An object stays at rest unless acted upon by an unbalanced force:

How pension auto-enrolment and the force of inertia cause debt

Adrian Brown

Student ID: 14273318

Supervisor: John Gathergood

Abstract

Does being automatically enrolled into a pension lead to an increase in debt? This is vital to investigate since the scheme's purpose is to give people more money for retirement and not for them to have a large pension that is offset by the unintended consequence of more debt. Literature in this area has given seemingly contradictory results. This paper seeks to address this by proposing an explanation for these differences – inertia. This paper explains how using a data set used by Gathergood et al. (2024), it would be possible to test if the theoretical argument of inertia may explain the literature's findings. This paper does not have access to Gathergood et al. (2024) data, so it discusses how the data and hypothesised results could be analysed. If the theoretical argument of this paper is upheld by the data, then this research would provide policymakers with the insights needed to design auto-enrolment pension schemes that effectively avoid the unintended side-effect of debt.

Table of Contents

1. Introduction

Inertia, also referred to as status quo bias in the literature, causes people to have a strong tendency to continue acting as they previously have simply because they cannot break out of their habits. The power of inertia can be crippling to an individual, a community, a nation, and the world; from trying to get started to going to the gym (Charness and Gneezy, 2009; Milkman, Minson, and Volpp, 2013) to climate change (Brulle and Norgaard, 2019), inertia plays a critical role in why people do not change their behaviour.

Inertia in pensions has previously been studied (Madiran and Shea, 2001; Thaler and Benartzi, 2004). In the UK, this research led to the introduction of auto-enrolment nationwide into a pension, which is the focus of this paper.

The UK introduced auto-enrolment in stages starting in 2012 until it was fully rolled out by 2018. The rollout created a natural experiment where some people were enrolled in a pension while others were not. This allowed for research into the effects of auto-enrolment using a difference-in-difference approach. Gathergood et al. (2024) used this approach to investigate the impact of this policy on debt, finding that auto-enrolment is associated with higher debt.

Gathergood et al. (2024) studied firms' employees, with new and existing employees studied as one. I propose that these two groups be investigated separately because of

inertia. The increase in debt found may only result from existing employees accumulating this debt, not new employees. If this is the case, this would be explained by inertia since existing employees would not have taken account of their wage being lower the month of auto-enrolment than the previous month (when they were not auto-enrolled). This would not happen for new employees since they have not been accustomed to a higher monthly salary. This would also explain the seemingly contradictory findings of Beshears et al. (2022) and Choukhmane and Palmer (2023), which will be discussed in detail in Chapter 2.

The significance of investigating these two groups separately is that if only the existing employees accumulate debt, it implies that auto-enrolment only causes a transitory increase in debt, with no long-term effect for new or existing employees. The reason it is transitory is due to liquidity constraints as is further discussed in Chapter 7.

This paper thus seeks to investigate whether the debt accumulation found by Gathergood et al. (2024) is likely to be a permanent or merely a transitional problem arising from auto-enrolment into a pension and inertia.

This paper will first explore the existing literature on inertia and pensions before discussing the background of the auto-enrolment policy and the experiment. It will then describe the data, outline the econometric model, describe what the results would look

like, and discuss the findings that would be correct according to this paper's theory. The paper will then conclude.

2. Literature Review

Decisions made by real people are different from those in standard economics texts, with humans seemingly violating or not acting per standard assumptions. One such area is inertia, as people have been shown to disproportionally stay with the status quo in experiments (Samuelson & Zeckhauser, 1988). This inertia has also been found to affect pensions, with Ameriks and Zeldes (2000), using a 10-year panel of data, finding nearly half of the participants in TIAA-CREF, the large retirement plan that catered to university employees, made no changes to their pension plan over a 10-year period.

Having found evidence of inertia in pension savings, the literature sought ways to use inertia to benefit individuals. Madrian and Shea (2001) investigated how changing the default of pension savings from an opt-in to an opt-out system (meaning they are auto-enrolled) affected pension savings. They found significantly higher participation in the pension scheme by doing this. Beyond just the impact on participation through auto-enrolment, they also found that having a well-designed default system is vital because the people who joined under auto-enrolment stayed with the default rate and fund allocation even though few employees pre-auto-enrolment chose this set-up. They assign the reason for this being inertia.

Thaler and Benartzi (2004) created their Save More Tomorrow (SMarT) scheme to address the importance of a well-designed default saving rate and fund allocation. In this design, people commit in advance to save part of their future salary increase into their pension savings. This was to overcome the problem of people having low savings rates and the phenomenon of people not wanting to have their current take-home income decrease. Their study found that 78% of those offered the scheme joined it, 80% remained in it, and the savings rate increased from an average of 3.5% to 13.6% over 40 months.

This type of research, along with the extremely low rates of employees opting into a pension in the UK (DWP, 2023), led to the government implementing the Pensions Act 2008. The details of the Act will be discussed in the next section, but in summary, it meant the default for pensions was switched to an opt-out system. The Act has successfully increased the number of people saving into pensions (DWP, 2023).

The introduction of this Act allowed for further research into the effects of auto-enrolment, including by Gathergood et al. (2024), which, as discussed, is the focus of this research. Gathergood et al.'s (2024) research links in with two other papers that look at the effect of auto-enrolment and debt: Beshears et al. (2022) and Choukhmane and Palmer (2023).

Beshears et al. (2022) studied the effect of auto-enrolment using a natural experiment when the US Army started enrolling newly hired civilian employees into a pension. They find no statistically significant evidence that auto-enrolment affects debt balances, creditworthiness or increases financial distress. All these findings held for subpopulations.

Choukhmane and Palmer (2023) studied the effect of auto-enrolment in the UK using the same policy that enabled Gathergood et al. (2024) and this paper's work. However, they focus on the effect of the increase in the default rate that occurred. The minimum contribution amount increased from 2% to 5% in April 2018 to 8% in April 2019. They compared the behaviour of those paying the minimum (so they had to pay more due to the increase) to those already paying above (so they did not have an increase). They found for a £1 reduction in take-home pay due to higher employee contributions, £0.34 was accounted for by a reduction in spending, while the remainder was financed through lower deposit balances and higher credit card debt.

At first glance, the findings of these three papers seem inconsistent since Choukhmane and Palmer (2023) and Gathergood et al. (2024) find that auto-enrolment does cause an increase in debt, while Beshears et al. (2022) find that it does not. However, these findings are precisely what would be expected from the argument of this paper, which is that the effect of auto-enrolment is different depending on whether you are a new or existing employee due to inertia. This is because Beshears et al. (2022) studied new employees (and found no effect), Choukhmane and Palmer (2023) studied an increase

in contribution on existing employees (and found an effect), and Gathergood et al. (2024) studied existing employees being auto-enrolled with a small sample of new employees included in the sample but not studied separately. Hence, this paper aims to separate these two groups to see if they respond differently to auto-enrolment. This is what is implied from the inertia argument and would be consistent with the findings of the other two papers.

This paper will now discuss the policy that makes this research possible and the experimental design that enables testing the effect of auto-enrolment on debt, differentiating between new and existing employees.

3. Policy Background and Experimental Design

As previously discussed, there were low numbers of people saving into a pension, which led to the United Kingdom introducing the Pension Act 2008. This required all firms with at least two eligible employees to automatically enrol all their eligible employees into a workplace pension. An eligible employee is one aged between 22 and state pension age (currently age 66 but increasing to 67 for those born after 5th April 1960; however, for the data sample studied here, it was age 65 for men and 62-64 for women, subject to their date of birth) and earning at least £10,000 a year.

The policy involved over 10 million new people being enrolled into a pension. To help with implementation, the UK created the National Employee Savings Trust (Nest) to

provide low-cost pensions with a public service mandate to serve all eligible firms. Due to the vast numbers, the roll-out occurred in stages by firm size, with the largest starting on 1st April 2012 and going until 1st April 2017 for firms in existence before April 2012 and between 1st April 2017 and 1st February 2018 for firms formed afterwards. The Pensions Regulator (TPR) assigned firms a 'staging date' on which auto-enrolment of all eligible employees must typically occur for firms not already offering an employer-provided pension. Employers could postpone this for up to three months after the staging date, but they had to inform employees of the delay and allow opt-in enrolments in the interim. It was illegal to auto-enrol employees before the staging date. However, employers could apply to TPR to move the staging date earlier or allow employees to opt into making Nest contributions.

Staging dates were assigned by size. The largest firms, those with 120,000 or more employees, were obliged to start the scheme by 1st October 2012; for those with 50,000-119,999, it was 1st November 2012; for those with 30,000-49,000, it was 1st January 2013; and this continued down to firms with 30-39 employees, with a staging date of 1st October 2015.

This research studies firms that existed before April 2012 with 29 employees or less. This is because the staging worked differently for these firms as they were randomised due to the large number of new enrolments for this group. TPR assigned staging dates to these firms between 1st June 2015 and 1st April 2017 (See Appendix A) using the last two digits of their Pay-As-You-Earn (PAYE) number. PAYE numbers are the unique payroll tax

identifier for firms in the UK, and the last two digits a firm has is as-good-as-random. See Appendix B for details of PAYE assignment as defined by Gathergood et al. (2024). Due to the randomisation of PAYE numbers, the staging date a firm got was random, creating a natural experiment since, at any point after June 2015, employees had been exogenously subject to auto-enrolment for longer than others based on the firm they worked for.

All employees were notified that they would be auto-enrolled, starting with a public information campaign with TV and radio adverts. Employees who were auto-enrolled into the Nest pension were also given written communication from the employer and then by Nest with a letter and brochure sent to their home address (as provided by the employer). The information included information about the pension and how to opt out. An evaluation report by the UK government for the automatic enrolment policy indicates an awareness of 74% of the introduction of auto-enrolment amongst the targeted population (DWP, 2014).

The data for this research will be drawn from a linked data file of individual-level pension contributions, employer data from Nest, and credit file data from the credit reference agency Experian. Nest is the UK's largest provider of auto-enrolment pensions. They offer a defined contribution scheme with a choice of investment funds and a default target retirement date fund. It is free for employees to use and has a public service obligation, so any employer can use them to meet their mandatory auto-enrolment obligation. There is no requirement to use Nest, but the vast majority of small firms use

it due to its low cost and public service mandate. As of 31st March 2021, Nest managed 9.9 million pensions for 881,000 employers, representing around a third of working-aged individuals in the UK. Experian is one of the so-called 'Big Three' credit-reporting agencies, which makes it an ideal data source.

A concern about this research may be that it does not identify inertia when separating new and existing employees. Individuals may have different characteristics, causing selection bias, which explains the debt difference. There are two possible cases here: people choose the job due to being able to get a pension, and people who switch jobs are fiscally different. Both mean that new employees would not take on extra debt.

The argument for the first case would be that people may have chosen to join a firm that offered auto-enrolment over one that did not because they would get a pension. If this were the case, they would want to be enrolled into a pension and would have accounted for lower take-home income because they are saving into a pension. Thus, they would not take on debt since they have actively chosen to put some of their earnings into a pension rather than use it for consumption or other savings products. Thus, this paper would not be identifying inertia but two groups: one who wants to save into a pension and one who does not.

The first issue with this argument is that if people did not want to save into the pension, then they could opt out of the scheme rather than take on debt while saving into the

pension. It could be argued that opting out is more cumbersome than mere borrowing, so borrowing is more effortless than opting out. However, if this were the case, the amount borrowed should equal the amount paid into the pension.

A second issue is to consider the purpose of auto-enrolment: to increase people's pension contributions. As discussed previously, the number of employees who opted into a pension was meagre. Given the extremely low uptake of a pension, it seems implausible that job selection was influenced by a firm's staging date.

A third issue is the effort required to determine when a firm offers an auto-enrolment pension. To know a firm's staging date, a person would need to know the firm’s PAYE number and the date at which that number was being staged, as well as ensure that the firm had not requested their date to be changed. While this can be found, it would require someone to know the information just listed, which would be cumbersome, so an individual would unlikely have done this.

Therefore, assessing the argument that there is selection bias due to firm selection based on offering a pension holds little substance and is unsound.

The second case is that people who switch jobs are systematically different from those who do not, and this systematic difference includes a difference in managing their finances. The argument is people who switch jobs have a reason to switch jobs, and one

of those reasons is to get a higher salary. People who are more fiscally prudent or money-motivated are, therefore, more likely to switch jobs, but they are also more likely not to acquire debt. Debt here refers to that which is accumulated due to overspending (this type of debt is expected to be seen in the form of credit card or overdraft debt since this debt is typically consumption spending) and not other types of debt, which may be investments (which would likely be debt in the form of a mortgage). These people do not acquire debt because they ensure their outgoings are within their earnings. This would mean that if the results showed that debt was lower for new employees than existing employees, it would not result from inertia but merely measure one group of people being more fiscally prudent than another.

However, if there is a difference between new and existing employees caused by fiscal character differences rather than inertia, this will be able to be seen from the data. Since it would be in their character to be fiscally prudent, they would have been this way before auto-enrolment. This would be revealed by examining their credit files before starting their new job and comparing this to the existing employees' credit files before auto-enrolment. While controls for characteristics such as age and gender are needed, if the data shows no significant difference between the credit files of new and existing employees, then it would mean the argument of fiscal prudency holds no substance. Ergo, the difference would have to be the result of inertia.

An alternative argument could be that the difference between new and existing employees may be that when starting a new job, individuals reassess their financial

situation to ensure that their spending is within their means so as not to take on debt. This, though, supports the argument of this paper since what causes the new employees to reassess their finances is that they have broken free of their inertia by starting a new job. Ergo, this argument is still an argument of inertia but one in which people have managed to overcome it.

It should be noted that before the Pension Act 2008, the UK required employers with five or more workers to offer an opt-in 'stakeholder' pension. The Act removed this obligation, although existing accounts had to be maintained, and employees contributing to one of these were not subject to automatic enrolment. Failure of an employer to comply led to fines and potential imprisonment for company directors.

The data this research would use does not show contributions made to stakeholder pensions. However, the 2015 Annual Survey of Hours and Earnings (ASHE) data showed that only 2.1% of workers in firms with one to 12 employees and 5.2% of workers in firms with 13 to 99 employees had a stakeholder pension.

This paper has been through the policy background and how it enables this paper's experimental design. It then establishes that the identification and selection bias argument is flawed, meaning if there is a difference between new and existing employees, this results from inertia. It will now discuss the data used to investigate if inertia is the reason for the accumulation of debt found by Gathergood et al. (2024).

4. Data

As stated, this paper does not have the data. Still, since it would use the same data set as Gathergood et al. (2024), the data will be explained in line with how the original authors used it and explain how this paper would investigate the difference between new and existing employees.

4.1 Sample Selection

The empirical strategy uses the random assignment of staging dates of firms with 2-29 employees that existed prior to 1st April 2012, the same as Gathergood et al. (2024). The primary analysis would use eligible workers determined by birth date to ensure eligibility for auto-enrolment and be enrolled in a Nest pension. However, it would separate the existing and new employees (those employed up to three months after the firm-reported staging date). This means the same analysis will be done as in Gathergood et al. (2024), apart from separating the new and existing employees.

Since the same data set will be used as Gathergood et al. (2024), it is known that some (14% in the sample) chose to opt out of the Nest pension before the one-month deadline. Nest holds basic information on those who opt out, including age. Opt-out increases with age (7.6% for those under 30 to 31.4% for those aged 60). There is a decline in the opt-out rate over the sample period for different staging cohorts (16% in June 2015 to 13% in April 2015). Those who opt out are excluded from the analysis.

Like Gathergood et al. (2024), four sample restrictions will apply to the employee records provided by Nest. First, the employees of firms who registered for a PAYE reference before 1st April 2000 were dropped since this system did not give as-good-as-random assignments for the last two digits. Second, only employees for whom a credit file could be matched were kept. Third, any firm that reported a staging date to Nest that was outside the feasible staging, as listed in Appendix A, was dropped. Finally, only those individuals whose birth date made them eligible for auto-enrolment at every possible firm staging date were kept in the sample.

The initial sample from Nest has 712,818 employee records across 173,570 firms, and after the four restrictions, the number of employees is reduced to 91% of the starting sample and 93% for firms. Thus, a baseline sample of 649,747 employees was made across 161,707 firms.

For this paper to investigate the difference between new and existing employees, the two need to be separated. Unfortunately, the data set does not have employment records stating when an employee started the job. For this, the data from His Majesty's Revenue and Customs (HMRC) would be needed as they keep employment histories with start and end dates. Thus, the proposal is to create the variable by taking the date at which a firm enrols the majority of its employees, which will be 'existing employees'. The reason for this is it is feasible that firms will have wanted to test the auto-enrolment system and so initially enrolled just one or two of the firm's employees and then do the rest altogether, knowing that the system runs smoothly. The data would be checked to

see if this is likely, and if it is not, then the research would use the date the first group were enrolled. However, for the rest of the paper, the date when the majority were enrolled will be assumed to be the better method. It will consider any employee enrolled into the pension on a date after the majority were enrolled as a 'new employee'. Therefore, individuals enrolled after this are likely to be new employees and so are classified as one. It is estimated that by this method, around 4% would be new employees based on an exploratory look at the data by Dr Christopher Firth. This means around 26,000 employees would be new.

4.2 Variables

The age, gender, individual-level monthly observations of pensionable pay, employer and employee contributions, tax relief, and accumulated pension balances for each individual from the Nest data. Employee addresses are obtained from employer records.

From the Experian credit file, there are the following variables: the Experian UK general-purpose banking and financial credit score; debt measures for mortgage debt, monthly mortgage payment due total vehicle loan debt, and total unsecured debt, which is divided into revolving and non-revolving debt; Experian's estimate of an individual's gross annual income; and financial distress as measured by a flag for whether the individual filed for bankruptcy within the past six years and a flag for whether the consumer entered default within the past six years.

4.3 Summary Statistics

The summary Statistics will be the same as those of Gathergood et al. (2024) but will also include summary statistics for new and existing employees rather than just the combined stats. It can thus hypothesise the likely differences between new and existing employees compared to the whole data set.

The average age is likely to be lower for new employees than existing ones due to younger people being more likely to start new jobs than older people. Thus, the average age of 43, as found by Gathergood et al. (2024), will be higher than the average for new employees but around the same for existing ones. The 41% of the sample that are female are likely to remain about the same for new and existing employees since there is no strong case for a substantial difference.

The average monthly Nest pension contribution of £29 and median of £21 will likely differ for new and existing employees. This is because new employees tend to be younger, and so are earlier in their careers. By being earlier in their careers, their earnings are typically lower than older employees, resulting in lower contributions to their pension for new employees compared to existing ones. The mean credit score is 935 with standard deviation of 183 are hypothesised to be similar for both groups since being a new employee does not inherently seem to indicate a difference in credit score. The mean income is £35,916, and the median income is £29,889 (this is close to the UK population median employee earnings of £29,588 in 2018), but this will likely be

different for new and existing employees since new employees are likely to have moved for a pay rise, and so may have higher averages. However, given that they are likely to be younger, the figure may be lower due to younger people being likely to be paid less due to factors such as experience. 98.6% of individuals in the sample have bankruptcy and delinquency information, with 1.4% filing for bankruptcy in the past six years and 16% defaulting within the past six years. This is likely to be similar for new and existing employees, given that starting a new job does not directly link to these issues. Still, if they are getting higher income by starting a new job, then it may be the case that they have more money, and so are less likely to default or go bankrupt than an existing employee.

Around a third of individuals have a positive mortgage balance with a mean value of £52,095 and mean monthly payments of £305. This will likely be different for new and existing employees, especially if there is a big difference in age since younger people are less likely to have bought a house and so less likely to have a mortgage. For vehicle loans, 11% have one with a mean outstanding debt of £1,253. There is not a strong case for the difference between new and existing employees, given that a new employee may be likely younger, so they may be more likely to be paying off debt. While an existing employee may have been more likely to get offered the loan in the first place meaning they have more debt on average since more of them have a loan. For unsecured debt, 67% of the sample hold this with a mean total balance of £3,836 and a median of £482. For revolving debt, the mean is £1,630, while for non-revolving debt, it is £2,206. It will be interesting to see the difference between new and existing employees as the inertia

argument explains why there may be a difference once enrolled. Still, if the debt is substantially different in the summary statistics, it may be due to a difference in willingness to lend.

4.4 Compliance with Staging Dates

A fundamental assumption of the research is that firms enrolled employees on or close to the staging date they were assigned. Gathergood et al. (2024) showed that 90% of employers have self-reported staging dates that are the same as those they were assigned. Due to their work, there is high confidence that the randomisation of assigned staging dates does give rise to exogenous variation in outcome.

4.5 Nest Contributions Prior to the Staging Date

An employee may have changed firms, and if they were auto-enrolled in a firm with more than 30 employees in April 2012, they would not be shown in the sample. This would lead to an underestimation of the effect of automatic enrolment since treated individuals are in the untreated control group.

Gathergood et al. (2024) looked at this by exploring if individuals had positive Nest balances prior to their staging date, with only 3.5% in April 2017 (the last staging date) and 1.4% in June 2015 (the earliest staging date). This is very low and is not a concern when comparing new and existing employees if the levels are similar.

4.6 Tests of Employee Characteristic Balance Across Staging Dates

The randomisation of when firms would auto-enrol should lead to a balance in employee characteristics across the staging dates. Gathergood et al. (2024) checked to ensure this using three sets of tests: balance in frequencies, balance of birth and gender, and balance of characteristics. All of these suggest that it is random. This also provides further evidence counter to an earlier point about job selection to get a pension, which suggests this is not the case. To build further on this, it will be useful to check the balance of new employees at each stage to ensure they do not all happen at the beginning of the staging dates, which would suggest switching jobs to get the pension, thus removing the randomisation that is presupposed for this work.

A potential complication may also be an individual being employed by more than one sample company on the staging dates. Thus, like Gathergood et al. (2024), this research will assign individuals to the earliest possible staging date applicable to them. This potentially creates a more stringent selection for later assigned staging dates. However, it was found that this was a tiny issue, with the fraction of employees who could be assigned to the last three staging dates being around 99% for all three.

5. Econometric Model

This paper will use the same two econometric models as Gathergood et al. (2024) but would run them first with existing employees and then with new employees to allow for a direct comparison with Gathergood et al. (2024).

Outcomes are observed for November 2016, November 2017, and November 2018. Let Y_{it} be the outcome for individual i observed at month t. The regressor of primary concern is the number of months that have elapsed between i's staging date, and the observation $m_{it} = min\{0, t - j_i\}$, where j_i is the staging of the firm that employs i. This variable will be called 'months post-enrolment', as most employees are enrolled on their staging date.

The first model flexibly estimates the effect of time since auto-enrolment using a set of months-post-enrolment dummies, with controls for age, gender, and observation date:

$$Y_{it} = \sum_{m=1}^{M} \lambda_m \mu_{mit} + \delta_t + X_i'\beta + \epsilon_{it} \ (1)$$

λ_m multiplies the dummies μ_{mit} for whether the number of months i is post-enrolment as of t is equal to m. M is the number of unique values of m_{it} in the sample. δ_t is a calendar date of observation fixed effect. X_i is a vector of individual covariates (dummies for one-year-wide age groups measured as of November 2017 and a gender dummy). ϵ_{it} is the residual.

The second model has a linear functional form on the effect of m_{it}, generating a summary measure of auto-enrolment's effect over time:

$$Y_{it} = \gamma m_{it} + \delta_t + \epsilon_{it} \ (2)$$

Both regressions have clustered standard errors at the employer level. Due to calendar date fixed effect controls, the regressions use cross-sectional variation to identify λ_m

and γ, combining estimates from the 2016, 2017, and 2018 cross-sections. For each of the three observation dates, the value of m_{it} is exogenously determined, allowing for the identification of the effect of auto-enrolment.

The models would be estimated on employees who did not opt out, meaning the coefficient estimates are a treatment effect on the treated, comparing the exogenous difference of being enrolled later than earlier.

6. Expected Results

This paper would estimate equations (1) and (2) using the baseline sample of 649,747 individuals. To allow for a direct comparison, the same outcome variables as Gathergood et al. (2024) will be used, grouped into debt outcomes (unsecured debt, mortgage debt, and vehicle debt), creditworthiness outcomes (bankruptcy, default, and credit score), and savings outcomes (cumulative Nest pension contributions). However, Equations (1) and (2) will be run separately for new and existing employees to allow for the comparison that is the concern of this paper. Due to the likely differences in characteristics mentioned, particularly age, the regressions would also be rerun with a subset of the sample where key characteristics, such as age and pay, are matched. This would allow for a direct comparison between like-for-like groups meaning the difference can only be explained by being auto-enrolled into a pension. The paper will now discuss the expected results based on what Gathergood et al. (2024) found and the hypothesised difference created by inertia.

6.1 Debt

6.1.1 Unsecured Debt

Gathergood et al. (2024) found an approximately linear increase in unsecured debt months after enrolment, driven by non-revolving debt. There was a £7.30 increase in total unsecured debt per month post enrolment, with £5.66 being non-revolving debt. For revolving debt, there was a £1.79 increase in credit card balances per month post-auto-enrolment, which is only 12 to 14% of the monthly increase in employee pension contributions and was statistically insignificant.

Non-revolving debt includes unsecured loans, overdrafts, sales agreements, and other products. Of the increase in non-revolving debt, 62% was from unsecured loans, which was statistically significant, while 23% was from overdrafts, which was statistically insignificant.

In this setting, it would be expected that the new employees' unsecured debt would not increase, while the debt found by Gathergood et al. (2024) would be seen for existing employees. This will be because of overspending resulting from inertia.

6.1.2 Secured Debt

Mortgage debt balances increase approximately linearly by £118 per month post enrolment. There was no change in monthly mortgage payments or vehicle debt

balances. The probability of having a mortgage increases with months post enrolment and was statistically significant, with no effect on vehicle debt except for the pilot group of June 2015 at 41 months. The probability of having a mortgage increases by 0.046 percentage points per month or a cumulative 1.9 percentage points at month 41 from a baseline of 38%.

The results may show that new employees drive this increase when investigating new and existing employees. The reason for this is the life cycle argument that life has stages, and once you have a pension, you then move on to getting a house. While this would affect both groups, the difference would be driven by inertia since new employees are reviewing their finances and starting a new job which gets them to think about their life cycle. This is because it would be logical that new employees would be looking to buy a home more than existing employees who would be stuck in their ways by inertia.

6.2 Default and Credit Score

Defaults were found to decrease by 0.04 percentage points per month, meaning after 41 months, there was a 10% decrease in defaults compared to the baseline rate. Credit scores increased by 0.3 points per month. There was no statistically significant effect on bankruptcy filings. Overall, this shows a modest improvement in creditworthiness resulting from auto-enrolment. This could be a mechanical result of the increase in

mortgage holdings, while the decrease in defaults suggests a reduction in financial distress.

If mortgages are driving the improvements in creditworthiness for new employees, then this group should see a more significant improvement. As for financial distress, both would likely see this improvement if they are both maintaining payments, as the overall data suggests.

6.3 Pension Contributions

The total pension contribution was estimated at a lower bound of £32 and an upper of £38, with £16-19 from the employer, £13-15 for the employee, and £3-4 in tax relief. This means that unsecured debt offsets 19-27% of total pension contributions induced by auto-enrolment. Given that a mortgage is used to purchase a house, which is an asset, the extent to which the increase in mortgage origination caused by auto-enrolment crowds out net savings cannot be assessed.

The total contributions for new and existing employees are expected to be the same, but the percentage lost to unsecured debt is expected to be higher for existing employees.

6.4 Heterogeneity in Treatment Effects by Income, Credit Score, and Age

Gathergood et al. (2024) found evidence that the effect of auto-enrolment on debt depends on the availability of credit and the possession of significant financial assets. The increase in mortgage holdings is concentrated among those with below-median income, credit scores, or age. Pension contributions are higher for individuals with higher incomes, credit scores, and age. Ergo, auto-enrolment benefits are more favourable for higher-income and older individuals.

When the difference between new and existing employees is analysed, it is expected that these findings will not hold for new employees because they have overcome inertia and are benefitting fully from auto-enrolment. This means that the people who benefit most would be new employees, and then it would be the existing employees, in line with Gathergood et al.'s (2024) findings.

7. Discussion

If the results found are in line with this paper's hypothesis, it would explain the difference in the literature between Choukhmane and Palmer (2023) and Gathergood et al. (2024), who found that auto-enrolment led to an increase in debt held, and Beshears et al. (2022), found no effect.

The broader implications of the implied inertia from this research are both a cause for optimism and concern. The optimism is that auto-enrolment does not lead to new employees accumulating debt; instead, they will reduce their consumption or other savings products (or a combination of the two) to fund the pension contribution (compared to if they had not been automatically enrolled). This is good news as concerns that individuals are acquiring assets and liabilities in parallel do not appear to result from auto-enrolment.

The concern of these findings is the effect on existing employees. As shown here, existing employees have not had the full benefit of being auto-enrolled into a pension (while a new employee has) since they have accumulated some debt, partially offsetting the gain of pension savings. This has a more minor transitional problem and a more significant future implication. The more minor transitional problem is that existing employees will have acquired debt, harming their future ability to consume as they have to service this debt, especially when it is on higher-interest products such as credit cards and overdrafts. This will have caused some hardship for people in the interim. However, the debt is relatively small, and upon realising the growth in debt, individuals will overcome their inertia, reassess their finances, and start spending within their new budget. This will have to occur at some point due to liquidity constraints at the very latest point, and most likely, it will occur before this due to interest payments to service the debt or simply a realisation that the debt is there when individuals review their accounts. Hence, this issue will be transitory and will correct itself.

The more significant concern here is for future implications and is drawn from the transitory issue. If the government wants to get people saving a higher proportion of their income into a pension, and so decides to raise the minimum amount to 10% or 12%, there will be harmful side-effects of this policy – an increase in debt, just as was found by Choukhmane and Palmer (2023) when the minimum was raised. Therefore, there will need to be a smart policy design for increasing the minimum saving amount if the government wants to avoid the danger of debt accumulation. The solution to this already exists – SMarT.

As previously discussed, the policy developed and studied by Thaler and Benartzi (2004) exploited many behavioural insights we know of, such as loss aversion and inertia, to design a policy whereby you commit to increasing your pension contributions in the future, and this increase comes from future pay rises rather than an individual's existing salary. The government can adopt this policy by mandating an increase in the minimum deposit amount but allowing firms, for example, a year to implement it. They can then use employee pay rises over that period to increase the amount that needs to be saved up to the required minimum without it coming from their existing salary. This should result in no (or negligible) debt accumulation from the increase in the mandated minimum contribution by avoiding the inertia problem discussed in this paper. Thus, removing the transitory issue of debt accumulation results in the policy having maximum benefit to employees.

Having discussed what the predicted results would mean for policy, this paper will now move on to conclude.

8. Conclusion

The title of this paper is a paraphrase of Newton's first law of motion, with the idea that an object, or in this case a person, will stay at rest, remain in inertia, until acted upon by an external force. In this case, it has been argued that a person will continue to spend their monthly wage even when it has been slightly reduced since the force is too small for them to overcome their inertia and so continue to spend as before. However, a new employee does not suffer from this inertia and so does not acquire debt when automatically enrolled into a pension.

This paper proposed this as an explanation for seemingly contradictory findings in the literature, where it has been found that auto-enrolment does lead to an increase in debt (Choukhmane and Palmer, 2023; Gathergood et al., 2024) but also has no effect (Beshears et al., 2022). Using the data set of Gathergood et al. (2024), the paper explained how it would investigate the possibility of inertia being the cause of debt accumulation – by comparing new and existing employees. It justified why a difference between these groups would show the cause is inertia rather than other possible explanations. The paper then moved on to how the data would be analysed, followed by a discussion of what the hypothesised results would mean for individuals and governments looking to introduce or increase the contribution rate of auto-enrolled

pensions. The next stage of this research would be to access the data so that the data analysis can be undertaken to investigate the difference between new and existing employees, as laid out in the data section of this paper.

9. References

Beshears, J., J. J. Choi, D. Laibson, B. C. Madrian, and W. L. Skimmyhorn, 2022. 'Borrowing to save? The impact of Automatic enrollment on debt'. *The Journal of Finance*, 77:1, 403-447.

Brulle, R. J. and K. M. Norgaard, 2019. 'Avoiding cultural trauma: climate change and social inertia'. *Environmental Politics*, 28:5, pp. 886-908.

Charness, G., and U. Gneezy, 2009. 'Incentives to exercise'. *Econometrica*, 77:3, pp. 909-931.

Choukhmane, T. and C. Palmer, 2023. 'How do consumers finance increased retirement savings?'. *Working Paper*.

Department of Work and Pensions and Pensions' Automatic Enrolment Team, 2014. 'Automatic enrolment evaluation report 2014'. Available at: https://www.gov.uk/government/publications/automatic-enrolment-evaluation-report-2014.

Department of Work and Pensions, 2023. 'Workplace pension participation and savings trends of eligible employees: 2009 to 2022'. Available at: https://www.gov.uk/government/statistics/workplace-pension-participation-and-savings-trends-2009-to-2022/workplace-pension-participation-and-savings-trends-of-eligible-employees-2009-to-2022#:~:text=Across%20the%20whole%20economy%2C%20there,(10.7%20million%20eligible%20employees).

Gathergood, J., J. Beshears, M. Blakstad, J. J. Choi, C. Firth, D. Laibson, R. Notley, J. D. Sheth, W. Sandbrook, N. Stewart, 2024. 'Does pension automatic enrollment increase debt? Evidence from a large-scale natural experiment'. *National Bureau of Economic Research*, no. w32100.

Madrian, B. C., and D. F. Shea, 2001. 'The power of suggestion: Inertia in 401(k) participation and saving behaviour'. *The Quarterly Journal of Economics*, 116:4, pp. 1149-1187.

Milkman K. L., J. A. Minson, K. G. M. Volpp, 2013. 'Holding the Hunger Games hostage at the gym: An evaluation of temptation bundling'. *Management Science*, 60:2, pp. 283-299.

Samuelson, W., and R. Zeckhauser, 1988. 'Status quo bias in decision making'. *Journal of Risk and Uncertainty*, 1, pp. 7-59.

Thaler, R. H., and S. Benartzi, 2004. 'Save more tomorrow: Using behavioural economics to increase employee saving'. *Journal of Political Economy*, 12:S1, pp. S164-S187.

10. Appendix

Appendix A: Staging Dates for Firms with 2-29 Employees

Final two PAYE digits	Staging Date
92	1st June 2015
02-04	1st January 2016
00, 05-07	1st February 2016
01, 08-11	1st March 2016
12-16	1st April 2016
17-22	1st June 2016
23-29	1st July 2016
30-37	1st August 2016
38-46	1st October 2016
47-57	1st November 2016
58-69	1st January 2017
70-83	1st February 2017
84-91, 93-99	1st April 2017

Appendix B: Assignment of Employer PAYE Reference Numbers – As defined by Gathergood et al. (2024) in their Appendix A1

Our empirical design exploits the staggered rollout of automatic enrollment across firms. Staging dates were assigned to firms based upon the final two digits of their PAYE reference number. These numbers are assigned to firms in a quasi-random manner, which we explain in this section.

Employer Pay-As-You-Earn (PAYE) references are a unique reference for each firm in the UK used by His Majesty's Revenue and Customs (HMRC) to identify an employer for the purposes of employment reporting and compliance. References are combinations of numbers and letters assigned at firm birth.

Since April 2000, PAYE references take the format of 3 numbers, followed by two letters, followed by 5 numbers, e.g. 123/AB45678. The first 3 numbers correspond to a HMRC office number. The remaining letters are assigned in sequence by incrementing the first letter until it reaches a value Z, at which point the final number is incremented and the initial letter is reset back to A, and so on. For example, the first firm to register is assigned AA00001, the next firm BA00001, and so on, until the series reaches ZA00001, at which point the series continues to AA00002, BA00002, etc. Only 20 of the 26 letters in the alphabet are used. The last two digits therefore increment after each 20 firms register. When the final number reaches 99999, the second letter is incremented, the first letter goes back to A, and the final number is reset to 00001. Hence, the sequence would be ZA99999, AB00001, BB00001…

Staging dates were assigned to firms based upon the final two digits of this PAYE reference. If the number of registrations were low, there could potentially be economically meaningful correlations between PAYE digits and seasonalities. Given the very large number of newly registered employers in any given year (in 2016, approximately 414,000, or 1,636 per day), the last two digits change approximately 81 times per day. Because the system cycles through the complete set of

last two digits of the employer PAYE reference almost daily, we are confident that the assignment of the last two digits generates no economically relevant sequence patterns in the data. Therefore, we conclude that these digits are as good as randomly assigned, and hence regard the TPR's assignment of employers to staging dates based on the last two digits of employer PAYE references to be as-good-as-random.