The impact of housing subsidy cuts on the labour market outcomes of claimants

Borbely, Daniel

Published in:
Journal of Housing Economics

DOI:
10.1016/j.jhe.2022.101859

Publication date:
2022

Journal of Housing Economics 57 (2022) 101859

Contents lists available at ScienceDirect

Journal of Housing Economics

journal homepage: www.elsevier.com/locate/jhec

The impact of housing subsidy cuts on the labour market outcomes of claimants: Evidence from England

Daniel Borbely[1]

Economics Group, School of Business, University of Dundee, 1-3 Perth Rd, DD1 4JW, Dundee, United Kingdom

ARTICLE INFO

JEL classification:
H31
H42
H53

Keywords:
Housing subsidies
Welfare programmes
Labour market behaviour

ABSTRACT

Housing subsidies are aimed at helping low-income individuals afford appropriate housing, but are costly to offer and, in the view of some experts and policy makers, reduce incentives for claimants to participate in the labour market. This paper investigates the labour market impacts of recent housing subsidy cuts in England that were aimed at encouraging labour market participation and increasing work effort among claimants. I utilise variation in exposure to the subsidy cuts within a difference-in-differences framework and find limited evidence at the individual-level that claimants increased employment and labour force participation in response to the subsidy cuts. Nonetheless, these findings lack robustness and aggregate-level evidence suggests that the subsidy cuts did not succeed in encouraging employment or participation among claimants. Overall, my results show that labour market responses to the subsidy cuts were likely preempted by a strong mobility response, whereby claimants moved into other parts of the rental market to maintain subsidy coverage.

1. Introduction

Housing subsidies assist low-income households with their rent obligations to help them afford appropriate housing. Critics note that these subsidies might depress work incentives among claimants leading to a reduction in labour supply. Following this line of reasoning, cuts to housing – and other in-kind subsidies – are often justified by governments as measures to encourage work or search effort by claimants (Taylor-Gooby, 2012). In England, recent housing subsidy cuts were justified in a similar manner, the main government objective being the creation of a 'fair' system where claimants are encouraged to work in order to afford quality housing.[2]

Despite the common justification, economic theory yields largely ambiguous predictions on the sign (and size) of labour supply impacts from housing subsidies (Murray, 1980; Schone, 1992; Moffitt, 2002; Shroder, 2002). The standard neoclassical model of labour supply predicts that housing subsidy provision will lead to a reduction in labour supply through both substitution and income effects (Moffitt, 2002). Conversely, alternative models predict that labour supply effects may vary depending on whether housing is a substitute or a complement to leisure (Murray, 1980); whether housing is a complement to other consumption goods (Schone, 1992); or could even be positive if subsidies lead to reduced housing uncertainty and allow claimants to spend more time seeking employment (Collinson et al., 2015).

While theoretical predictions of housing subsidy effects on labour supply are ambiguous, empirical evidence mostly confirms the stipulations of the neoclassical model. Recent quasi-experimental studies from the U.S. find conclusive evidence that housing subsidy provision has a negative effect on labour supply, although this effect diminishes with time and is mostly rather small (Mills et al., 2006; Jacob and Ludwig, 2012; Carlson et al., 2012). Observational studies report similar (though larger) negative effects (Olsen et al., 2005; Susin, 2005). In contrast, empirical evidence from outside the U.S. is inconclusive and rather sparse (Shroder, 2010). Furthermore, the case studies in the literature tend to focus on the effects of housing subsidy provision, and very few studies focus on the effects of a reduction or withdrawal of subsidy entitlements from existing claimants. More specifically, when housing subsidies are withdrawn or reduced, does this induce claimants to increase labour supply either along the extensive (having a job) or intensive (hours of work) margin? In this paper, I aim to answer this question through an analysis of recent housing subsidy cuts in England.

In England, through the Housing Benefit (HB) system, housing subsidies constitute a significant share of welfare expenditures.[3] In the first decade of the 21*st* century, expenditures on housing subsidies

E-mail address: dborbely001@dundee.ac.uk.

[1] I would like to thank Otto Lenhart, Jonathan Norris, Julia Darby, Graeme Roy, Gennaro Rossi and two anonymous referees for helpful comments and suggestions.

[2] See House of Commons, Work and Pensions Committee (2010) for an overview of the policy consultation.

[3] In what follows, I will use the terms 'Housing Benefit' and 'housing subsidy' interchangeably.

https://doi.org/10.1016/j.jhe.2022.101859
Received 12 July 2021; Received in revised form 6 May 2022; Accepted 12 August 2022
Available online 24 August 2022

increased by 46%, with annual expenditures in 2010/11 totalling at £21.4 billion (Wilson et al., 2016). Along with a range of austerity measures aimed at reducing the public deficit, the 2010 Coalition Government introduced several changes to the Housing Benefit system.[4] Specifically, the June 2010 Budget announced changes to the way Local Housing Allowance (LHA) rates, which determine Housing Benefit entitlements for claimants in the private rented sector (PRS), are set.[5] These changes, rolled out in 2011–12, resulted in substantial cuts to the Housing Benefit entitlements of private rented sector claimants.

The government's justification for the housing subsidy cuts relied on two underlying objectives. First, they intended to curb expenditures from the Housing Benefit system. Second, they wanted to ensure that households on subsidies do not occupy more expensive housing than low-income working families. This second objective was an attempt to induce claimants to seek employment, and to create the means for housing through labour and not benefit income (see Wilson et al., 2016).

The Housing Benefit cuts came under significant scrutiny upon their implementation. Critics of the new system commented on the hardships the subsidy cuts created for claimants by reducing rent affordability leading to increased risk of eviction by landlords. A briefing report released by the Shelter Foundation claimed that while the Housing Benefit cuts led to a deterioration in housing conditions for private rented sector claimants, they did not lead to a noticeable increase in employment (Shelter, 2015). There is also evidence from the academic literature that the subsidy cuts led to a substantial reduction in rent affordability (Brewer et al., 2014, Fetzer et al., 2019). On the other hand, evidence on the reforms' impacts on claimant mobility is mixed (Brewer et al., 2014; Braakmann and McDonald, 2018) and no analysis to date has focused explicitly on their effect on labour market outcomes.

My analysis draws on individual-level panel data on Housing Benefit claimants in England from the Understanding Society (US) longitudinal survey for the period 2009–2017. The empirical strategy makes use of a panel difference-in-differences approach that examines the effects of the subsidy cuts through a comparison, over time, of groups affected (treated) and unaffected (comparison) by the policy. Since subsidy cuts are not assigned randomly to claimants, my empirical approach relies entirely on within-claimant variation in outcomes over time. To control for diverging trends arising from structural differences between treated and untreated individuals, I also make use of the synthetic difference-in-differences approach (Arkhangelsky et al., 2021) along with propensity score matching (PSM) to optimise comparison group selection and ensure the robustness of my results to the choice of comparison group. Moreover, I use data collected at the local authority district level to investigate the aggregate level labour market effects of the Housing Benefit cuts.

My results show that the Housing Benefit cuts had a strong positive effect on claimant mobility, as claimants moved away from the private rented sector potentially in order to avoid (or mitigate) the effects of the subsidy cuts. On the other hand, I find mixed evidence on labour market responses to the reforms. My analysis of survey data suggests a positive effect on employment and participation rate, but these results are sensitive to changing the specification of the comparison group, and are not supported by evidence from district-level data on the same outcomes. Overall, my findings provide no aggregate level evidence that the government's objective to get claimants back to work (or even get them to seek work) was accomplished, or at least that positive labour market responses to the subsidy cuts were not large enough to materialise in an aggregate level effect.

The remainder of this paper is organised as follows. Section 2 outlines the recent changes to the UK housing subsidy system. Section 3 describes the data. Section 4 outlines the empirical strategy and presents my results. Section 5 concludes.

2. Policy background

In England, all individuals can apply for Housing Benefit (HB),[6] a means-tested housing subsidy that provides assistance with the rental costs of housing, as long as claimants: (1) live in rental accommodation (2) are on a low-income or are claiming benefits and (3) possess savings lower than £16,000. The amount of Housing Benefit received by each individual/household is determined as follows:

$$HB = \min\left\{rent,\ HB_{max}\right\} \quad if\ Y \leq Y^T$$

or

$$HB = \min\left\{rent,\ HB_{max}\right\} - 0.65\left(Y - Y^T\right) \quad if\ Y > Y^T$$

where HB_{max} is the maximum eligible housing subsidy amount, Y is household income, and Y^T is the threshold income for Housing Benefit eligibility. Above the threshold income, Housing Benefits are withdrawn at the taper rate of 65%, i.e. the subsidy amount is reduced by 65 pence for every £1 increase in income. Note, that increased subsidies will lead to higher marginal tax rates by increasing the opportunity cost of labour (and increasing earnings in the absence of labour). In the UK case, due to the high taper rate (65%) the income increases of claimants from increased work effort are 'punished' by a particularly high marginal tax rate. Increased housing subsidies therefore make labour relatively 'costly' in comparison to leisure, inducing a substitution effect towards the latter. Conversely, a reduction in housing subsidies should have the opposite effect: a reduction in the opportunity cost of labour and increased incentives to provide work effort.

Housing Benefit can be claimed by households renting their accommodation from either the private or social rented sectors. In the private rented sector (PRS), low-income households can access Housing Benefit through the Local Housing Allowance (henceforth LHA) system. The LHA system was introduced on April 7, 2008 and provides a set of rules that determine the amount of Housing Benefit low-income households are eligible to claim (see Wilson et al., 2016). The extent to which LHA will cover rental costs is determined by the private market rent distribution in the specific geographical area (referred to as Broad Rental Market Area); the size of the household; and the earnings and income from other benefits claimed by household members. Upon their introduction, LHA rates were set to cover housing costs for properties with values below the local median house price (the cheapest half of local properties). Rates were then adjusted (monthly) to reflect inflation in rents.

Starting in 2011, the UK Government introduced (as part of the Welfare Reform Act 2012) several Housing Benefit cuts and eligibility rule changes to the LHA system, applicable to claimants renting from the private rented sector:

- LHA rates were capped so that they only cover the bottom $30th$ percentile of local rental properties instead of the bottom $50th$ percentile. This means that after the reforms, Housing Benefits can only cover at most an amount equal to the $30th$ percentile of local rents, as opposed to an amount equal to the 50th percentile;

[4] The reforms were rolled out UK wide, however the changes discussed in this paper concerned the English subsidy system to the greatest extent, as devolved administrations in other parts of the UK had some discretion over related housing policies and in some cases decided to mitigate the impacts of housing subsidy reform. For this reason, my analysis of the recent reforms to the housing subsidy system only concerns England.

[5] Claimants in the private rented sector are renting accommodation from private landlords.

[6] This system is currently (in 2021) being replaced by Universal Credit, a means-tested benefit system that pulls together different types of benefits, including Housing Benefit. Claimants continue to receive Housing Benefits under the Universal Credit system.

- the Shared Accommodation Rate (a lower rate for claimants not living in shared accommodation) was extended to cover a wider age group;
- the £15 per week excess, the amount claimants could keep when their rent was below the LHA rate, was removed;
- LHA rates no longer adjust to the inflation of rental costs — they are currently uprated in line with CPI inflation.[7]

The introduction of the reform package was staggered: it was announced through the June 2010 budget and rolled out starting from April 2011, however, due to transitional protection periods some claimants were not rolled in until late 2012 (see Brewer et al., 2014). Transitional protection also meant that the earliest enrolment date for existing claimants was January 2012, and most claimants were rolled in during the first half of 2012.

The changes to the LHA system only applied to the private rented sector, and not the social rented sector.[8] In the social rented sector (SRS), accommodation is rented from local authorities or housing associations, rents are kept at affordable levels, and long-term rental security is ensured for tenants (Wilson and Barton, 2021). Regardless, social rented sector claimants can still receive Housing Benefit to help with the already lower social sector rents, but their entitlements are not determined by the LHA rules that apply to private sector claimants.[9] How much Housing Benefit households are eligible for in the social sector is instead determined by factors such as the overall rent payable (plus service charges), the number of rooms in the rented property, household income and household circumstances (such as disabilities or age). For the social sector, there were recent plans to introduce the same LHA rules as in the private sector, but this policy was first deferred and then scrapped altogether by the UK Government (see Wilson et al., 2016). Instead, a tax on 'spare' bedrooms was introduced to encourage downsizing among claimants (Gibbons et al., 2018). This policy, often referred to as the 'bedroom tax', was introduced in April 2013, and constituted a small monetary 'punishment' for households occupying properties with more rooms than they are entitled to based on policy rules. The 'bedroom tax' therefore targeted a different policy base in comparison to the LHA reforms: it only led to a subsidy cut for those in the social rented sector not adhering to the specific subsidy eligibility rules, whereas the LHA reforms applied to all Housing Benefit claimants in the private rented sector.

The objectives of the government with the Housing Benefit cuts were related to two concerning aspects of the housing subsidy system: its cost, and its effects on the housing and labour markets (Tunstall et al., 2015). First, a concern was that spending on Housing Benefit constituted a large share of total welfare expenditures in the pre-reform period and yet did not contribute to an investment in the housing stock. This was because the subsidies were paid to private rented sector landlords for existing property. Second, as Housing Benefit helps pay rental costs for those at low wages, they effectively subsidise low wage (labour) income.

Considering these concerns, on one hand, the government's objective was to reduce expenditures from the Housing Benefit system and make the system simpler. On the other hand, they wanted to encourage labour market participation amongst claimants by providing more incentive to withdraw from benefits and seek work. As the Department for Work and Pensions (DWP) argued:

> Providing some customers, mainly in London, with the ability to live in very high cost rented properties makes it extremely unlikely they would ever move completely off Housing Benefit because of the very high income levels required. Moving to more affordable accommodation could therefore encourage households to take up employment and move completely off benefits.[10]

During the consultation period, the main government justification of the subsidy cuts was based on the assertion that Housing Benefit claimants occupy more expensive housing than working individuals not in receipt of benefits. The Minister of Pensions put it this way:

> Low-income households rent at about 90% of what the Housing Benefit recipients are renting at. So they are renting at a lower level. [...] The facts are that low-income people who are not taking Housing Benefit are having to live in cheaper housing.[11]

The government argued that housing subsidy cuts were necessary to eliminate an 'unlevel' playing field, so that subsidy recipients would not enjoy higher quality housing than low-income working families, and would be encouraged to seek work. Some experts found this justification unsatisfactory. For example, a study by the Cambridge Centre for Housing and Planning Research (Fenton, 2010) claimed that the Housing Benefit cuts were likely to have no impact on the labour market activity of claimants as most subsidy recipients who were able to work were already in employment before the reforms. The author also claimed that the reason unemployment levels were high in some areas was due to structural weaknesses in regional economies, and not due to individuals' lack of willingness to work. Qualitative assessment of the recent reform by Shelter (2015) also found that it had no noticeable impact on the labour market activity of claimants, but led to worse housing conditions and even homelessness in some cases. Later sections of this paper will aim to identify these labour and housing market impacts empirically.

3. Data

To assess the labour market impacts of recent private rented sector housing subsidy cuts in England, I use data on Housing Benefit claimants from the Understanding Society (US) survey, covering the time period 2009 to 2017. The US survey, also known as the UK Household Longitudinal Study (UKHLS), is a representative longitudinal survey of 40,000 UK households.[12]

The data can be described as an unbalanced panel where individuals are observed in waves (these need not overlap with years). I use data from the first nine waves of the US survey. Some individuals are not measured in consecutive waves. I track individuals using the cross-wave person identifier ('*pidp*'). Adult individuals that share the same household all receive claimant status even if only one of them is indicated as a housing subsidy claimant — whilst from an administrative point of view only one person claims the subsidies, those accrue to the entire household. Under these circumstances, household outcomes are considered, with two individuals forming a benefit unit. When a single person occupies a household, she is the only benefit unit. I exclude observations living outside of England before the Housing Benefit changes. I focus on England to make sure that heterogeneity in policy responses by devolved governments in Scotland and Wales do not bias my results. I also exclude those claimants eligible for the bedroom tax (see Section 2) – these individuals make up 3.2% of my baseline sample – so that responses to this policy change do not

[7] This is important because according to analysis by Shelter (2015) rents have been rising more sharply than LHA rates, particularly in London.

[8] In England in 2019/20, roughly 19% of households lived in the private rented sector, while around 17% of households lived in the social rented sector (Cromarty, 2021; Wilson and Barton, 2021).

[9] See UKGovernment---HousingBenefits for an overview of Housing Benefit rules for the social rented sector.

[10] See House of Commons, Work and Pensions Committee (2010).

[11] See House of Commons, Work and Pensions Committee (2010).

[12] In Petersen et al. (2013), the respondents surveyed in the General Population Survey (which includes most of the households surveyed in US) were found to be representative of the census population at the neighbourhood level.

Table 1
Summary statistics — full sample.

	Mean	SD	Count
Treatment			
HB recipient in PRS before the reform	0.03	0.17	229107
Housing Benefits			
On Housing Benefit	0.10	0.30	229107
Covariates			
Female	0.54	0.50	229107
Lives in urban area	0.80	0.40	229107
Married	0.55	0.50	229107
Age	46.34	17.40	229107
Household size	2.92	1.48	229107
Number of siblings	0.11	0.50	229107
Number of children aged 0–15	0.62	0.99	41700
Incapacity benefit	0.02	0.15	228700
Employment and support allowance	0.00	0.07	228700
Severe disablement allowance	0.00	0.06	228700
Carers allowance	0.01	0.10	228700
Disability living allowance	0.04	0.19	228700
Attendance allowance	0.01	0.09	228700
Retired	0.20	0.40	229107
Disabled	0.03	0.17	229107
Student	0.06	0.25	229107
Outcomes			
Behind with rent	0.07	0.26	229107
Number of beds in HH	1.25	0.71	229071
Would like to move from property	0.33	0.47	214538
Self-employed	0.08	0.28	229107
Employed	0.48	0.50	229107
Active	0.61	0.49	229107
Log hours of work	3.47	0.56	125249
Mental health (Caseness)	1.75	2.97	196889
Changed tenure	0.27	0.45	229107
Moved to SRS	0.01	0.11	229107
No longer on HB	0.03	0.18	229107
HB Coverage of rent	0.23	0.39	53565

influence my results. Summary statistics on key variables are provided in Table 1.[13] Moreover, Table A.1 provides summary statistics from only the pre-reform period.

4. Empirical evidence

4.1. Individual-level evidence

To assess the effects of Housing Benefit cuts on private rented sector claimants, I estimate a simple pooled difference-in-differences model, which takes the following form:

$$y_{i,q,w,r} = \alpha_i + \beta \times Ant_{i,q} \times PRS_i + \gamma \times Post_{i,q} \times PRS_i + \theta_{q,w,t} + \epsilon_{i,q,w,r} \quad (1)$$

where our main coefficient of interest is 'γ', which captures the change in outcome variables (see below) for each individual 'i', in quarter 'q', in survey wave 'w' and district 'r', after the housing subsidy cuts became effective (indicated by $Post_{i,q}$) for private rented sector claimants (PRS_i). I include individual fixed effects 'α_i', meaning that the estimation of the effects of Housing Benefit changes will rely entirely on within-individual variation. The term $Ant_{i,q} \times PRS_i$ is meant to capture potential anticipation effects during the transition period, when the new subsidy system was already in place, but claimants were not yet enrolled due to transitional protection periods (see Section 2). Following the approach of Fetzer (2019), who used the US survey to analyse the effects of austerity measures on voting preferences, I include time and survey wave fixed effects specific to each 331 local authority district in the sample ($\theta_{q,w,t}$).[14] These fixed effects should capture any district-specific shocks that vary over time. In some specifications, I also add quarter-wave fixed effects specific to each qualification type, to control for labour market shocks specific to different skill levels, and quarter-wave fixed effects specific to claimants of other types of benefits (incapacity benefits, employment benefits, etc.), to control for shocks related to other sources of benefit income.[15]

Treatment group. The treatment group – the group of individuals affected by the Housing Benefit cuts – consists of claimants of Housing Benefit who resided in the private rented sector when last observed in the US survey prior to the reforms. I am fixing treatment group membership to the pre-reform time period so that responses to the reforms (for example, moving away from the PRS) do not induce selection bias. The cut-off time for being included in the treatment group is December 2011 — if a survey respondent is claiming Housing Benefit while residing in the private rented sector when we last observe them pre-reform in, say, September 2011, they are in the treatment group.[16] In the baseline specification, the comparison group – the group of individuals who were unaffected by the reforms – are the individuals who make up the rest of our analytical sample. This group consists of individuals who did not claim Housing Benefit while residing in the private rented sector when last observed prior to the reforms; individuals (both claimants and non-claimants of Housing Benefit) residing in the social rented sector; and individuals who own their own housing units or rent accommodation from family members.

Overall, the treatment group contains 6,450 individuals, who make up 3% of the full sample. Looking at the local authority district level, some districts have a higher share of private rented sector claimants compared to the rest of the sample, with the Greater London Area having a larger concentration of these claimants both before and after the reforms (see Figure A.1). These regional differences, and related changes in outcomes, are controlled for in the individual-level specifications through the use of time fixed effects specific to each local authority district.

Importantly, private rented sector claimants of Housing Benefits might be different from the rest of the sample in terms of their observable characteristics. This, in itself, is not an issue for our difference-in-differences strategy, as level differences can be controlled for through individual fixed effects or control variables. Nonetheless, if structural differences correlate with both treatment assignment and trends in outcomes, this could bias our results. To assess whether structural differences are present across treated and comparison groups, I regress treatment group membership on key observable demographic variables using data from the pre-reform period only. The results of these 'balance tests' are summarised in Fig. 1. The coefficient plots in Fig. 1 show that the survey respondents who belong to the treatment group are different from those in the comparison group along a number of important dimensions. Claimants of Housing Benefit in the private rented sector are by and large less likely to be retired from work, have more children, and live in larger accommodation compared to the individuals in the comparison group. In Eq. (1), I control for time-invariant structural differences between the treated and comparison groups through individual fixed effects, while time-variant differences (labour market shocks or changes to other benefit income) should be captured by time fixed effects interacted with qualification types and other benefit categories. Nonetheless, given these structural differences and the lack of quasi-random assignment into treated and comparison

[13] This baseline sample contains all individuals from the first nine waves of the US survey who were not excluded for the reasons mentioned in this section.

[14] These fixed effects combine each local authority district with each 36 year-quarter and each 9 survey wave. There are 333 local authority districts in England, but there are no observations in our sample from two of them.

[15] I fix claimant status for these benefits in the pre-reform period, considering the information available at the last survey wave an individual is observed before the reforms.

[16] I chose the cut-off to be December 2011 because most claimants were protected from the effects of the subsidy cuts until 2012 (see Section 2).

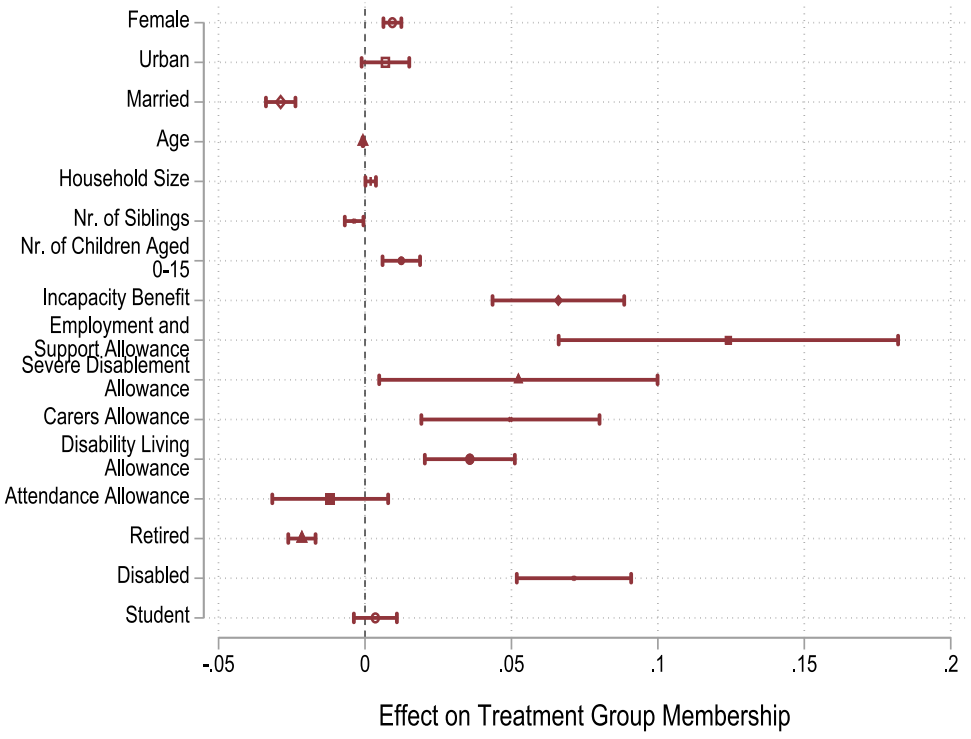

Fig. 1. Balance tests — Correlates of treatment group membership before the reforms.
Notes: These figures plot the point estimates from OLS regressions of the treatment group indicator (in the private rented sector and on HB before the reforms) on covariates. The confidence intervals not spanning zero indicate significance at the 5%-level.

groups, the question arises as to how appropriate the baseline comparison group specification is. To ensure the robustness of my results to the comparison group specified, I carry out a number of sensitivity checks in Section 4.1.3.

Finally, my way of specifying the treatment and comparison groups could be problematic if many individuals who did not claim Housing Benefit prior to the reform become claimants in the post-reform period, positively affecting their work incentives. If this coincides with the reform roll-out, such effects on the comparison group could bias our estimates as both treated and comparison units would face similar post-reform work incentives. In Fig. 2, I show that this is unlikely to be the case as there does not seem to be any evidence of a post-reform change in the shares of those on Housing Benefit in the comparison group, suggesting that these individuals did not start to claim housing subsidies in large numbers after the reforms. Members of the treatment group are naturally more likely to be on Housing Benefits in the years leading up to the reforms – treatment group membership is precipitated on being on Housing Benefit when last observed before the reforms – but seem likely to lose subsidy entitlements in later years, perhaps due to changes in work incentives or subsidy rules resulting from the reforms.[17] In the Appendix A, Figure A.2, I also plot changes in Housing Benefit shares in the PRS and SRS, and show that these shares reduced over time in both parts of the rental sector. Note, that since rents are already lower in the SRS (see Section 2), tenants moving to the SRS to avoid the subsidy cuts may have no longer been eligible for Housing Benefit there, and yet maintained or improved rent affordability. The next sections discuss these issues in more detail.

The Size of the Subsidy Cuts. Another question that could arise is whether the size of the subsidy cuts was sufficient to trigger changes in claimants' labour market behaviour. To assess this, I examine how the reforms affected the Housing Benefit entitlements of treated claimants. Based on the available information on median rents in local authority districts, the change from housing subsidies covering up to 50% of local median rents to up to 30% has imposed a maximum cost on claimants of roughly £58 per month. This is 15.5% of the mean net (excluding Housing Benefit payments) monthly rent in our sample. The maximum cost is estimated using a scenario where the claimants' Housing Benefit entitlements before the reforms were equal to the 50th percentile of local rents, and are reduced to equal the 30th percentile after the reforms (see Section 2). Concurrent policy changes, such as capping the housing subsidies, removing the £15 per week excess that claimants could keep, and reducing rates for claimants in shared accommodation, likely lead to higher reductions in subsidies for some claimants, although I cannot directly identify the individuals who these policies applied to.[18] Nonetheless, the US data do allow me to look at how the coverage of rents by housing subsidies – i.e. the percentage of gross rental payments covered by Housing Benefit payments – has changed over time for claimants. In Fig. 3, I show the Housing Benefit coverage of rents over time for those claimants remaining in the private rented sector throughout the sample period and also for those initially in the private rented sector who later moved away. Those who moved away clearly have higher post-reform coverage compared to those remaining in the private rented sector, suggesting that (1) the reforms did lead to some reduction in Housing Benefit coverage for those who stayed in the private sector and (2) that there were strong incentives to move away from the private sector after the reforms, as moving away has potentially resulted in higher benefit coverage. In line with this, in the Appendix, I also show that there was a sharp post-reform increase in

[17] Note, that in Fig. 2, the treatment group has a share of 1 (everyone is claiming Housing Benefits) in the reform year (2011) by construction, as the treatment group is defined as those individuals on Housing Benefit residing in the private rented sector (PRS) when last observed before the reform year. For this reason, everyone on Housing Benefit in the PRS observed in 2011 is in the treatment group.

[18] See Section 2 for an overview of these policy changes.

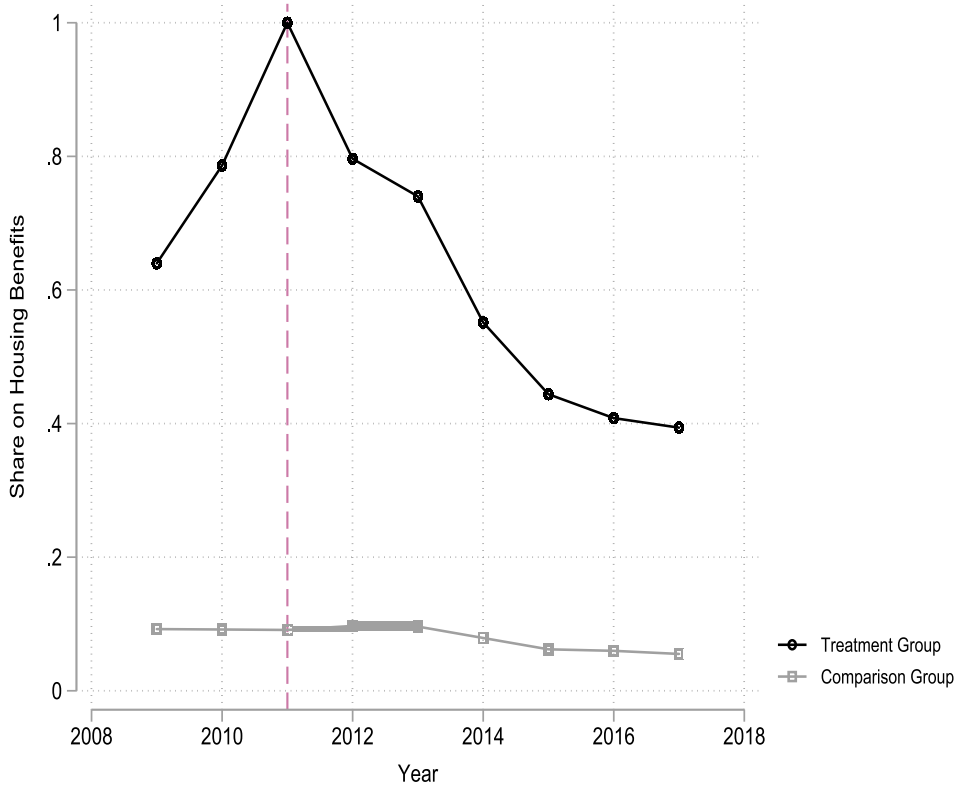

Fig. 2. Housing benefit shares over time in treatment and comparison groups.
Notes: This figure plots the shares of those on Housing Benefit in the treated (black line) and comparison group (grey line), over time. The treatment group has a share of 1 (everyone on Housing Benefits) in the reform year (2011) by construction, as the treatment group is defined as those individuals on Housing Benefit in the private rented sector when last observed before the reform year. The vertical dashed line indicates the reform year. Three percent of our full sample is in the treatment group (black line) while the rest of the sample is in the comparison group (grey line).

PRS tenants who moved to the SRS, when compared to tenants in non-SRS parts of the housing market (Figure A.3). Although it is possible that these effects are mechanical due to too few periods in my sample before the reforms – the probability of moving away should increase over time, as one has to be in the PRS in at least one period beforehand – these figures do imply that moving away was a possible behavioural response to the subsidy cuts for PRS tenants. Below, I aim to confirm this by assessing the effects of the reforms on Housing Benefit coverage empirically.

4.1.1. Baseline results

My baseline analysis focuses on two main labour market outcomes of interest: an individual's likelihood of being employed, and the likelihood of participating in the labour market. Both of these outcomes are measured through dummy variables that indicate each survey respondent's employment and labour force participation status.[19] Results from the baseline specification detailed in Eq. (1) are summarised in Table 2. The coefficients for the treatment variable clearly indicate positive effects on both employment and participation. More specifically, my findings for the more demanding specification indicate that experiencing the reforms can be associated with a 8.9 percentage point (0.18SD) increase in the probability of being employed, along with a 5.5 percentage point (0.11SD) increase in the probability of actively participating in the labour market. My results for anticipation effects suggest that we can largely rule these effects out, the point estimates being close to zero and precisely estimated. Nonetheless, as mentioned earlier, the treated and comparison groups are structurally different – allocation of treatment is non-random – and if these differences correlate with time-varying changes in labour market trends, Eq. (1) need not necessarily identify the effects of the subsidy cuts. To ensure this is not the case, I evaluate the validity of the parallel trends assumption, along with the sensitivity of my results to different comparison groups specifications, in the next two sections.

4.1.2. Parallel trends

Here, I evaluate the presence of pre-trends in the difference-in-differences specifications. The underlying assumption of the difference-in-differences approach is that of parallel trends between treatment and comparison groups prior to the policy change (Angrist and Pischke, 2008). To lend validity to this assumption, and to assess the effects of the subsidy cuts over time, I estimate an event-study specification of Eq. (1) where the treatment group indicator is interacted with a set of year fixed effects.[20] In this specification, I use 2011 as the reference (baseline) period — divergence in time trends between the treatment and comparison groups each year are compared to the level difference between the two groups in this reference period. If there are noticeable pre-reform trends (point estimates significantly different from zero before 2011) this should make us doubt the validity of the parallel trends assumption.[21] The event study estimates for our main labour market outcomes are plotted in Fig. 4.

[19] Information on employment and labour force participation status is from the Understanding Society survey (see Section 3). Participation is defined as those either seeking employment or being employed.

[20] I use year fixed effects instead of year-quarter (or year-month) fixed effects because individual year-quarters in the US sample will sometimes have very few observations in the treatment group, leading to reduced power. On the other hand, interacting with year dummies means that I only have two pre-reform point estimates to test the parallel trends assumption on.

[21] Note, that level differences in outcomes for the treated and comparison groups are not a threat to the difference-in-differences estimation but diverging trends compared to the baseline (reference period) level difference are.

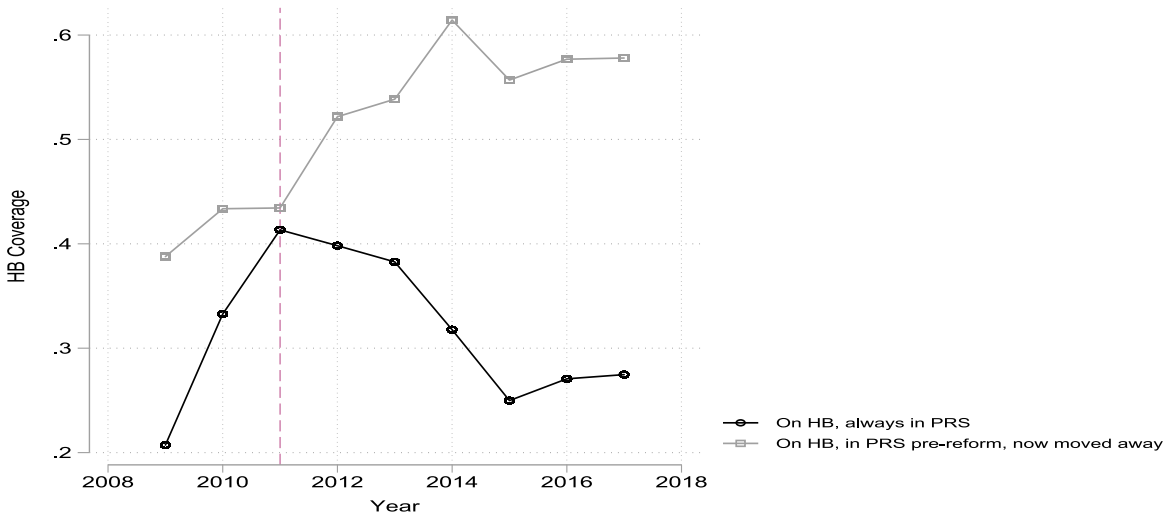

Fig. 3. HB recipients in private rented sector over time.
Notes: This figure plots the Housing Benefit coverage of rents – i.e. the share of gross monthly covered by Housing Benefit payments – for two groups of survey respondents: (1) those claiming housing subsidies in the private rented sector who remained in the private rented sector after the reforms and (2) those claiming housing subsidies in the private rented sector who moved away from the private rented sector after the reforms. The vertical dashed line at 2011 indicates the reform year.

Table 2
OLS results — baseline specification.

	Employment		Participation	
	(1) Employed	(2) Employed	(3) Active	(4) Active
Anticipation effect	0.001	0.007	0.006	0.014
	(0.017)	(0.016)	(0.018)	(0.018)
Treatment effect	0.091***	0.089***	0.056***	0.055***
	(0.016)	(0.015)	(0.015)	(0.014)
Observations	229107	227609	229107	227609
R^2	0.777	0.785	0.798	0.810
Mean DV (Pre-treatment)	0.48	0.48	0.61	0.61
SD DV (Pre-treatment)	0.50	0.50	0.49	0.49
Individual FE	Yes	Yes	Yes	Yes
LA × Wave × Quarter FE	Yes	Yes	Yes	Yes
Qual × Wave × Quarter FE	No	Yes	No	Yes
Other Benefits × Wave × Quarter FE	No	Yes	No	Yes

Notes: This table summarises the estimates from the baseline difference-in-differences specification in Section 4. The model is estimated using ordinary least squares (OLS). Odd numbered columns show estimates for specifications with individual fixed effects and quarter-wave fixed effects specific to each local authority district. Even numbered columns show estimates for specifications that also include quarter-wave fixed effects specific to each qualification type and quarter-wave fixed effects interacted with dummies for other benefit membership. The coefficient corresponding to the 'anticipation effect' measures the effects of the reforms during the transition period between April 2011 and December 2011. The coefficient corresponding to the 'treatment effect' measures the effects of the reforms after the end of the transition period from January 2012. Standard errors are in parentheses and clustered at the local authority district level. $p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

The event study plots show no evidence of significant pre-reform point estimates, and diverging trends starting after the reform are evident when looking at our outcomes. A notable concern however is that pre-reform trends, while not significant, are present and simply continue into the post-reform period: we see some evidence of this on the event study plots for both of our main labour market outcomes, which seem to suggest that the post-reform trends are a clear continuation of pre-reform ones (see Panel a) of Fig. 4, for example. Since these findings are predicated on the inclusion of survey wave and time fixed effects specific to each local authority district – meaning that time-varying changes in outcomes not specific to treated and comparison groups should be well accounted for – this could be suggestive of issues driven by structural differences between our treated and comparison groups that could bias trends in key outcomes. Based on this, the question arises: how sensitive are the baseline estimates to the likelihood that the parallel trends assumption does not hold?

To answer this question, I implement the robustness procedure developed by Rambachan and Roth (2019). This procedure allows for the likelihood that the parallel trends assumption does not hold exactly and checks the sensitivity of estimates to this possibility. In Fig. 5, I compare 95% confidence intervals obtained – for each outcome – from my baseline difference-in-differences model to those obtained after allowing for deviations from a linear trend (in each period) of up to an 'M' amount. I select 'M' arbitrarily and then check how quickly the confidence intervals indicate a null effect (or an effect of the opposite sign to the baseline one) as 'M' increases. If we observe a null effect already at $M = 0$, this indicates that my results are sensitive to small violations of the parallel trends assumption. Based on Fig. 5, our baseline estimates for participation rate are sensitive to small violations of the parallel trend assumption, but our estimates for employment are not. Nonetheless, taken together with the visible pre-trends on the event-study plots, there is still some suggestive evidence that using the baseline comparison group specification could be a significant concern for this study.

4.1.3. Alternative specifications and additional outcomes

To address concerns about the suitability of the comparison group, I estimate a number of alternative specifications. Balancing tests for covariates are provided for all alternative (and baseline) specifications in Table A.2.

First, to optimise the selection of a comparison group, I employ the synthetic difference-in-differences approach developed by Arkhangelsky et al. (2021). This approach effectively re-weights the regression in Eq. (1) to rely on treated and comparison units that are more similar in the pre-treatment period in terms of observable characteristics. The unit weights in the regression are selected so that pre-treatment trends in outcomes are approximately parallel for treated units and the weighted average of comparison units. According to Arkhangelsky et al. (2021), the advantages of using this approach are two-fold: the re-weighting procedure weakens the reliance on the parallel trend assumption, while incorporating this in a difference-in-differences framework allows for valid large panel inference. A limitation of this approach in the context of my analysis is that it can only be used on a balanced panel of individuals, I therefore restrict the sample to those individuals who are observed in every survey wave. Due to considerable attrition in the US survey, this reduces the sample size to about a fifth of the

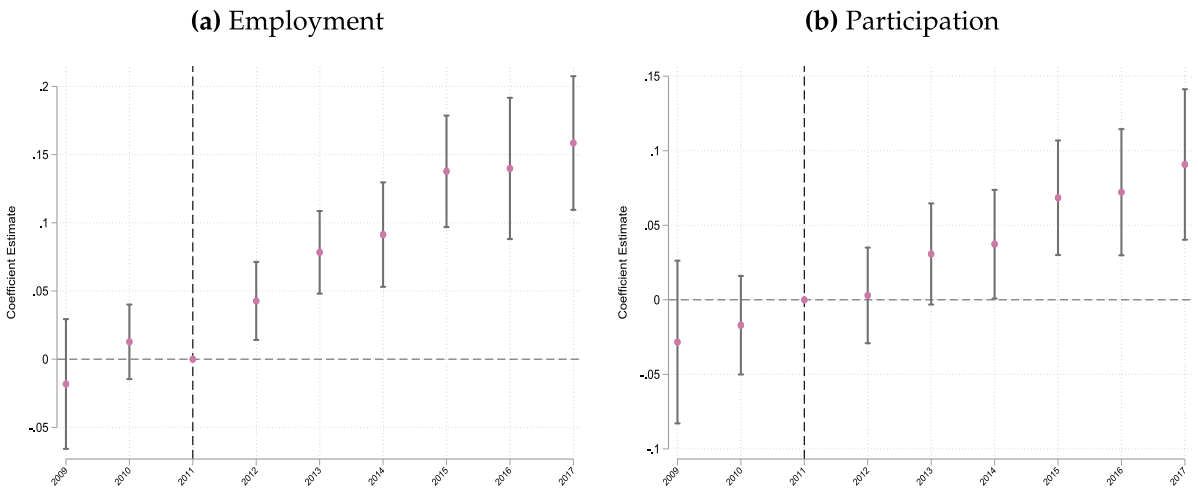

Fig. 4. Event studies – Labour market outcomes – Full Sample
Notes: These figures plot point estimates from the full sample event-study specification outlined in Section 4.1.2. Each point estimate corresponds to the specific year fixed effect interacted with the treatment indicator. The reference period for the event-study specifications is 2011. Confidence intervals not spanning zero indicate significance at the 5% level. The vertical dashed line at 2011 marks the start of post-reform period.

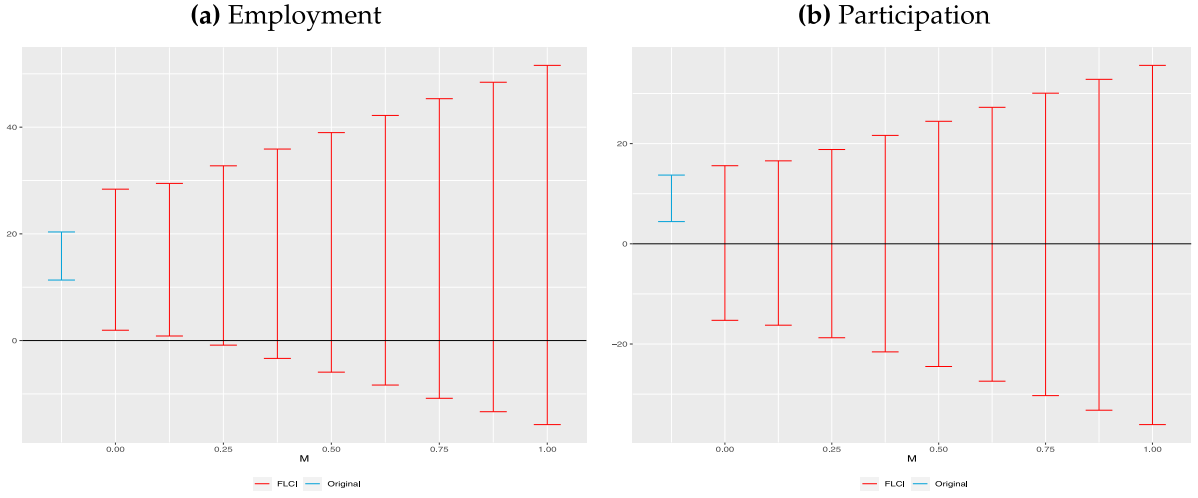

Fig. 5. Sensitivity to parallel trends — Labour market outcomes
Notes: These figures show the results, for each outcome variable, from the robustness procedure developed by Rambachan and Roth (2019). The first 95% confidence interval from the left (marked blue) is obtained from my baseline difference-in-differences model. The other confidence intervals (marked red) are obtained after allowing for deviations from a linear trend (in each period) of up to an $\}M\}$ amount, labelled on the horizontal axis. If the confidence intervals indicate a null effect, or an effect of the opposite sign to the baseline one, at $M = 0$, this suggests that my results are sensitive to small violations of the parallel trends assumption. The y-axis is re-scaled so that numbers correspond to percentage point change in outcomes. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

baseline sample, although this reduced sample remains fairly similar to the baseline one in terms of covariate balance across treated and untreated individuals (see Table A.2).

Second, I also attempt to optimise comparison group selection using propensity score matching as an alternative method.[22] Matched pairs are created based on the following individual characteristics: age, gender, location of residence (urban or rural), household size, dummy variables for economic activity status, dummy variables for tenure type, and dummy variables for different types of benefit income (including Housing Benefits). Matching based on these variables ensures that individuals in the treated and comparison groups are similar in terms of their observable characteristics — this is largely the case for the matching variables in our matched sample, see Table A.2. I then estimate Eq. (1) using the matched sample. Note, that using the matched comparison units results in a considerable reduction in sample size, which could lead to a loss of statistical power for the models that rely on high-dimensional fixed effects (Fetzer, 2019). Nonetheless, unlike in the case of the synthetic difference-in-differences approach, where my sample is reduced due to attrition, this reduction in sample size is a consequence of a closer match between treatment and comparison units.

Finally, following the strategy outlined in Fetzer (2019), I refine the baseline specification in Section 4 in a way that the comparison group only includes those individuals who, at any point during the sample period, were in receipt of Housing Benefit. Restricting the sample this way leads to a pool of survey respondents who are arguably more likely to be similar in terms of both observable characteristics and the housing and labour market decisions they are likely to face. This specification, while *ad hoc*, is still informative in examining how sensitive our results are to changing the comparison group.

Results from alternative specifications, along with baseline results, are summarised in Table 3. Detailed results for each specification

[22] For an overview, see Abadie and Imbens (2016).

are presented in the Appendix A.[23] Here, I include various additional outcome variables of interest from the US survey:

- **Mobility and benefit status:** This set of outcome variables are meant to assess the mobility responses to the subsidy cuts. Claimants residing in the private rented sector who were hit by the Housing Benefit cuts could still avoid a reduction in their housing subsidies by moving away from the private sector, to, for example, continue receiving housing subsidies whilst renting from another part of the market. I test whether this was the case by estimating Eq. (1) using two mobility outcomes: (1) a dummy for whether an individual has changed tenure type (for example, became a homeowner) and (2) a dummy for whether an individual moved to the social rented sector after the reforms. Another possibility is that some claimants immediately lost eligibility to Housing Benefit after the reforms under the new, less generous, system. To test this, I use an outcome which is a dummy variable indicating whether someone who previously (at any point during the survey) claimed Housing Benefit is no longer in receipt of these. Finally, a crucial question is how the reforms affected rent affordability for claimants. To assess this, I use two variables: (1) a dummy variable indicating whether an individual is behind on rent payments and (2) Housing Benefit as a percentage share of gross rents (Housing Benefit coverage).
- **Housing Conditions and Mental Health:** According to the initial report by Shelter (2015), the Housing Benefit cuts led to a deterioration of housing conditions for private rented sector claimants. Some claimants ended up struggling to pay rent, had to move to lower quality housing units, and even had a higher risk of becoming homeless after the reforms (see Fetzer et al., 2019). To assess the changes in housing conditions induced by the Housing Benefit cuts, I make use of four outcome variables: (1) the number of bedrooms in the household to assess effects on accommodation size (2) a dummy variable indicating whether a respondent would like to move out from their current housing unit and (3) a variable measuring the mental health of individuals on a 'caseness' basis. The latter variable counts cases of serious mental health issues over the last year for each respondent, and is included to assess the possible link between increased housing insecurity and mental health (Phinney et al., 2007).
- **Labour Market Outcomes:** Finally, I include two additional labour market outcomes on top of the main outcome variables in Section 4.1.1: (1) a dummy for whether an individual is self-employed to assess whether claimants moved into self-employment in response to the subsidy cuts and (2) log hours worked to assess labour market responses along the intensive margin (i.e. increased work effort).

Overall, the results summarised in Panel A of Table 3 suggest clear effects in terms of mobility responses to the Housing Benefit cuts. Claimants of housing subsidies affected by the reforms were less likely to change tenure (become homeowners), more likely to move to the social rented sector, and were more likely to move off Housing Benefits after the reforms. The effect on the extent to which Housing Benefits cover rental payments – for those claimants who continue to receive these – is negatively affected by the reforms in most specifications but not when I optimise comparison group selection using the synthetic difference-in-differences approach. On the other hand, the reforms have had no significant effect on the incidence of survey respondents being behind on their rental payments. It is possible that moving to the social rented sector, or moving off Housing Benefit by moving away or seeking (and finding) work have mitigated the effects of reduced subsidies on rent affordability for some claimants. Note, however, that

[23] The results summarised in Table 3 are from the specification where all fixed effects are included.

Table 3
Results — all specifications.

	Baseline	Synth - DID	PSM	AC
Panel A: Mobility and housing benefit status				
Changed Tenure	−0.042**	−0.048	−0.094**	−0.066**
	(0.019)	(0.031)	(0.041)	(0.024)
Moved to SRS	0.080***	0.102***	0.073***	0.086***
	(0.012)	(0.006)	(0.024)	(0.017)
No longer on HB	0.200***	0.289***	0.174***	0.088***
	(0.013)	(0.011)	(0.040)	(0.018)
Behind with rent	−0.012	−0.018	0.011	−0.009
	(0.017)	(0.017)	(0.041)	(0.027)
HB Coverage	−0.046**	0.030***	−0.133***	−0.027
	(0.022)	(0.012)	(0.039)	(0.026)
Panel B: Housing conditions and mental health				
Nr of Beds	0.023**	−0.060**	0.171***	−0.016
	(0.011)	(0.025)	(0.051)	(0.019)
Would like to move	−0.003	0.019	0.026	−0.003
	(0.020)	(0.033)	(0.052)	(0.028)
Mental health	−0.222	−0.153	−0.115	−0.066
	(0.135)	(0.193)	(0.398)	(0.199)
Panel C: Labour market outcomes				
Employment	0.089***	0.079**	−0.045	0.040*
	(0.015)	(0.032)	(0.038)	(0.021)
Participation	0.055***	0.041	0.051	0.020
	(0.014)	(0.027)	(0.043)	(0.022)
Log hours of work	0.129***	0.012	0.432***	0.069
	(0.028)	(0.028)	(0.088)	(0.059)
Self-employment	0.019***	0.003	0.042	0.016
	(0.007)	(0.018)	(0.021)	(0.010)

Notes: This table summarises treatment effect estimates for each of our outcomes for all of our main specifications. Column 1 shows results for the baseline sample. Column 2 shows results for the synthetic difference-in-differences analysis. Column 3 shows results for the sample that was matched using propensity score matching. Column 4 shows results for the sample that uses an alternative comparison group.

while the descriptive evidence (see Figure A.2, Figure A.3, or the discussion on the size of the subsidy cuts above, for example) support these findings, mobility responses to the subsidy cuts are difficult to measure in the individual data as both being able to move away from the PRS and being in the treatment group are predicated upon being in the PRS before the reforms. This could lead to the issue of mean reversion, whereby lower post-reform levels of mobility outcomes (moving away from the PRS, for example) are a consequence of mechanically higher pre-reform levels due to the way in which the measures were constructed. To corroborate these results, I therefore look at aggregate (district) level measures of the shares living in the SRS and Housing Benefit spending in Section 4.2 below.

The effects of the subsidy reform on housing conditions and mental health are mostly inconclusive (see Panel B of Table 3). There is a positive effect on accommodation size (number of bedrooms) in the baseline specification but this is not robust to changes in the specification of the comparison group. I also find no evidence of a significant effect on claimants' wanting to move away from their current accommodation, or on their mental health.

Finally, Panel C of Table 3 suggests positive effects on labour market outcomes which however lack robustness and are not significant in most specifications. For example, the synthetic difference-in-differences specification confirms the baseline finding on employment, but this effect is no longer significant in the alternative control group specification and changes sign when I use the PSM approach to optimise comparison unit selection. The baseline (positive) effect on participation remains positive in all specifications, but is smaller and/or less precisely estimated in alternative comparison group specifications, and the same is true for the log hours of work and self-employment variables. Some of these effects remain similar to baseline effects in terms of size, and the lack of significance is likely a result of a lack of power due to a reduced sample. In many cases, the upper bounds of the 95% confidence intervals for these effects would still yield

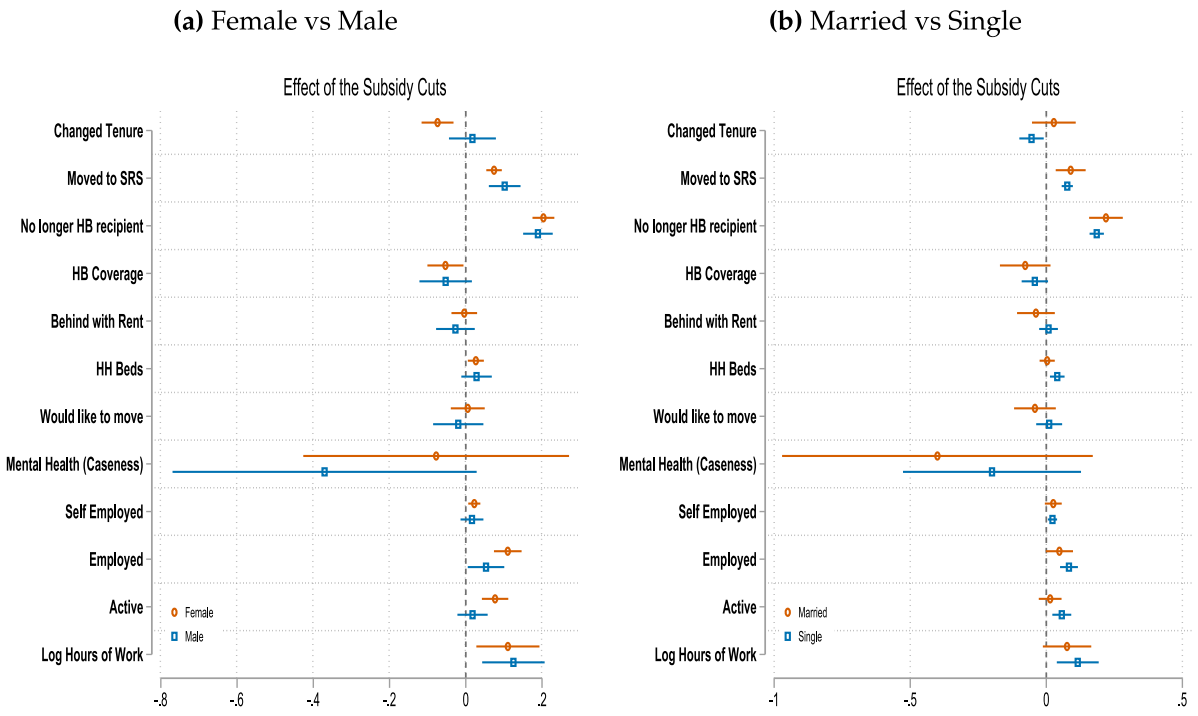

Fig. 6. Heterogeneous effects
Notes: These figures plot the point estimates from OLS regressions of the baseline difference-in-differences models, restricted to specific individual categories. Panel (a) summarises the point estimates from regressions using a female/male only sample. Panel (b) summarises the point estimates from regressions using a married/single sample. Confidence intervals not spanning zero indicate significance at the 5% level.

economically significant labour market effects, suggesting that some of the null effects presented here are not precisely estimated. I attempt to find further evidence of labour market effects at the aggregate (district) level in Section 4.2 below.

4.1.4. Heterogeneous effects

In this section, I check whether the effects of the subsidy cuts were different along observable individual characteristics. I assess the heterogeneity of the baseline effects along the gender, marital status, area of residence (urban or rural) and age (older or younger than median) dimensions. The point estimates of the reform's effects for each of these categories on each outcome are plotted in Figs. 6 and 7. The effects on mobility seem to be similar across all observable dimensions. On the other hand, the coefficient plots suggest that labour market effects are larger and more likely to be positive for females, single claimants, and claimants living in urban areas. Possibly, this indicates that the labour market decisions of these groups are more elastic to changes in income from Housing Benefits. For example, according to the literature on welfare programmes, women are more likely to be the target of welfare schemes and are also more sensitive to changes in associated labour supply incentives than men (see Meghir and Phillips, 2010; Blundell et al., 2016). It is also possible that, quite simply, claimants belonging to these demographic groups have an easier time adjusting to income shocks. Nonetheless, most of the differences between these subcategories are not statistically significant and the evidence on heterogeneous effects is therefore entirely suggestive.

4.2. District-level evidence

Finally, I collate some data on employment and participation rates aggregated at the local authority district level to see if there is evidence of labour market responses to the subsidy cuts at the aggregate (regional) level. This section is also meant to assess whether the positive effects found for employment and participation in the individual-level baseline analysis are confirmed using aggregate data. The data on employment and labour market participation rates in each local authority district are from the Office for National Statistics (ONS).[24] I build a panel of English local authority districts from 2005 to 2018,[25] and estimate the following model:

$$y_{r,t} = \alpha_r + \gamma \times Post_t \times Exp_{r,2010} + \theta_t + \beta' X_{r,t} + \epsilon_{r,t} \quad (2)$$

where $y_{r,t}$ denotes labour market outcomes (employment and participation rate) at the local authority district level ('r') in each year ('t'). I include a local authority district fixed effect (α_r) to control for time-invariant differences in outcomes across districts, and a time fixed effect (θ_t) to control for random shocks. The main coefficient of interest is γ, estimated on the interaction between the exposure to the private rented sector cuts ($Exp_{r,2010}$) in a given district and a dummy for the post-reform period ($Post_t$). I also include regional controls (information on local shares of benefit holders and demographic variables) that vary over time to control for localised shocks to labour markets. Exposure to the benefit cuts is calculated based on data on the costs of different austerity measures from Beatty and Fothergill (2016), and is given by the normalised (by working age population) product of the number of claimants in each district affected by private rented sector housing subsidy cuts and the average financial loss per claimant. I run several different specifications at the district level. First, in some specifications, I restrict the sample to observations before 2014, to ensure that benefit system changes around this year do not influence the estimates (see Fetzer et al., 2019). Second, I exclude the Greater London area from some specifications as this region might be an outlier in terms of its rental market (see Gibbons et al., 2018). The results are summarised in Table 4. Parallel trends in outcomes in the pre-reform period across districts with high and low exposure to the Housing Benefit cuts are visible in Fig. 8.

[24] The data can be accessed through the NOMIS website: https://www.nomisweb.co.uk/.

[25] These data are missing for labour force participation for the year 2018.

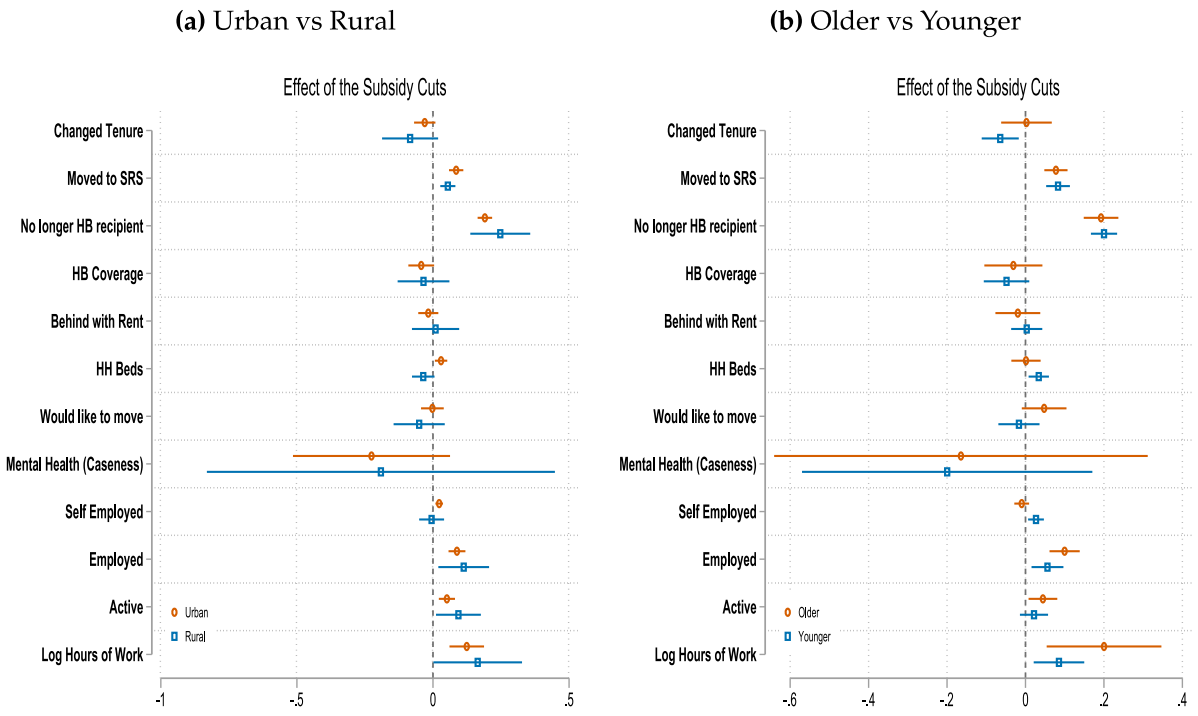

Fig. 7. Heterogeneous effects
Notes: These figures plot the point estimates from OLS regressions of the baseline difference-in-differences models, restricted to specific individual categories. Panel (a) summarises the point estimates from regressions using a urban/rural only sample. Panel (b) summarises the point estimates from regressions using a older/younger only sample. Confidence intervals not spanning zero indicate significance at the 5% level.

The mean amount lost from the Housing Benefit cuts in local authority districts is roughly £160k per year (with SD of £439.7k), and the maximum amount lost is £5.34 million per year, in the London Borough of Brent. In Table 4, the point estimates correspond to a percentage change in participation/employment in response to a 1SD deviation in district-level exposure to the Housing Benefit cuts. Translating this exposure measure into quantifiable figures, a 1SD change corresponds to losses from the cuts being roughly £440k higher per year at the district-level considering the average losses from the Housing Benefit cuts and the number of households affected. The point estimates presented in Table 4 all have a negative sign, but are close to zero and are not significant at any reasonable level of statistical significance. These null effects are consistent across different specifications. Even the upper bound of the confidence interval for the largest effect we observe would only suggest a 0.2% increase in the local employment rate, a rather small effect.

In a back-of-the-envelope calculation, I re-scale these estimates using the district shares of the population affected, given that labour market outcomes are aggregated across districts but, on average, only 12.3% of a district's working age population is affected by the subsidy cuts.[26] When re-scaled this way, and applying the upper bound of the estimated effect to the entire working age population, my estimates would still suggest, at maximum, a 1.3% increase in employment, and a 1.4% increase in labour force participation, but equivalently the lower bounds of these estimates would suggest small negative effects.[27] On the local authority district level, there is therefore no evidence of an aggregate level change in employment or participation rates in response to the Housing Benefit cuts.

To corroborate the results shown in Section 4.1, I also look at changes in the shares living in the SRS, and changes in (log) Housing Benefit expenditures in response to the reforms at the district-level.[28] While I cannot track transitions from different segments of the rental market in the aggregate-level data, an increase in the share living in the SRS after the reforms would be suggestive of tenants responding to the benefit cuts in the PRS by moving to a segment of the rental market where these cuts did not apply (see Section 2). Results for these outcomes are summarised in Table 5 and Table 6. For SRS shares, I only have data for a single pre-reform (2001) and a single post-reform period (2018). For this reason, I modify my baseline model to a simple two-period difference-in-differences specification where change in the SRS share is the outcome variable. The significant point estimates presented in Table 5 indicate a positive aggregate level effect on the share of the district population living in the SRS. My results suggest that a 1SD change in exposure to the Housing Benefit cuts can be associated with a positive change in the share living in the SRS of around 3–4 percent. We can observe the opposite effect for (log) Housing Benefit expenditures at the district-level, where a 1SD change in exposure is associated with a between 1.8% (before 2014) and 7.5% (full sample period but outside London) reduction in Housing Benefit expenditures. Taken together, these findings corroborate both the first-stage effect of the reforms on Housing Benefit spending and entitlements, and the mobility response seen in the individual-level data.

5. Conclusions

This paper looked at the labour market impacts of a reform to the English housing subsidy system that has led to a substantial reduction

[26] This is calculated using the number of households affected by the Housing Benefit cuts in each district from Beatty and Fothergill (2016), and using the UK average household size of 2.4 to calculate the number of individuals affected.

[27] Using this calculation, if we observe a large employment effect of 10%, for example, the observed effect size in Table 4 would have to be at least 0.013 (1.3%).

[28] Data on Housing Benefit expenditures at the local authority level are from the ONS (Nomis). Data on SRS shares are from Shelter (WebLink).

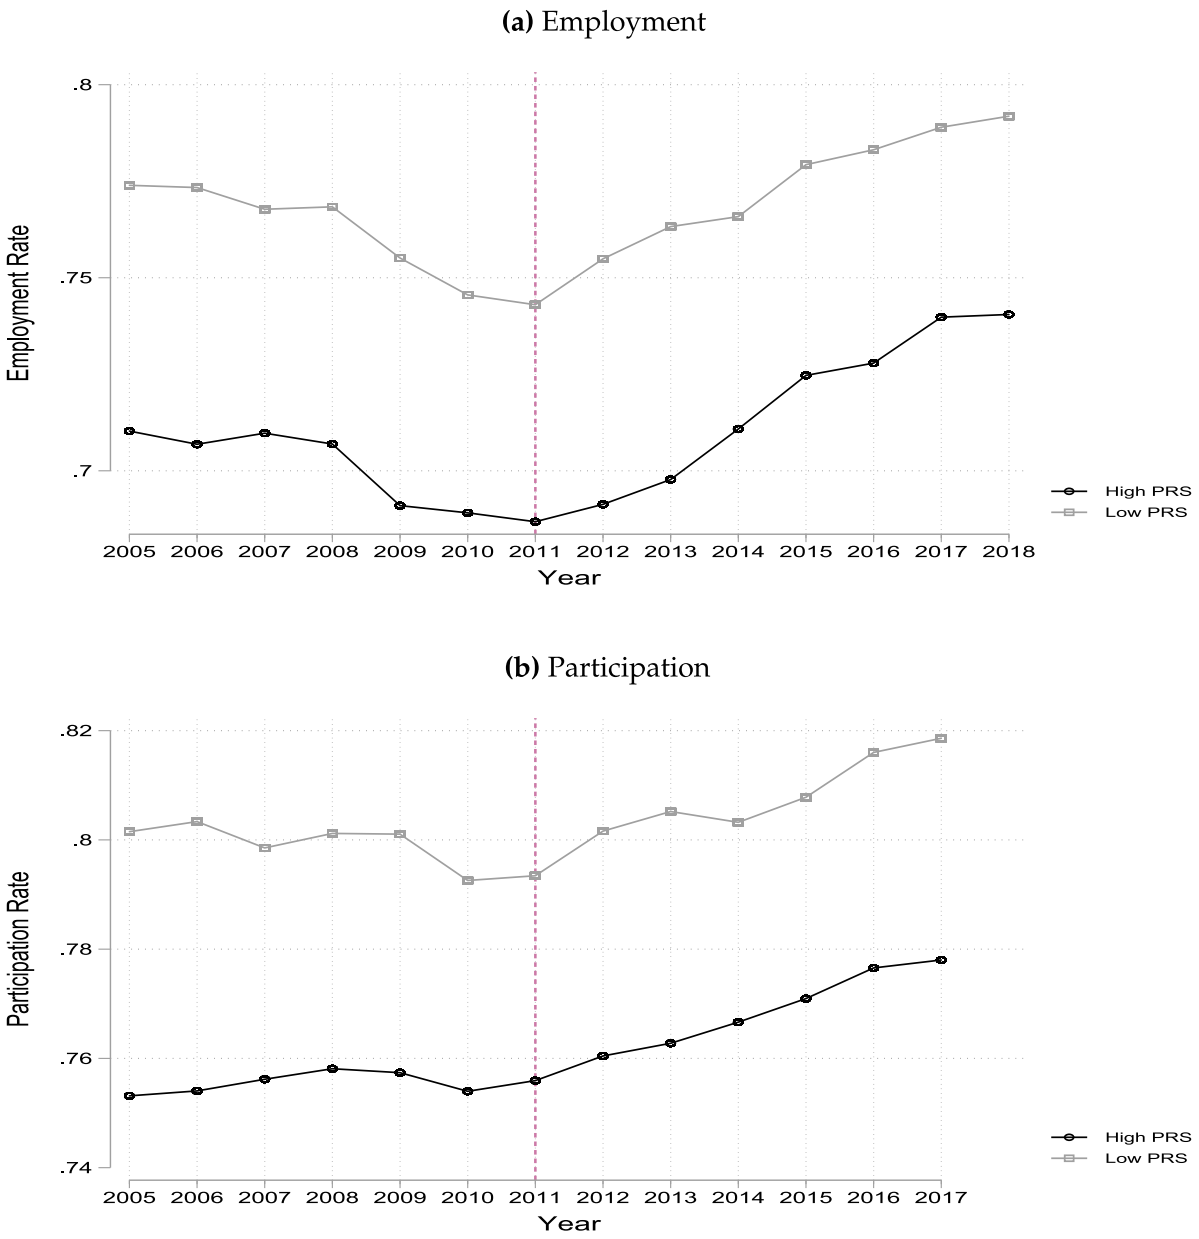

Fig. 8. Time trends – Labour market outcomes – District-level analysis
Notes: These figures plot time trends for district-level employment rate and participation rate for districts with low (below the median) exposure to the Housing Benefit cuts (grey) and districts with high (over the median) exposure to the Housing Benefit cuts (black). The vertical dashed line at 2011 indicates the implementation of the Housing Benefit cuts.

in subsidy entitlements for private rented sector (PRS) claimants. These subsidy cuts were intended to encourage labour market participation and increased work effort by claimants. My empirical strategy made use of a pooled panel difference-in-differences (difference-in-differences) approach where I compared outcomes across groups affected (treated) and unaffected (comparison) by the policy changes. In order to make sure that my results are robust to the choice of comparison group, I estimated a number of different specifications and used both synthetic difference-in-differences and propensity score matching (PSM) methods to optimise comparison group selection. Moreover, I assessed the aggregate level impact of the subsidy cuts on labour market outcomes through a specification that made use of data on participation and employment rates at the local authority district level.

Using individual survey data, I find some evidence of increased employment and participation in response to the subsidy cuts, but these findings lack robustness, and are not supported by results using aggregate labour market data at the district-level. Overall, my findings indicate that, on aggregate, the reforms likely did not have a robust impact on the labour market outcomes of affected claimants, who predominantly responded to the subsidy cuts by moving away to other (unaffected) parts of the rental market. Nonetheless, the precise mechanisms behind my findings are not identified in this paper, and more detailed analysis is needed to understand the factors that drive the labour market decisions of subsidy claimants.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.jhe.2022.101859.

Table 4
OLS Results — district level analysis.

Panel A: Labour force participation				
	Full sample period		Before 2014	
	(1)	(2)	(3)	(4)
	All districts	London excluded	All districts	London excluded
Post × Treatment	−0.002	−0.004	−0.001	−0.003
	(0.002)	(0.003)	(0.003)	(0.005)
Observations	3740	3359	2622	2365
R^2	0.637	0.613	0.667	0.638
Mean DV (Pre-treatment)	0.78	0.78	0.78	0.78
SD DV (Pre-treatment)	0.05	0.04	0.05	0.04
Panel B: Employment rate				
	Full sample period		Before 2014	
	(1)	(2)	(3)	(4)
	All districts	London excluded	All districts	London excluded
Post × Treatment	−0.001	−0.002	−0.003	−0.003
	(0.002)	(0.003)	(0.003)	(0.005)
Observations	3698	3286	2411	2154
R^2	0.731	0.714	0.768	0.745
Mean DV (Pre-Treatment)	0.73	0.73	0.73	0.73
SD DV (Pre-Treatment)	0.06	0.05	0.06	0.05
LA FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Covariates	Yes	Yes	Yes	Yes

Notes: This table summarises the estimates from the difference-in-differences specification in Section 4.2. This specification uses aggregate level data on labour force participation (panel A) and employment (panel B) in a difference-in-differences framework where the coefficient of interest is estimated on the interaction term between a post-reform indicator and a treatment intensity measure. The treatment intensity measure is based on the local authority district level average financial losses from the Housing Benefit cuts multiplied by the local number of affected individuals and normalised by population. The model is estimated using ordinary least squares (OLS). All specifications include local authority district and year fixed effects along with a number of covariates aggregated at the local authority district level. Columns (1) and (2) include the full sample period, while Columns (3) and (4) only include years before 2014. Odd numbered columns include all local authority districts, while even numbered columns exclude districts in the Greater London Area. Standard errors are in parentheses and clustered at the local authority district level. $p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 5
OLS results – district level analysis – SRS shares.

SRS shares		
	(1)	(2)
	All districts	London excluded
Post × Treatment	0.040***	0.033***
	(0.006)	(0.006)
Observations	609	547
R^2	0.295	0.227
Pre-Treatment SRS Share Mean	0.14	0.14
Pre-Treatment SRS Share SD	0.07	0.05
Post PRS Dummy	Yes	Yes
Covariates	Yes	Yes

Notes: This table summarises the estimates from the difference-in-differences specification in Section 4.2. This specification uses aggregate level data on district-level shares of residents living in the SRS, in two-period (before and after) difference-in-differences framework where the coefficient of interest is estimated on the interaction term between a post-reform indicator and the treatment intensity measure described in Section 4.2. The model is estimated using ordinary least squares (OLS). Standard errors are in parentheses and clustered at the local authority district level. $p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 6
OLS results – district level analysis – HB expenditure.

Log HB expenditures				
	Full sample period		Before 2014	
	(1)	(2)	(3)	(4)
	All districts	London excluded	All districts	London excluded
Post × Treatment	−0.049***	−0.075***	−0.018***	−0.031***
	(0.004)	(0.005)	(0.004)	(0.006)
Observations	4019	3607	2622	2365
R^2	0.997	0.997	0.998	0.998
Mean DV (Pre-Treatment)	3.37	3.22	3.37	3.22
SD DV (Pre-Treatment)	0.88	0.77	0.88	0.77
LA FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Covariates	Yes	Yes	Yes	Yes

Notes: This table summarises the estimates from the difference-in-differences specification in Section 4.2. This specification uses aggregate level data on district-level (log) Housing Benefit expenditures in a difference-in-differences framework where the coefficient of interest is estimated on the interaction term between a post-reform indicator and the treatment intensity measure described in Section 4.2. The model is estimated using ordinary least squares (OLS). Standard errors are in parentheses and clustered at the local authority district level. $p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

References

Abadie, A., Imbens, G.W., 2016. Matching on the estimated propensity score. Econometrica 84 (2), 781–807.

Angrist, J.D., Pischke, J.-S., 2008. Mostly Harmless Econometrics: An Empiricist's Companion. Princeton University Press.

Arkhangelsky, D., Athey, S., Hirshberg, D.A., Imbens, G.W., Wager, S., 2021. Synthetic difference-in-differences. Amer. Econ. Rev. 111 (12), 4088–4118.

Beatty, C., Fothergill, S., 2016. The Uneven Impact of Welfare Reform: The Financial Losses to Places and People. Sheffield Hallam University.

Blundell, R., Costa Dias, M., Meghir, C., Shaw, J., 2016. Female labor supply, human capital, and welfare reform. Econometrica 84 (5), 1705–1753.

Braakmann, N., McDonald, S., 2018. Housing subsidies and property prices: Evidence from England. Reg. Sci. Urban Econ..

Brewer, M., Emmerson, C., Hood, A., Joyce, R., 2014. Econometric Analysis of the Impacts of Local Housing Allowance Reforms on Existing Claimants. Department for Work and Pensions.

Carlson, D., Haveman, R., Kaplan, T., Wolfe, B., 2012. Long-term earnings and employment effects of housing voucher receipt. J. Urban Econ. 71 (1), 128–150.

Collinson, R., Ellen, I.G., Ludwig, J., 2015. Low-income housing policy. In: Economics of Means-Tested Transfer Programs in the United States. Vol. 2. University of Chicago Press, pp. 59–126.

Cromarty, H., 2021. Housing Conditions in the Private Rented Sector (England). House of Commons Library, London.

Fenton, A., 2010. How Will Changes to Local Housing Allowance Affect Low-Income Tenants in Private Rented Housing? The Cambridge Centre for Housing and Planning Research, University of Cambridge.

Fetzer, T., 2019. Did austerity cause brexit? Amer. Econ. Rev. 109 (11), 3849–3886.

Fetzer, T., Sen, S., Souza, P.C., 2019. Housing Insecurity, Homelessness and Populism: Evidence from the UK. CEPR Discussion Paper No. DP14184.

Gibbons, S., Sanchez-Vidal, M., Silva, O., 2018. The bedroom tax. Reg. Sci. Urban Econ..

House of Commons, Work and Pensions Committee, 2010. Changes to Housing Benefit Announced in the June 2010 Budget. House of Commons Library, London.

Jacob, B.A., Ludwig, J., 2012. The effects of housing assistance on labor supply: Evidence from a voucher lottery. Amer. Econ. Rev. 102 (1), 272–304.

Meghir, C., Phillips, D., 2010. Labour supply and taxes. In: Dimensions of Tax Design: The Mirrlees Review. Oxford University Press Oxford, pp. 202–274.

Mills, G., Gubits, D., Orr, L., Long, D., Feins, J., Kaul, B., Wood, M., Jones, A., et al., 2006. Effects of Housing Vouchers on Welfare Families. Vol. 8. US Department of Housing and Urban Development, Office of Policy Development and Research, Washington, DC, p. 2010, Retrieved October.

Moffitt, R.A., 2002. Welfare programs and labor supply. In: Handbook of Public Economics. Vol. 4. Elsevier, pp. 2393–2430.

Murray, M.P., 1980. A reinterpretation of the traditional income–leisure model, with application to in-kind subsidy programs. J. Public Econ. 14 (1), 69–81.

Olsen, E.O., Tyler, C.A., King, J.W., Carrillo, P.E., 2005. The effects of different types of housing assistance on earnings and employment. Cityscape 163–187.

Petersen, J., Rabe, B., et al., 2013. Understanding Society–A Geographical Profile of Respondents. Understanding Society Working Paper Series No..

Phinney, R., Danziger, S., Pollack, H.A., Seefeldt, K., 2007. Housing instability among current and former welfare recipients. Am J Public Health 97 (5), 832–837.

Rambachan, A., Roth, J., 2019. An Honest Approach to Parallel Trends. Harvard University.[99], Unpublished Manuscript.

Schone, B.S., 1992. Do means tested transfers reduce labor supply? Econom. Lett. 40 (3), 353–357.

Shelter, 2015. Cuts and changes to housing benefits for private renters (LHA). URL https://england.shelter.org.uk/_data/assets/pdf_file/0020/1142183/7567.01_LHA_Briefing_Report_v4.pdf.

Shroder, M., 2002. Does housing assistance perversely affect self-sufficiency? A review essay. J. Hous. Econ. 11 (4), 381–417.

Shroder, M.D., 2010. Housing subsidies and work incentives. Available At SSRN 1691112.

Susin, S., 2005. Longitudinal Outcomes of Subsidized Housing Recipients in Matched Survey and Administrative Data. JSTOR, Cityscape, pp. 189–218.

Taylor-Gooby, P., 2012. Root and branch restructuring to achieve major cuts: The social policy programme of the 2010 UK coalition government. Soc. Policy Adm. 46 (1), 61–82.

Tunstall, R., et al., 2015. The Coalition's Record on Housing: Policy, Spending and Outcomes 2010–2015. Vol. 18. Social Policy in A Cold Climate Working Paper.

Wilson, W., Barton, C., 2021. Social Rented Housing (England): Past Trends and Prospects. House of Commons Library, London.

Wilson, W., Barton, C., Keen, R., 2016. Housing Benefit Measures Announced Since 2010. House of Commons Library, London.