

Crawford, Rowena; Simpson, Polly

Working Paper

The impact of house prices on pension saving in early adulthood

IFS Working Paper, No. W20/38

Provided in Cooperation with:

The Institute for Fiscal Studies (IFS), London

Suggested Citation: Crawford, Rowena; Simpson, Polly (2020) : The impact of house prices on pension saving in early adulthood, IFS Working Paper, No. W20/38, Institute for Fiscal Studies (IFS), London, <https://doi.org/10.1920/wp.ifs.2020.3820>

This Version is available at:

<https://hdl.handle.net/10419/242897>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Rowena Crawford
Polly Simpson

20/38

Working paper

The impact of house prices on pension saving in early adulthood

The impact of house prices on pension saving in early adulthood

Rowena Crawford
Institute for Fiscal Studies

Polly Simpson
London School of Economics

December 2020

Abstract

In this paper, we estimate the effect of house prices on whether or not young adults actively save in a private pension. We use job-level data from a survey of employers, matched to average house prices at the level of an individual's location of employment, exploiting geographical variation in local house price movements in England over the decade 1997 to 2007. We find that after controlling for individual and job characteristics there is no statistically significant effect, on average, across all employees. There is a negative effect for public-sector workers and those in the middle of the earnings distribution. However, the effects are small – for example, among public-sector workers, if house prices are £100,000 higher, then this is associated with a 3 percentage point lower probability of contributing to a pension. The effect is larger among employees in the NHS, education and non-uniformed services, who face a higher employee contribution than employees in the civil service.

JEL: D14, D15

Acknowledgements

This research has been funded by the IFS Retirement Saving Consortium – which comprises Age UK, Association of British Insurers, Aviva UK, Canada Life, Chartered Insurance Institute, Department for Work and Pensions, Investment Association, Legal and General Investment Management, and Money and Pensions Service – and by the Economic and Social Research Council through the Centre for the Microeconomic Analysis of Public Policy at the Institute for Fiscal Studies (grant reference ES/M010147/1). All interpretation and any errors are those of the authors.

This work is based on data from the Annual Survey of Hours and Earnings, produced by the Office for National Statistics (ONS) and accessed through the ONS Secure Research Service. The data are Crown Copyright and reproduced with the permission of the Controller of HMSO and the Queen's Printer for Scotland.

This work has been produced using statistical data from ONS. The use of the ONS statistical data in this work does not imply the endorsement of the ONS in relation to the interpretation or analysis of the statistical data. This work uses research data sets that may not exactly reproduce National Statistics aggregates.

1. Introduction

House prices in England have increased substantially over the past two decades: the median house price increased by 173% between 1997 and 2007 (Cribb and Simpson, 2018). The ratio between average house prices and average individual gross earnings increased from 3.7 to 7.6 over the same period. This trend has been linked in public debate to the lower rate of homeownership amongst the most recent generations. Compared with the cohorts born a decade earlier, UK residents born in the early 1980s were 18 percentage points (ppt) less likely to own a home at age 30 (Cribb, 2019).

In this paper, we seek to understand the impact of house prices on another aspect of the balance sheet of young adults: pension membership. We take an empirical reduced-form approach, estimating whether house prices have an impact on pension membership in early adulthood by exploiting geographical variation in local house price movements over the decade 1997 to 2007.

There are several mechanisms through which house prices plausibly affect the pension saving of young adults. If mortgage lenders require buyers to pay a minimum proportion of the property price as a down-payment, then higher house prices will mean that aspiring homeowners have to save more in order to purchase a house of a given ‘quality’. In theory, households could respond in numerous ways: by saving a larger proportion of their income, saving for longer, purchasing a smaller or lower-quality house, or even permanently forgoing owner-occupation. Cribb and Simpson (2018) find that mortgage repayments for young homeowners in England increased over this period despite substantial falls in interest rates, implying that individuals are not just buying smaller or lower-quality properties. Bottazzi, Crossley and Wakefield (2015) find that historic cohort gaps in homeownership rates at age 30, which are associated with differences in UK house prices, have mostly closed by age 40, suggesting that higher house prices in the past have not led to large proportions permanently forgoing owner-occupation. Sheiner (1995) finds that higher house prices are associated with greater wealth accumulation among young renters – direct evidence, albeit in the US context, that higher house prices imply greater down-payment saving. To the extent that higher house prices do drive greater saving for down-payments, with a fixed budget constraint this could be accommodated by reduced pension saving rather than (or in addition to) lower spending.

Other mechanisms may also be at play. For example, higher house prices could increase housing expenditures in the form of higher spending on mortgage repayments for owners – as shown empirically to be the case in Cribb and Simpson (2018) – and higher rents (if these are positively correlated with house prices) for renters. This could cause individuals to postpone more of their pension saving until later in life when earnings are typically higher.

Higher housing costs could also affect how much pension wealth individuals seek to accumulate over their lifetimes. For homeowners, if higher house prices cause a shift in the optimal allocation of

consumption between housing and non-housing, then less pension saving may be needed in order to finance the lower levels of non-housing consumption in retirement (with housing consumption in retirement already having been ‘paid for’ through mortgage payments during working life). Those downsizing their housing in retirement will also have more housing wealth to access if house prices are higher, implying that less pension saving is required to finance non-housing consumption in retirement. For those who are renters for their entire lifetime, however, higher house prices will likely imply higher rents – in both working life and retirement. To the extent that higher house prices shift individuals from being owners to being lifetime renters, this could substantially increase housing expenditures in retirement. This could require an increase in pension saving to finance these costs, but the picture is complicated by the provisions of the UK welfare system (which includes a component specifically for renters).

In summary, aside from those who are shifted from expecting to be owner-occupiers by retirement to expecting to be lifetime renters, higher house prices would be expected to reduce pension saving.¹ This could be either a temporary effect, if the timing of pension saving is shifted until later in working life, or a permanent effect, if optimal expenditure shares are shifted more towards housing versus non-housing consumption.

In this paper, we estimate the reduced-form effect of house prices on pension membership. While formally we abstract from the different mechanisms discussed above, our interest in the potential effect through increased down-payment saving leads us to focus on individuals at the start of their working lives, who will be making decisions about whether and how intensely to save for the upfront costs of homeownership. Specifically, we analyse the pension saving behaviour of individuals aged 21–35, an age range over which cohort-specific homeownership rates historically increased rapidly.

We use job-level data from a survey of employers – the Annual Survey of Hours and Earnings (ASHE) – matched to average house prices at the level of an individual’s location of employment, for employees in England over a period of rising house prices (1997–2007). Our ordinary least-squares (OLS) estimates show that an increase of £100,000 (in 2018 prices) in local average house prices is, on average, associated with a 2 ppt reduction in the probability of being enrolled in a workplace pension, after controlling for time, age, earnings, occupation, industry, sector, employer size and gender. We estimate that this effect is more negative for younger individuals, as would be expected were the down-payment saving mechanism to be the driving factor.

These small (and statistically insignificant) results, however, hide heterogeneity across individuals. In particular, we find evidence of a ‘U-shape’ when we examine how the effect of house prices varies

¹ We abstract in this theoretical discussion from the potential correlation of house prices with other asset returns, which would affect overall incentives to save. We also abstract from the role of intergenerational transfers, which could potentially act to ameliorate the effects of higher house prices and deposit requirements.

with current earnings. Those with the lowest and highest earnings are least sensitive to house prices, while for those in the middle of the earnings distribution higher house prices have a small but statistically significant negative effect on pension membership. For those in the middle earning quintile, an increase of £100,000 in house prices is associated with a 2 ppt lower probability of contributing to a pension. This result is intuitive. Those on the lowest incomes are least likely to save in a private pension irrespective of house prices. Those on the highest incomes are more likely to be able to afford to save for a down-payment and contribute to a pension simultaneously, and they are more likely to have become homeowners already when observed (and therefore they will no longer be saving for a down-payment). The greater sensitivity to house prices among middle earners also accords with the evidence that recent falls in homeownership rates have been sharpest for young adults with middle incomes (Cribb, Hood and Hoyle, 2018).

We also examine heterogeneity in the estimated effect between public- and private-sector employees. It has been well documented that employees in the public sector benefit from relatively more generous private pension provision (Disney, Emmerson and Tetlow, 2009; Cribb and Emmerson 2016) – in particular, public-sector employees are still typically offered access to a defined benefit (DB) pension, where pension benefits in retirement are calculated using a formula that explicitly depends on years in the pension scheme, while such schemes are rarely available to private-sector employees. We find a small negative effect of house prices on pension membership, on average, across the public sector and little effect, on average, across the private sector. At first sight, this is surprising as those opting out of saving in a public-sector DB scheme are forgoing large (implicit) employer pension contributions. However, the likely explanation is that participation in public-sector pension schemes involves higher employee contributions on average than participation in private-sector pension schemes. Public-sector employees choosing not to join the private pension offered to them would therefore increase their capacity to save for a house to a greater extent than private-sector employees making a similar decision. This explanation is supported by the greater estimated effect among employees in the NHS, education and other non-uniformed services than among employees in the civil service – the average effect of an increase of £100,000 in house prices is –4 ppt as opposed to –2ppt – because the former are required to make higher employee contributions to their pension.

Taken together, our results suggest that higher house prices do have some crowd-out effects on the extensive margin of pension saving. Specifically, higher house prices reduce the probability that young adults – particular those in the public sector and those in the middle of the earnings distribution – enrol in a private pension. The estimated effects are, however, small. We do not attempt in this work to unpick the mechanisms driving this result – whether pension saving is being postponed until individuals are on the housing ladder or until later in working life, or whether retirement consumption is being weighted more towards housing consumption than non-housing consumption – but this would be an interesting avenue for future research. Whether or not pension saving is crowded out on the

intensive margin, with individuals' decisions about the size of pension contributions affected by house prices, is also an important question for future research.

This paper builds on a large existing body of literature on the impact of house price shocks on household financial decisions. House prices have been shown causally to increase consumption (Aladangady, 2017), to increase spending on residential investment (Crossley, Levell and Low, 2020), to increase leverage (Mian and Sufi, 2011) and to affect portfolio choice (Chetty, Sándor and Szeidl, 2017). While most of this literature focuses on the response of existing homeowners to shocks to their housing equity, a small subset considers differential effects by age and tenure status. For example, using data from the UK between 1988 and 2000, Campbell and Cocco (2007) find an effect close to zero of house prices on consumption for young renting households, which is consistent with our finding of no significant effect (on average) of house prices on pension membership for the younger adults that we study. We build on this work, analysing financial behaviour in early adulthood, before many have made an initial housing purchase.

Our work is quite related to Sheiner (1995), who examined the effect of house prices on the liquid net wealth of young renters and found evidence of increased saving in forms that could be accessed to make down-payment on a property. Our interest though is the effect of house prices on pension membership through mechanisms such as increased down-payment saving, rather than the estimation of the effect of house prices on down-payment saving itself. Our work complements the literature on the impact of housing and house prices on portfolio choice over the lifecycle. Cocco (2005) finds in the US that down-payment constraints and mortgage contracts reduce the liquid wealth available to younger investors to invest in stocks, lowering participation in the equity market. We seek to understand whether a similar mechanism reduces pension saving in the UK, although we are unable to observe directly the entirety of the household portfolio, and we take a reduced-form rather than a structural approach.

The rest of this paper proceeds as follows. In Section 2, we provide context by describing the institutional framework for private pensions in the UK, and describing key trends in homeownership rates. In Section 3, we explain the data sources we use and we set out key facts about the geographical variation in house prices and pension membership that we exploit in our analysis. In Section 4, we outline our empirical strategy, including our baseline OLS specification, how we allow for heterogeneous effects and the two-stage least-squares (2SLS) strategy that we use to address concerns about omitted variables, exploiting geographical variation in supply constraints interacted with declining real interest rates. We discuss our results in Section 5. We conclude in Section 6 with implications for policy and avenues for future research.

2. Context

We start by briefly setting the context, describing the institutional setting and some key facts for private pension saving and owner-occupation in the UK.

Private pensions in the UK

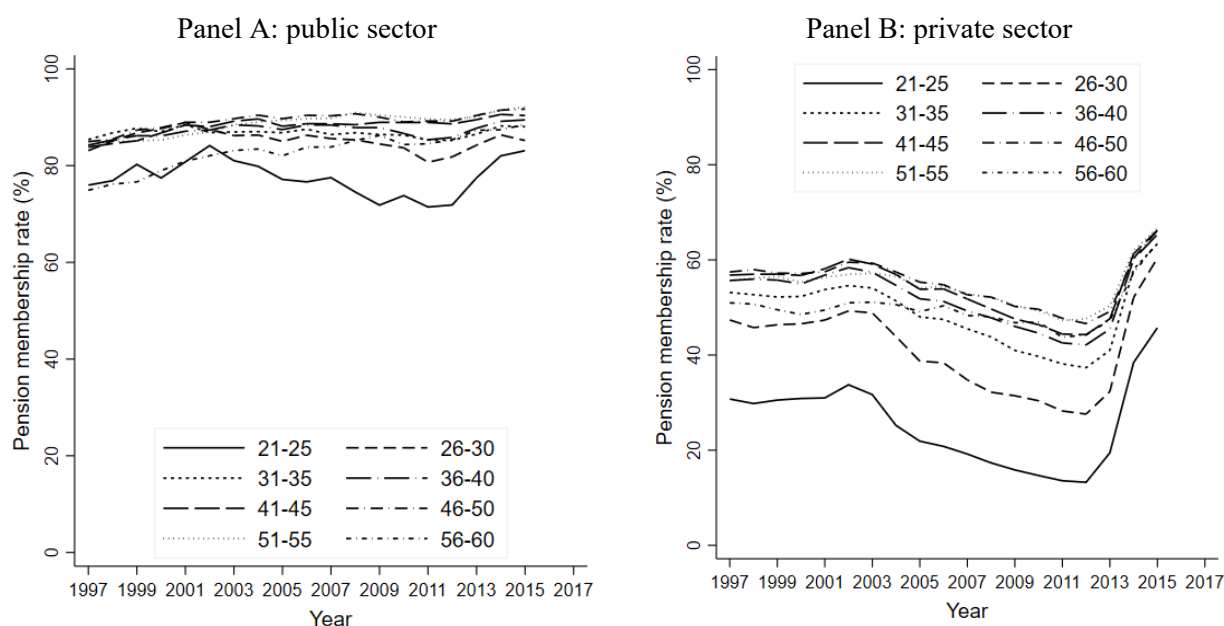
In the UK, private pensions play an important role in enabling individuals to smooth their living standards into and through retirement. The public pension paid to older individuals by the government (known as the ‘state pension’) is not particularly generous relative to average working-age earnings. In 2020–21, the ‘new state pension’ provides £175.20 per week (£9,110 per year) to those with full contribution records – that is, 35 years of national insurance contributions (payroll tax) or credits for certain caring responsibilities or receipt of certain out-of-work benefits. This is equivalent to around one-third of median earnings. The level of pension income received is not related to individuals’ earnings, or to years of activity in excess of 35 years. This means that for middle and higher earners the state pension provides only a relatively low replacement of their working-life income, and additional private saving is required by many to smooth consumption.

Saving in private pensions is encouraged through tax incentives: contributions are made out of pre-tax earnings and investment returns are untaxed, with pension income being taxed when it is drawn in retirement. Furthermore, up to one-quarter of accumulated pension saving can be withdrawn tax-free. However, private pension saving is illiquid, and cannot be accessed until age 55 except under exceptional circumstances (such as terminal illness).

Private pensions include those provided by public and private employers (‘occupational pensions’), and those provided by insurance companies. Employers may facilitate access to a pension provided by an insurance company. Together occupational pensions and these facilitated pensions are known as ‘workplace pensions’. These account for the vast majority of private pensions, and are what we capture in our data.

Figure 1 shows how the proportion of employees who are members of a workplace pension has changed over time for employees in the public sector (panel A) and in the private sector (panel B). In the public sector, membership of private pensions has been consistently high. In 1997, over 80% of public-sector employees aged 26–55 were members of a workplace pension, and this proportion has been relatively stable over time. In the private sector, where private pensions are not always provided by employers, and a declining proportion of pensions provided are DB pensions, pension membership rates are lower. In 1997, only around 56–57% of private-sector employees aged 36–55 were members of a pension, and this membership rate declined between 2003 and 2012, reaching 42–48%.

Figure 1. Membership rates in workplace pensions over time, by age and sector



Source: ASHE, 1997–2015.

In late 2012, the government started rolling out its ‘automatic enrolment’ policy. This reform required employers to provide access to a workplace pension, to enrol most of their employees automatically into it, and to pay minimum levels of contributions for those who did not choose to opt out. A discussion of this reform, and a formal evaluation of its effects on pension membership, can be found in Cribb and Emmerson (2020). The substantial boost to pension membership rates that resulted, particularly among younger employees, is clear from Figure 1.

It can also be seen from Figure 1 that there is an age gradient in pension membership, with membership rates being markedly lower among younger private-sector employees than older private-sector employees. For example, in 1997, only 31% of those aged 21–25 were members of a workplace pension, compared with 47% of those aged 26–30 and 53% of those aged 31–35. Among those aged 36–55, the proportions saving in a pension are much more similar across the five-year age groups. Among public-sector employees, it is only those aged 21–25 among whom the pension membership rate is appreciably lower. These differences in pension membership by age have reduced since 2012. Bourquin, Cribb and Emmerson (2020) found that automatic enrolment substantially reduced differences in pension membership across individuals with different characteristics – such as age – largely as a result of the switch to membership by default.

The key facts on private pensions for the purposes of our paper are therefore that, while saving in a private pension is important for smoothing living standards into retirement for middle and higher earners, before automatic enrolment most individuals were not saving in a pension in all years of working life. In particular, employees were less likely to save in a workplace pension at younger ages, particularly in the private sector, and the proportion of young employees saving in a pension declined

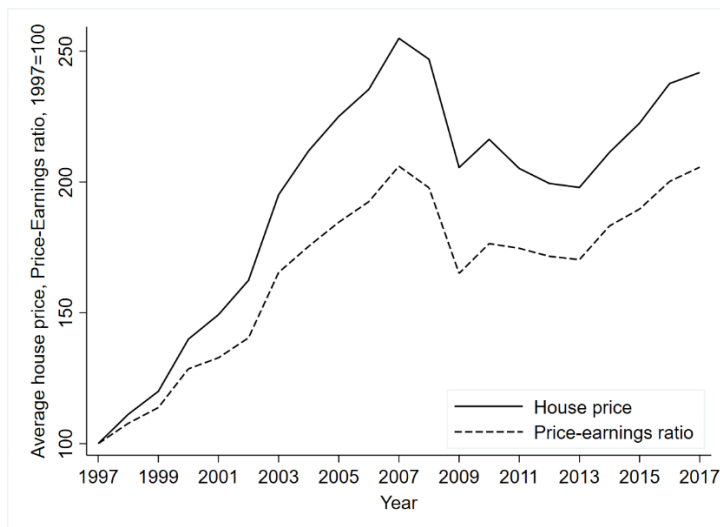
over time (before 2012). Individual decisions over this period are being taken against a backdrop of less generous pension provision in the private sector, associated with a shift from DB to defined contribution (DC) pension provision, which will vary in systematic ways with the characteristics of employers.

Housing and homeownership in the UK

House prices increased substantially in the UK between 1997 and 2007, and far more rapidly than average earnings. Figure 2 shows that the ratio between the UK average house price and average weekly earnings of full-time employees more than doubled over this period. After 2007, house prices fell sharply as a result of the financial crisis and associated recession, and they remained at a depressed level until around 2013, since when prices and the ratio of house price to earnings has again increased.

This dramatic increase in the ratio of house prices to earnings over the past two decades as a whole has been linked in the public debate to the marked decline in homeownership among more recently born generations (Cribb, 2019). Figure 3 shows how homeownership rates at different ages have changed over time (left-hand panel), and how the ownership rates of particular birth cohorts have changes as the cohort has aged (right-hand panel). Three important facts from these graphs are important for the context of this paper.

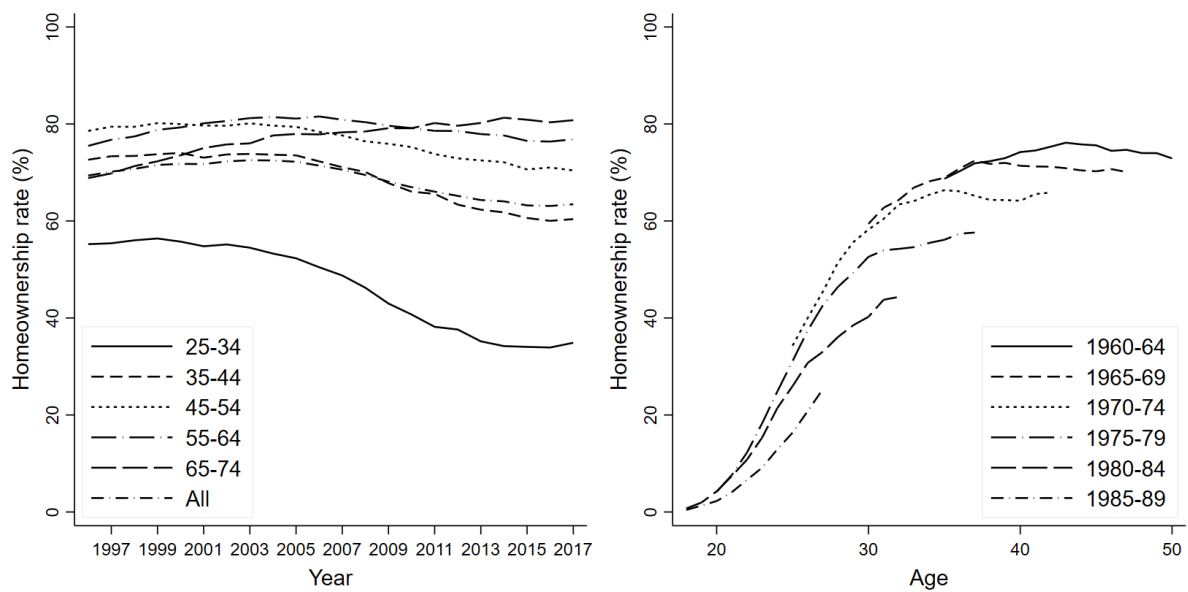
Figure 2. Average house price and price–earnings ratio in UK, 1997–2017



Note: Data are for whole of the UK. House price is the UK average from the UK House Price Index (UKHPI). Earnings are gross weekly earnings of full-time employees from the ASHE (multiplied by 52). House prices are in real terms in 2018 prices.

Source: UKHPI, ASHE.

Figure 3. Homeownership in the UK, by age–year (left-hand panel) and cohort–age (right-hand panel)



Source: Left-hand panel: figure 1 from Cribb and Simpson (2018). Right-hand panel: figure 1 from Cribb, Hood and Hoyle (2018).

First, homeownership rates increase rapidly at younger ages. For example, among those born in the late 1970s, homeownership rates were near zero at age 20 but over 40% by age 30. This motivates our interest of the sample of individuals aged 21–35, to capture those most likely to be making decisions over whether and how to save for a housing down-payment.

Second, in recent decades, the majority of those reaching retirement have done so owning their own property. Housing has therefore been an important resource in retirement, providing a flow of housing consumption, or alternatively a stock of wealth that can be drawn on to fund consumption.

Third, cohorts born more recently are noticeably less likely to own a home in their 20s and 30s than cohorts that came before them. At age 30, 40% of adults born in 1980–84 owned their own home, compared with 53% of those born in the five preceding years, and 58% of those born between 1970 and 1974. Whether this is a temporary divergence in ownership rates, or whether these rates will persist through to individuals’ retirement, is an important open question with implications for individuals’ retirement income needs.

3. Data

Annual Survey of Hours and Earnings

We undertake our analysis using job-level microdata on a 1% sample of employees in England from the Annual Survey of Hours and Earnings (ASHE).² This is an annual survey that has been conducted each April since 1997. The data are collected directly from employers and, as such, contain accurate information on firm and job characteristics as well as earnings, but only a limited set of individual characteristics. Particularly relevant for our analysis, we observe: earnings, membership of workplace pensions, industry (five-digit SIC code), occupation (four-digit SOC code), employer size, geographical location of work, age and gender.³ In addition, from 2005 onwards, we consistently observe pension contributions and geographical location of residence. We do not observe any other individual characteristics, such as education or whether an individual is part of a couple, or details of an individual's financial situation other than their earnings, such as whether or not they are an owner-occupier (or other details such as benefit income, self-employment income or wealth).

We select our sample from the ASHE on the basis of several criteria. First, we use only individuals aged 21–35. This age range, our definition of 'early adulthood', captures the window during which homeownership rates have historically increased rapidly with age for each cohort (before broadly stabilising), therefore making this the age range in which most individuals make decisions about whether and how intensely to save for the upfront costs of homeownership.

Second, we focus primarily on data from 1997–2007. This covers from the first year of ASHE until the financial crisis and property price bust that happened in 2007–08 in the UK. As such, our estimation focuses on a time period with a relatively stable macroeconomic environment, and rising housing demand and prices. This is particularly important in our 2SLS analysis, as the asymmetry of housing supply means that the intuitive link between housing supply constraints and house prices holds in a time of growing, but not falling, demand. We do, however, present some results for how the estimated relationship between house prices and pension membership changes after 2007 (when the financial crisis dominated the economic landscape), and after 2012 (when automatic enrolment significantly altered individuals' private pension saving).

Third, we restrict analysis to individuals employed in England outside of London to avoid our results being driven by the particularities of the London labour and housing market. We measure 'local' house prices that individuals are exposed to, on the basis of the local authority in which they work. For young employees across most of England, around half work and live in the same local authority.

² Office for National Statistics (2020).

³ The ASHE data contain measures of average hourly earnings and average weekly earnings, both with and without overtime. Throughout, we use, as our measure of earnings, average gross weekly earnings including overtime. Our results are not sensitive to this choice of earnings measure.

For those who do not, we think of using their work local authority as a way of avoiding bias induced by residential sorting. The vast majority of employees in London, however, work in a small number of locations (Westminster, the City, Camden, and Tower Hamlets) whilst often commuting from much further afield, including from other regions of the UK. We therefore drop all individuals employed in London to avoid complicating the interpretation of our results, and to ensure that the house prices we assign to individuals are in fact relevant to their housing choices.

Finally, we drop the small proportion of employees (around 4%) working for employers with fewer than five employees. Many of these individuals will be owner-managers, whose private saving decisions are influenced by different factors to other employees, and for whom private pensions may be a less appropriate method for saving for retirement.

Table 1 presents summary statistics for our ASHE sample. The first column pools data from 1997–2007, the second column uses only 1997 data, and the third column uses just 2007 data.

Table 1. Summary statistics for ASHE variables used in analysis, in 1997, 2007 and all years pooled

Description	Sample mean		
	1997–2007	1997	2007
Age	28.6	28.5	28.3
Gross weekly earnings (2018 prices)	470.4	407.6	509.7
Female	48.2%	47.9%	48.6%
Pension member	50.2%	51.8%	43.4%
DB pension member	34.9%	42.5%	25.0%
DC pension member	14.7%	9.4%	16.3%
Public sector	20.9%	21.0%	20.8%
Full time	79.6%	76.6%	78.7%
Occupation 1: managers and senior officials	11.6%	11.1%	11.4%
Occupation 2: professional occupations	11.9%	10.2%	13.8%
Occupation 3: associate professional and technical	15.1%	13.9%	16.9%
Occupation 4: admin and secretarial	15.1%	15.1%	13.1%
Occupation 5: skilled trades	10.0%	10.7%	8.9%
Occupation 6: personal service	7.7%	7.6%	8.2%
Occupation 7: sales and customer service	8.9%	9.1%	9.2%
Occupation 8: process, plant and machine operatives	8.1%	9.7%	6.5%
Occupation 9: elementary occupations	11.6%	12.5%	12.0%
Industry 1: Agriculture	0.6%	0.8%	0.4%
Industry 2: Mining	0.2%	0.2%	0.1%
Industry 3: Manufacturing	17.3%	22.5%	11.5%
Industry 4: Electricity and Gas	0.5%	0.7%	0.4%
Industry 5: Waste	0.5%	0.5%	0.5%
Industry 6: Construction	4.2%	3.1%	5.1%
Industry 7: Retail and Wholesale	16.6%	16.6%	17.6%
Industry 8: Transport+Storage	4.4%	4.6%	4.1%
Industry 9: Accom+FoodServices	4.2%	3.7%	5.4%
Industry 10: Info+Comms	4.4%	3.3%	4.6%
Industry 11: Fin+Insur	5.8%	5.9%	5.5%
Industry 12: Real Estate	0.8%	0.6%	1.0%
Industry 13: Prof,Scien+Tech	5.6%	5.1%	6.2%
Industry 14: Admin+Support	5.6%	4.6%	6.7%
Industry 15: Public Admin+Defence	5.3%	6.1%	4.4%
Industry 16: Education	10.8%	8.4%	12.5%
Industry 17: Health	10.1%	10.5%	10.7%
Industry 18: Arts+Rec	1.6%	1.2%	1.8%
Industry 19: Other Services	1.6%	1.4%	1.7%
Employer size: 1–4	4.0%	4.2%	3.4%
Employer size: 5–49	17.7%	17.8%	16.2%
Employer size: 50–159	9.7%	9.9%	9.4%
Employer size: 160–349	6.6%	6.7%	6.5%
Employer size: 350–5,999	30.6%	32.9%	30.5%
Employer size: 6,000–29,999	19.0%	16.0%	21.5%
Employer size: 30,000+	12.4%	12.5%	12.6%

Note: The sample is 21–35 year olds, employed in England (excluding London).
Source: ASHE, 1997–2007.

Additional data sources

We combine ASHE microdata with information on house prices from the UK House Price Index (UKHPI). The UKHPI is produced by the Office for National Statistics (ONS), and measures monthly average house prices at the local authority level back to 1995. It is estimated from transaction data collected by HM Land Registry, a public register of all housing transactions in the UK, with hedonic adjustments applied by the ONS in order to control for monthly variation in the composition of property transactions.

In our instrumental variables approach, we use data on local housing supply conditions and the yield on 10-year government bonds to construct our instruments for house prices. Our measures of constraints on housing supply, explained in detail at the end of Section 4, come from replication data published alongside Hilber and Vermeulen (2016). Data on the real-terms yield on 10-year government bonds are published by the Bank of England.

Variation in house prices and pension membership

In this paper, we seek to use geographical variation in house prices to identify the effect of house prices on the probability that an employee aged 21–35 is enrolled in a workplace pension. Our chosen unit of geography is the lower tier of English local authorities, of which there were 293 across the country (outside London) during our period of interest. These vary in both geographical size and population, from around 35,000 to around 1.1 million people.

In 1997, the start of our study period, the average house price in England was £86,000 (in 2018 prices, adjusted for consumer price inflation excluding housing costs). Table 2 summarises the variation across local authority areas in 1997, 2002 and 2007 (the start, middle and end of our study period). The first column contains the average figure for England. Subsequent columns show points in the distribution of these variables across local authority areas.⁴ For example, the 10th percentile here means that in 10% of local authority areas the average house price in 1997 was less than £61,000 (in 2018 prices). In 1997, the average house prices in an area at the 90th percentile were more than double the 10th percentile (£128,000 or more, compared with £61,000 or less).

House prices in part reflect differences in local labour market opportunities, but geographical dispersion in average wages in 1997 was relatively modest, and large differences in prices relative to earnings remained. The second panel in Table 2 illustrates the geographical variation in median earnings of employees aged 21–35. In 1997, the median employee aged 21–35 outside London earned £373 per week (in 2018 prices). Median earnings in an area at the 90th percentile were £447, 45% higher than in an area at the 10th percentile (£309). Because variation in earnings was modest relative

⁴ Note that there is a slight discrepancy here for reasons of data availability – the first column contains the mean house price for England as a whole, whereas the subsequent columns exclude London boroughs.

to the variation in house prices (and far from perfectly correlated), average house price to earnings ratios varied almost as much around the country as house prices, as detailed in the third panel of Table 2. In 1997, the median house price to earnings ratio outside London was 4.1. In an area at the 90th percentile, it was 6.3, almost double the 10th percentile (3.3).

Table 2. Geographical variation in house prices, price–earnings ratios and pension membership

	England average	Percentiles of local authority distribution					Std dev.
		10 th	25 th	50 th	75 th	90 th	
<i>Average house price (£)</i>							
1997	86,043*	60,917	68,408	83,285	101,593	128,400	27,437
2002	143,619*	77,239	102,408	141,897	180,390	238,629	61,481
2007	241,773*	165,925	190,039	236,547	277,961	337,326	67,210
<i>Median weekly earnings</i>							
1997	373	309	338	367	408	447	53
2002	444	363	397	432	483	541	74
2007	455	380	409	448	493	536	64
<i>Price–earnings ratio</i>							
1997	4.1	3.3	3.7	4.3	5.4	6.3	1.2
2002	5.6	3.8	4.8	6.2	7.9	9.2	2.1
2007	9.4	7.3	8.4	10.1	12.2	13.5	2.5
<i>Pension membership</i>							
1997	52%	39%	44%	50%	56%	62%	9%
2002	55%	40%	48%	52%	58%	63%	9%
2007	43%	28%	34%	41%	48%	52%	10%

Note: *Average house price is mean hedonically adjusted house price from the UKHPI. Weekly earnings, price–earnings ratio and pension membership are all based on the sample of employees aged 21–35 in the ASHE, employed in England (excluding London) with non-zero gross pay and non-missing data. Weekly earnings are gross of tax. The price–earnings ratio is the median ratio of local mean house price to gross weekly pay for individuals in our sample in that local authority. Pension membership is the proportion of employees enrolled in a workplace pension. Earnings and prices are in 2018 prices using consumer price inflation after housing costs. Due to sampling error, dispersion across local authorities in this table is likely to be larger than the true dispersion.

Source: UKHPI, ASHE.

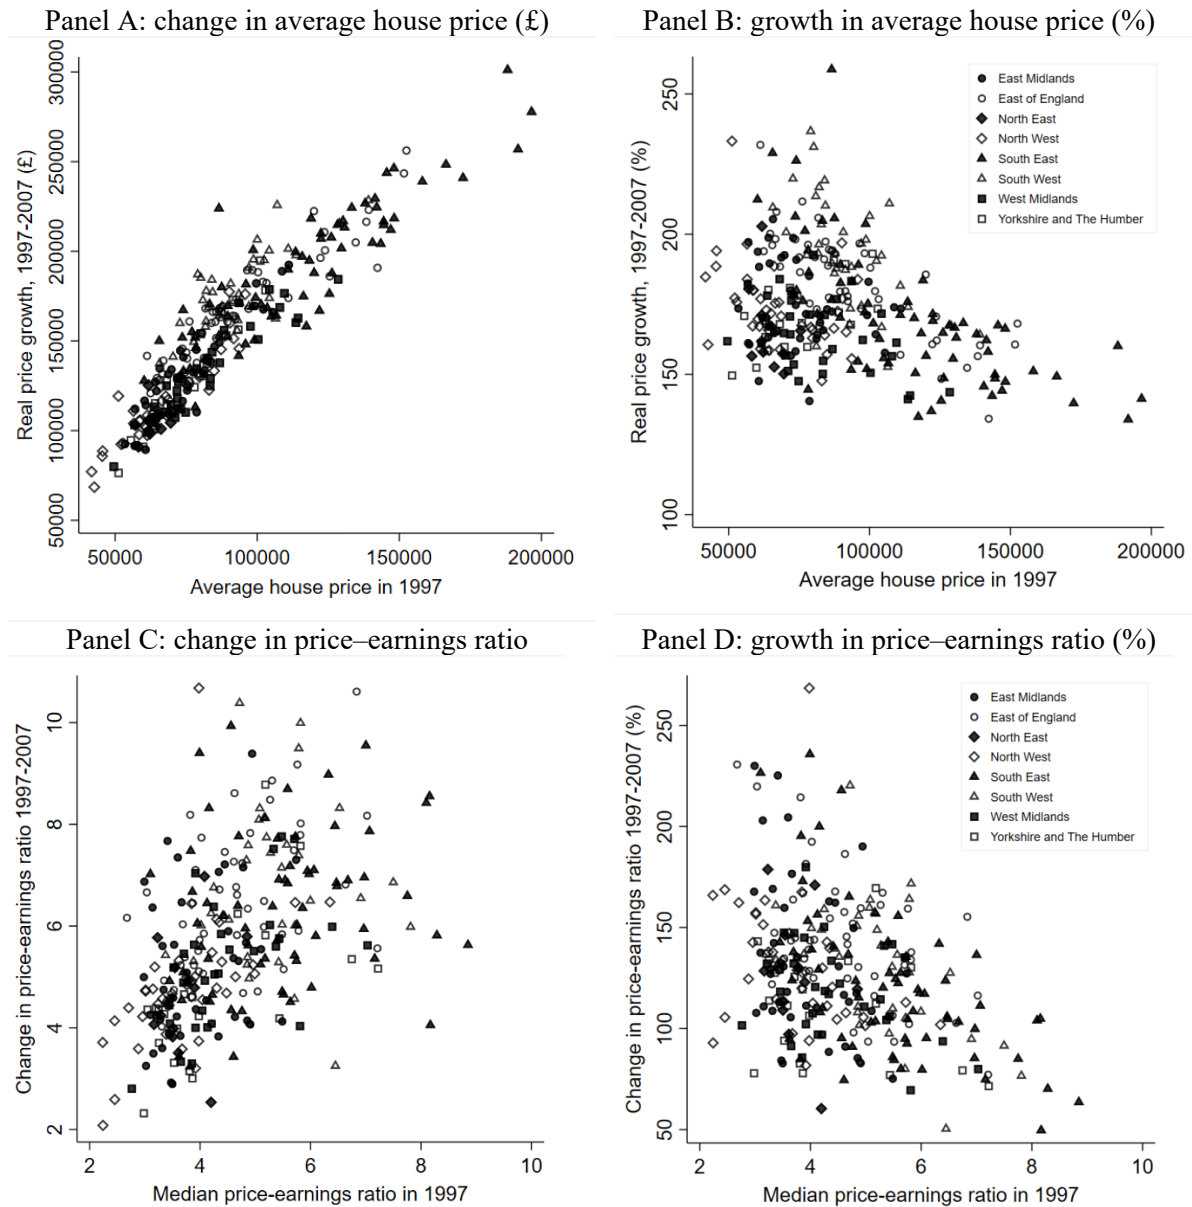
The average house price and price–earnings ratio in England more than doubled between 1997 and 2007, as shown in Table 2. Across England as a whole, the average house price increased in real terms from £86,000 to £242,000 (an increase of £156,000 or 181%), while the price–earnings ratio for employees aged 21–35 outside London increased from 4.1 to 9.4 (+129%).⁵

Interestingly, house price growth differed across the country, even when comparing local authorities in the same region. Figure 4 shows the variation in house price (Panels A and B) and price–earnings ratio (Panels C and D) across local authorities. Areas with higher house prices and price–earnings ratios in 1997 saw larger increases in these values over the subsequent decade, as demonstrated by the positive correlation in Panels B and D. However, in percentage terms, the growth of house prices and price–earnings ratios was, in fact, lower in areas with initially higher levels. Whilst the pattern is

⁵ The price–earnings ratio is calculated as the median ratio of local mean house price to gross weekly earnings.

similar for both house prices and price–earnings ratios, the correlation between the initial price–earnings ratio and subsequent changes is weaker, because earnings and house price growth were not perfectly correlated over this period.

Figure 4. Geographical variation in house price and price–earnings ratio 1997–2007



Note: Average house price is mean hedonically adjusted house price from the UKHPI, in 2018 prices using consumer price inflation after housing costs. Price–earnings ratio is defined as the median ratio of local mean house price to gross weekly pay for individuals in our sample in that local authority, where the sample is employees aged 21–35 in the ASHE, employed in England (excluding London) with non-zero gross pay and non-missing data. Due to sampling error, dispersion across local authorities in this table is likely to be larger than the true dispersion.

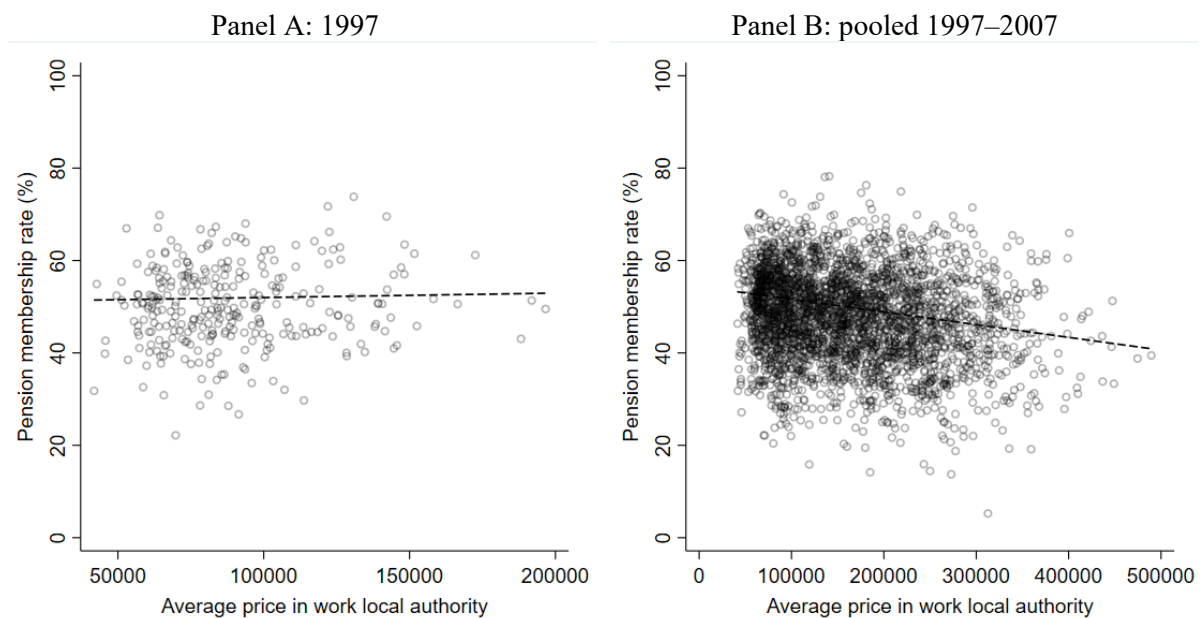
Source: ASHE, 1997–2007 and UKHPI.

The final panel of Table 2 describes the variation in pension membership rates across local authorities. There is geographical variation: in 1997, across England, 52% of employees aged 21–35 were members of a pension, but in one-quarter of local authorities this proportion was 44% or lower,

while in 25% of local authorities this proportion was 56% or greater. By the end of our period in 2007, fewer young adults were members of a pension – the equivalent figures are 43%, 34% and 48%.

Figure 5 plots the simple correlation between pension membership rates and local average house prices in 1997 (left-hand panel) and pooled across all years 1997–2007 (right-hand panel). This suggests a negative association. Fewer young individuals who work in an area with higher local house prices save in a private pension. However, we must be cautious of inferring from unconditional correlations, particularly those exploiting variation across time given the strong time trends in private pension membership. Therefore, we turn now to our empirical approach where we can explicitly control for time trends and the association between individual and job characteristics on pension membership.

Figure 5. Correlation between pension membership rates and local (real) average house prices at the local authority level



Source: ASHE, 1997–2007 and UKHPI.

4. Empirical approach

Ordinary least-squares approach

We estimate the reduced-form effect of average house prices on the probability that an employee aged 21–35 is enrolled in a workplace pension. We do so by exploiting geographical variation in local house prices, and by using individual-level data such that we can control for important individual and job characteristics that are important for pension saving decisions and that may also vary geographically.

Our basic specification is a linear probability model:

$$PM_{it} = \gamma HP_{jt} + \alpha_{st} + X'_{it} \beta + \varepsilon_{it}. \quad [1]$$

Here, PM_{it} is a binary variable taking the value 1 if individual i at time t is enrolled in a workplace pension, and HP_{jt} is the median house price at time t in the local authority j in which individual i works. As controls, we include a vector of individual-level controls, X_{it} , which includes: single year of age interacted with gender (so 29 dummy variables), occupation (8 dummy variables), industry (18 dummy variables), a quadratic in log earnings and employer size (5 dummy variables). We also include sector-specific year fixed effects (α_{st}), to capture flexibly the different trends in pension membership among those working in the public and private sectors, as discussed in Section 2 (21 dummy variables). Standard errors are clustered at the local authority level as this is the level of variation in average house prices in our data. The parameter of interest is γ . Our estimates of this will be reported as the percentage point increase in the probability of being enrolled in a workplace pension associated with an increase of £100,000 (in 2018 prices) in local house prices.

Heterogeneous effects

One limitation of our data is that we do not observe individuals' housing tenure. We are therefore limited in our ability to unpick the mechanisms through which house prices affect pension saving – in particular, the relative contribution of higher house prices increasing down-payment saving, which crowds out pension saving, compared with higher housing costs delaying pension saving until later in life when earnings are greater, or higher housing expenditure reducing non-housing consumption over the life cycle and therefore reducing the need for pension saving.

To unpick our estimated average effect and to provide some indication of the circumstances under which house prices may 'crowd out' pension saving, we do however extend our analysis to examine heterogeneous effects by age, current earnings and sector of employment.

We consider age because, as illustrated in Figure 2, the youngest employees are the least likely to be homeowners already. The 'down-payment saving' channel would therefore be expected to be more evident among younger employees.

We consider current earnings because, intuitively, the likelihood that greater down-payment saving crowds out pension saving is likely to vary across the earnings distribution. We hypothesise that the lowest earners are most likely to have the lowest capacity to save in any form and to be more likely to be lifetime renters, while the highest earners are likely to be able to save in both a pension and for a deposit simultaneously and are more likely to be owners already at any given age (and therefore the down-payment channel would no longer be relevant for them). Middle earners are perhaps most likely to be marginal pension members, whose saving choices may be influenced by house prices.

The final dimension of heterogeneity we consider is sector of employment. The extent to which house prices affect individuals' decisions about private pension membership is likely to be related to the generosity of the private pension an individual is offered. Those who would receive a large employer contribution to their pension would have more to lose by not joining the pension, and so their pension membership would be expected to be less sensitive to house prices. In the ASHE data, we do not observe anything about the pension offer made to those who do not join a pension, we only observe the pension type and contribution levels for those who have joined. However, given the well-documented differences in the generosity of pensions between the public and private sectors (Disney, Emmerson and Tetlow, 2009; Cribb and Emmerson, 2016), we examine heterogeneity by sector of employment as a proxy for the pension offer made.

We implement estimation with heterogeneous effects by modifying equation [1] as follows:

$$PM_{it} = \gamma HP_{jt} + HP_{jt} * \mathbf{Z}'_{it} \boldsymbol{\delta} + \mathbf{Z}'_{it} \boldsymbol{\varphi} + \alpha_{st} + \mathbf{X}'_{it} \boldsymbol{\beta} + \varepsilon_{it}. \quad [2]$$

Here, \mathbf{Z}_{it} represents, in turn, a set of dummies for age in five-year bands (21–25, 26–30, 31–35), a set of dummies for earnings quintile (where the quintiles are calculated conditional on a single year of age and data year) and a dummy for sector of employment.

Identification

Identification in our estimation strategy comes from the assumption that geographical variation in house prices is exogenous as far as individuals' pension saving decisions are concerned. There are two potential concerns with this assumption: residential sorting and geographical variation in local labour market opportunities.

Residential sorting

Individuals can choose where to live, and therefore house prices in an individual's area of residence cannot be considered exogenous. In particular, we may be concerned that individuals sort across areas with different house prices in a way that correlates with unobserved preferences for saving or portfolio composition, biasing our OLS estimates. For example, individuals with a preference for a balanced portfolio may move to an area with low house prices in order to become a homeowner and still contribute to their pension, whereas those who do not wish to become a homeowner may be more content to move to an area with expensive house prices as they can still save in a pension.

We minimise these concerns by 'treating' individuals with house prices in the local authority in which they are employed instead of the local authority in which they are resident. Whilst location of work is still a choice, the possibility and prevalence of commuting gives us confidence that house prices in the location of work are less likely to be correlated with individual preferences for housing and saving. In

2005, when both location of work and location of residence are available in the ASHE data, half of employees aged 21–35 who worked outside of London lived and worked in different local authorities.

Local labour market opportunities

Having defined ‘local’ as local to an individual’s location of employment, a second concern arises: expected earnings growth may differ across locations, even conditional on current earnings, age, gender, sector, occupation and industry (all of which we control for). Unobserved future earnings are likely to affect both house prices and pension savings decisions, so omitting these from our estimation could bias our results.

The idea that locations differ in the future opportunities that they offer was recently formalised in Bilal and Rossi-Hansberg (2020), who argue for the existence of dynamic gains from location on the basis that ruling them out implies implausibly large migration costs given the geographical differences in lifetime income and outcomes observed in the US. Also in the US context, Martellini (2019) shows that the well-established ‘city-size wage premium’ grows with years of experience, as the benefits of knowledge spill-overs become compounded over time. This suggests that young adults working in larger cities may expect more rapid income growth than those working elsewhere.

To address concerns about differences in local labour market opportunities (and, by extension, other unobserved factors that may correlate with both local house prices and individuals’ pension saving decisions), we have two strategies.

First, we test the robustness of our results to the inclusion of local authority fixed effects. This approach eliminates the risk of a time-invariant impact of unobserved characteristics across areas. However, this comes at the cost that we must then rely only on within-area variation in house prices over time to identify our effect of interest. This considerably reduces the precision with which we can estimate the impact of house prices on pension saving.

Second, we use an instrumental variables strategy to isolate exogenous variation in house prices arising from differences in local housing supply conditions.

Instrumental variables approach

We implement our instrumental variables strategy using standard 2SLS estimation. We employ two alternative sets of instruments, in both cases consisting of the interaction between measures of constraints on local housing supply SC_j (from geographical and regulatory sources), and long-run interest rate r_t . Our estimating equations are set out as

$$HP_{jt} = \alpha_{st} + X_{it}'\beta + S_j'\delta + (r_t * SC_j)'\vartheta + v_{it} \quad [3]$$

and

$$PM_{it} = \gamma \widehat{HP}_{jt} + \alpha_t + X'_{it} \beta + SC'_j \lambda + \varepsilon_{it} \quad [4]$$

for the first and second stages, respectively, with other variables defined as in equation [1]. (The level of the interest rate r_t is not explicitly included in equation [3] separately from the interaction with supply constraints, as this will be captured by the single-year dummies α_{st} .)

The intuition for our instruments is as follows. Geographical and regulatory constraints on housing supply affect the cost of building the marginal home, changing the slope of the supply curve. In the short run, a shock that increases demand for housing should induce a larger price response in locations where housing supply is constrained relative to those where it is less so. This creates temporary price differentials, unrelated to underlying differences in the quality of a location or property, which can be used to identify the impact of prices on pension decisions. Over time, the instruments should naturally become less relevant, as migration should reallocate demand geographically and eliminate these price differentials.

Estimated housing supply elasticities at the city level have been used extensively in the literature to instrument for house prices in US settings (e.g. Mian and Sufi, 2011; Chaney, Sraer and Thesmar, 2012; Adelino, Schoar and Severino, 2015; Aladangady, 2017; Chetty et al., 2017). Davidoff (2015) calls into question the exogeneity of these instruments, on the basis that the most commonly used measures of US city-level supply elasticities (Saiz, 2010) are correlated with observable demand factors.

Rather than using the cross-sectional (time-invariant) variation in supply constraints as our instruments, we use an interactive approach following Aladangady (2017), instrumenting for the level of house prices in year t with the interaction between supply constraints in area j and long-run interest rate in year t (r_t), which we treat as a national housing demand shifter. This allows us to control for the direct effect of supply constraints, minimising our exposure to the concerns that Davidoff raises.

For our measures of supply constraints, we follow Hilber and Vermeulen (2016). They propose the following measures of local supply constraints: (i) the average refusal rate on major planning applications between 1978 and 2007, (ii) the share of developable land already developed in 1990, and (iii) the range of elevation. Recognising the possible correlation of these supply constraint measures with productivity or demand conditions, Hilber and Vermeulen (2016) also propose the following instruments for supply constraints: (i) the change in the major project delay rate, (ii) the share of votes for the Labour Party at the 1983 general election, and (iii) population density in 1911. The first set of constraint measures have the benefit of intuitiveness, measuring directly certain constraints that we would expect to influence construction costs, whilst the second are plausibly more

exogenous to demand and productivity trends. We opt to estimate each of our 2SLS specifications twice using these two different sets of supply constraint measures.

5. Results

Headline OLS results

We begin by looking at the correlation between pension membership and house prices for adults aged 21–35 in the years from 1997 to 2007. The first pair of columns in Table 3 shows the result of regressing a binary indicator of pension membership on average house prices in the individual's local authority of work and a constant. Working in an area with higher house prices is associated with a lower propensity to be enrolled in a workplace pension. Specifically, a £100,000 greater average local house price is associated with a 3 ppt lower probability of being enrolled in a workplace pension. This quantifies the correlation demonstrated graphically in Figure 3.

In the second pair of columns, we add controls for an individual's age and gender, gross weekly earnings, hours (part-time or full-time), occupation, industry, and employer size. We control flexibly for time trends using a set of year dummies interacted with sector of employment. The inclusion of these controls together weakens the association between house prices and pension membership, reducing the estimated magnitude of the coefficient to -2 ppt, an association that is not statistically significantly different from zero at the 5% level.

As previously discussed, one potential concern with our OLS specification is endogeneity caused by factors that we do not control for being correlated with both pension membership and house prices. We can control for the impact of such factors, if they are constant over time, through the inclusion of local authority fixed effects. However, this comes at the cost that we must then rely only on within-area time variation in pension membership and average house prices to identify our effect of interest. The results of this specification are shown in the final pair of columns in Table 3. Adding fixed effects for local authority of work does not appreciably change the estimated correlation between local house prices and pension membership. As expected, however, the standard error of the estimate is increased considerably.

Table 3. OLS regression results of pension membership on individual, firm and area characteristics

	(1)		(2)		(3)	
	Coeff.	Std err	Coeff.	Std err	Coeff.	Std err
Local house price (£00,000s)	-0.027***	(0.004)	-0.020	(0.036)	-0.012	(0.084)
Firm size: 50–159			0.085***	(0.005)	0.085***	(0.005)
Firm size: 160–349			0.136***	(0.006)	0.135***	(0.006)
Firm size: 350–5,999			0.204***	(0.005)	0.202***	(0.005)
Firm size: 6,000–29,999			0.225***	(0.006)	0.222***	(0.006)
Firm size: 30,000+			0.293***	(0.008)	0.294***	(0.007)
<i>Industry</i>						
2: Mining			0.167***	(0.031)	0.157***	(0.031)
3: Manufacturing			0.125***	(0.017)	0.119***	(0.017)
4: Electricity and Gas			0.263***	(0.021)	0.257***	(0.021)
5: Waste			0.106***	(0.024)	0.102***	(0.024)
6: Construction			-0.002	(0.018)	-0.006	(0.018)
7: Retail and Wholesale			-0.001	(0.017)	-0.006	(0.016)
8: Transport+Storage			0.039**	(0.019)	0.033*	(0.019)
9: Accom+FoodServices			-0.120***	(0.017)	-0.127***	(0.017)
10: Info+Comms			0.099***	(0.017)	0.090***	(0.017)
11: Fin+Insur			0.239***	(0.017)	0.234***	(0.017)
12: Real Estate			-0.067***	(0.023)	-0.072***	(0.023)
13: Prof,Scien+Tech			0.039**	(0.017)	0.031*	(0.017)
14: Admin+Support			-0.112***	(0.017)	-0.115***	(0.017)
15: Public Admin/Defence			0.114***	(0.018)	0.105***	(0.018)
16: Education			0.097***	(0.017)	0.091***	(0.017)
17: Health			0.054***	(0.017)	0.047***	(0.017)
18: Arts+Rec			-0.043**	(0.019)	-0.049***	(0.019)
19: Other Services			0.046**	(0.019)	0.040**	(0.018)
Log real earnings			-0.069***	(0.010)	-0.070***	(0.010)
sqr(Log real earnings)			0.021***	(0.001)	0.021***	(0.001)
<i>Occupation</i>						
1: manager/senior official			0.047***	(0.005)	0.046***	(0.005)
2: professional			0.021***	(0.005)	0.021***	(0.005)
3: associate professional			0.034***	(0.006)	0.034***	(0.006)
4: admin and secretarial			-0.048***	(0.006)	-0.048***	(0.006)
5: skilled trades			-0.072***	(0.006)	-0.071***	(0.006)
6: personal service			-0.059***	(0.006)	-0.059***	(0.006)
7: sales/customer service			-0.071***	(0.007)	-0.070***	(0.007)
8: process, plant and machine operatives			-0.086***	(0.006)	-0.085***	(0.006)
Full-time employee			-0.016***	(0.004)	-0.016***	(0.004)
Local authority fixed effects			Not included		Included	
Observations	374,584		374,584		374,584	

Note: Regressions (2) and (3) additionally control for gender-specific age profiles (using single year of age dummies) and sector-specific time trends (using year dummies). ***, ** and * denote statistical significance at the 1%, 5% and 10% levels, respectively.

Heterogeneous effects

We now examine whether there are heterogeneous effects. The first dimension of heterogeneity we examine is with respect to age. As discussed in Section 4, our interest here is that if the channel through which higher house prices affect pension saving is through greater saving being required to accumulate a down-payment, then we would expect to find a greater effect at younger ages (where a greater proportion of individuals are not yet homeowners). The results of estimating equation [2] (with Z_{it} being a set of dummies for the five-year age bands) are presented in the first panel of Table 4. The estimated effect of house prices on pension membership is indeed more negative at younger ages, but the differences are not statistically significant and for none of the age groups do we estimate an average association that is significantly different to zero.

Table 4. OLS regression results allowing heterogeneity in estimated effect of local house prices

Marginal effect of local average house price	Without local authority fixed effects		With local authority fixed effects	
	Marginal effect	Std error	Marginal effect	Std error
<i>Heterogeneity by age:</i>				
Aged 21–25	–0.003	(0.004)	0.001	(0.009)
Aged 26–30	0.000	(0.004)	0.003	(0.009)
Aged 31–35	0.002	(0.004)	0.005	(0.009)
<i>Heterogeneity by current earnings:</i>				
Lowest earning quintile	0.009**	(0.005)	0.016*	(0.008)
Quintile 2	–0.011**	(0.004)	–0.005	(0.009)
Quintile 3	–0.016***	(0.004)	–0.011	(0.009)
Quintile 4	–0.006	(0.004)	–0.001	(0.009)
Highest earning quintile	0.010**	(0.004)	0.0015*	(0.009)
<i>Heterogeneity by sector:</i>				
Public sector	–0.030***	(0.006)	–0.026***	(0.010)
Private sector	0.007	(0.004)	0.010	(0.009)

Note: Each panel represents the results of estimating equation [1] with additional interaction terms to examine the heterogeneity of interest. All other covariates are as described in Table 3. The number of observations in all regressions is 374,584, and the number of clusters is 292. *, ** and *** denote statistical significance at the 1%, 5% and 10% levels, respectively.

We then explore whether local house prices have a different association with pension membership depending on individuals' current earnings. We again estimate equation [2], now with Z_{it} being dummies for quintiles of the (age-year specific) earnings distribution; the results are shown in the second panel of Table 4. The estimated effect of local house prices on private pension membership is U-shaped in current earnings. Without local authority fixed effects, we find that the estimated effect is positive for the lowest and highest earnings quintiles, respectively, but negative for the second, third and fourth earnings quintiles. The effect is most negative for those in the middle of the distribution (–1.6 ppt). The statistical significance of our estimates is reduced by the addition of local authority fixed effects, but the magnitude of the estimated coefficients is similar and the U-shaped pattern remains.

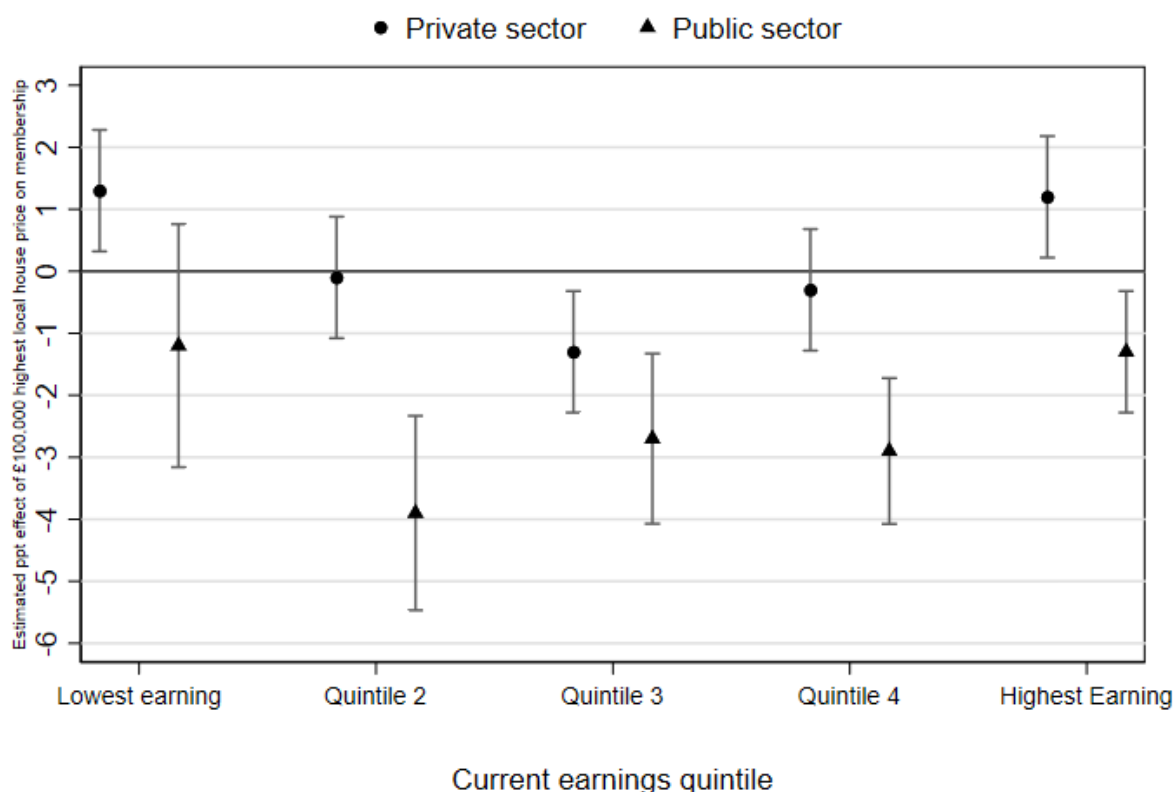
While we cannot pin down the mechanisms for this result with our reduced-form approach, one possible intuitive explanation is that those on the lowest earnings are unlikely to be saving privately for retirement irrespective of local house prices – if earnings are temporarily low, then they would be better placed to save when earnings are higher, while if earnings are permanently low, then they can expect the state pension to compare favourably with their working-life earnings and might prefer to hold any working-life savings in more liquid forms. At the other end of the spectrum, those on the highest earnings are more likely to be able to save for a house and in a private pension simultaneously, and so again the level of local house prices might matter less for the extensive margin decision of whether or not an individual contributes anything to a private pension.

The final dimension of heterogeneity we consider is sector of employment, which serves as a proxy for the type of workplace pension (if any) offered by an individual's employer. The results of estimating equation [2] with an additional interaction between local average house price and sector of employment are presented in the bottom panel of Table 4. We estimate a larger negative association between house prices and pension membership in the public sector than in the private sector: house prices being £100,000 higher is associated with a 3 ppt reduction in the probability that a public-sector employee is enrolled in a workplace pension. In contrast, on average across employees in the private sector there is no significant association between pension enrolment and house prices. This negative association, on average, among public-sector employees is strongly statistically significant, and is robust to the inclusion of local authority fixed effects (shown in the final two columns of Table 4).

Given this dramatic difference in the estimated effect between the public and private sectors, we also examine whether the effect of local house prices on pension membership varies with age or current earnings conditional on sector of employment. To do this, we estimate equation [2] separately for employees working in the public and private sectors, including as Z_{it} dummies for five-year age bands and dummies for quintile of earnings in turn. As is the case for the sample as whole, there are no statistically significant differences in the estimated size of the effect by age for employees working in either sector. The results for earnings quintiles are summarised in Figure 6. In both sectors, there is evidence of the U-shaped pattern, with the negative association between local house prices and pension membership being greater for individuals in the middle of the earnings distribution.

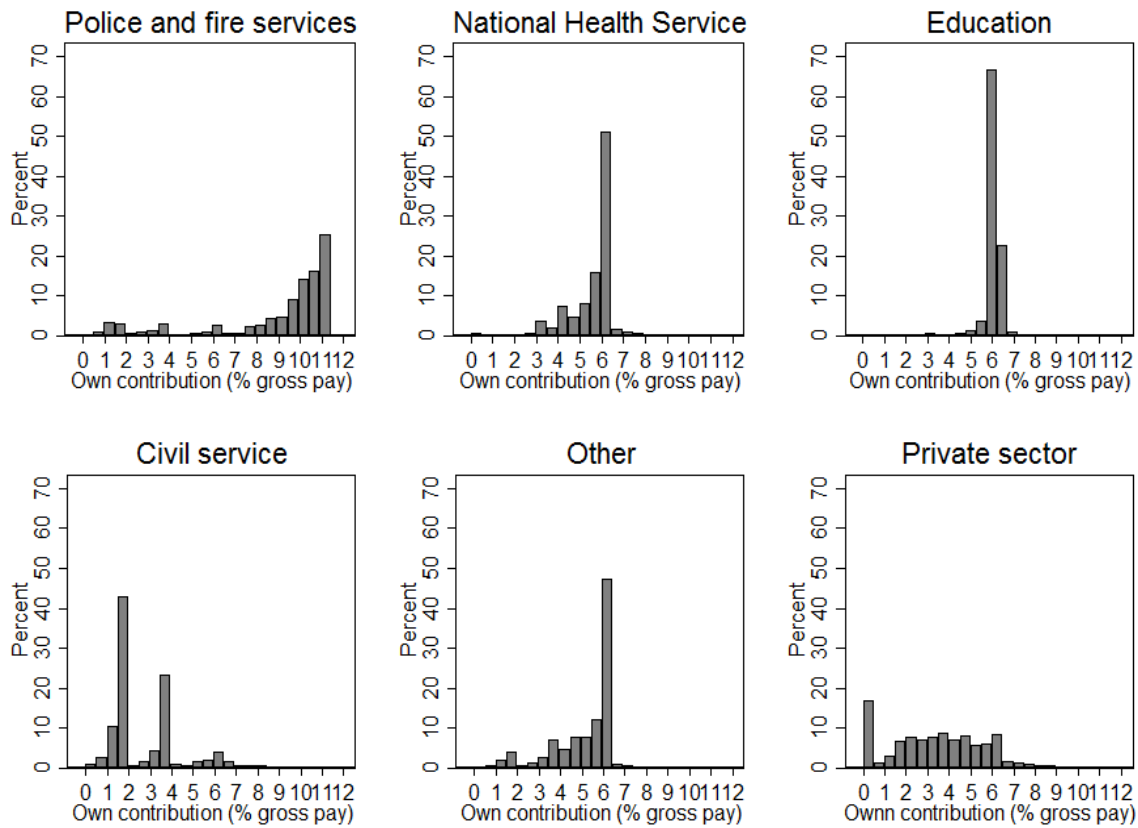
The more negative effect of house prices on pension membership estimated for the public sector, compared with the private sector, is at first sight surprising. Employees in the public sector are forgoing a significant employer contribution by not joining a DB scheme as soon as they are able to. However, there are likely two factors at play. First, individuals in the private sector are less likely to be offered a pension by their employer than those in the public sector. This means that, all else equal, the pension saving choice of private-sector employees may be less sensitive to local house prices because a greater proportion are not going to save in a private pension regardless.

Figure 6. Heterogeneity in estimated marginal effect by quintile of earnings and sector



Second, most public-sector pension schemes require fixed employee contributions. During this period, the local government, education and NHS pension schemes had employee contribution rates of 6% of pensionable pay, while the police and fire services required contributions of 11% and the civil service required 3.5% (Independent Public Service Pensions Commission, 2010). In contrast, employees in the private sector typically made lower contributions than these, either because their employer required lower fixed contributions, or because they had greater flexibility to choose their own contribution level and opted to contribute a lower amount. This is illustrated in Figure 7, which shows the distribution of employee pension contributions (as a percentage of total gross pay) for public- and private-sector employees in our data who are members of a private pension. Public-sector employees' pension membership may therefore be more sensitive to house prices because, on average, public-sector workers would free up more disposable income for down-payment saving by choosing not to join their employer pension than private-sector employees would.

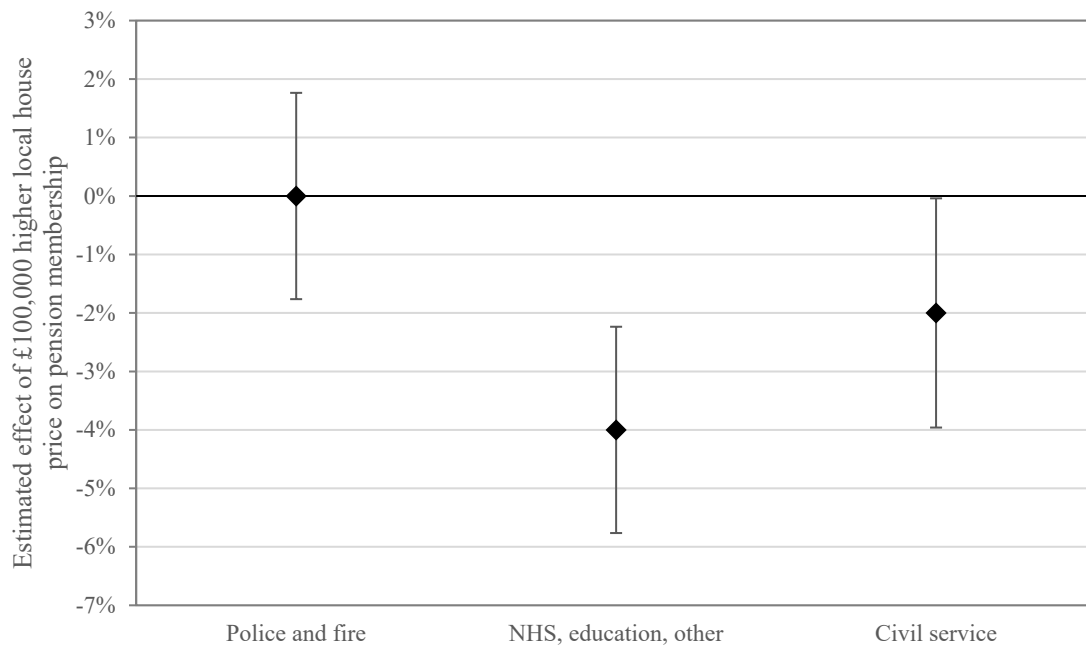
Figure 7. Distribution of employee pension contributions (as a percentage of total gross pay) for public- and private-sector employees



Note: Public-sector workers are grouped according to their standard occupation classification that approximates which public-sector pension scheme an individual is likely to be eligible to join.

To test this explanation, we group public-sector workers into three categories according to their occupation: (i) police and fire service, (ii) civil service administration, and (iii) education, health service and other public-sector workers. Given the distribution of employee pension contributions documented in Figure 7, these can be thought of as areas of the public sector with high, low and medium contribution requirements, respectively. We then estimate one final specification – we estimate equation [2] for public-sector workers, with additional interactions between house prices and these three occupation groups. The estimated marginal effect of house prices on pension membership for these different groups is illustrated in Figure 8. House prices do indeed have a less negative effect on pension membership rates among employees in the civil service, for whom the required contribution rates are lower, than among employees in the health service, education, and other areas of the public sector outside the uniformed services (–2 ppt rather than –4 ppt, a difference that is statistically significant). Perhaps surprisingly, the estimated effect for those in the uniformed services, for whom contribution rates are very high, is not significantly different from zero. This could be due to the greater salience of pension benefits in these occupations, which allowed retirement at much younger ages than in other areas of the public sector.

Figure 8. Heterogeneity in estimated marginal effect by quintile of earnings and occupation within the public sector



Instrumental variables analysis

The results from our OLS analysis are subject to the concern that, while house prices vary geographically, this variation may not be exogenous. In particular, if omitted variables such as local labour market opportunities not only push up house prices but also directly affect pension saving decisions, then our OLS estimates will be biased.

Much of this concern is assuaged by the robustness of our main results to the inclusion of local authority fixed effects. This approach eliminates the risk of time-invariant differences across areas. While the standard errors of our estimates are increased markedly, removing our power to identify estimates as being statistically significantly different from zero, it is reassuring that the size of the estimated effects is generally quantitatively unaffected by the inclusion of fixed effects.

However, there may be residual concern about omitted variables that do not have a fixed impact across time for a given local authority. To address this concern, we here employ our instrumental variables approach to isolate plausible exogenous geographical variation in house prices. As previously described in more detail, our approach exploits the fact that there is variation across the country in how easy it is to build new housing, as a result of geographical and regulatory factors, as well as the density of pre-existing development. Because these measures do not vary over time, we interact them with a demand shifter (in this case, long-run interest rates). The intuition is that in supply-constrained areas, an increase in housing demand resulting from falling interest rates will lead to a larger price increase than in areas where it is cheaper to build new housing.

We use two different sets of supply constraint measures. Set 1 (referred to as ‘HV’) contains ‘direct’ measures of supply constraints – the refusal rate on planning applications, variation in elevation, and pre-existing development in 1990. Hilbert and Vermeulen (2016), who originally used these supply constraint measures to study the impact of regulation on house prices, pointed out that measures such as regulatory strictness or pre-existing development may be correlated with the same local labour market/productivity factors that we are trying to avoid. They proposed instead using instruments for these supply constraint measures – including the Labour Party’s vote share in the 1983 general election, changes in regulatory strictness around the time of a policy reform, and population density in 1911. We use these measures as our second set of supply constraint measures (referred to as ‘HV-I’). In all of the tables that follow, the header ‘HVxR’ means that we use as instruments the first set of supply constraint measures interacted with interest rates, whereas ‘HV-IxR’ indicates that we use the second set interacted with interest rates.

The main first-stage results from our baseline 2SLS specification are presented in Table 5. As described above, the two columns correspond to the two different sets of instruments that we use in our analysis. In the top panel, we report the estimated coefficients on the interaction terms that we use as instruments. In the bottom panel, we report estimates for the direct effect of supply constraints – controls that appear in both the second- and first-stage equations. In both cases, we standardise the supply constraint measures so a one-unit change corresponds to one standard deviation. We define the interest rate measure such that a one-unit change is a 1 ppt change in interest rates. The interpretation of the coefficient on the interaction terms is therefore the impact on house prices (measured in £00,000s) of a one standard deviation difference in supply constraints when interest rates change by 1 ppt.

According to our intuition for the instruments, we expect all the estimated coefficients in the top panel, with the exception of the 1983 Labour Party vote share and change in delay rate, to be negative. When interest rates fall, prices should increase by more in supply constraint areas. Several of the estimated coefficients do not go in the direction we expect. In particular, our results suggest that when interest rates fall, prices go up by *less* in areas with more pre-existing development, higher population density in 1911 or a larger range in elevation – contrary to our expectation. However, the coefficients on the refusal rate, the 1983 Labour Party vote share and the change in the delay rate do go in the anticipated direction. In all cases, the estimated coefficients are statistically significant at the 10% level or more, and the F -test on our instruments far exceeds the rule of thumb value of 10 for a relevant first-stage.

Table 5. First-stage results from 2SLS estimation

	HVxR		HV-IxR	
	Coefficient	Std error	Coefficient	Std error
<i>Instruments</i>				
1990 development $\times r$	0.027**	(0.008)		
Elevation range $\times r$	0.037***	(0.008)	0.022**	(0.007)
Refusal rate $\times r$	-0.110***	(0.007)		
1911 population density $\times r$			0.081*	(0.036)
1983 Labour Party votes $\times r$			0.106***	(0.006)
Change in delay rate $\times r$			0.026***	(0.006)
<i>Controls (direct effects)</i>				
1990 development	-0.154***	(0.046)		
Elevation range	-0.212***	(0.039)	-0.126***	(0.037)
Refusal rate	0.663***	(0.035)		
1911 population density			-0.326	(0.193)
1983 Labour Party vote share			-0.605***	(0.033)
Change in delay rate			-0.160***	(0.037)
Number of observations	353,621		353,621	
Number of clusters	284		284	
<i>F</i> -test on instruments in first stage	150.5		126.56	

Note: First-stage regressions include additional controls as set out in equation [1] and Table 3. Variables are standardised so that a one-unit change corresponds to one standard deviation. Interest rate measure is defined such that a one-unit change is a 1 ppt change in interest rates. Interpretation of the coefficient on the interaction term is the impact on house prices (measured in £00,000s) of a difference of one standard deviation in supply constraints when interest rates change by 1 ppt.

The second-stage results from our 2SLS specification are summarised in Table 6. Despite the apparent strength of our instruments in the first stage, the standard errors in the second stage are large and the estimates somewhat unstable. Quite different estimated coefficients are obtained from our two different measures of supply constraints, though given the size of the standard errors involved, while the differences may be economically significant (and in some cases of opposite sign) they are not statistically different from each other or from zero. Some of the key results from our OLS specification are similarly estimated in the 2SLS estimation. In particular, the larger negative effect for public-sector workers, compared with private-sector workers, and the U-shaped size of the effect across the earnings distribution. However, we lack precision, and the standard errors on these estimates are large.

Table 6. Second-stage results from 2SLS estimation

Marginal effect of local average house price	HVxR		HV-IxR	
	Marginal effect	Std error	Marginal effect	Std error
All	-0.003	(0.014)	0.016	(0.014)
<i>Heterogeneity by age:</i>				
Aged 21–25	-0.004	(0.015)	0.016	(0.015)
Aged 26–30	-0.002	(0.014)	0.015	(0.015)
Aged 31–35	0.004	(0.014)	0.018	(0.014)
<i>Heterogeneity by current earnings:</i>				
Lowest earning quintile	0.070	(0.040)	-0.013	(0.016)
Quintile 2	-0.058*	(0.031)	-0.017	(0.015)
Quintile 3	-0.022	(0.044)	0.000	(0.014)
Quintile 4	-0.029	(0.032)	0.022	(0.014)
Highest earning quintile	0.030	(0.020)	0.015	(0.013)
<i>Heterogeneity by sector:</i>				
Public sector	0.008	(0.014)	0.025	(0.014)
Private sector	-0.031	(0.017)	-0.019	(0.018)

Note: Each panel represents the results of estimating equation [1] by 2SLS with additional interaction terms to examine the heterogeneity of interest. All other covariates are as described in Table 3. *, ** and *** denote statistical significance at the 1%, 5% and 10% levels, respectively.

Extending the time period

Our analysis has focused on the period 1997–2007 because this was a period over which the macroeconomic environment was relatively stable, and house prices in the UK increased rapidly. After 2007, house prices fell sharply as a result of the financial crisis, and they stayed depressed for some years (as shown in Figure 2). This abrupt change in the level and trajectory of house prices might be expected to alter the relationship between house prices and individuals' pension saving decisions, at least temporarily. Furthermore, individuals' saving behaviour may be affected for other reasons. Other asset prices also fell sharply as a result of the financial crisis, which may have affected incentives to save more generally, and labour market prospects were weaker, with low earnings growth and higher unemployment rates for younger individuals.

Given this dramatically changing environment, it is reasonable to restrict the focus of this paper to the 'pre-crisis' years. However, we also separately estimate the relationship between house prices and pension membership after 2007. We examine two separate additional time periods: 2008–12 and 2013–15. The reason for this division is that in October 2012 the government started rolling out its automatic enrolment policy (discussed in Section 2), which led to a substantial increase in workplace pension membership rates among private-sector employees. Even absent the different macroeconomic

context, we would therefore expect a changing relationship between house prices and pension membership as individuals become exposed to automatic enrolment.

The results of this analysis are presented in Table 7. This shows that the small negative relationship between house prices and pension membership estimated for the period 1997–2007 disappears in the subsequent period. This is not surprising given the increased pension membership as a result of automatic enrolment, particularly among groups such as young adults who were less likely to be actively saving in a private pension before the reform (Bourquin, Cribb and Emmerson, 2020).

Table 7. OLS regression results of estimated effect of local house prices in different time periods

Effect of local average house price (£00,000s)	1997–2007	2008–12	2013–15
All	–0.000 (0.004)	0.005 (0.004)	0.004 (0.003)
<i>Heterogeneity by age:</i>			
Aged 21–25	–0.003 (0.004)	–0.000 (0.004)	0.001 (0.004)
Aged 26–30	–0.000 (0.004)	0.006 (0.005)	0.006 (0.004)
Aged 31–35	0.002 (0.004)	0.007 (0.005)	0.004 (0.004)
<i>Heterogeneity by current earnings:</i>			
Lowest earning quintile	0.009** (0.005)	0.004 (0.005)	–0.003 (0.005)
Quintile 2	–0.011** (0.004)	–0.002 (0.005)	0.005 (0.005)
Quintile 3	–0.016*** (0.004)	–0.002 (0.005)	0.007 (0.004)
Quintile 4	–0.006 (0.004)	0.010 (0.006)	0.006 (0.004)
Highest earning quintile	0.010** (0.004)	0.005 (0.007)	0.002 (0.004)
<i>Heterogeneity by sector:</i>			
Public sector	–0.030*** (0.006)	–0.010 (0.006)	0.004 (0.007)
Private sector	0.007 (0.004)	0.008* (0.004)	0.004 (0.003)

Note: Results for column 1 are as presented in Table 4. Columns 2 and 3 report the results of estimating equivalent regressions for the time periods 2008–12 and 2013–15, respectively. Each panel represents the results of estimating equation [1] with additional interaction terms to examine the heterogeneity of interest. Regressions additionally control for gender-specific age profiles, quadratic log earnings, occupation, industry, firm size, hours, and sector-specific time trends (as described in Table 3). The number of observations for regressions for 1997–2007 is 374,584, for 2008–12 is 175,008 and for 2013–15 is 116,886. *, ** and *** denote statistical significance at the 1%, 5% and 10% levels, respectively.

6. Discussion and conclusion

In this paper, we set out to estimate the impact of house prices on pension membership in early adulthood. In doing so, we sought to complement existing research on homeownership decisions, providing a broader perspective on the impact of house prices on savings decisions across an age range in which homeownership and pension membership historically both increase sharply.

Our empirical approach exploits time-varying geographical differences in house prices and controls for both employer characteristics (which influence the offer of workplace pension schemes) and individual characteristics, such as age and earnings. On average, we find that the effect of house prices is negative, but small and statistically insignificant. However, allowing the effect to vary by

current earnings reveals a ‘U-shape’, with a small but statistically significant negative effect concentrated amongst middle earners. Examining heterogeneity by sector of employment shows stronger negative effects for public-sector workers, among whom an increase of £100,000 in house prices is, on average, associated with a 3 ppt lower probability of contributing to a pension. Within the public sector, those working in the civil service are estimated to be less responsive than those working in the NHS, education or other non-uniformed services, which is consistent with the fact they are required to make smaller employee contributions to their pension scheme.

These results suggest that high house prices have a small crowd-out effect on the extensive margin of pension saving for middle earning employees. We do not seek to unpick the mechanisms for this result in this work – in particular, whether the higher house prices increase down-payment saving that crowds out pension saving, whether higher ongoing housing costs cause pension saving to be delayed until later in life when earnings are greater, or whether higher housing expenditure reduces non-housing consumption over the lifecycle and therefore reduces the need for pension saving. However, further research to unpick the relative importance of these different effects would be beneficial. The effect of house prices on the intensive margin of pension saving – how much an individual contributes to a pension, rather than just whether or not they contribute at all – would be another important avenue for future research.

References

- Adelino, M., Schoar, A. and Severino, F. (2015), ‘House prices, collateral, and self-employment’, *Journal of Financial Economics*, 117, 288–306.
- Aladangady, A. (2017), ‘Housing wealth and consumption: evidence from geographically-linked microdata’, *American Economic Review*, 107 (11), 3415–3446.
- Bilal, A. and Rossi-Hansberg, E. (2020), ‘Location as an asset’, Working paper, Princeton University.
- Bottazzi, R., Crossley, T. F. and Wakefield, M. (2015), ‘First-time house buying and catch-up: a cohort study’, *Economica*, 82, 1021–1047.
- Bourquin, P., Cribb, J. and Emmerson, C. (2020), ‘Who leaves their pension after being automatically enrolled?’, Institute for Fiscal Studies (IFS) Briefing Note BN272, <https://www.ifs.org.uk/publications/14742>.
- Campbell, J. Y. and Cocco, J. F. (2007), ‘How do house prices affect consumption? Evidence from micro data’, *Journal of Monetary Economics*, 54, 591–621.
- Chaney, T., Sraer, D. and Thesmar, D. (2012), ‘The collateral channel: how real estate shocks affect corporate investment’, *American Economic Review*, 102 (6), 2381–2409.

- Chetty, R., Sándor, L. and Szeidl, A. (2017), 'The effect of housing on portfolio choice', *Journal of Finance*, 72, 1171–1212.
- Cocco, J. F. (2005), 'Portfolio choice in the presence of housing', *Review of Financial Studies*, 18, 535–567.
- Cribb, J. (2019), 'Intergenerational differences in income and wealth: evidence from Britain', *Fiscal Studies*, 40, 275–299, <https://doi.org/10.1111/1475-5890.12202>.
- Cribb, J. and Emmerson, C. (2016), 'Workplace pensions and remuneration in the public and private sectors in the UK', *National Institute Economic Review*, 237, R30–R37.
- Cribb, J. and Emmerson, C. (2020), 'What happens to workplace pension saving when employers are obliged to enrol employees automatically?', *International Tax and Public Finance*, 27, 664–693.
- Cribb, J., Hood, A. and Hoyle, J. (2018), 'The decline of homeownership among young adults', Institute for Fiscal Studies (IFS) Briefing Note BN224, <https://www.ifs.org.uk/publications/10505>.
- Cribb, J. and Simpson, P. (2018), 'Barriers to homeownership among young adults', in C. Emmerson, C. Farquharson and P. Johnson (eds), *IFS Green Budget 2018*, Chapter 9, London: Institute for Fiscal Studies, <https://www.ifs.org.uk/publications/13475>.
- Crossley, T., Levell, P. and Low, H. (2020), 'House price rises and borrowing to invest', Institute for Fiscal Studies (IFS) Working Paper W20/02, <https://www.ifs.org.uk/publications/14684>.
- Davidoff, T. (2015), 'Supply constraints are not valid instrumental variables for home prices because they are correlated with many demand factors', available at SSRN, <http://dx.doi.org/10.2139/ssrn.2400833>.
- Disney, R., Emmerson, C. and Tetlow, G. (2009), 'What is a public sector pension worth?', *Economic Journal*, 119 (541), F517–F535.
- Hilber, C. A. and Vermeulen, W. (2016), 'The impact of supply constraints on house prices in England', *Economic Journal*, 126 (591), 358–405.
- Independent Public Service Pensions Commission (2010), *Interim Report*, https://www.ucu.org.uk/media/4328/Independent-Public-Service-Pensions-Commission---interim-report-7-Oct-10/pdf/hutton_pensionsinterim_071010.pdf.
- Martellini, P. (2019), 'Essays on labour markets and cities', University of Pennsylvania, ProQuest Dissertations Publishing.
- Mian, A. and Sufi, A. (2011), 'House prices, home equity-based borrowing, and the US household leverage crisis', *American Economic Review*, 101 (5), 2132–2156.

Office for National Statistics (2020), Annual Survey of Hours and Earnings, 1997–2020: Secure Access. [data collection]. 17th Edition. UK Data Service. SN: 6689, <http://doi.org/10.5255/UKDA-SN-6689-16>.

Saiz, A. (2010), ‘The geographic determinants of housing supply’, *Quarterly Journal of Economics*, 125, 1253–1296.

Sheiner, L. (1995), ‘Housing prices and the savings of renters’, *Journal of Urban Economics*, 38, 94–125.