

The Effects of Pension Reforms on Physician Labour Supply: Evidence from the English NHS

Carol Propper, George Stoye, Max Warner*

May 2024

Abstract

We examine the labour supply response of senior doctors in England following a reform of the public sector pension system that moved employees from a final salary to a career average pension plan, thereby increasing the return to work. Exploiting the staggered rollout of the reform across narrowly defined age groups, we find that doctors increased labour supply in response to this greater return to work by just under 4% four years after exposure, driven by increases on the extensive margin. This implies an extensive margin labour supply elasticity with respect to the return to work of 0.09.

Keywords: Doctor labour supply; Labour supply elasticity; Defined Benefit pensions; Public pension reform.

JEL Classification: H55, J22, J26, I10.

*We thank Jonathan Cribb, Carl Emmerson, Eric French, Nils Gutacker, Giuseppe Moscelli, Cormac O’Dea, Victoria Serra-Sastre, Jon Skinner, Jannis Stöckel, David Sturrock, Marcos Vera-Hernández, Ben Zaranko and seminar participants at IFS, Monash University, University of Melbourne, ESRI, CEPR Paris Symposium, Winter 2023 Health Economics Study Group, International Health Economics Association 2023 Congress, European Health Economics Associate Spring 2023 Seminar Series, and European Association of Labour Economists 2023 Annual Conference for helpful discussions and comments. We thank the Department of Health and Social Care for access to the Electronic Staff Record. Figures published using the ESR may be different from the official workforce statistics published by NHS Digital. The views expressed are those of the authors and are not necessarily those of the Department of Health and Social Care. We gratefully acknowledge financial support from the UK Economic and Social Research Council through the Centre for the Microeconomic Analysis of Public Policy (CPP) at IFS (ES/M010147/1) and from ERC Advanced Investigator Grant 788529 (HealthcareLabour). Any errors are our own. Author affiliation: Propper (Imperial College London, Monash University, Institute for Fiscal Studies and CEPR. c.propper@imperial.ac.uk); Stoye (Institute for Fiscal Studies, 7 Ridgmount Street, London, WC1E 7AE, UK. george.s@ifs.org.uk); Warner (Institute for Fiscal Studies. max.w@ifs.org.uk).

1 Introduction

Concerns over the financial sustainability of pension plans resulting from rising life expectancy and low interest rates have meant that generous defined benefit (DB) plans have either been reformed or scrapped in both the public and private sectors of many countries (Hinrichs, 2021; Whiteford and Whitehouse, 2006). Such reforms are also intended to strengthen labour supply incentives at younger ages, by moving away from benefits that are strongly linked to final years in the labour market to those linked to earnings across the life cycle. But the responsiveness of labor supply to changes in pension incentives for individuals who are not yet at retirement age is an open question. Individuals often appear either not to be fully aware of the details of pension systems (Mitchell, 1988; Arenas de Mesa et al., 2006; Crawford and Karjalainen, 2020), over-discount the future (Laibson, 1997) or do not process financial information well (Banks and Oldfield, 2007; Banks and Crawford, 2022). This potentially impacts the responses to pension incentives (Chan and Stevens, 2008; Bottazzi et al., 2006; Mastrobuoni, 2011; Liebman and Luttmer, 2015).

In this paper we study the impact of a reform to public sector pensions in the UK. This reform, in common with reforms enacted in many other OECD countries, retained the DB nature of the pension scheme but moved accrual from final salary to career average earnings.¹ These pension changes were announced in 2012 and implemented in 2015 for all public sector workers. We focus on the responses of a tightly defined group of highly skilled public sector workers: senior doctors employed in the English National Health Service (NHS).

Our setting has several advantages for the study of the labour supply impacts of pension reforms of this nature. First, because the pension reform we examine was part of similar reforms across the UK public sector and applied to all 1.4 million NHS employees, it was salient and well publicised. As an indication of this, following the announcement of the reforms to come, doctors in the NHS went on strike to protest against the proposed changes and the doctor unions predicted a large reduction in NHS participation by doctors.² Second, the way the reform was

¹Examples of countries introducing such reforms to their public pension plans in the 2000s include Austria, Finland, France and Portugal. More broadly, 20 of the 25 OECD countries with earnings-related public pension plans in 2005 used a lifetime (or near-lifetime) average of earnings to calculate pension benefits (OECD, 2005).

²See, for example, <https://www.bbc.co.uk/news/uk-16618194>

implemented meant that it affected individuals sufficiently far away from retirement to be able to alter their labour supply. Third, the senior doctors we examine have flexibility in their labour supply to the NHS, may have generous outside options from working in the private healthcare market and are a highly educated group whose financial literacy is likely to be considerably higher than that of other employees (Banks and Oldfield, 2007). Thus they are potentially able to alter their labour supply in response to changes in the incentives embodied in the pension reform. Finally, senior doctors play a central role in the delivery of healthcare, in the training of younger doctors and face increasing demand for their time due to rising population needs for healthcare. Understanding their response is therefore central to the delivery of healthcare in the UK and, more generally, to understanding how skilled labour may respond to pension changes.

Our empirical approach exploits the fact that the reform was announced and introduced for all public sector employees at a single point in time, but its implementation affected employees, including senior doctors, in a staggered fashion. The new pension plan was announced in 2012 and introduced in 2015. On introduction, most senior doctors were immediately moved to the new pension plan while those close to retirement under the old plan were never moved. An intermediate group were moved onto the new plan over time in a staggered manner every two months based on their month of birth. Using detailed administrative monthly payroll data, we exploit this staggered roll-out, restricting our sample to senior doctors born within a seven year time span, all of whom were treated by the reform. To avoid the well known issues when estimating models with two-way fixed effects and staggered treatment timings (De Chaisemartin and d'Haultfoeuille, 2018; Goodman-Bacon, 2021; Sun and Abraham, 2021; Borusyak et al., 2021), we use the imputation method proposed by Borusyak et al. (2021). We control for time, age and cohort effects, and subject our results to a wide range of robustness tests for cohort trends.

The move to a career average plan was also accompanied by a change in the retirement age of the plan, with the normal retirement age rising from 60 to 67. The reform therefore had two main financial impacts on senior doctors, both of which have ambiguous effects on labour supply. First, the move from a final salary to a career average plan increased the returns to work in a given year as the pension value in the new plan depends on current rather than

final income. Second, the reform also increased the normal retirement age of the plan. This both reduced the returns to work in a given year as the pension could now be claimed for fewer years and reduced the total value of the pension, so decreasing pension wealth. The announcement of the new scheme three years prior to its implementation and our focus only on those affected by the staggered roll out allows us to isolate the labour supply responses to the change in the return to work from any response to the reduction in pension wealth brought about by the change in pension age. The announcement of the reforms three years prior to implementation, the opposition of doctor unions to the changes and the doctor strikes against the reforms meant that doctors would have been well aware of the overall reductions in their pension wealth, regardless of the date that they were actually moved on to the new scheme. Our design means we control for such anticipatory changes in labour supply arising from the reduction in pension wealth due to the later retirement age of the new plan. This means our findings are also applicable to reforms which change pension accrual rules without altering the pension age.

We find that the change in the return to work associated with being moved onto the new pension plan increased average full-time equivalent labour supply by 3 percentage points after four years, equivalent to a 3.5% increase on baseline labour supply. This was driven by an increase on the extensive margin: moving to the new plan increased the probability of working in the NHS four years later by 3.4 percentage points (3.7%). These effects persist and grow over longer horizons, suggesting that the change in the return to work caused by the reform extended the number of years doctors worked in the NHS. By contrast, there was no impact on the intensive margin, with no significant effect on hours conditional on participation.

We use our labour supply estimates combined with the estimated financial impacts of the reform to derive an estimated extensive margin labour supply elasticity with respect to work of 0.09. This is at the lower end of other estimates of doctor labour supply (e.g. Lee et al. (2019)). However, as we estimate a return to delayed remuneration, a lower elasticity than for current wages would be expected. The elasticity is also smaller than estimates of similar pensions responses in wider populations (French et al., 2022), suggesting that changes in financial incentives from pension reforms have less impact on doctors than on the general population.

Our contributions are the following. First, we contribute to the literature on the impact of pensions on labour supply. Most of the literature focuses on retirement incentives for those at or near retirement (Blundell et al., 2016). These studies suggest that changes in retirement incentives have a large impact on retirement decisions and large implicit taxes often hinder older age employment (Gruber and Wise, 1999). But we might expect different responses for those who are further from retirement, as a move from a final salary to a career average or defined contribution plan generates stronger work incentives at younger ages. Two recent papers have examined the average responses of public sector employees not yet close to retirement to similar reforms to those we study in two European countries (French et al., 2022; Bovini, 2019). Both studies derive an average for employees across all occupations and sectors of the economy. We build on these, and other previous studies, by focusing on the responses of a set of highly skilled workers, who are also key to the health and production of the economy.

Second, we contribute to the literature on doctor labour supply. Perhaps surprisingly, there is a rather limited literature on the labour supply of doctors and much comes from the US context, which is characterised by very different employment contracts and incentives for doctors than many other healthcare systems (Nicholson and Propper, 2011; Lee et al., 2019). Much of this literature also has difficulties establishing causality. The UK evidence is particularly scarce (Ikenwilo and Scott, 2007; Lee et al., 2019) despite the importance of the NHS as the dominant provider of healthcare and the large share of public expenditure accounted for by healthcare spending (21% in 2019-20, Institute for Fiscal Studies (2021)). Our contribution is to provide causal evidence on the labour supply of doctors in England using novel detailed administrative data by exploiting a reform that generates exogenous variation in pension compensation. We thus contribute to the understanding of the labour supply of doctors in countries where doctors are hospital or healthcare facility employees, as distinct from being self-employed (e.g. (Brekke et al., 2017; Andreassen et al., 2013; Broadway et al., 2017)). Most of this literature focuses on current remuneration for doctors: our contribution is to show the importance of pensions, and more generally delayed remuneration, for labour supply.

Finally, as the reform we examine was for almost all public sector employees, it also affected other high skilled individuals in the UK public sector (judges, academics, senior public

sector managers) for whom there is no evidence on their response to the changes. Our results may be useful for understanding the responses of these employees.

The paper is organised as follows. Section 2 describes the institutional background, the pension reforms that we study, and how we might expect them to affect labour supply. Section 3 describes the data. Section 4 explains our empirical strategy. Section 5 presents the results of the labour supply analysis and Section 6 uses these results to estimate labour supply elasticities. Section 7 concludes.

2 Our setting

2.1 Healthcare in England

Almost all hospital healthcare in England is provided by the National Health Service (NHS). Care is free at the point of use and funded out of general taxation. Elective (pre-planned) and emergency care is provided by large, public hospitals with a small amount of provision of NHS funded elective care by private hospitals (Stoye, 2019). Patients do not pay for care when treated under the NHS and hospitals are reimbursed by the government through a set of nationally agreed tariffs for a given treatment, using a system close to Diagnosis-related Groups (DRGs) in the US. Elective care is rationed through waiting times and by the requirement to have a referral from a primary care doctor. Emergency care is accessed through hospital emergency departments, which provide urgent care and admit patients for further treatment if required. There is also a small private-pay elective healthcare sector, specialising in elective care for which there is long NHS waiting lists. For this, patients pay out-of-pocket or through medical insurance.

Staff in public hospitals are salaried public sector employees, with short-term shortages filled by external agency staff. Staff are contracted to an NHS ‘Trust’, comprised of one or a small number of geographically close hospital sites.³ We focus here on senior doctors in NHS hospitals, known as consultants. These senior doctors made up 42% of the total doctor

³For simplicity, we refer to these trusts as ‘hospitals’ throughout the text. In rare cases, individuals may be employed by two Trusts.

workforce and 8% of the total qualified clinical workforce in the NHS in 2021 (NHS Digital, 2022). They lead medical teams to provide care in their given speciality and every patient has a named consultant. They may also have other management duties within the hospital.

All doctors earn a basic salary depending on their role and experience, with this salary set by national pay scales. Senior doctors can receive additional pay, for example by working extra shifts or taking management responsibilities. Some of this additional pay is pensionable, including pay for permanent extra responsibilities, and some of it is not, including payment for hours worked over full-time contracted hours. Senior doctors can also apply for clinical excellence awards, based on achievements in their work, which can substantially increase their pay. Over most of the period we study, all such awards were pensionable.⁴

Senior doctors can also supplement their NHS salaries by doing additional work in private hospitals. To do this, senior doctors who are on full-time contracts are expected to provide 10 percent additional hours to the NHS, though recording of this increase in contracted hours varies between NHS hospitals. Senior doctors may alternatively choose to leave the NHS and work entirely within the private sector. Data on private earnings for doctors is difficult to obtain but, historically, private earnings have been a significant percentage of total income for some senior doctors because pay-rates in the private sector are often considerably higher than in the NHS. The availability of private work varies considerably across both region and clinical specialities, being particularly high in specialities for which there are long NHS waits and in regions where patient income is higher (Morris et al., 2008).

2.2 The 2015 Pension Reform

In March 2011, the government's Independent Public Service Pensions Commission recommended that UK public sector pensions move away from their traditional final salary form (which typically averaged earnings of the last three years of employment) towards a career average revalued earning (CARE) plan. This was still a defined benefit plan, but with different features to the previous plan. These recommendations led to changes across all public sector

⁴Clinical excellence awards can be either local (allocated and funded by the employer trust) or national (allocated and funded nationally). National awards were pensionable during all of our sample period. From April 2018, local awards were no longer pensionable.

pension plans, including those for civil servants, teachers, police, judges, firefighters and NHS employees.

Prior to 2015, the NHS had two pension plans: the 1995 and 2008 plans. We focus on the former as almost all senior doctors we examine were in that plan.⁵ This was a final salary defined benefit pension plan, where each year of service accrued 1/80th of the member's best salary in their final three years before retirement. Under the new 2015 plan the pension depended on average earnings and a member would accrue 1/54th of their earnings in each year they worked in the NHS. Before retirement, each year's accrual would also be uprated at Consumer Price Index (CPI) inflation plus 1.5%. Thus the value of the 2015 plan depended much more on career rather than final (three year) earnings.

The 2015 plan had a number of other differences from the 1995 plan. The normal retirement age was 60 in the 1995 plan and this was raised to the state pension age in the 2015 plan, which was 67 for the senior doctors in our analysis. So the retirement age was substantially higher under the new plan, though in practice many senior doctors continued to work beyond their retirement age under the old plan (Appendix Figure A2). Both plans allowed early retirement in return for reduced benefits during retirement. Under the 1995 plan, those who joined before 2006 (most of those that we study) had an early retirement age of 50. Under the 2015 pension plan, the earliest early retirement age was 55. Thus early retirement at the same age in the 2015 plan meant a larger reduction in benefits than in the 1995 plan. Finally, the 1995 plan included a lump sum worth three times the annual pension value upon retirement. The 2015 plan did not include a lump sum, but members could choose to take one in return for reduced subsequent pension payments.

Despite these differences in generosity, both plans had the same employee and employer (the NHS Trusts) contribution rates.⁶ Changes to the pension plan therefore did not directly affect either the take-home pay of senior doctors or the costs for hospitals of employing senior

⁵The 2008 plan was introduced in 2008 for new joiners, so any senior doctor that worked continuously in the NHS since before 2008 would still be in the 1995 plan. All of our cohort worked in the NHS in 2012. Our data does not go back further, but the mean age of our cohort in 2008 was 42 and so it is likely that almost all our senior doctors were working in the NHS prior to 2008. The 2008 plan was also a defined benefit plan, but with slightly different accrual rules and a later retirement age (65).

⁶Contribution rates for employees depended on income and ranged from 5% to 14.5% for those earning more than £111,377 and contribution rates for employers were 14.3% until March 2019 and 20.6% from April 2019.

doctors.

The new pension plan was announced in March 2012 and implemented from 2015 in a staggered fashion for all NHS employees. In April 2015 all of those more than 13 years and 5 months away from retirement age (under their previous plan) in April 2012 were immediately moved to the new plan. Rather than transferring the value of their pension in the old plan to the new pension plan, employees maintained their old pension and started a new pension in the new plan. The final salary used to calculate benefits from the old plan remained linked to the employee's current salary, although no further years of service could be accumulated on the old plan. The benefits from the old plan could be taken from the retirement age of that plan if the employee no longer worked for the NHS. All new joiners to the NHS started in the 2015 plan.

Those within 13 years and 5 months of their retirement age under the old plan (age 60) in April 2012 received either 'tapered' or 'full' protection. Those within 10 years of their retirement age under the old plan (aged 50 and above in April 2012) received 'full' protection, and were not moved onto the 2015 plan, but remained in the old plan. Those between 10 years and 13 years and 5 months (aged between 46 years and 7 months and 49 years and 11 months in April 2012) received 'tapered' protection, which meant they would spend additional time in the old plan before being moved to the 2015 plan. In June 2015, those 13 years and 5 months away from retirement were moved to the 2015 plan, and in August 2015, those 13 years and 4 months away were moved. This pattern continued over time, with each subsequent month age group shifted every two months, until all were shifted by February 2022.⁷

The pension reform was followed by two national changes to the taxation of pensions in 2016. The first involved changes to the annual allowance (the amount a pension is allowed to increase tax-free in a year) and the second to the lifetime allowance (the amount a pension is allowed to increase tax-free over a lifetime). While there were reports that the changes to the annual allowance negatively affected doctors because of volatility in their NHS income when they undertake extra duties, and the lifetime allowance could have affected those who

⁷The staggered rollout was subsequently ruled to have discriminated against younger members. As a result, the government has announced that members at retirement can choose whether they want their benefits from earnings between 2015 and 2022 to come from the old or the new pension plans. This was not announced until 2021 and is still subject to some uncertainty. For our sample period, the senior doctors we study would have believed their pension benefits would come from the new 2015 plan upon being moved.

had accrued larger pension pots, neither change is a significant concern for our analysis given our empirical design. These changes, particularly the change in the annual allowance, would have affected relatively few doctors for one to two years over our sample period as they would have affected doctors on both the old and new pension plans, and our analysis compares the response to the *change* in plan.

2.3 How the reform may affect doctor labour supply

The reform had two main financial impacts on senior doctors. The first was a change to total pension wealth, for any given level of labour supply. The higher retirement age of the new plan lowers pension wealth because the pension is now paid out over fewer years in retirement. This lump sum change in pension wealth may increase or decrease labour supply. The second was a change to the return to work. The change from final salary to career average accumulation increases the return to work, for those sufficiently far from retirement, because the pension value now depends on all NHS earnings rather than final 3 years of earnings. In contrast, the change in the retirement age decreases the return to work because any increase in pension value is paid out over fewer years.

As we discuss in more detail in Section 4, our data and empirical design allow us to cleanly identify the impact of the reform's change to the return to work while controlling for the impact of the change in pension wealth. We therefore focus here on how the reform changed the return to work. The return to work for a pension plan is the impact of labour supply (on either the extensive or intensive margin) on the discounted value of the pension. A higher return to work means that labour supply increases this value by more than a lower return. The change in the return to work caused by any reform is the percentage change in the pension return from earning income y_{age} at age age under the previous and current schemes. This will depend on the differences in accrual rates, which income is used to calculate accruals, expected income trajectories, the discount rate and the expected age of death.

In Appendix B we show that for the reform we study this change in the return to work can be expressed, for an individual at age age who is more than three years from their retirement,

as:

$$\frac{80}{54} \times \frac{y_{age}}{\max\{\hat{y}_{r-1}, \hat{y}_{r-2}, \hat{y}_{r-3}\}} \times \frac{(1 + \pi + 0.015)^{66-age}}{(1 + \pi)^7} \times \frac{\sum_{a=67}^{100} P_a \frac{1}{(1+\delta)^{a-60}}}{3P_{60} + \sum_{a=60}^{100} P_a \frac{1}{(1+\delta)^{a-60}}} - 1 \quad (1)$$

where \hat{y}_{r-t} is the expected income t years before retirement, π is the expected inflation rate, δ is the real discount rate and P_a is the probability of living to age a .

Equation (1) depends on four terms. The first term represents the differences in accrual rates between the two pension plans. The new pension plan has a higher accrual rate (1/54) than the old pension plan (1/80). This increases the return to work because a greater share of income is accrued into the pension. The second term represents the change by which income is used to calculate accruals. The new pension plan uses current income (y_{age}) while the old pension plan used the best year of the final three years of income. The change in the return to work will depend on the expected income trajectory over the accrual period relative to the expected final salary.

The third term represents the change in how contributions are uprated. The reform altered this such that current year earnings were uprated by inflation plus 1.5% in each year prior to retirement. The motivation was to match the expected increase in wages over time, but in practice the uprating under the new plan was much greater than nominal wage growth of senior doctors over the period we study. The final term represents the later retirement age of the new pension plan and the fact that it does not include a lump sum upon retirement. This reduces the return to working because pension benefits are paid out over fewer years, do not include a lump sum, are paid later and are received when the individual has a higher probability of death.

Despite one goal of the reform being to increase labour supply incentives, ex-ante the sign of Equation (1) is unclear. The first term is greater than one, while the final term is less than one. The second and third terms may be greater or smaller than one as they depend on the income profiles and age of senior doctors and the uprating of contributions in the previous and current plans. In practice, the second term is likely to be less than one if earnings rise with age, and the third term is always greater than one in our analysis given the age of senior doctors affected by the reform. The reform therefore has an ex-ante ambiguous impact on the return to

work, and thus an ambiguous effect on labour supply.

Equation (1) also shows that the return to work changes on both the extensive and intensive margin. On the extensive margin, working at a constant y_{age} has a different return than under the old scheme because of changes to the accumulation rate, uprating and the retirement age. Increasing y_{age} also affects the increase in pension value because the new pension plan is a career average plan. This increases the return to work on the intensive margin.

For those within three years of their actual retirement (which may be earlier or later than their pension plan's normal retirement age), the change in the return to work is more complex. The previously accumulated benefits in the old pension plan still depend on the actual final salary of the senior doctor, and so higher earnings in the final three years of work give an additional return to work because they increase the value of the old pension plan. The decision to retire or leave the pension scheme, which determines the salary used as the final salary, is also affected.

3 Data

Our primary data source is the Electronic Staff Record (ESR), the monthly payroll data for all staff directly employed by the NHS. It includes occupational codes, demographic characteristics, a breakdown of monthly pay and the number of hours or shifts worked for each employee. Our sample period is April 2012 to August 2021. In our analysis years are fiscal years (April to March). Using these data we create a cohort of senior doctors. This contains any doctor who was employed as a senior doctor in an acute (short term general) hospital in April 2012 who was born between April 1962 and December 1969. All these doctors were affected by the staggered rollout. The vast majority of those in this plan were likely on the 1995 pension plan prior to the reform.⁸ In robustness tests, we also use a cohort of senior doctors who were never moved onto the new pension plan. We define this group as senior doctors born between January 1959 and March 1962.

⁸This matters because the rollout was based on normal retirement age which is different in the 1995 and 2008 plans, and so different date of births were shifted at the same time depending on which plan they were in. Footnote 5 explains why almost all senior doctors in our cohort were likely in the 1995 plan. Any senior doctors in our cohort that were in the 2008 plan would not have been affected by the staggered rollout as they would have been too far from retirement. Instead, they would have all been immediately moved onto the new plan.

We define three labour supply outcomes. First, we measure total labour supply as the full-time equivalent (FTE) contracted hours of each senior doctor, including zeroes for those who have left the NHS. Second, we measure the extensive margin of labour supply using a dummy variable for whether the senior doctor is still employed by the NHS at time t . Third, we measure the intensive margin of labour supply as FTE contracted hours conditional on being employed by the NHS.

Our cohort contains 11,872 senior doctors. Table 1 presents summary statistics for 2014, focusing on 2014 as this is the year immediately before the new pension plan was introduced. 32.6% of the senior doctors in our cohort were female and the average age in 2014 was 48. 93.9% were still working in the NHS (i.e. 6.1% had left the NHS between 2012 and 2014). 98.8% of those who were still working in the NHS were in the NHS pension plan. This means that almost all senior doctors in our cohort would have been affected by the transition to the new pension plan. Appendix Figure A1 shows this very high participation rate is similar to that of younger senior doctors prior to the reform.

Table 1: Consultant cohort summary statistics in 2014

	Mean	SD
Age	48.0	2.24
Female	32.6%	46.9%
NHS employment	93.9%	23.0%
NHS pay conditional on employment	£121,000	£33,600
NHS FTE conditional on employment	0.972	0.122
NHS pension plan membership conditional on employment	0.988	0.0982
N	11,872	

4 Empirical strategy

To identify the impacts of the new pension plan on labour supply we exploit its staggered rollout. In our cohort, those born between October 1966 and December 1969 were immediately moved to the new pension plan in April 2015. Those born between April 1962 and September

1966 were treated in the staggered fashion described in Section 2.2, with the youngest treated first from June 2015.⁹ The oldest in this group had yet to be moved onto the new plan by the end of our sample period, but would eventually be moved. Restricting our analysis cohort to this narrow group of birth years has the following advantages. First, it means we only compare the labour supply of senior doctors born at most seven years apart. This means many other characteristics of these doctors are likely to be similar, as are their labour supply trends in the absence of the reform. Second, our analysis only includes senior doctors who would be affected eventually by the reform. By comparing only those who would eventually be treated, we eliminate any time-varying anticipation effects that were independent of treatment timing. For example, all of those in our cohort knew their retirement age would increase from 60 to 67, irrespective of when they were actually moved onto the new pension plan. Third, defining the cohort in 2012 allows us to consider labour supply outcomes for three years prior to the reform.

To estimate the impact of the reform we model the labour supply of individual i in birth month group j (e.g. those born in October 1963) in month t as

$$y_{ijt} = \sum_{h=0}^{74} \beta_h \mathbf{1}(t = E_j + h) + \alpha_j + \delta_t + \gamma \text{age}_{it} \times \text{gender}_{it} + u_{ijt} \quad (2)$$

where E_j is the month that each group is first treated, and so each β_h measures the effect of the new pension plan h months after being moved onto it. We include birth month fixed effects α_j and month fixed effects δ_t , as well as age (dummy variables for each year) by gender fixed effects.

We do not use OLS to estimate Equation (2) because of the well known issues when estimating models with two-way fixed effects and staggered treatment timings (De Chaisemartin and d'Haultfoeuille, 2018; Goodman-Bacon, 2021; Sun and Abraham, 2021; Borusyak et al., 2021). Instead, we use the imputation estimator proposed by Borusyak et al. (2021). This uses the not-yet-treated observations to estimate the values of each fixed effect in Equation (2). These fixed effects are then used to impute a counterfactual for each treated observation and the difference between the actual outcomes and these counterfactuals are aggregated to

⁹The rollout actually grouped those born between the 2nd of one month and the 1st of the next month. For simplicity we refer to the month that the vast majority of those in the group were born in.

produce estimates of the treatment effects. The Borusyak et al. (2021) imputation estimator is more computationally feasible in our context than other proposed methods which aggregate many pairwise comparisons, such as that proposed by Callaway and Sant’Anna (2021), because of the large number of time periods our treatment is staggered over.¹⁰

The key identifying assumption is that conditional on our controls, all other factors determining each senior doctor’s labour supply are uncorrelated with their treatment timing. We control for age and gender, since they are correlated with both labour supply and treatment timing via date of births. We assume that all other trends in labour supply are the same for all those in our cohort, since they were all born within a similar period. The key threat is that there are time-varying anticipation effects that vary with treatment timing. We might expect considerable responses in advance of being moved onto the reform. But conditional on group, time and age-gender fixed effects in Equation (2), these are unlikely to be correlated with treatment timing. We therefore make the assumption that there are no such effects. To examine this we test for pre-trends using the methodology proposed by Borusyak et al. (2021) and in robustness tests also re-estimate our model with the inclusion of cohort time fixed effects.

Our estimates do not identify the effect on labour supply of changes to pension wealth. As the reforms were announced three years prior to implementation, all doctors in our sample would have been aware of the changes in scheme rules and the large reduction in pension wealth. Standard theory suggests that doctors may change their labour supply in advance of actually being moved onto the new pension scheme in response to this lifetime reduction in wealth. But any such anticipatory changes in labour supply that are either constant over time (but may vary by group) or depend on age and/or gender will be absorbed by our fixed effects. Instead, our estimates capture the impact of the actual change in financial incentives that occur when a doctor is moved onto the new pension scheme. This is the change in the return to work discussed in Section 2.3 above.¹¹

¹⁰We repeated our primary analysis with OLS. The results are similar but less precise. Intuitively, this is possible if the forbidden comparisons made by OLS are noisier than the correct comparisons used by the imputation estimator.

¹¹Technically it is not being moved onto the new pension plan, since not all senior doctors are members of the plan, and therefore closer to an intention to treat (ITT) treatment effect. But since 99% were members pre-reform, for simplicity we interpret the estimated treatment effects as the effect of being moved.

5 Results

In this section, we present our baseline labour supply results which we subject to a number of robustness checks. We also examine heterogeneity between different groups of senior doctors based on their age and gender. In Section 6, we then use our results to estimate a labour supply elasticity with respect to the expected increase in the present value of pension wealth from work.

5.1 Baseline results


We first consider total labour supply. Figure 1 shows the estimated effect of the change in returns to work on contracted FTE hours, including zeros. There is a clear upward trend in labour supply after being moved onto the new plan, which is statistically significant after three years. The confidence intervals grow over time because Borusyak et al. (2021)'s imputation method uses the not-yet-treated group to impute the time fixed effects, and this group gets smaller over time as more doctors are treated. Four years after being moved onto the new plan, senior doctors work 3% more of an FTE, and after six years they work 8% more. The average labour supply of those in our cohort is 0.91 of an FTE in 2014, prior to the introduction of the 2015 plan. So this is equivalent to a 3.5% increase after four years, and a 8.6% increase after six years.

Importantly, Figure 1 also shows that there is no statistically significant pre-trend in labour supply, with the point estimates all close to zero. The standard errors grow for pre-trend estimates closer to the treatment date, the opposite of the pattern often found in standard OLS event studies. This is because the Borusyak et al. (2021) imputation estimator estimates pre-trends relative to those even further from treatment (those more than 36 months away in our case) rather than the period before treatment as in standard OLS event studies.¹² The pre-trend estimates suggest that there are no time-varying anticipation effects between those moved earlier and later onto the plan that would bias our estimated treatment effects, or any differences

¹²This pattern is often found in other papers implementing the Borusyak et al. (2021) imputation estimator. When we estimate our results with OLS, we find the same pattern if our base period is the same as in Borusyak et al. (2021) (e.g. those more than 36 months away from treatment) and the reverse pattern if our base period is the period before treatment (e.g. the standard even study approach). Results are available on request.

in labour supply trends. This provides evidence against the concern that doctors change their labour supply immediately prior to being moved on to the new plan or that there are different labour supply trends for different cohorts of doctors in our sample.


Figure 1: Total labour supply



Note: The lines for each coefficient represent 95% confidence intervals.

We next decompose these changes in total labour supply into changes on the extensive and intensive margin. Figure 2 shows the extensive margin, which we measure using a dummy variable for working in the NHS as the outcome. There is again a sustained increase in labour supply after senior doctors are moved onto the new pension plan. After four years, senior doctors are 3.4 percentage points more likely to be working in the NHS, and after six years they are 7.9 percentage points more likely. As almost all (94%) of those in our cohort were working for the NHS in 2014 prior to the introduction of the plan, this is equivalent to a 3.7% increase after four years and a 8.4% increase after six years. Since almost all of those in our cohort were working for the NHS prior to the reform, we interpret this increase in NHS employment as a reduction in leaving the NHS relative to the control group of doctors, leading to a *relative* increase in participation, rather than an *absolute* increase.


Figure 2: Extensive margin



Note: The lines for each coefficient represent 95% confidence intervals.

Figure 3 shows the intensive margin, measured as FTE contracted hours conditional on working in the NHS. There is no significant change in contracted hours at any point after being moved on to the new pension plan, suggesting that the main response was on the extensive margin.

Figure 3: Intensive margin (FTE contracted hours)



Note: The lines for each coefficient represent 95% confidence intervals.

For both total labour supply and labour supply on the extensive margin, our estimated treatment effects grow over time. This is perhaps surprising given the change in incentives is instantaneous after being moved. One likely explanation for the gradual increase in treatment effects is that because most of our cohort are working in the NHS when treated, there is limited scope to *increase* their labour supply when first treated, particularly on the extensive margin. However, over time and with age, the control group reduces their labour supply. Appendix Figure A2 shows the older doctors unaffected by the reform gradually reduced their labour supply from age 50. If similar patterns occur in our slightly younger analysis cohort, this reduction in labour supply by the control group over time could mean that our treatment effects grow over time because there is more scope for the treated senior doctors to respond to the change in financial incentives.

Table 2 summarises our labour supply results by aggregating our event study estimates into average treatment effect estimates by year from treatment. On average, the change introduced by the new pension plan increased the labour supply of the senior doctors in our cohort, driven by an increase in labour supply on the extensive margin. While there are significant pre-trends on the intensive margin three years prior to the reform, these are relatively small in magnitude (Figure 3) and we find no significant impacts on this margin. In comparison, there are no significant pre-trends for total labour supply or the extensive margin, where we do find major impacts of the reform.

Table 2: Summary of primary results

	(1) Total Labour Supply	(2) Extensive Margin	(3) Intensive Margin
3 years before	0.000600 (0.00275)	-0.00227 (0.00270)	0.00303** (0.00129)
2 years before	0.00116 (0.00428)	-0.00127 (0.00420)	0.00255 (0.00199)
1 year before	0.00317 (0.00598)	0.000790 (0.00591)	0.00263 (0.00265)
0 years after	0.00231 (0.00225)	0.00207 (0.00215)	0.000370 (0.00119)
1 year after	0.00812* (0.00449)	0.00675 (0.00437)	0.00179 (0.00226)
2 years after	0.0130* (0.00751)	0.0113 (0.00737)	0.00228 (0.00356)
3 years after	0.0233** (0.0115)	0.0237** (0.0114)	0.000588 (0.00536)
4 years after	0.0437*** (0.0169)	0.0460*** (0.0167)	0.0000361 (0.00806)
5 years after	0.0647*** (0.0234)	0.0656*** (0.0231)	0.00371 (0.0109)
Outcome Mean	0.880	0.910	0.970
N	1,317,792	1,317,792	1,198,717

Standard errors clustered at the doctor level in parentheses


* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Appendix Figure A3 estimates the total labour supply impact of the change in returns to work on our cohort. This total impact increases over time both because the estimated individual treatment effects grow over time (Figure 1) and because the staggered roll-out means that more senior doctors are moved onto the new plan over time. By the start of 2021, our estimates suggest that the change in returns to work brought about by the reform had increased the number of FTE senior doctors working in the NHS by 666, out of a cohort of 11,900. Given the high pre-reform labour supply of senior doctors in our cohort, this increase was mainly driven by fewer senior doctors leaving the NHS than would have happened had the reform not occurred.

5.2 Robustness


We begin by testing the robustness of our primary results to varying the definition of the treated cohort. For simplicity, we focus on results for total labour supply. Our results above use all senior doctors born between April 1962 and December 1969. Figure 4 shows the results if we shorten the definition of our cohort and only include those born between April 1962 and December 1967. Figure 5 shows the results if we extend the definition of our cohort and include those born between April 1962 and December 1971. In both cases, the results are very similar to our primary results, showing that they are not dependent on the precise boundaries of our cohort. This also suggests there are not large differences in the responses of senior doctors by age, and we examine this further below.

Figure 4: Total labour supply with a smaller cohort



Note: The lines for each coefficient represent 95% confidence intervals.


Figure 5: Total labour supply with a larger cohort



Note: The lines for each coefficient represent 95% confidence intervals.

Another potential concern is that we exclude senior doctors that were never treated from our analysis. As explained in Section 4, we exclude these doctors from our primary analysis to ensure we are comparing a similar group of doctors and to eliminate some common pre-treatment anticipation effects. To test robustness to this choice, Figure 6 repeats our primary analysis including never-treated senior doctors born between January 1959 and March 1962. Inclusion of the never-treated senior doctors changes the interpretation of our results, in particular the extent to which our fixed effects absorb the labour supply effect of the change in pension wealth. Constant group-specific changes in labour supply in response to the change in pension wealth are still absorbed, but the age by gender fixed effects are now estimated in part using the never-treated doctors. These fixed effects will therefore not fully absorb age- and gender-specific changes in labour supply in response to the change in pension wealth. Our estimated treatment effects in this case may therefore include some of the effect of the change in pension wealth.

Figure 6: Total labour supply with never treated senior doctors




Note: The lines for each coefficient represent 95% confidence intervals.

The estimated treatment effects are similar to our primary results, although estimated somewhat less precisely: for example, four years after being moved onto the new pension plan, Figure 6 suggests that senior doctors are working 1.8% more after 4 years, compared with 3.5% more in our primary analysis. This suggests that there were age-gender specific reductions in labour supply in response to the changes in pension wealth from the new plan that are absorbed by the fixed effects in our primary analysis. However, there is a small upwards pre-treatment trend in labour supply in this model which, while not statistically significant, is consistent with our methodological concerns about including these never-treated doctors.

Second, to allow for potential trends in labour supply behaviour that are unrelated to the change in returns to work introduced by the reform but that vary across cohorts, we include date of birth year by age fixed effects. As with our previous robustness test, we include never-treated senior doctors born between January 1959 and March 1962 in our sample. Figure 7 shows that the inclusion of these trends eliminates these weak pre-treatment trends while still finding a qualitatively similar pattern of results. Importantly, though, these additional controls also absorb some of the genuine treatment effect of the change in the return to work, since treatment is allocated based on date of birth. Moreover, due to the requirements of Borusyak et al. (2021)'s imputation estimator we can only estimate treatment effects for the first three

years after being treated. This means the estimates are slightly smaller in magnitude while the standard errors are clearly much larger. This is consistent with the additional fixed effects absorbing some of the genuine response to the change in the return to work.


Figure 7: Total labour supply with never treated senior doctors and additional fixed effects



Note: The lines for each coefficient represent 95% confidence intervals.

Third, another potential concern is that our results may be influenced by the hospital or region in which senior doctors work which may be subject to time-varying shocks correlated with exposure to the reform. Figure 8 shows the impact of the changes in returns to work on total labour supply conditional on including hospital by year fixed effects. These capture both differences in the types of doctors working in each hospital and differences by local geography and region over time. The results are very similar to our primary results.


Figure 8: Total labour supply with hospital FEs



Note: The lines for each coefficient represent 95% confidence intervals.

Two common macro-shocks could be potential concerns. The first is the impact of the Covid-19 pandemic as our sample period extends to August 2021. Figure 9 therefore repeats our primary analysis excluding March 2020 onwards from our sample. The results are qualitatively similar to our primary results, although slightly less precisely estimated. This suggests that our results are not being driven by changes in senior doctor labour supply during the Covid-19 pandemic. The second is Brexit (the decision to leave the European Union), which occurred in 2016. For example, European Union doctors may have chosen to leave the NHS, while British doctors may have worked more to compensate. We consider this unlikely to affect our results. National trends over time in labour supply are captured by our time fixed effects and it is not clear why Brexit would affect doctors differently based on their month of birth.

Figure 9: Total labour supply excluding Covid-19 period



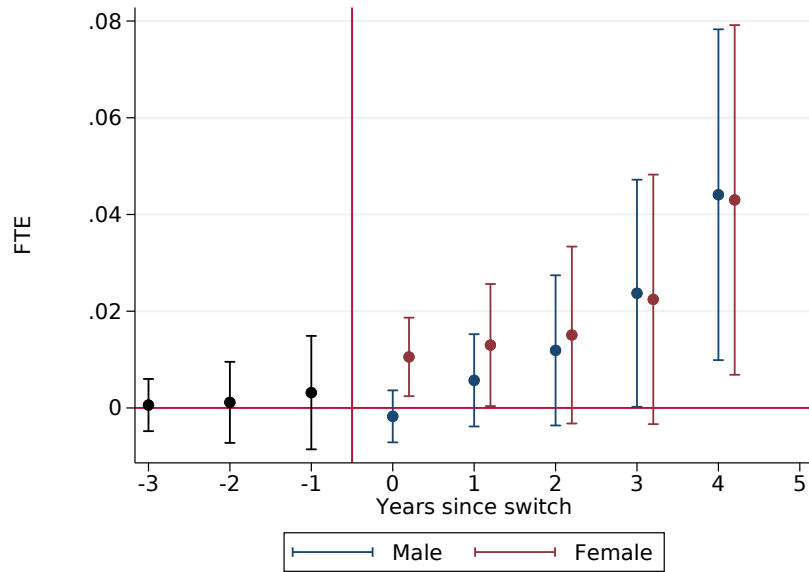
Note: The lines for each coefficient represent 95% confidence intervals.

5.3 Heterogeneity

We consider differences by gender and differences by age of doctors when they were treated. Figure 10 presents the estimates for male and female senior doctors. These are estimated jointly for males and females so the pre-reform estimates are the same as in our main analysis. To present as clear a picture as possible, we average the estimated β_h into the average treatment effect for each year from treatment rather than each month. There is little evidence of a meaningful difference in response by gender, except perhaps less than a year after being moved. The point estimates for the two groups are very close and are always well within the standard errors for each group.

The labour supply responses of men and women to financial changes in other contexts are often very different (Meghir and Phillips, 2010). But in our setting, where most individuals work full time and have relatively similar outside options, we find no differences in responses to the pension reform. This is consistent with evidence that female doctors in England reduce their labour supply after maternity leave by much less than female workers in the rest of the economy (Kelly and Stockton, 2022).

Figure 10: Total labour supply by gender




Note: The lines for each coefficient represent 95% confidence intervals.

Figure 11 presents estimates separately for whether doctors were moved onto the new pension scheme in the 40s or in their 50s. The point estimates for doctors in their 50s are consistently above the point estimates for doctors in their 40s, but the estimates are always within the standard errors for each group. This suggests there were no significant differences in response by age group. While the benefits of a career average scheme relative to a final three year scheme may be larger for those who are younger rather than those who are older, the younger group have very high full time participation rates (perhaps for career concerns). Thus there is less scope to increase their hours than for older senior doctors who may have more flexibility.

These results also suggests that the response to the new plan does not differ much by pension wealth, since those moved onto the new pension scheme in their 50s had previously accumulated greater pension wealth than those moved in their 40s. Likewise, these results suggest that there was little difference by treatment timing, since those in their 40s were moved onto the new scheme earlier in time than those in their 50s.

Figure 11: Total labour supply by age when treated



Note: The lines for each coefficient represent 95% confidence intervals.

Taken together, our heterogeneity results suggest that there was little difference in response to the change in returns to work caused by the reform by either gender or age, and therefore also suggests there was little difference by treatment timing or prior pension wealth.

6 Labour supply elasticity estimates

We now use our estimates of the labour supply changes to derive a labour supply elasticity with respect to the return to work by treating the reform as a source of exogenous change to the pension return to work. To do this, it is necessary to quantify the financial impact of the reform. We estimate elasticities by relating this to our estimate of the change in labour supply in response to the reform.

Our labour supply estimates capture the change in the return to work driven by the pension reform. Because our estimates of the increase in total labour supply were driven by the extensive margin, it is very likely that our labour supply results were only driven by the change in the return to work on the extensive margin. Given this, we focus on measuring the change in the extensive margin return to work, making the assumption that there is no impact of intensive margin incentives.

As discussed in Section 2.3, we can measure the financial change in the return to work as the change in the discounted pension value of working an additional year. Given that all senior doctors in our sample were far from the new plan’s retirement age when moved onto it, we assume that all were more than three years from retirement, and so the change in the return to work is given by Equation (1). To estimate an elasticity, we therefore need to estimate this financial change. This depends on income expectations, age and life expectancy. We estimate income expectations using the age and gender income profiles of older doctors who were not affected by the reform. For age, we use the post-treatment distribution of ages among the treated doctors in our sample. For life expectancy, we use gender-specific life expectancy data for the whole population (rather than just senior doctors). Appendix B has full details.

We estimate that the average change in the return to work for treated doctors in our sample is an increase of 32.7%. This is a large increase in financial incentives. To understand what drives this change, Table 3 shows the average value of each component of Equation (1). The higher accrual rate and the uprating of the new pension scheme increase the returns to working another year. This is partially offset by the later retirement age and the change in accrued income.

Table 3: Average value of each component of Equation (1)

Component	Description	Average Value
80/54	Higher accrual rate	1.45
$y_{age} / \max\{\hat{y}_{r-1}, \hat{y}_{r-2}, \hat{y}_{r-3}\}$	Change in accrued income	0.91
$(1 + \pi + 0.015)^{66-age} / (1 + \pi)^7$	Uprating of new plan	1.88
$\sum_{a=67}^{100} P_a \frac{1}{(1+\delta)^{a-60}} / (3P_{60} + \sum_{a=60}^{100} P_a \frac{1}{(1+\delta)^a})$	Later retirement age and no lump sum	0.53

Our estimates in Section 5.1 suggest that the change in the returns to work from being moved to the new pension scheme increased labour supply on the extensive margin by 2.9%. Dividing this change in labour supply by the change in financial incentives provides an estimate of the labour supply elasticity with respect to the financial change. This extensive margin labour supply elasticity estimate with respect to the return to work is 0.090, with a 95% confidence interval of 0.024 to 0.157.¹³

¹³To derive the confidence interval, we calculate the elasticity using the same estimate of the change in return

Previous work has estimated doctor wage elasticities in the range of 0.1 to 0.5 (Lee et al., 2019). In a comparable NHS system, for example, Baltagi et al. (2005) estimated a short-run wage elasticity of 0.3 for male doctors in Norway. Our extensive margin labour supply elasticity (0.09) is low relative to this and to other studies that have examined the labour supply responses of doctors to the return to current income. It is likely that doctors are less responsive to changes in delayed remuneration, even after discounting, relative to current wage changes. In addition, it may also be that, because of the complex nature of the pension plans and the reform, many doctors did not realise or underestimated the extent to which their current incentives had changed. In particular, whilst doctors would have been well aware of the change in accrual due to the change in accrual rates (term 1 in Equation (1)), the change to the career average (term 2) and the fact that the pension is paid over fewer years (term 4), they are less likely to have understood that the uprating in the new pension plan was very generous relative to their previous plan (term 3). As this drives a large part of the relative financial benefit of the new pension plan, underestimating this would have meant underestimating the return to work.

French et al. (2022) estimate a labour supply elasticity with respect to the net return to work of 0.44 for Polish employees aged 51-54 using a reform to the public pensions system. Their definition of net return to work is similar but not directly comparable to our elasticity. Their estimated elasticity (0.44) is much larger than our elasticity (0.09), suggesting that doctors are less responsive to pension incentives than the whole population of (Polish) workers. This is not surprising as, within a cohort composed of all workers, there are more individuals who can easily adjust their labour supply than amongst a cohort of senior doctors.

Taken together, our estimated labour supply elasticity is close to the lower bound of labour supply elasticities in the literature, suggesting that our cohort of senior doctors are relatively unresponsive to changes in pension financial incentives. However, as our estimates of the impact of the reform show, if changes in financial incentives are large enough, there can still be material changes in labour supply.

to work but replace our labour supply estimate with its 95% confidence interval (an increase of 0.8% to 5.1%).

7 Conclusion

This paper examines the consequences of a change in financial incentives brought about by a pension reform in the English NHS on the labour supply of senior doctors. The new pension plan involved a number of changes, including both a change in the return to work via pension accrual and a change to the retirement age of the plan, and the reform has a theoretically ambiguous impact on labour supply. To isolate the causal impact of the change in the return to work induced by the reform changes to pension accrual, we exploit the staggered rollout of the new plan and focus on a narrow group of senior doctors born within 7 years of each other. This allows us to difference out common labour supply anticipation effects that may have been driven by the change in pension wealth and allows us to cleanly identify the labour supply impact of the change in returns to work.

We find that the effect of the change in returns to work after being moved onto the new scheme was to increase the labour supply of senior doctors. This was driven by an increase on the extensive margin. Our findings are robust to a large set of tests. Combining these labour supply results with estimates of the size of the financial change in the return to work brought about by the reform, we estimate a labour supply elasticity of 0.09 with respect to returns to work given by the pension change.

Our focus on identifying causal estimates of the pension change means we do not estimate the labour supply response to the large change in pension age (except via responses to the changes this made to the return to work). However, this has been studied extensively in other contexts for workers closer to retirement, as summarised by Blundell et al. (2016). Nor can we estimate the effect on younger doctors further down the career ladder or those considering entry into the profession. For these individuals, the pension changes could induce less entry, less effort to attain high salaries at the end of their careers or more movement out of the NHS earlier in their careers, any of which could reduce total labour supplied to the NHS. Or it may be that myopic behaviour means that pension reforms have little effect on the behaviour of these younger individuals.

Our results also indicate that reforms to public sector pension plans of the kind that have been adopted in many countries to increase the return to work can increase - at least in the

short term - the labour supply of key senior public sector workers, many of whom have had generous pension schemes to encourage them to enter and stay. However, as a tool to increase labour supply in the public sector, our findings that the senior doctors in our analysis are less responsive to the change in the return to work from delayed remuneration than for current wages may suggest that pension reforms may be a less effective method to increase labour supply in public services than direct wage increases. However, there is also the cost side to be considered: direct wage increase will certainly increase the costs of service provision, whilst pension reforms are often intended to decrease these costs.

References

- Andreassen, L., M. L. Di Tommaso, and S. Strøm (2013). Do medical doctors respond to economic incentives? *Journal of Health Economics* 32(2), 392–409.
- Arenas de Mesa, A., D. Bravo, J. Behrman, O. S. Mitchell, and P. E. Todd (2006). The Chilean pension reform turns 25: lessons from the social protection survey.
- Baltagi, B. H., E. Bratberg, and T. H. Holmås (2005). A panel data study of physicians' labor supply: the case of Norway. *Health Economics* 14(10), 1035–1045.
- Banks, J. and R. Crawford (2022). Managing retirement incomes. *Annual Review of Economics* 14, 181–204.
- Banks, J. and Z. Oldfield (2007). Understanding pensions: Cognitive function, numerical ability and retirement saving. *Fiscal Studies* 28(2), 143–170.
- Blundell, R., E. French, and G. Tetlow (2016). Retirement incentives and labor supply. In *Handbook of the economics of population aging*, Volume 1, pp. 457–566. Elsevier.
- Borusyak, K., X. Jaravel, and J. Spiess (2021). Revisiting event study designs: Robust and efficient estimation.
- Bottazzi, R., T. Jappelli, and M. Padula (2006). Retirement expectations, pension reforms, and their impact on private wealth accumulation. *Journal of Public Economics* 90(12), 2187–2212.
- Bovini, G. (2019). *Essays in applied economics*. Ph. D. thesis, London School of Economics and Political Science.
- Brekke, K. R., T. H. Holmås, K. Monstad, and O. R. Straume (2017). Do treatment decisions depend on physicians' financial incentives? *Journal of Public Economics* 155, 74–92.
- Broadway, B., G. Kalb, J. Li, and A. Scott (2017). Do financial incentives influence GPs' decisions to do after-hours work? A discrete choice labour supply model. *Health Economics* 26(12), e52–e66.

- Callaway, B. and P. H. Sant'Anna (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics* 225(2), 200–230.
- Chan, S. and A. H. Stevens (2008). What you don't know can't help you: Pension knowledge and retirement decision-making. *The Review of Economics and Statistics* 90(2), 253–266.
- Crawford, R. and H. Karjalainen (2020). Awareness of state pension entitlements. *The Dynamics of Aging: Evidence from the English Longitudinal Study of Aging 2002–2019*, 10–52.
- De Chaisemartin, C. and X. d'Haultfoeuille (2018). Fuzzy differences-in-differences. *The Review of Economic Studies* 85(2), 999–1028.
- French, E., A. S. Lindner, C. O'Dea, and T. A. Zawisza (2022). Labor supply and the pension-contribution link. Technical report, National Bureau of Economic Research.
- Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics* 225(2), 254–277.
- Hinrichs, K. (2021). Recent pension reforms in europe: More challenges, new directions. an overview. *Social Policy & Administration* 55(3), 409–422.
- House of Commons Library (2019). Public service pensions - employer contributions. <https://researchbriefings.files.parliament.uk/documents/CBP-7539/CBP-7539.pdf>.
- Ikenwilo, D. and A. Scott (2007). The effects of pay and job satisfaction on the labour supply of hospital consultants. *Health Economics* 16(12), 1303–1318.
- Institute for Fiscal Studies (2021). Ifs spending composition sheet. <https://ifs.org.uk/taxlab/taxlab-data-item/ifs-spending-composition-sheet>.
- Kelly, E. and I. Stockton (2022). Maternity and the labour supply of nhs doctors and nurses.
- Laibson, D. (1997). Golden eggs and hyperbolic discounting. *The Quarterly Journal of Economics* 112(2), 443–478.

- Lee, T., C. Propper, and G. Stoye (2019). Medical labour supply and the production of health-care. *Fiscal Studies* 40(4), 621–661.
- Liebman, J. B. and E. F. Luttmer (2015). Would people behave differently if they better understood social security? evidence from a field experiment. *American Economic Journal: Economic Policy* 7(1), 275–299.
- Mastrobuoni, G. (2011). The role of information for retirement behavior: Evidence based on the stepwise introduction of the social security statement. *Journal of Public Economics* 95(7-8), 913–925.
- Meghir, C. and D. Phillips (2010). Labour supply and taxes. *Dimensions of tax design: The Mirrlees review*, 202–74.
- Mitchell, O. S. (1988). Worker knowledge of pension provisions. *Journal of Labor Economics* 6(1), 21–39.
- Morris, S., B. Elliott, A. Ma, A. McConnachie, N. Rice, D. Skåtun, and M. Sutton (2008). Analysis of consultants' nhs and private incomes in england in 2003/4. *Journal of the Royal Society of Medicine* 101(7), 372–380.
- NHS Digital (2022). Nhs workforce statistics - december 2021. <https://digital.nhs.uk/data-and-information/publications/statistical/nhs-workforce-statistics/december-2021/>.
- Nicholson, S. and C. Propper (2011). Medical workforce. In *Handbook of health economics*, Volume 2, pp. 873–925. Elsevier.
- OECD (2005). Pensions at a glance: Public policies across oecd countries.
- Office for National Statistics (2024). National life tables: England. <https://www.ons.gov.uk/peoplepopulationandcommunity/birthsdeathsandmarriages/lifeexpectancies/datasets/nationallifetablesenglandreferencetables>.


Stoye, G. (2019). Recent trends in independent sector provision of nhs-funded elective hospital care in england. *London: Institute for Fiscal Studies*.

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

Whiteford, P. and E. Whitehouse (2006). Pension challenges and pension reforms in oecd countries. *Oxford review of economic policy* 22(1), 78–94.


A Appendix: Additional results

Figure A1: Average pension plan participation by age 2012-2014 conditional on NHS participation




Note: Average share of senior doctors in the NHS pension by year of age in each month between 2012-2014. Excludes ages with less than 1,000 senior doctor month observations.

Figure A2: Average NHS participation by age for never treated doctors



Note: Estimated on senior doctors born between 1959 and March 1962 and working in the NHS in April 2012.

Figure A3: Estimated total impact of policy



Note: Total effect in each period is calculated as the sum of each treatment coefficient multiplied by the number of senior doctors that number of months after treatment.

B Appendix: Calculating financial impacts of the reform

B.1 Measuring the change in the return to work

We assume that the senior doctor chooses between working this year or not, rather than choosing whether to leave permanently or not. Under the old pension plan, working an additional year (at any income) increases the annual pension value by $1/80$ of the best final three years of income, given by Equation B1.

$$\frac{1}{80} \times \max\{\hat{y}_{r-1}, \hat{y}_{r-2}, \hat{y}_{r-3}\} \quad (\text{B1})$$

Under the new pension plan, working an additional year (at income y_{age}) increases the annual pension value by $1/54$ of y_{age} , uprated for each year of working.¹⁴ Assuming that the doctor retires upon turning 67, this means the increase in the annual pension value is Equation B2.

$$\frac{1}{54} \times y_{age} \times (1 + \pi + 0.015)^{66-age} \quad (\text{B2})$$

¹⁴This uprating is based on CPI inflation + 1.5%. We do not model the in-year timing of this annual uprating.

The annual value of the pension is paid each year from retirement and is maintained in real terms.¹⁵ We measure the return to work as the change in the discounted value of the pension at age 60 (in age 60 prices), as this is the retirement age of the old scheme.

Assuming the real discount rate δ , under the old pension plan the value of working an additional year (at any income) is Equation B3, where P_a is the probability of living to age a . This accounts for the additional lump sum in the first year of retirement.

$$\frac{1}{80} \times \max\{\hat{y}_{r-1}, \hat{y}_{r-2}, \hat{y}_{r-3}\} \times \left(3P_{60} + \sum_{a=60}^{100} P_a \frac{1}{(1+\delta)^{a-60}} \right) \quad (\text{B3})$$

For the new pension plan, the value of working an additional year at income y_{age} aged 67 (in age 67 prices) is Equation B4.

$$\frac{1}{54} \times y_{age} \times (1 + \pi + 0.015)^{66-age} \times \sum_{a=67}^{100} P_a \frac{1}{(1+\delta)^{a-67}} \quad (\text{B4})$$

We move this to age 60 (in age 60 prices) by dividing by $(1 + \pi)^7(1 + \delta)^7$, which gives Equation B5.

$$\frac{1}{54} \times y_{age} \times \frac{(1 + \pi + 0.015)^{66-age}}{(1 + \pi)^7} \times \sum_{a=67}^{100} P_a \frac{1}{(1 + \delta)^{a-60}} \quad (\text{B5})$$

We calculate the percentage change in the return to work as Equation X divided by Equation Y minus 1. This equals Equation B6.

$$\frac{80}{54} \times \frac{y_{age}}{\max\{\hat{y}_{r-1}, \hat{y}_{r-2}, \hat{y}_{r-3}\}} \times \frac{(1 + \pi + 0.015)^{66-age}}{(1 + \pi)^7} \times \frac{\sum_{a=67}^{100} P_a \frac{1}{(1+\delta)^{a-60}}}{3P_{60} + \sum_{a=60}^{100} P_a \frac{1}{(1+\delta)^{a-60}}} - 1 \quad (\text{B6})$$

which is Equation 1.

B.2 Estimating the change in the return to work

We assume that the (expected) inflation rate equals the Bank of England target of 2.0%. We assume that the real discount rate is the UK government's SCAPE rate at the time of the reform, 3.0%, which is used to discount public sector pension plans (House of Commons Library,

¹⁵Note we abstract away the in-year timing of this annual increase because it affects both pension schemes equally and so does not affect the percentage difference.

2019). We use official life tables for the whole population between 2017 and 2019 to estimate P_a separately by gender (Office for National Statistics, 2024).

The second term of Equation 1 requires us to estimate expected earnings profiles. We do this using the cohort of older senior doctors born before 1970 who were not affected by the reform over our whole sample period (2012 to 2021). We calculate the average pensionable income by age and gender in 2014 prices using pay awards to adjust for nominal earnings growth. Figure B1 shows the age-profiles in earnings separately for male and female senior doctors.

We then assume that nominal pay was expected to grow at the same rate as in 2012-13 to 2014-15, which was 0.5% per year. We use these counterfactual earnings to estimate $\hat{y}_{r-1}, \hat{y}_{r-2}, \hat{y}_{r-3}$, but also y_{age} as we calculate the change in the extensive margin to work (i.e. we calculate the change in the return to work holding earnings unchanged).

We do not assume that $r = 60$ and instead account for the strategic dropout of the old pension plan. Figure B2 shows the age-profiles in pension participation separately for male and female senior doctors, calculated on the same sample we calculate earnings profiles. For simplicity, we assume that all senior doctors were in the pension scheme until 50 and then use the actual decline in participation. We calculate $\hat{y}_{r-1}, \hat{y}_{r-2}, \hat{y}_{r-3}$ weighted by the share who retire by each age after 50.

Figure B1: Pensionable earnings profile of never treated senior doctors in 2014 prices




Figure B2: Pension participation profile of never treated senior doctors in 2014 prices

