi-experimental results hold in the model without selection and simulate long-run effects. Section 8 concludes with a summary of our findings and discusses implications for policy. A data appendix describes the sample selection criteria that we used to arrive at our final sample.

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. (2022). 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

[2]Outside of the retirement saving setting, Medina and Pagel (2021) also fail to reject that increased deposits from text message nudges are financed with unsecured debt.

γ of government retirement savings subsidies $s(ret_i, \pi_i, \gamma)$ and taxes $\tau_i(liq_i, \pi_i, \gamma)$:

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

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

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

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. 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:

$$W(\gamma) = \int_i \omega_i \left[u(c_i(\gamma)) + \beta_i(1 + p_i)V_i(ret_i(\gamma), liq_i(\gamma))\right] di + \mu \int_i (\tau(liq_i(\gamma), \gamma) - s(ret_i(\gamma), \gamma)) di \tag{3}$$

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

$$\frac{dW(\gamma)}{d\gamma} = \int_i \omega_i \left\{ \underbrace{\frac{dc_i}{d\gamma} u'(c_i)}_{\text{cons. response}} + \beta_i (1+p_i) \left[\underbrace{\frac{dret_i}{d\gamma} V_1'}_{\text{ret. savings response}} + \underbrace{\frac{dliq_i}{d\gamma} V_2'}_{\text{crowd-out liq. savings}} \right] \right\} d_i$$
$$+ \mu \int_i \left\{ \underbrace{\frac{d\tau_i(\gamma)}{d\gamma} - \frac{ds_i(\gamma)}{d\gamma}}_{\text{fiscal effect}} \right\} di$$

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:

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

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:

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

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. (2022), 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. (2022)).

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 consumption 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, savings accounts, and credit cards.[3] Unlike the flow data, month-end balances are rounded to the nearest £100. Finally, the data contain individual characteristics available at an annual frequency: age, gender, and a consumer marketing segment category known as a Fresco segment. The raw data is available up to mid-December 2019, with a few consumers having data available as early as January 2011. Pension contribution data start in January 2011, as do account balance data; account flows data start in August 2011, and demographic data start in 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.

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

[3]These categories also include different product variants within these account types, such as "added value accounts," which are current accounts packaged together with other services such as insurance.

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. Appendix Figure A3 plots the monthly averages of several key variables in our data. There are clear time trends, with each category generally increasing over time, subject to persistent seasonality.

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 A4 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.[4] 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.[5] 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 2% to 5%.[6] 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%.[7] Employees can choose to opt-out (but cannot contribute less than the

[4]Qualifying 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.

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

[6]Optionally, 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.

[7]Technically, 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.

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 connot contribute below the new minimum) and changes the financial incentives for contributing (by raising employer contribution levels).

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%.[8] We set the lower bound of the lowest contribution rate to be at 1.5% to focus 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

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.

[8]Increasing the upper bound to a 20% contribution rate would add an additional 10,000 individuals to the sample—only a 1.6% increase in sample size at the potential expense of reduced comparability across contribution-rate groups.

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 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 average 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 average contributions compared with the relatively flat time trends for the other groups. As expected, the lower contribution rate groups have their average pension contributions increased significantly by the increase in default minimum contribution rates in both periods. For example, the 2% contribution rate group's average 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 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.

As a final visualization of how contrasting the differential response of these contribution rate groups isolates variation in pension contributions coming from the increase in autoenrollment minimums, Figure 6 plots monthly average take-home pay by contribution rate group. As predicted, these net wage income series mirror the contribution series in Figure 4. In the left-hand panel, the two most-affected groups (2 and 3) show a similar evolution of net wage income over time, except for noticeable drops in net wage income levels right at the first increase in default contributions in April 2018. In the right-hand panel, the two groups with contributions that were near or above the default contribution rates as of March

2018 are plotted. Here, too, incomes are increasing over time, with group 8 being unaffected by the April 2018 increase in minimum contribution rates and group 5 experiencing a slight drop in take-home pay before continuing to increase as overall incomes grow.

Lastly, we use the life-cycle model in section 7 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—analogous 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

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

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

$$\begin{aligned} Outcome_{it} &= \beta \cdot PensionContributions_{it} + \alpha_i + \delta_t + \varepsilon_{it} \\ Pension\,Contributions_{it} &= \sum_{s \in \{1,2\}} \sum_{k \in \{2,3,5\}} \pi_{ks} Group_i^k \times Post_t^s + \psi_i + \phi_t + v_{it} \end{aligned}$$

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-hone 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.

In the data, we find that increasing pension contributions by £1 indeed reduced take-home-pay by around £0.67 (panel I of Table 3).

Our comparison of groups of customers with different contribution rates as of March 2018 has shown that employees who were more affected by the contribution rate step-ups of April

2018 and 2019 experienced a significant increase in their monthly pension contribution and a drop in their take-home-pay. Next, we evaluate how individuals offset the reduction in their take-home pay by either reducing their consumption, increasing their debt levels, or reducing their non-retirement savings.

Table 3 displays the results of the regression analysis on the effect of increased pension contributions on spending (panel II), debt payments (panel III), and net flows into current accounts and out of credit card accounts (panel IV). The results are expressed in terms of the effect of increasing pension contributions by £1 (which is associated with an approximate reduction in take-home pay of £0.67) for each of the outcome variables.

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 8 reports 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 respond to the drop in take-home pay with 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). This changes in monthly flows translate into significant reduction in current account balances and an increase in credit-card balances. Figure 9 summarizes the effects on balances, showing that for every £1 increase in monthly take-home pay, checking account balances decrease by £4.2, internal credit card balances increase by £0.8 (with the caveat that unlike monthly flows, the delivered balance data have been rounded to the nearest £100) Panel IV of Figure 3 presents regression estimates of these effects on flows. Finally, to characterize the timing of these balance changes, Figure 10 plots event study estimates of the change in checking account (panel I) and credit card (panel II) balances. In contrast to the more immediate effects of the policy on retirement savings, the effect on credit-card debt and checking accounts is more gradual, suggesting that workers draw on these sources of financing to cope with decreased take-home pay gradually over time.

6 Heterogeneous Effects by Liquidity Status

Our results suggest that, on average, only one-third 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 11 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 inital 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.

7 Life-Cycle Model

In this section, we build on the estimated life-cycle model of Choukhmane (2023) to explore alternative assumptions about the incidence of the policy, simulate the response of spending in the long-run, and assess the targeting of retirement incentives. We briefly sketch the model here and refer the reader to Section 4 of Choukhmane (2023) for further details on the model and Section H.3. for the calibration using micro-data from the U.K. Annual Survey of Hours and Earnings.

The model features a rich economic environment along with a parsimonious specification of preferences. During their working life, agents choose how much to consume, how much to contribute to an illiquid tax-favored employer-sponsored retirement account, and how much liquid savings or unsecured debt to hold. Agents face income and employment risk that varies with age and tenure. After age 65, agents can withdraw ressources from their retirement savings account (subject to income taxation) and receive public pension benefits calibrated to match the UK State Pension benefit level. The government maintains a progressive tax schedule (calibrated to match the UK income tax system) and funds unemployment insurance and a public pension system (calibrated to match the UK State Pension benefit level). Households face mortality risk and changes to household composition that both vary with age. Preferences include an elasticity of intertemporal substitution (set at 0.52) and an exponential discount factor ($\delta^4 = 0.96$). Households choosing to actively change their contribution rate (including those opting out of participation entirely) pay a £171 utility

cost. In an extension, we explore the model under a specification with naive present bias (as described in Section E.1. of Choukhmane (2023)).

The state variables consist of age, employment status, labor productivity, the employer retirement contribution formula, tenure, the current ratio of liquid wealth to debt, current retirement wealth, the default contribution rate and the aggregate policy environment (pre- vs post- contribution step-ups). Agents maximize their lifetime utility by each period choosing their next period ratio of liquid wealth to debt and their savings contribution rate to their defined contribution retirement account, which is negative during retirement to capture withdrawals.

Using the model, we simulate the effect on total spending of increasing the default employer and employee contribution rates. We increase employee contribution rates to 3% and then 5% a year later and employer contribution rates to 2% and then 3% a year later. Figure 12 reports the results of this exercise. For comparison, the red bars with confidence intervals repeat the quasi-experimental estimates using the data plotted in Figure 11. The dark blue bars show the model-simulated effects at a comparable two-year horizon. The model matches the pattern seen in the data of strong spending effects of autoenrollment at higher contribution rates for agents in the bottom third of initial level of liquid savings. As in the data, the spending effects in the model fade out as initial deposits increase. This comparison demonstrates that the quasi-experimental effects are not merely an artifact of unobserved differences across workers with varying baseline contribution rates. Using the model, we simulate effects at a longer horizon to gauge persistence. The light blue bars in Figure 12 show that while there is some fade-out of the initial spending effects at a twenty-year horizon, spending still seems permanently lower. Over time, workers with higher initial deposits eventually spend down their savings buffer and ultimately also finance ongoing higher retirement savings with lower spending.

We use the calibrated model to evaluate other policies aimed at increasing retirement contributions. As shown in Section 2 a well-targeted policy ensures positive covariances between the change in retirement contributions, and both the elasticity of spending and the degree of undersaving bias. We simulate the effect of a financial incentive for retirement saving, in the form of a one-time matching subsidy for increasing one's contribution rate by 1% of salary. As shown in Figure 13, we find that such financial incentives generate negative covariances: the take-up of the matching incentives is highest among those with substantial liquid savings (left panel) who are least likely to cut their spending (middle panel) and tend to be the least present biased individuals (right panel).

8 Conclusion

In this paper, we study the effects of increasing retirement savings on consumption, total savings, borrowing, and welfare. For identification, we use the increases in minimum default pension contributions mandated by the Pension Act of 2008 and newly compiled data combining pension contributions data with spending, saving, and borrowing data. Using a simple framework for assessing effects on welfare in the presence of undersaving, we show that the consumption response is a sufficient statistic for welfare. However, the average treatment effect on consumption is not sufficient. Instead, the covariance between the contribution response, the elasticity of consumption, and any bias in consumers' savings decision-making determines social welfare. Empirically, we find that, after the increases in minimum contributions in 2018 and 2019, affected consumers' net wage income declined by approximately £0.67 for every £1 in additional pension contributions, suggesting that roughly two-thirds of the increase in pension contributions was funded by employee contributions and the remainder by employer contributions as we would expect. For every £1 decrease in their monthly take-home pay, consumers respond by cutting their spending by £0.35 and financing the remainder by dissaving in their checking accounts and accumulating a modest amount of additional credit card debt. We find that the relatively discretionary areas of spending, such as restaurant and leisure spending, are the most elastic to the decrease in income net of pension contributions. Younger customers and customers with lower current account balances cut their total spending the most, with customers seeming to follow a pecking order where they prefer to spend down current account balances before cutting their habitual spending if possible.

Returning to our welfare framework, what lessons do these results have for the design of retirement policy? Consistent with the theoretical framework of Allcott et al. (2022), the average increase in retirement savings from an intervention is a poor guide for welfare. Instead, a well-targeted policy is one that that increases the retirement contributions of those with the larger consumption responses—such that $Cov\left(\frac{dret_i}{d\gamma}, \frac{dc_i}{dret_i}\right) > 0$—and larger levels of bias—such that $Cov\left(\frac{dret_i}{d\gamma}, p_i\right) > 0$. Empirically, we find the elasticity of spending to retirement contributions to be largest for low-liquidity individuals. Moreover, models of present bias predict that a low level of liquid savings is also an indicator of a larger bias and undersaving. This result highlights a potential downside to making retirement savings more illiquid. While raising penalties on early withdrawals can provide a source of commitment, such penalties may deter low-liquidity households—with the largest spending responses and the most to gain from saving—from participating, worsening the quality of targeting.

Our results further suggest that tax incentives and employer matching subsidies are

often poorly targeted given that the largest share of fiscal and employer expenditures appear to be taken up by those with the highest liquidity (Choukhmane et al., 2023). These households have the smallest spending response and are likely the least biased, such that $Cov\left(\frac{dret_i}{d\gamma}, \frac{dc_i}{dret_i}\right) < 0$ and $Cov\left(\frac{dret_i}{d\gamma}, p_i\right) < 0$. High-liquidity households can be more responsive to financial incentives precisely because they can take advantage of the subsidy without changing their spending patterns—simply by shifting existing savings from non-retirement to a retirement accounts. Consequently, even absent any redistributive motive ($g_i = 1$), our results suggest that an asset (or income) cap on tax incentives or on forced savings may be efficient.

References

Allcott, Hunt, Daniel Cohen, William Morrison, and Dmitry Taubinsky, "When do" Nudges" Increase Welfare?," Technical Report, National Bureau of Economic Research 2022.

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.

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*, 2021.

_ **, James J Choi, David Laibson, Brigitte C Madrian, and William L Skimmyhorn**, "Borrowing to save? The impact of automatic enrollment on debt," *The Journal of Finance*, 2022, *77* (1), 403–447.

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.

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.

Choukhmane, Taha, "Default Options and Retirement Saving Dynamics," 2023. Working Paper.

_ , 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 Obliged 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.

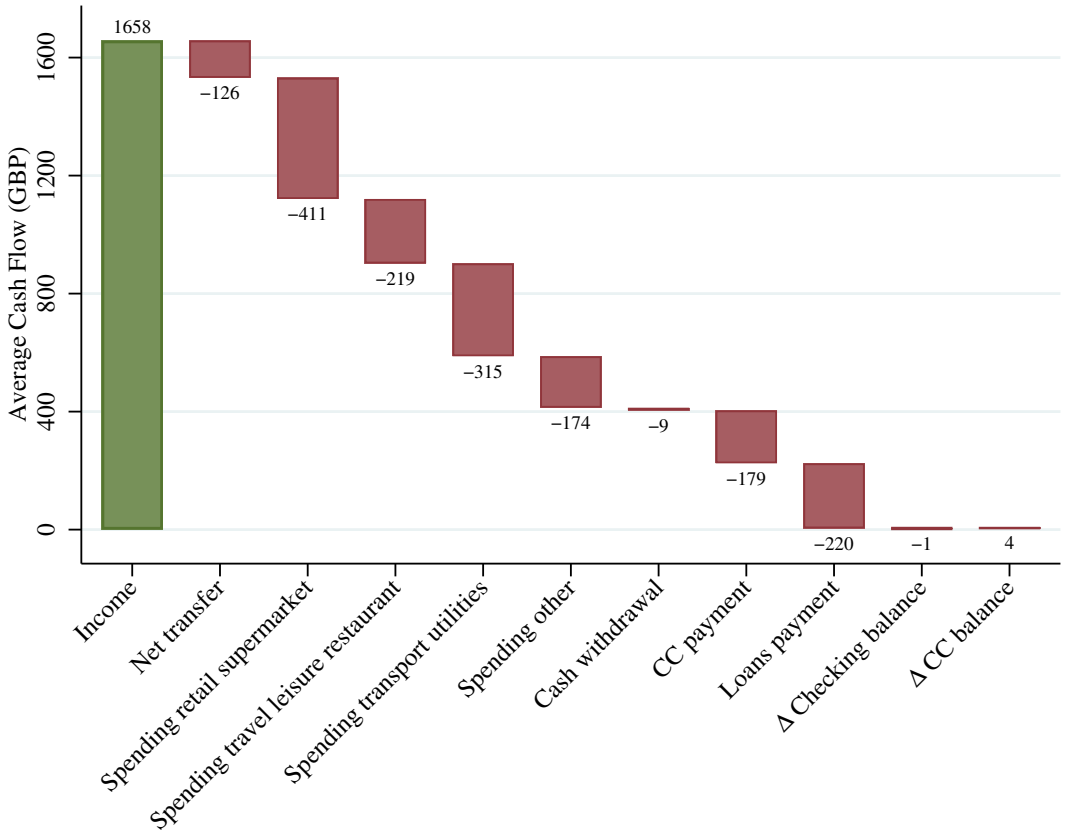
Kolsrud, Jonas, Camille Landais, Daniel Reck, and Johannes Spinnewijn, "Retirement Consumption and Pension Design," 2021.

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?," 2021. NBER Working Paper No. 28956.

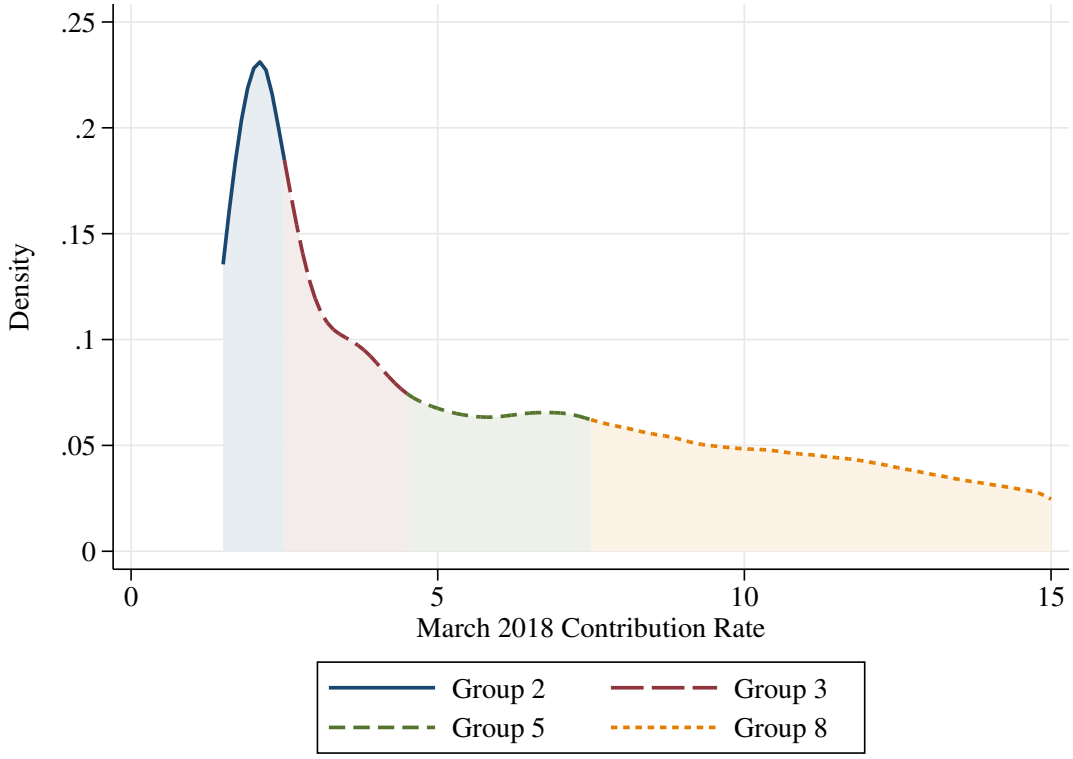
Sun, Liyang and Sarah Abraham, "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects," *Journal of Econometrics*, 2021, *225* (2), 175–199.

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



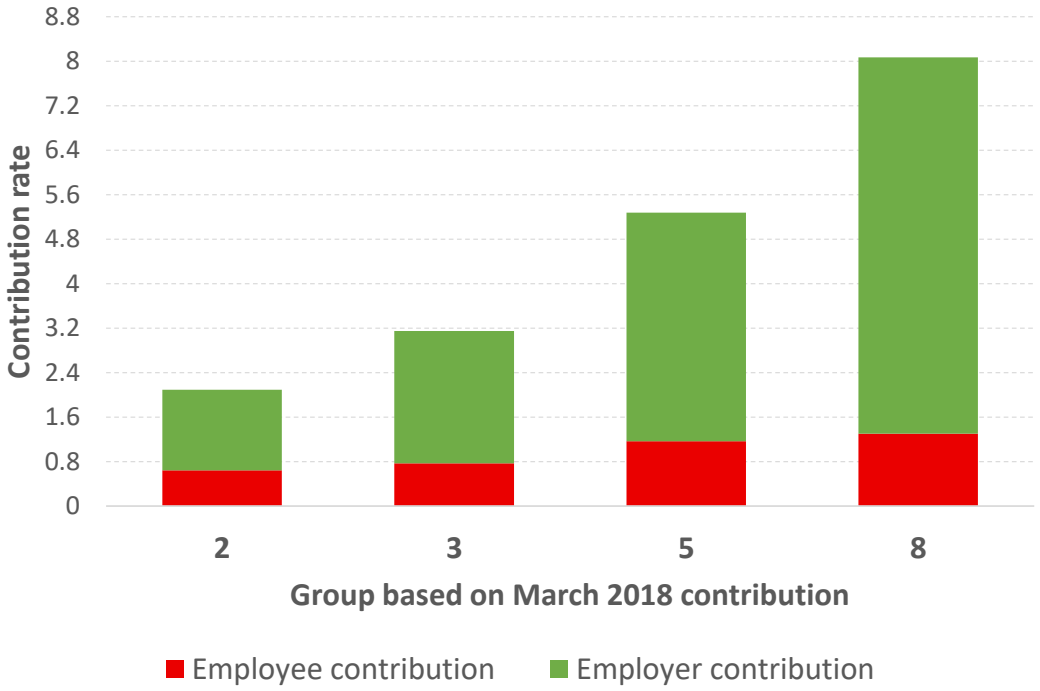
Notes: The figure plots the average cash flows in and out of checking accounts for workers in the middle tercile of income. Bar heights report the net cash flow for the indicated category.

Figure 2: Distribution of March 2018 Contribution Rates by Group



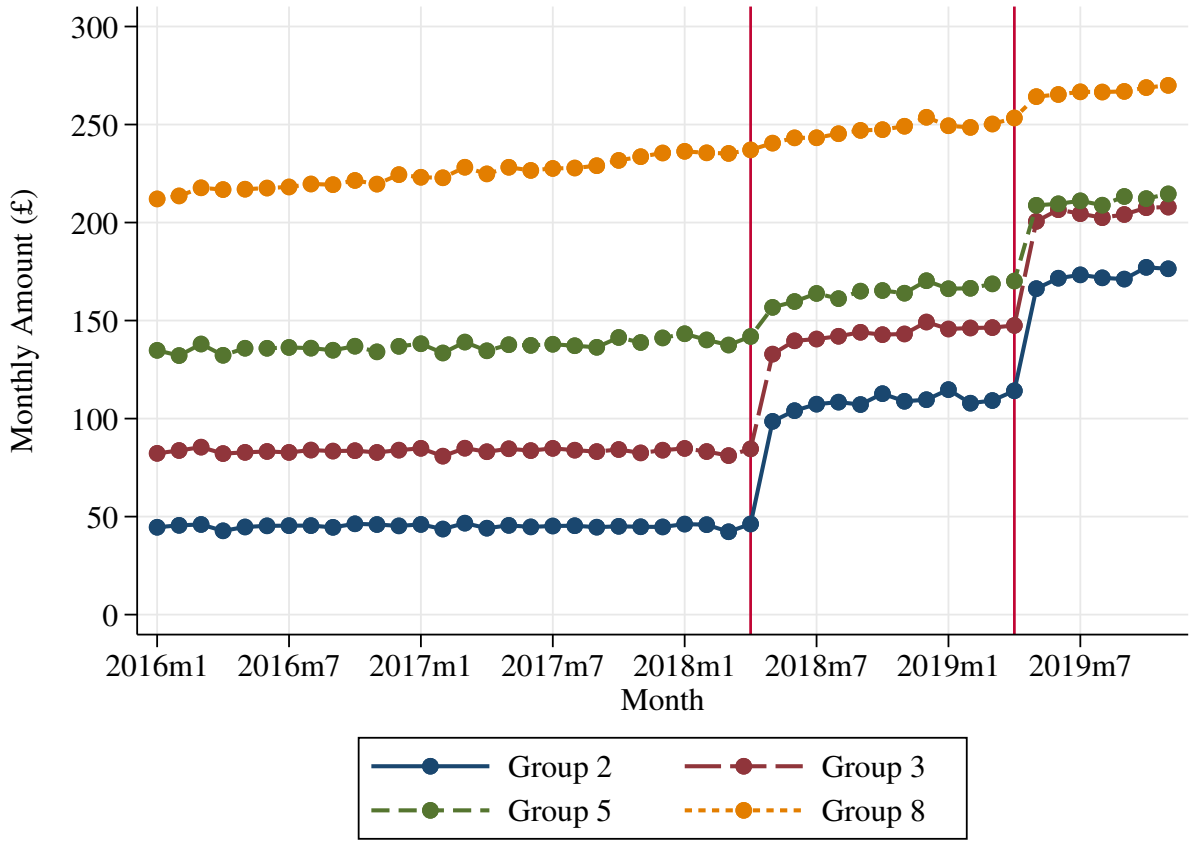
Notes: The figure plots the probability density of contribution rates in March 2018, with each contribution rate group plotted in a different color. See section 4 for an explanation of the contribution rate groups.

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



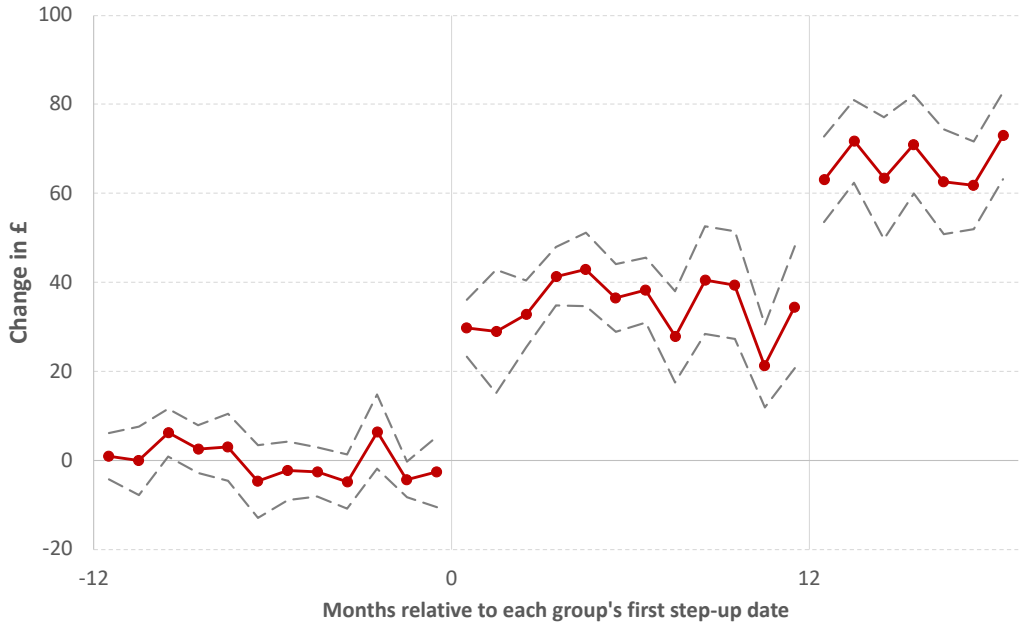
Notes: Figure plots the average employee (red) and employer (green) contribution rate for each contribution rate group for the subsample of workers with data differentiating employer and employee contributions. Contribution rate groups are defined based on each employee's March 2018 total contribution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

Figure 4: Average Monthly Pension Contributions by Contribution Rate Group



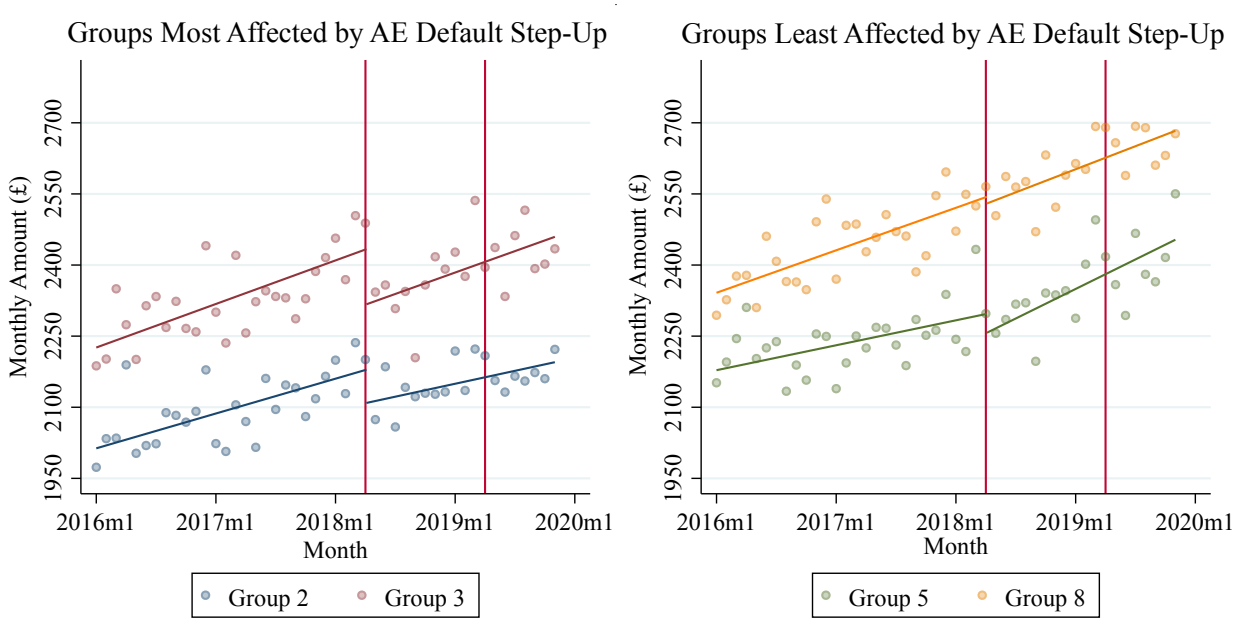
Notes: Figure plots the average monthly contributions by contribution rate group. Apart from the criteria in Appendix A, we also exclude months where an individual changed jobs, had multiple employers, or where the contribution amount was zero or in the top 5% of contributions. The vertical lines indicate the increases in the default pension contributions on 6 April 2018 and 6 April 2019.

Figure 5: Event Study of Change in Monthly Contributions



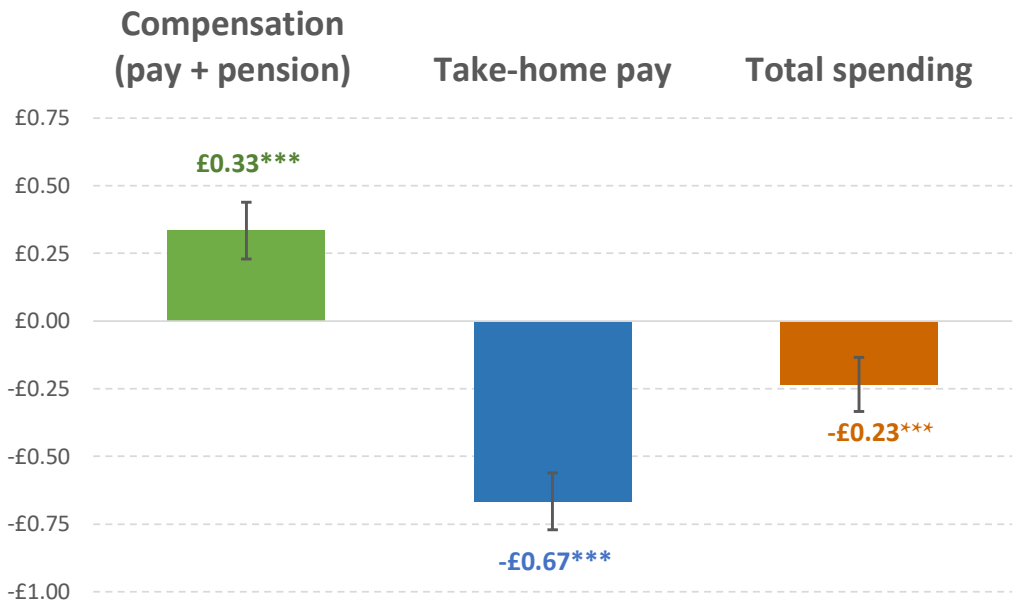
Notes: Figure plots the average change in monthly retirement contributions for treated contribution-rate group workers relative to control-group workers using the Sun and Abraham (2021) estimator, normalizing March 2018 to zero. Dashed lines plot 95% confidence intervals.

Figure 6: Average Monthly Net Wage Income by Contribution Rate Group



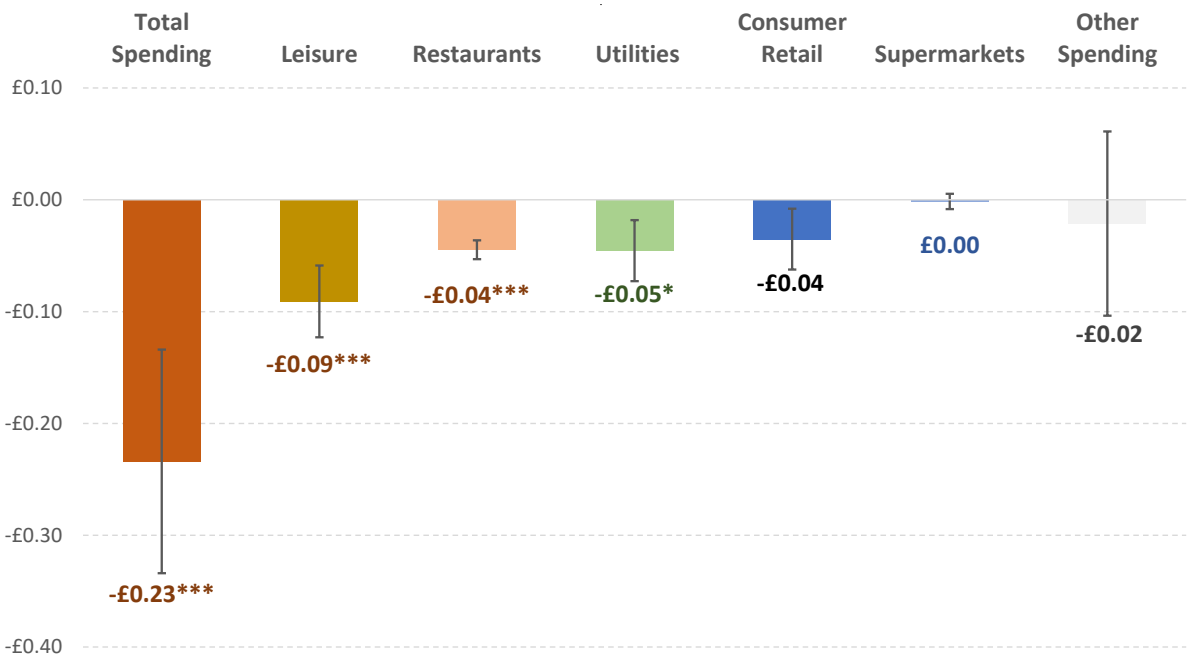
Notes: Figures plot the average monthly net wage income (take-home pay) by contribution rate group using the same data as Figure 4.

Figure 7: Efffect on Total Compensation, Take-home Pay and Total Spending



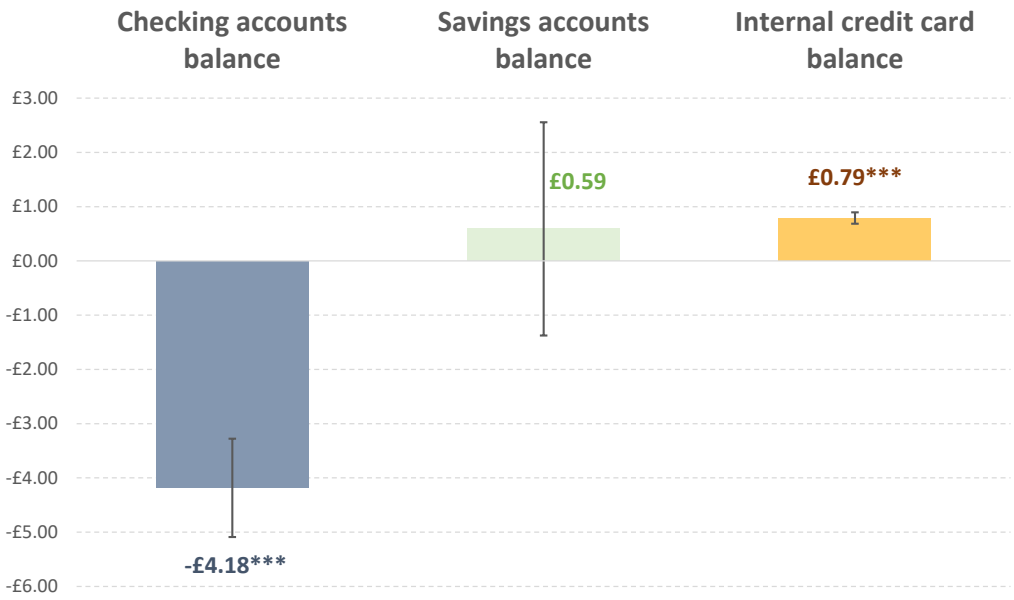
Notes: Plotted coefficients are the instrumental variables coefficients from a regression of monthly total compension (i.e., the sum of take home pay and employee and employer pension contributions), take-home pay and total spending on instrumented pension contributions along with error bars indicating 95% confidence intervals. .

Figure 8: Effects on Spending by Expenditure Category



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

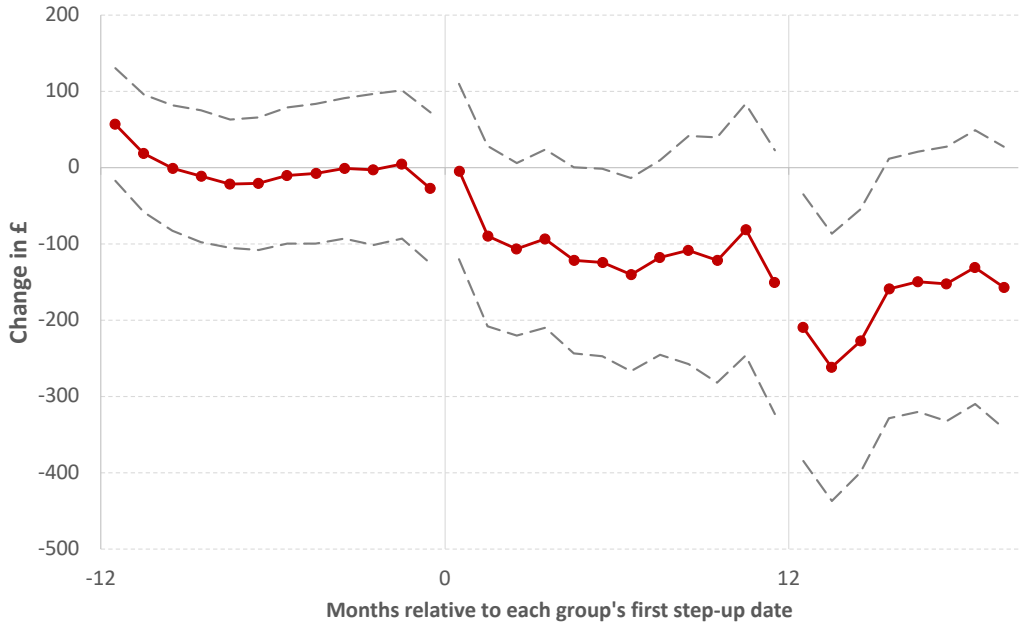
Figure 9: Effects on Deposit and Credit Card Balances



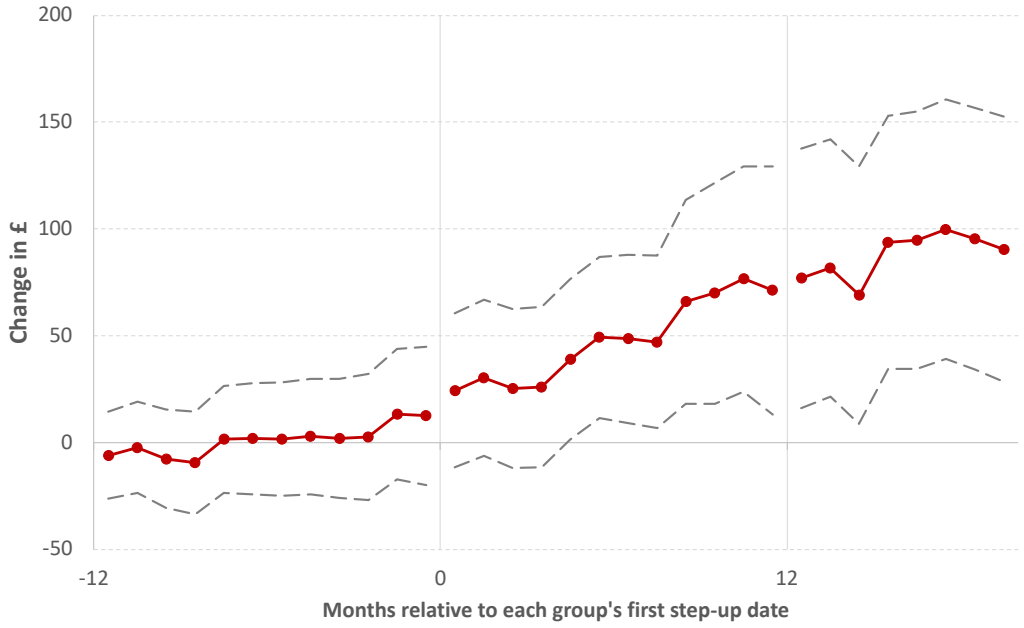
Notes: Plotted coefficients are the instrumental variables coefficients from a regression of account balances on take-home pay along with error bars indicating 95% confidence intervals.

Figure 10: Event Study of Change in Checking Account and Credit Card Balances

I. Effects on Checking Account Balances

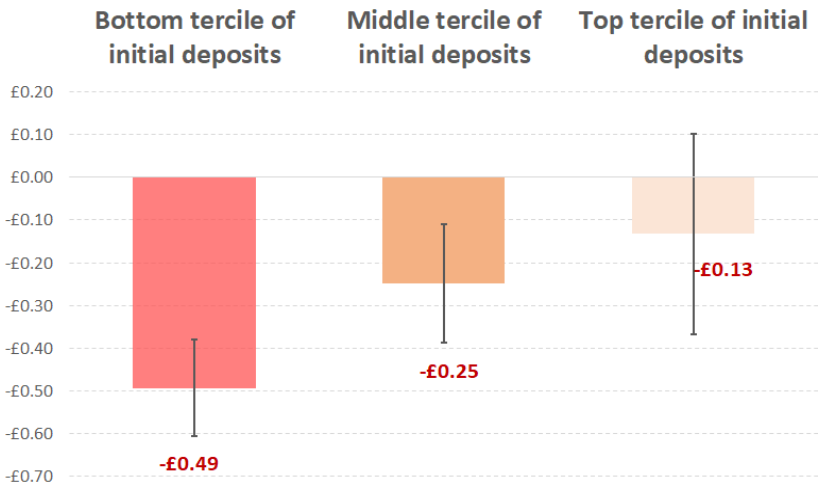


II. Effects on Credit Card Balances



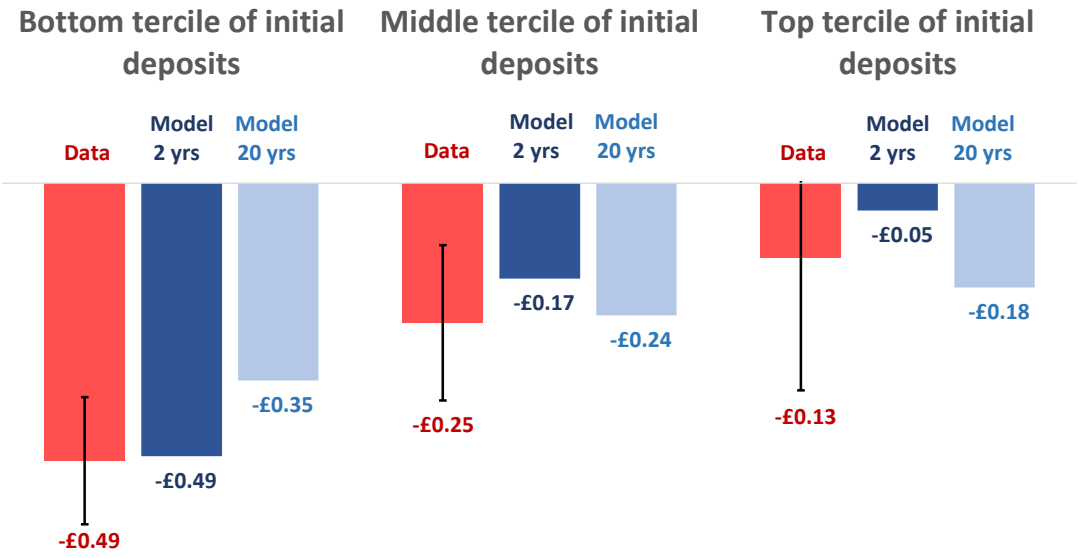
Notes: Figures plot the average change in checking account balances (panel I) and credit card balances (panel II) for credit cards issued by our partner financial institution for treated contribution-rate group workers relative to control-group workers using the Sun and Abraham (2021) estimator, normalizing March 2018 to zero. Dashed lines plot 95% confidence intervals.

Figure 11: Effects on Total Spending by Deposit Tercile



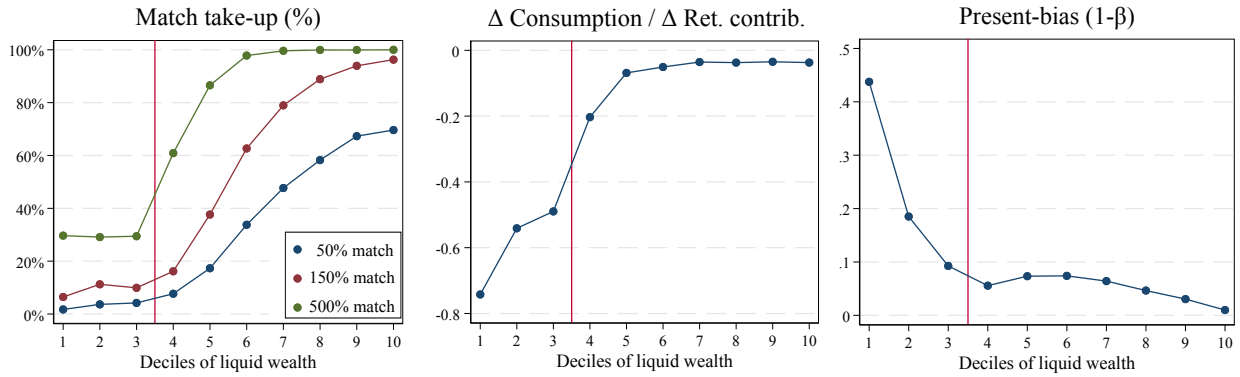
Notes: Plotted coefficients are the instrumental variables coefficients from a regression of total spending on pension contribution amounts for each tercile of March 2018 deposits level along with error bars indicating 95% confidence intervals.

Figure 12: Simulated Effects on Total Spending by Deposit Tercile



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 11. Dark blue 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. Light blue bars report simulated effects after 20 years.

Figure 13: 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 veritcal line are deciles with net negative liquid wealth (i.e., unsecured debt holders). The middle (right) panel corresponds to the average spending reduction per addionatl 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. Income and Spending			
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
II. Debt and Balances			
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
III. Individual Attributes			
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	3.4	6.0	11.0
	(0.28)	(0.57)	(0.88)	(2.24)
Net Wage Income	2101.1	2478.8	2567.8	2471.6
	(2322.2)	(3089.0)	(3000.6)	(1990.3)
Pension Contribution Amount	41.5	84.9	153.3	270.7
	(46.2)	(110.4)	(181.1)	(218.8)
Total Spending	1248.8	1387.6	1389.2	1447.4
	(1831.0)	(1767.4)	(2083.7)	(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. Income	
Net Wage Income	-0.67*** (0.10)
II. Spending	
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)
III. Debt Payments	
Credit Card Payments	-0.22*** (0.06)
Mortgage Payments	-0.02 (0.05)
Other Loans Payments	0.02 (0.04)
IV. Net Account Flows	
Net Current Account Inflows	-0.34 (1.07)

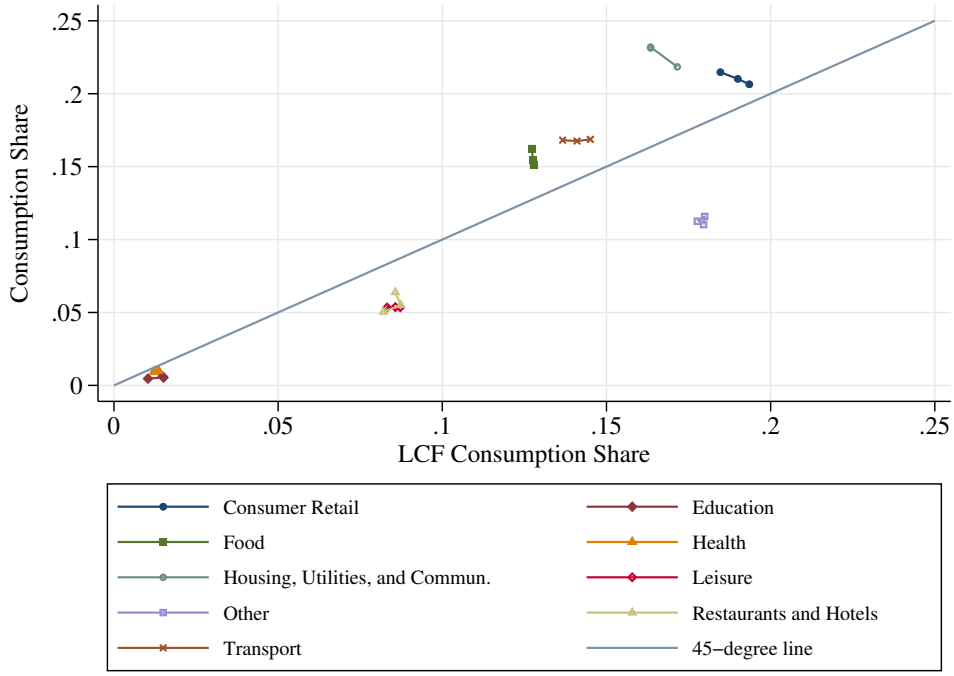
Notes: Table reports instrumental variables regression estimates of the effect of pension contribution increases on income (panel I), spending (panel II), debt payments (panel III), and net account flows (panel IV). The estimation sample is described in Table 1 and contains 1,534,654 monthly observations from 44,122 individuals. Total Spending equals Spending via Current Accounts plus Spending via Credit Cards. Loan Payments represent debt payments that are not included in total spending. Robust standard errors in parentheses. Individual and calendar-month fixed effects included in all models. * $p<0.10$ ** $p<0.05$ *** $p<0.01$

A Data Appendix

In the raw data, observations are at the account-month level. We aggregate information on pension contributions, net flows, deposit account balances, and demographics from the different accounts of each individual to get a panel of workers at the individual-month level as follows. First, we drop observations where pension contributions are missing or where net wage income is zero (since the ratio of pension contributions to income is therefore undefined). We treat "leading zeroes" in pension contributions, continuous sequences of pension contributions of zero for the first months of each pension account, as missing. This is because these zeroes can indicate misreported pension contributions – in particular, at times the zero pension contributions are outright impossible, such as when the pension balance jumps up from zero but the pension contributions remain at zero. We also note that some individuals have multiple pension accounts in the same month; this is often due to individuals moving from one pension to another, e.g., from switching jobs. When this happens, for that individual, we drop the month and the adjacent two months since the act of moving to another pension potentially introduces measurement error for the surrounding months, e.g., pension contributions misreported as positive despite the pension balance going to zero. We also only keep individuals who had a pension contribution of 1.5%-15% in March 2018. Typically, the contribution rate must be set as an integer percentage, so we round the contribution rates to the nearest integer; the contribution rate of 1.5% for the lower bound is chosen because it corresponds to a rounded rate of 2%, under the rationale that individuals whose pension contributions are below the default minimum of 2% have most likely opted out of the default contribution amounts and therefore would not been affected by the autoenrollment step-up. The upper bound is set at 15% in order to make group 8 more comparable to the other groups. Finally, the sample window is restricted to January 2016 – November 2019. The last month in the raw data is December 2019, but this is dropped because we only have observations up to mid-December. Whilst 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. Each of these steps is taken to form a more representative sample of individuals who would be potentially impacted by the treatment and a more representative control group that is likely to be otherwise similar on observable dimensions in the time period close to the step-up dates.

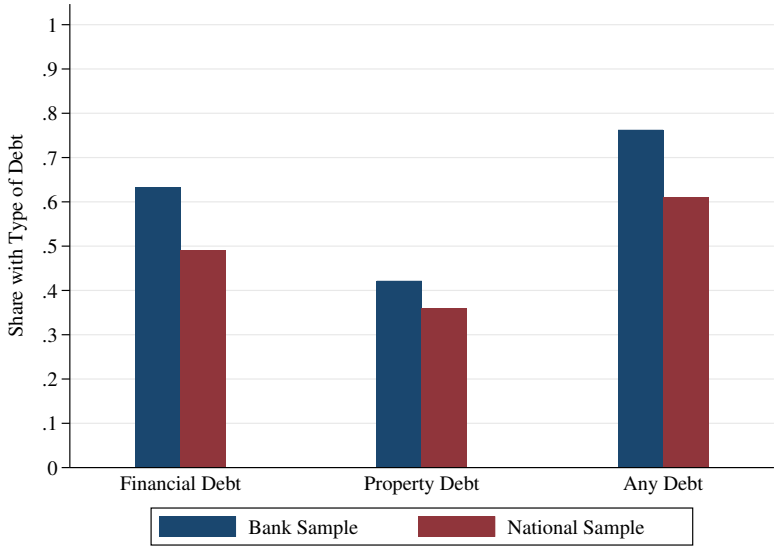
We then consider restricting the panel to a subsample of individuals for which the data provides a complete picture of finances. More specifically, we limit to the sample to individuals who use their current account for day-to-day transactions and are therefore adequately captured in the data as described in Table A1.

Figure A1: Benchmarking Sample Consumption Data with National Survey Data



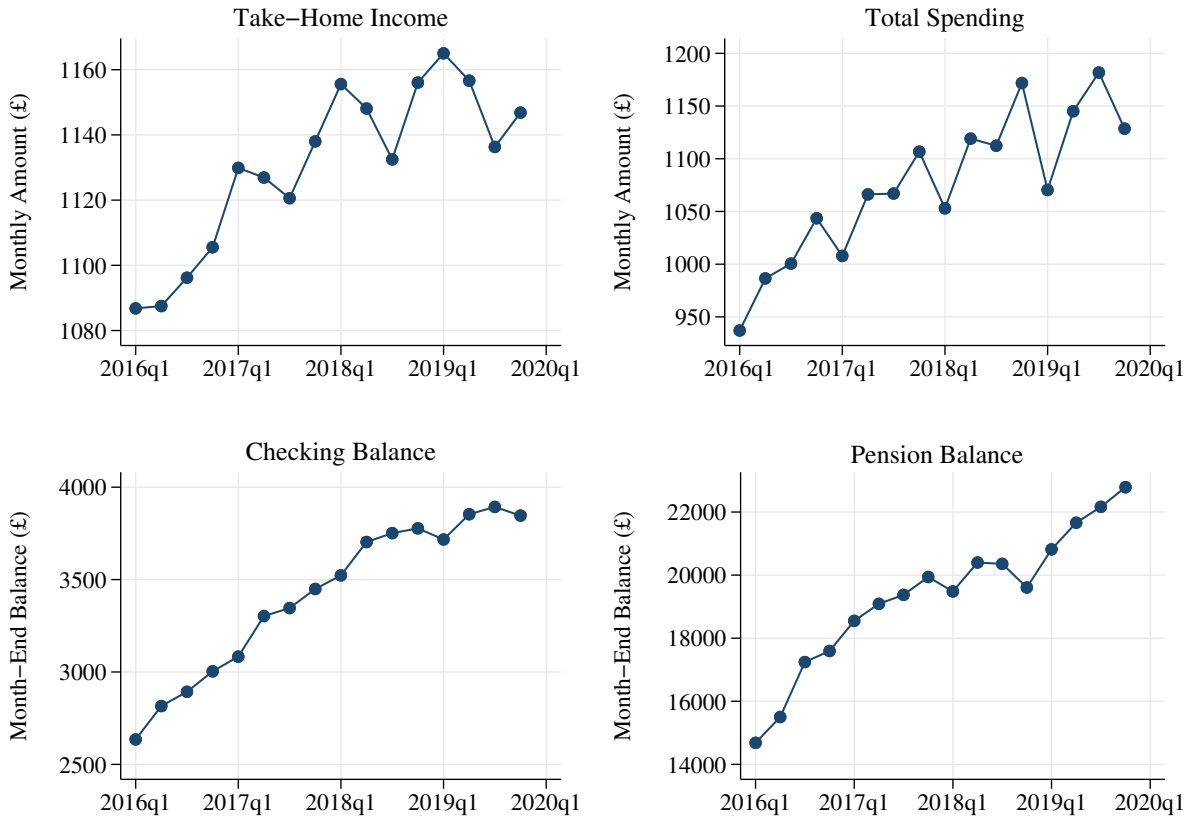
Notes: UK survey data comes from the Living Costs and Food Survey. Each point represents an annual average between the years 2016 and 2018, and the lines chronologically connect the points within each category.

Figure A2: Benchmarking Debt Data with National Survey Data



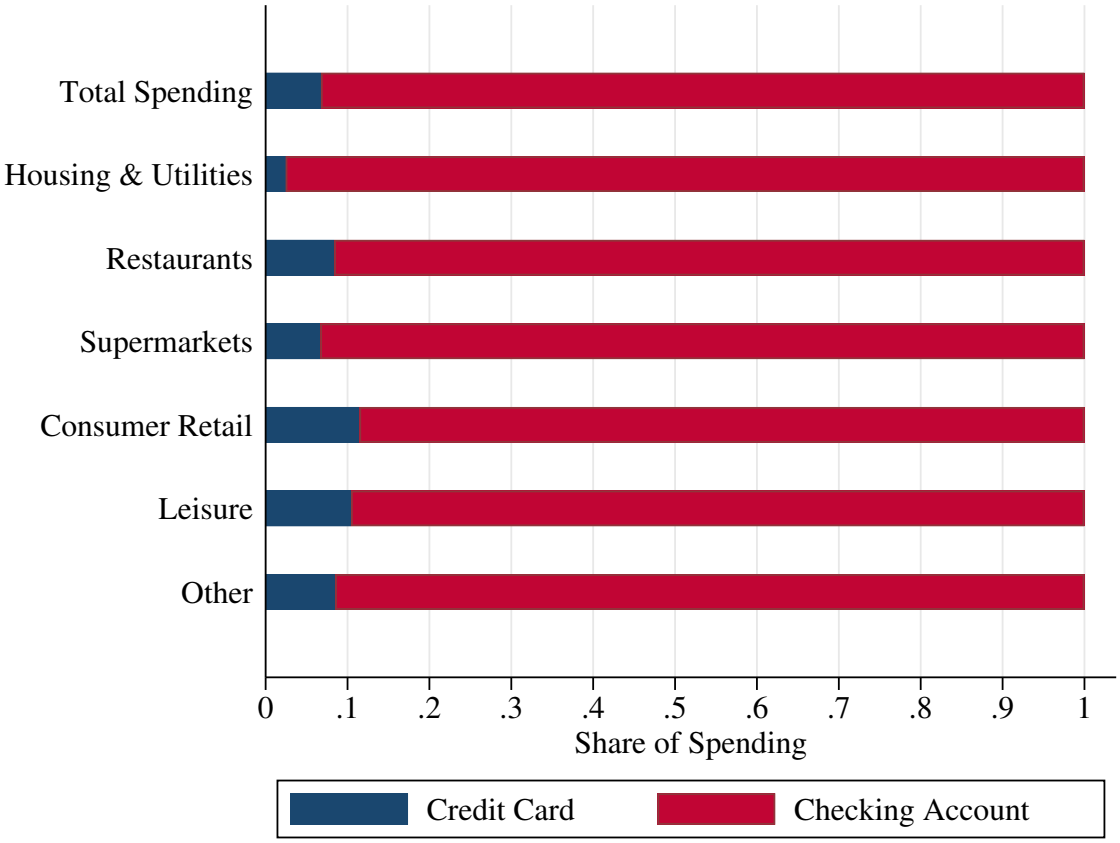
Notes: This figure compares the our data with UK survey data from the Office for National Statistics (https://www.ons.gov.uk/peoplepopulationandcommunity/personalandhouseholdfinances/incomeandwealth/datasets/householddebtwealthingreatbritain, section 7.1) for the April 2016–March 2018 period. For our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

Figure A3: Quarterly Averages of Key Variables



Notes: This figure graphs the mean of each variable at the quarterly frequency (monthly variables aggregated and averaged at a quarterly level). Net wage income (take-home income) is calculated as after-tax income deposited in a current account or savings account. Total spending is the sum of spending from current accounts, savings accounts, or credit cards.

Figure A4: 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.

Table A1: Data Cleaning Steps

	Step	Number of Individuals	Number of Employers
0	All Observations after Merging and Aggregation	614,313	107,609
1	Keeping Months with Non-Missing Contributions (Ex. Leading Zeroes) and Non-Zero Income	457,549	71,069
2	Dropping Months Associated with Multiple Pension Accounts	452,010	68,813
3	Keeping Customers with a Pension Contribution in March 2018 of 1.5%-15%	106,345	25,182
4	Time: Jan 2016–Nov 2019	106,345	20,690
5	Keeping Customers with a Strictly Positive Average Current Balance	85,135	18,478
6	Keeping Customers with a Strictly Positive Average Net Wage Income	85,132	18,478
7	Keeping Customers Who Ever Have an Observed Credit Card	50,409	14,134
8	Keeping Years with at Least One Strictly Positive Pension Contribution	50,409	13,531
9	Keeping Years with a Spending-to-Income Ratio of at Least 30%	47,688	13,027
10	Keeping Years with Median Spending on Food of at Least £50	44,722	12,630

Notes: "Months" and "years" refer to months and years for each individual. We interpret missing pension contributions as including leading zeroes: continuous sequences of pension contributions of zero for the first months of each pension account. Customer averages for current balance and net wage income are taken over all non-missing observations 2011–2019.