



The curious incidence of rent subsidies: Evidence of heterogeneity from administrative data

Mike Brewer^{a,b}, James Browne^{b,1}, Carl Emmerson^b, Andrew Hood^b, Robert Joyce^{b,*}

^a University of Essex, Wivenhoe Park, Colchester, CO4 3SQ, United Kingdom

^b Institute for Fiscal Studies, 7 Ridgmount Street, London, WC1E 7AE, United Kingdom

ARTICLE INFO

JEL classification:

H22
H53
I38

Keywords:

Housing
Incidence
Subsidy

ABSTRACT

This paper provides new evidence on the incidence of rent subsidies. We use administrative panel data on subsidy recipients in the UK and exploit a natural experiment in which entitlements were cut for about a million households. In the short-run, about 90% of the incidence of the cuts is found to be on tenants. We also uncover an important dimension of heterogeneity in the balance of incidence between tenants and their landlords. We find that the share of the incidence of the cut that falls on landlords, rather than tenants, is higher in cases where the previous system looked more generous relative to tenants' likely housing needs. This is informative about the likely incidence of alternative rent subsidy schemes.

1. Introduction

Targeted demand-side subsidies for rented housing are a major and growing element of modern welfare states. In this paper we provide new evidence on their incidence, exploiting a natural experiment provided by a substantial package of cuts to a means-tested cash transfer for renters in the UK, known as Housing Benefit (HB). In the short run, we find that a large majority of the incidence of the cuts fell on tenants, but we also uncover significant heterogeneity in the balance of incidence between tenants and their landlords. This yields new insights into the relationship between the design of rent subsidies and the extent to which the programmes actually help their intended beneficiaries.

The design and incidence of these subsidies is of great importance. In the US, on which much of the empirical literature on the incidence of rent subsidies is based, the federal government spent about \$18 billion (0.1% of GDP) in 2014 subsidising the rents of 2.2 million families through Housing Choice Vouchers, the largest such federal program. Demand-side subsidies are more substantial elsewhere. In Great Britain, the government spends around £9 billion per year, or 0.5% of GDP, subsidising the rents of 1.6 million families in privately rented accommodation through HB, which any renter with income and financial assets

low enough can receive.² If rent subsidies raise the price of rented accommodation, governments are transferring some of these resources to landlords rather than low-income tenants.

We estimate the incidence of a package of cuts to HB in the UK which was phased in during 2011 and 2012. The date at which existing recipients were affected depended, for the most part, on the calendar month in which their claim began. This means we can implement a difference-in-differences design to estimate the effect of the subsidy reduction on rents, using those not yet rolled onto the reformed system as a control group at each point in time, and can track the impact of the subsidy cut up to 11 months out. Importantly, the cuts were larger for certain groups which were thought to have previously had overly-generous entitlements. For new claims, the reformed system applied to all claims beginning after a particular date, so we use an interrupted time series

² The US government also provides housing assistance to low-income families through programmes such as Project-based Rental Assistance (which cost \$12 billion in 2014). In addition to HB in the private rented sector, the UK government subsidises tenants in public housing through a combination of sub-market rents and HB. US figures from Congressional Budget Office (<https://www.cbo.gov/publication/50782>); GB figures from Department for Work and Pensions (<https://www.gov.uk/government/statistics/benefit-expenditure-and-caseload-tables-2015>). The evolution of the US system is described in Susin (2002). Hills (2007) provides a detailed account of the UK case.

* Corresponding author at: Institute for Fiscal Studies, 7 Ridgmount Street, London, WC1E 7AE, UK.

¹ Present address: The Tony Blair Institute for Global Change, <https://institute.global/>.

E-mail address: robert.j@ifs.org.uk (R. Joyce).

approach. This analysis of new claimants addresses concerns that our estimates for existing recipients might have been confounded by a contemporaneous impact of the treatment (i.e. the reforms) on our control group: our analysis of new claimants relies on no such control group. Furthermore, new HB claimants are much more likely to be taking on new tenancies than existing HB recipients, so there is much less scope for adjustment costs to be limiting their response to the cuts. These are important differences between the analyses of new and existing claimants, since they help provide reassurance that the small rent adjustments found for existing claimants are not simply due to attenuation bias.

We estimate that the quality-adjusted rents of subsidy recipients were little changed, on average, by the reforms, meaning that, in the short-run, about 90% of the incidence of cuts to subsidies was on the tenants rather than their landlords. But we uncover significant heterogeneity in the balance of incidence between tenants and their landlords, with tenants for whom the original subsidy level was, arguably, high relative to their needs seeing significant falls in their quality-adjusted rents, meaning less than two-thirds of the cut was incident on them. A likely explanation for the significant heterogeneity in incidence is differences in the elasticity of demand for rented housing. The pre-reform HB system meant that certain groups were subsidised to consume an amount of housing that was high relative to their needs. As a result, their demand was relatively responsive when the subsidy was cut back: although basic housing is a necessity, it is more likely to be a luxury good at the margin if housing consumption is already high (relative to needs). As a result, more of the incidence of the subsidy cut on these subsidy recipients was shifted to their landlords. This has important policy implications. If, as cuts to subsidies bite further down the distributions of rent and housing quality, the average demand elasticity of affected tenants falls, then the incidence of any further reductions in generosity would tend to fall even more on subsidised tenants rather than their landlords.

The rest of the paper proceeds as follows. Section 2 discusses how this paper relates to the existing literature looking at the incidence of housing subsidies. Section 3 outlines the details of the reforms to HB that are key for our empirical strategy. Section 4 describes the data and explains our sample selection. Section 5 describes our empirical strategy and discusses identification. Section 6 presents our estimates of the incidence of rent subsidies including, importantly, heterogeneity in that incidence. Section 7 discusses our results and concludes.

2. Related literature

We build on evidence about the incidence of rent subsidies provided by a number of previous studies. These have used a variety of empirical strategies, data and institutional settings. As we will explain, the heterogeneity in incidence we uncover in our own analysis may help reconcile some of the variation in the findings of previous work.

For the United States, Susin (2002) effectively compared long run rent trends between areas where housing voucher supply has been expanded by different amounts, estimating that the existence of the voucher system had increased the rents of *non-recipients* by 16% – although Olsen (2003) argues that one might be concerned about a failure of common rent trends between areas that saw different increases in housing voucher supply. In this paper we focus on the rents of subsidy *recipients*, not non-recipients. But we are careful to assess (and rule out) the possibility that contemporaneous effects of rent subsidies on the rents of non-recipients, of the kind implied by Susin's results, might be confounding our own estimates – or more precisely, given our empirical setting, that changes in rent subsidies might affect the rents of recipients yet to be directly impacted by those changes.

More recent evidence from the US suggests that housing vouchers lead recipients to rent higher-quality properties, but have little impact on the average price of rental housing. Eriksen and Ross (2015) exploit an area-varying expansion in the US housing voucher program, and find that average rents are unaffected by the voucher expansion, but that this masks important heterogeneity: when the supply of housing vouchers is

increased, the price of low-quality housing falls, and the price of housing just below the subsidy ceiling increases, consistent with voucher recipients moving to more expensive units (a conclusion supported by Collinson and Ganong, 2018). The result that housing vouchers do not affect average rental prices towards the bottom of the rental market (and hence that the incidence falls on tenants) is consistent with our conclusion that, on average, the incidence of the cuts we analyse fell mostly on tenants. But a key contribution of this paper is to uncover another important dimension of heterogeneity: we show that the effect of targeted subsidies on rents varies according to the level of the rent ceiling. Among groups whose subsidy cap was cut from a level that looked high relative to their needs, the incidence of cuts fell on their landlords to a greater extent.

A small body of evidence from outside the US has found that a substantial share of the incidence of more generous rent subsidies falls on landlords. Laferrère and Le Blanc (2004), Fack (2006) and Grislain-Letremy and Trevien (2014), using different reforms and slightly different identification strategies, all find that the incidence of higher rent subsidies in France is mostly on landlords. Viren (2013) estimated that one-third to one-half of the incidence of a Finnish rent subsidy is on landlords. Sayag and Zussman (2015) found that a voucher scheme for students in central Jerusalem increased the rents of both recipients and non-recipients, such that landlords captured four-fifths of the value of the grants. The heterogeneity we uncover can provide a way of reconciling these results with those from the US and our own study. In some cases, the reforms these papers exploit are extensions of housing subsidies to specific groups that might be expected to have highly elastic housing demand (e.g. students). When we restrict attention to groups who plausibly share that characteristic (e.g. single adults aged under 35), we get more similar results.

The heterogeneity that we find is also important in understanding how our results relate to the only previous paper on this topic that uses UK data: Gibbons and Manning (2006), which looks at a cut to HB in the mid-1990s that applied only to new claimants, and affected those with the highest rents among households of their size in their local area. By effectively comparing the rent levels of new claimants and existing recipients, they estimated that 60% to two-thirds of the incidence of the cut was on landlords. This is consistent with our finding that the incidence on landlords was much higher where tenants were previously subsidised to rent properties further towards the top of the quality distribution for their type of household.

3. A natural experiment: reforms to UK housing benefit

Our estimates of the incidence of rent subsidies exploit a set of reforms to Housing Benefit (HB) in the UK in 2011 and 2012. Here we outline the features of the reform that are most important for our identification strategy; further details are available in Appendix A.

HB is an entitlement-based, means-tested, cash transfer for renters: any renter with income and financial assets low enough will receive it if they apply. We focus on tenants of private landlords; those in public housing were unaffected by the cuts that we exploit. In the majority of cases (80% in January 2011), HB is paid to the tenant (who remains responsible for paying rent to their landlord), rather than being paid direct to the landlord. There is no systematic direct contact between the landlords of subsidised tenants and any government agency. For recipients whose HB claim began before April 2011, maximum entitlement to HB (before the means-test is applied) was a function of actual rent and a cap known as the Local Housing Allowance (LHA) rate, given by $HB = \min(LHA \text{ rate}, \text{rent} + £15 \text{ per week})$. The applicable LHA rate varies geographically (according to Broad Rental Market Areas, BRMAs³) and by the subsidy recipient's family type (as detailed

³ There are 192 BRMAs in Great Britain, each containing around 140,000 households (330,000 individuals) on average, making them around three times the size of US counties.

in Appendix A). If subsidy recipients rented a property whose rent was below their applicable LHA rate, then they could keep the first £15 per week of the difference.

3.1. How reforms affected existing recipients

In our main empirical analysis we look at those who were already receiving HB before April 2011. These existing recipients saw their entitlements reduced in two ways. First, on the first annual anniversary of their claim after April 2011 (i.e. at some point between April 2011 and March 2012) the weekly ‘excess’ of up to £15 that subsidy recipients could keep if their rent was below the LHA rate was removed. Second, nine months later (i.e. at some point between January and December 2012) existing recipients saw their entitlements reduced by a number of changes to the applicable LHA rate. The overall impact of these cuts to LHA rates was to reduce the cap on HB payments, giving subsidy recipients an incentive to seek cheaper properties or to pay less for a given property. The biggest single change to LHA rates was to base them on the 30th percentile of local rent levels among the applicable property type, rather than the median, as previously. But there were additional cuts to LHA rates that were much more tightly targeted on three particular groups:

1. Those previously deemed to need a 5-bedroom property saw their LHA rate reduced from one based on the rental prices of local 5-bedroom properties to one based on the rental prices of local 4-bedroom properties. Affected families all had at least 4 children under 16 (see Appendix A).
2. Those in some areas of central London saw their LHA rate reduced by the imposition of ‘national caps’, which bound only in those areas.
3. Single adults aged 25–34 saw their LHA rate reduced from one based on the rental prices of local one-bedroom properties to one based on the local rental prices of single rooms in shared properties.

Our analysis of heterogeneity in the incidence of the subsidy reduction is based on these three groups of individuals. In each case, one can think of the group-specific reduction in LHA rates as being motivated by a belief that the existing system was particularly generous to these individuals, either in terms of the quality of housing it enabled them to purchase relative to their needs (those previously deemed entitled to a 5-bedroom property and single childless 25–34 year olds) or in terms of the area in which they were enabled to live (those in central London). One might therefore expect the elasticity of housing demand at the margin to be different (higher) for these particular groups.

Because of the phased roll-out of the reforms described above, the point in time at which individuals were affected depended on the month in which their claim anniversary fell. We refer to a group of individuals whose claim anniversary fell in a given month as a ‘claim-cohort’, and hence there are 12 claim-cohorts. Otherwise-identical recipients in different claim-cohorts faced different levels of subsidy at the same point in time. This motivates the difference-in-differences design that we specify in Section 5.

3.2. How reforms affected new claimants

Individuals begin a new claim to HB either because they move property and begin a new tenancy (around a third of cases) or because other changes in their circumstances – such as a fall in earnings or a change in household composition – render them eligible (the remaining two thirds of cases). New claims of HB were also affected by the reforms described above, but the way the changes were rolled out was different. Those individuals who began their claim to HB from April 2011 onwards were subject to the reductions in the subsidy immediately:⁴ they received no

‘excess’ if their rent was below the LHA rate, and their LHA rates were determined by the new, less generous rules.

4. Data and sample selection

We use administrative monthly panel data on HB claims in Great Britain, from the Single Housing Benefit Extract (SHBE). Details of the construction of key variables, including data cleaning, are given in Appendix B.

4.1. Existing recipients

Our estimates are based on a random one-in-three sample of individuals receiving HB, renting from a private landlord, and assessed under the LHA rules in January 2011, shortly before the reforms were implemented. These recipients were affected by the removal of the £15 excess on their first annual claim anniversary from April 2011, and were typically affected by the other elements of the reform package nine months after that. We use monthly observations for these subsidy recipients between January 2010 and November 2013 inclusive, meaning that we follow all recipients for 11 months after the cuts to LHA rates, or 20 months after their first annual claim anniversary following April 2011 (when they lost the ‘excess’ of up to £15 per week that was previously payable if their rent was below their LHA rate, as described above).

After dropping 15% of the sample with missing information, we have an estimation sample of 239,723 subsidy recipients. Table 1 shows basic demographic characteristics, employment status and a measure of local neighborhood deprivation for subsidy recipients, before the reform took effect, for the universe of HB recipients and our estimation sample. Almost 80% of recipients are single adults (of which slightly less than half have children), over half are aged between 25 and 44, in around 70% of cases neither the claimant nor any partner works, and around a third of recipients live in the most deprived fifth of neighborhoods.

Table 2 shows the distribution of weekly rents and HB entitlements for the two samples. The median rent paid by recipients is £115 a week,

Table 1
Demographic characteristics of existing recipients in January 2011.

Characteristic	Universe	Estimation sample
Household type		
Single man	29%	29%
Single woman	16%	16%
Couples without children	6%	6%
Single parents	33%	32%
Couples with children	16%	16%
Age		
Under 25	16%	16%
25–34	32%	32%
35–44	25%	25%
45–59	19%	19%
60 and above	8%	8%
Employment status		
In-work family	31%	31%
Out-of-work family	69%	69%
Neighborhood deprivation ^a		
Most deprived quintile	34%	33%
2nd overall quintile	27%	27%
3rd overall quintile	19%	19%
4th overall quintile	12%	13%
Least deprived quintile	8%	8%
N	850,249	239,723

Source: authors’ calculations using SHBE data.

^aFor more details on the construction of this index see <https://www.gov.uk/government/statistics/english-indices-of-deprivation-2010-technical-report>.

⁴ The only exception was the change in the entitlement of single 25–34 year olds without children, which took effect from January 2012.




Fig. 1. Average Housing Benefit entitlement of new claimants by date of claim. (seven-day moving average).

Table 2

Rents and HB entitlements of existing recipients in January 2011.

	Universe (£ per week)	Estimation sample (£ per week)
Rent		
Mean	134	134
Median	115	115
25th percentile	92	92
75th percentile	150	150
HB entitlement		
Mean	126	126
Median	111	110
25th percentile	91	91
75th percentile	144	144
N	850,249	239,723

Source: authors' calculations using SHBE data.

and the inter-quartile range is £92 to £150; this compares to median weekly household income in the UK population of around £500 a week.⁵ Median HB entitlements are £4 a week lower than median rents, meaning that at least half of recipients faced a shortfall (i.e. their rent exceeded their HB entitlement) before the reforms.

4.2. New claimants

Our estimates for new claimants are based on all claims that began between 1st June 2010 and 1st December 2011; we extract information on the claimants' circumstances at the time of the claim for HB. Fig. 1 shows a seven-day moving average of HB entitlements for new claims between June 2010 and December 2011, and Fig. 2 shows the equivalent for rents. A striking feature of both is the large spikes in mean entitlements and rents just before the reforms affected new claimants on 1 April 2011, which we discuss further in the Appendix. Because of this, we select our sample to exclude a window of data around the reform. We take a conservative approach and exclude all new claims made between 1st December 2010 and 31st May 2011, marked with vertical lines on Figs. 1 and 2.⁶ This means that our estimation sample contains new claims between June and November 2010 (giving us 336,486 ob-

servations from before the reforms took place) and between June and November 2011 (giving us 334,093 observations from after the reforms took place). It is then straightforward to see the patterns in the data that underlie our formal results for new claimants in Section 6. Having been quite flat in the pre-reform period (with, if anything, a gentle linear decline), entitlements to HB clearly settle at a lower level after the reforms, whereas there is little discernible break in the trend for rents across the pre- and post-reform windows.

5. Empirical strategy

5.1. Specification for existing recipients

For our main estimates of the extent to which entitlements to HB and rents changed as a result of the reforms, we exploit the variation in the point in time at which existing recipients were affected by the reforms. This variation is between 'claim-cohorts': the 12 groups whose claim anniversary fell in a different calendar month. The nature of the roll-out meant that otherwise-identical recipients in different claim-cohorts faced different levels of subsidy at the same point in time. In our analysis of the impact on all existing recipients and our examination of heterogeneity in the incidence of the reform, we use a difference-in-differences (DiD) linear regression, specified as:

$$y_{iact} = f_a(t) + \pi_c + z\beta + X'_{iact}\alpha + \varepsilon_{iact} \quad (1)$$

where i indexes subsidy recipients who are observed at multiple points in calendar time t , live in area (BRMA) a , and are members of claim-cohort c , defined by the calendar month in which their annual claim anniversary falls.⁷ The main outcome variables are rent, HB entitlement, and shortfall (the difference between the rent and HB entitlement), all measured in £s per week.⁸

⁷ As mentioned in Section 3, some individuals are affected by the reforms earlier than the rest of their claim-cohort as a result of having a change of circumstance that triggers a reassessment. Hence the difference-in-difference approach we employ could be thought of as an 'intent to treat' design – although rather than some individuals in our 'treatment' group avoiding treatment, we have some individuals in our 'control' group ending up as treated.

⁸ Our administrative data source records contractual rents, which may differ from actual rents either because tenants are in arrears or because landlords informally accept a rent that is lower than the contractual one. We are not able to tell whether the prevalence of these phenomena was affected by the reforms, but qualitative evidence suggests that rent arrears did increase after the reforms

⁵ Source: authors' calculations using 2010–11 Family Resources Survey.

⁶ We have conducted sensitivity analysis and, as the figure suggests, our estimates are robust to small shifts in the window of data excluded.




Fig. 2. Average rent of new claimants by date of claim. (seven-day moving average).

X is a vector of control variables. For our main outcomes, this includes dummies for the full set of interactions between BRMA and number of bedrooms in the property (which we have top-coded at 5), and quintiles of deprivation measured at the neighborhood level.⁹ This rich set of control variables helps to ensure we do not confound changes in property or neighborhood characteristics with changes in the price of accommodation. If, despite these controls, there remain unobserved falls in quality, then we will pick this up as a price change rather than a quality change, and this would lead us to over-estimate the incidence of the subsidy on landlords, and under-estimate the incidence on tenants. To anticipate our results, though, this would only mean that our conclusion about the incidence of the reforms holds even more strongly in reality. We also control for family type and age, in case these change over time for reasons unrelated to the reform in ways that are not adequately captured by our time trends: these make a negligible difference to our estimates. We also present results from regressions that do not include controls for property and neighborhood characteristics. A comparison of these results to those from our preferred covariate-adjusted estimates summarises the extent to which recipients responded to the subsidy cut by renting housing of lower quality or in less desirable areas.

The reforms we study included specific targeted components designed to reduce support for three particular subgroups, whose level of entitlement to support under the previous system was deemed to be especially high relative to their housing needs. A key part of our analysis focuses on whether we see evidence that the balance of incidence between tenants and landlords is different in these cases, as economic theory might suggest. To examine this potential heterogeneity in incidence as flexibly as possible, we estimate Eq. (1) separately for those subgroups. As part of our heterogeneity analysis we also consider as outcomes property size (as measured by number of bedrooms or

whether the property is shared), the probability of moving house, and the probability of moving out of central London. This allows us to examine whether the elasticity of housing demand appears greater among tenants who bear less of the incidence of the cuts, providing supporting evidence for the underlying mechanisms driving any heterogeneity in incidence. When looking at these outcomes, we remove controls for housing quality and replace them with fixed effects for BRMA at the start of the period of observation. When the outcome is whether the claimant moved house, we include controls for estimated rental contract and claim anniversaries.

We estimate Eq. (1) by OLS where the outcome variable is continuous, and as a probit when the outcome is dichotomous. Estimated standard errors allow for heteroskedasticity and for errors clustered at the BRMA level. Identification of β in Eq. (1) depends on the standard ‘common trends’ assumption: in the absence of the reforms, trends in the outcome variables would have been unrelated to treatment status (at the mean, conditional on covariates). It also depends on the assumption that the ‘control’ group is unaffected by the treatment. We discuss these assumptions in turn.

In our context, common trends means that there should be no systematic, unobserved, time-varying differences between subsidy recipients who are in different claim-cohorts. This is untestable, but there are at least two good reasons to believe it to be true. First, there is strong evidence of common trends across claim-cohorts before the reforms took place: Fig. 3 plots changes over time in mean HB entitlements of three example claim-cohorts (May, August and November), after controlling for the composition of properties being rented.¹⁰ It shows that, before May 2011, trends in HB entitlement are extremely similar across the three claim-cohorts. Then in May 2011, the May claim-cohort saw their HB entitlements fall, as they lost any excess. This cohort then saw a further fall in HB entitlement nine months later, in February 2012, as they were affected by the rest of the reform package. The same pattern holds for the August and November claim-cohorts, but with the drops in entitlements occurring three and six months later. The fact that differences between cohorts appear only as the reform takes effect for some cohorts and not others, and then disappear again once all cohorts are fully affected, is compelling evidence that in the absence of reforms there would

(Beatty et al, 2013). The ultimate implication of this for the incidence of the reforms is not clear, as it depends whether or not these arrears are eventually settled. It is also not clear whether household-survey-based measures of rents, as used in some other studies (including Gibbons and Manning, 2006), would be more likely to pick up actual rents paid or contractual rent.

⁹ The local deprivation measures were not matched in to the SHBE data at the time of writing of Beatty et al (2013). This means that the results presented here for new claimants are not numerically identical to those in Beatty et al (2013), though they are extremely similar and no qualitative conclusions are affected.

¹⁰ Specifically, we plot mean residuals from a regression of LHA entitlements on all interactions between BRMA and number of bedrooms.




Fig. 3. Average maximum entitlement of existing recipients by month.
(Residual from regression of entitlements on BRMA and number of bedrooms, £pw).

have been no differences in outcome variables across claim-cohorts. We provide a further assessment of common trends by allowing for treatment effects before the treatment was actually applied, and this reveals no economically important differences in trends in HB entitlements and rents across cohorts (after controls) prior to the roll-out of the subsidy cuts (see Section 6, and in particular the top panel of Table 3).

Appendix Table A1 described four example claim-cohorts (February, May, August and November) and shows that they look almost identical in terms of demographic characteristics, employment status, local-area deprivation, rents and housing benefit entitlements. Further evidence of the extremely high degree of similarity between different cohorts is provided by the fact that the inclusion of ‘claim-cohort’ fixed effects has a minimal impact on our estimated coefficients, as shown in column (3) of Table 3.

The second assumption required for identification – that the treatment does not affect the outcomes of the control group – might not hold if the reforms affected the rents of those not yet rolled onto the reformed system. In a market with no frictions and perfect competition, there would be a single (quality-adjusted) rental price at all times. In such a world, the rents paid by all claim-cohorts would change instantaneously, regardless of whether they had been directly affected by the cuts in subsidy, and estimation of Eq. (1) would therefore find that the reform had no impact on rents, regardless of the true impact. More generally, market-level effects of a less extreme nature would attenuate our estimates of the extent to which the reforms were incidence on the landlords of existing recipients. However, the rental market is not in reality characterised by spot prices: tenants’ rents are usually fixed for the duration of their contract (typically a year). Most crucially, though, our analysis of the impact on the rents of new claimants acts as a natural robustness check. The different nature of the roll-out for new claimants means we can use identifying assumptions for those claimants that would not be violated by market-level effects. We detail this below.

5.2. Specification for new claimants

New claimants were affected by the reformed HB system if their claim began in April 2011 or afterwards, motivating us to use an ‘Interrupted Time Series’ design, as follows:

$$y_{iat} = f_a(t) + 1(t \geq \text{April 2011})\beta + X'_{iat}\alpha + \epsilon_{iat} \quad (2)$$

i indexes new subsidy recipients starting a claim at time t who live in area a . The coefficient of interest is β . We specify $f_a(t)$ as a BRMA-specific

linear trend that is allowed to vary between the pre- and post-reform periods. X is a vector of control variables that includes dummies for the full set of interactions between BRMA and number of bedrooms in the property (which we have top-coded at 5), and family type and age. Eq. (2) is estimated using OLS, and standard errors are robust to heteroskedasticity and clustered at the BRMA level.

The key identifying assumption is that, in the absence of the ‘interruption’, the outcome variable would have been a smooth function of time around April 2011, when the reforms began applying to new claimants. In practice, our strategy is not a pure interrupted time series design because we exclude a 6-month window of data around the date of the reforms, so our treatment effects rely on some extrapolation, being based on any difference in the covariate-adjusted outcome between our pre- and post-treatment windows that cannot be explained by time trends. There are at least two good reasons to be confident about this identification strategy. First, the trends in HB entitlements and rents for new claimants both before and after the reforms were very simple, as shown in Figs. 1 and 2. Second, as Appendix Table A2 shows, the pre- and post-reform claimants in our estimation sample are extremely similar with respect to observed characteristics. Importantly, identification does not depend on a comparison group observed contemporaneously (as it does for the analysis of existing recipients), and so the estimates are not susceptible to bias if cuts to rent subsidies also impact the rents of those not (yet) affected by those cuts. We therefore view this analysis as an important robustness check on our main results.

5.3. Are our results informative of long-run effects?

In common with the majority of the empirical literature examining this question, we estimate effects over a relatively short time horizon. The nature of the identification strategy suggests we have accurate estimates of the incidence of the cut to HB on existing recipients up to 11 months after it took effect. Two main reasons why the longer-run effects might be different would be if there were rental contract rigidity (e.g. because of renegotiation costs) or fixed costs for tenants of moving property or for landlords of finding new tenants.¹¹

But there are several good reasons to believe that our short-run estimates are informative of the long-run effects of the subsidy reduction, although none is conclusive. First, by the time of the final observation

¹¹ See Genesove (2003) for empirical evidence on the importance of price stickiness in the rental market.

Table 3
Estimated impact of cuts to Housing Benefit on existing recipients, £/wk.

Model specification		(1)	(2)	(3)	(4)	(5)
12 months before main impact	Housing benefit	3.09 (0.69)	0.66 (0.35)	0.93 (0.35)	−0.53 (0.16)	−0.68 (0.15)
	Rent	3.43 (0.61)	0.95 (0.28)	1.29 (0.27)	−0.02 (0.12)	−0.14 (0.11)
	Rent net of HB	0.34 (0.24)	0.28 (0.25)	0.36 (0.26)	0.51 (0.15)	0.54 (0.14)
	Estimated pass-through (proportion)					0.20 (0.15)
Loss of excess	Housing benefit	−0.72 (0.73)	−3.40 (0.41)	−3.18 (0.39)	−4.90 (0.44)	−4.98 (0.42)
	Rent	3.73 (0.74)	0.93 (0.31)	1.25 (0.30)	−0.74 (0.29)	−0.81 (0.27)
	Rent net of HB	4.46 (0.23)	4.33 (0.27)	4.43 (0.28)	4.17 (0.34)	4.17 (0.34)
	Estimated pass-through (proportion)					0.16 (0.05)
Point of main impact	Housing benefit	−4.12 (1.93)	−6.67 (0.85)	−6.42 (0.82)	−8.31 (1.09)	−8.31 (1.01)
	Rent	5.17 (1.43)	2.78 (0.40)	3.08 (0.40)	−0.65 (0.78)	−0.73 (0.68)
	Rent net of HB	9.29 (0.71)	9.45 (0.79)	9.50 (0.80)	7.66 (0.84)	7.58 (0.85)
	Estimated pass-through (proportion)					0.09 (0.08)
11 months after main impact	Housing benefit	−1.35 (2.02)	−5.04 (0.82)	−4.87 (0.80)	−7.08 (1.07)	−6.84 (0.92)
	Rent	6.97 (1.84)	3.43 (0.56)	3.64 (0.56)	−0.79 (1.19)	−0.79 (1.09)
	Rent net of HB	8.32 (0.39)	8.47 (0.48)	8.51 (0.48)	6.29 (0.80)	6.06 (0.83)
	Estimated pass-through (proportion)					0.11 (0.15)
Controls for:						
Local area, number of bedrooms and interactions			X	X	X	X
Neighborhood deprivation			X	X	X	X
'Claim-cohort' fixed effects				X	X	X
BRMA-level linear time trend					X	X
Month dummies					X	X
Family type, age and interactions						X
N		239,576	239,576	239,094	239,094	238,782

Source: authors' calculations using SHBE data.

Note: Standard errors in brackets are robust to heteroskedasticity and clustering at the BRMA level. Standard errors on estimated pass-through parameters account for correlation in the errors from the two parent regressions. When controlling for local area, number of bedrooms and interactions we control for LA, BRMA, number of bedrooms in the property (shared accommodation, 1 bedroom, 2 bedrooms, 3 bedrooms, 4 bedrooms, 5 or more bedrooms), and interaction terms that capture all possible combinations of number of bedrooms and BRMA. When controlling for family type, age and interactions we define 37 mutually exclusive combinations of family type and age: families without children are split jointly by family type (single men, single women, couples) and age of claimant (under 25, 25–34, 35–44, 45–59, 60 or more); families with dependent children are split jointly by whether lone parents or couple parents, age of claimant (under 25, 25–34, 35–44, 45 or more), and number of children (1 or 2 or more for under 25s, and 1, 2 or 3 or more for other ages).

we have for each existing recipient, the vast majority have had (assuming they are on year-long contracts with their landlord) at least one opportunity to move or negotiate a lower rent since the point at which they were affected by the package of cuts, and some will have had two opportunities: our final observation is 20 months after the £15 excess was removed, and 11 months after being fully rolled on to the reformed system. Second, as Fig. 4 shows, we find no evidence that the incidence of the reforms changes between when they start to be applied and 11 months after they applied in full. On the contrary the estimated pass-through parameter remains flat throughout this period. As mentioned above, many recipients would have had an opportunity to renegotiate rent by the end of that 11-month window. Third, although we cannot rule out that recipients would renegotiate their tenancy agreements after the period covered by our data, for some subgroups we do see quick rent responses, including for a set of tenants with several children for whom

the fixed costs of moving are likely to be relatively high (we discuss the potential reasons for this in Section 6). Fourth, the UK rental market more generally is characterised by levels of mobility that suggests that the costs of moving are not particularly high for a large number of households. Survey data from 2013 to 2014 indicates that over a third of those in private rented accommodation had lived there for less than a year, and over half for less than two years.¹²

Finally, a key reason for believing that our estimates are informative about the long-run effects of the subsidy reduction is that our estimates for new claimants provide a valuable robustness check. As mentioned

¹² Source: 2013–14 English Housing Survey headline report (<https://www.gov.uk/government/statistics/english-housing-survey-2013-to-2014-headline-report>).




Fig. 4. Estimated pass-through parameter, by months since main impact of HB cuts.

Note: Figure uses estimates from the same models estimated in column 5 of Table 3. The pass-through parameter is calculated as the estimated impact on rent divided by the estimated impact on housing benefit.

(Impact on rent divided by impact on HB, with 95% confidence interval).

in Section 3, analysis of contemporaneous survey data indicates that around a third of new HB claims coincide with a new tenancy,¹³ where rigidities caused by renegotiation costs or fixed costs of moving are not applicable (unlike for the other two thirds, for whom the most likely reason for a new claim is a drop in income that makes them eligible, e.g. a movement out of paid work). If these kinds of rigidities were causing our estimated effects on rents for existing recipients to be close to zero, then one would not expect to obtain the same result for new claimants; but as we are about to show, we do obtain the same results for new claimants as for existing recipients.

6. Regression results

6.1. Results for existing recipients

Table 3 presents estimates based on Eq. (1) for the effect of the reforms on rents, entitlement to HB and the difference between them (the so-called “shortfall”). We add regressors to build up to our preferred specification. For each specification we show four sets of coefficients. The first set of coefficients captures whether we record any differences in rents and HB entitlements 12 months before the point of main impact of the reforms: the coefficients β_{-12} from Eq. (1). This effectively provides a ‘placebo test’ of our empirical strategy. The remaining sets of coefficients capture the impact of the reform at different stages: the point at which subsidy recipients reach their first annual claim anniversary after April 2011, at which time they lost the excess; the point nine months later when they were subject to all elements of the reform package; and eleven months after that, which is the latest point in our data at which we observe all claim-cohorts. These coefficients are β_{-9} , β_0 and β_{11} from Eq. (1).

In column 1 we have no control variables: the only variables on the right hand side are the vector of dummies corresponding to the number of months before or after the month in which the subsidy recipient was rolled onto the new system. In the absence of controls for underlying time trends or changes in our measures of housing quality, rents increase as the reform is rolled out, while HB entitlements fall only slightly. This is because there was an upwards underlying trend in rents and HB entitlements. In column 2, we add our controls for housing quality (local area, number of bedrooms and their interaction, along with a measure of neighborhood deprivation). Once these controls are included, the rise in rents over time is smaller, and the falls in HB entitlements are larger,

reflecting the fact that changes in quality were acting to increase rents (and hence entitlements) over this period, independent of the reform. Column 3 adds 12 fixed effects for our ‘claim-cohorts’: the inclusion of these fixed effects makes little difference to our estimates, providing evidence of no systematic differences between claim-cohorts. In column 4, we add our time trends (see Section 5). There was a secular upwards trend in rental prices and entitlements over time, and so the inclusion of time trends leads to a reduction in our estimated effects of the reforms on rents (which become negative after the loss of excess) and HB entitlements. Finally, our preferred specification (column 5) adds controls for family type and age; this makes little difference to the estimates.

Reassuringly, in our main specification there is almost no estimated impact on HB entitlements or rents 12 months before the main impact of the reforms, and this supports the suggestion provided by Fig. 3 that we do have common trends across claim-cohorts. As a result of the removal of excesses, our preferred specification suggests that recipients lost an average of about £5 per week, and rents were reduced slightly, by about £0.80 per week, which is statistically significant at the 1% level. We attribute this fall in rents to an anticipatory adjustment to the remaining, impending reforms; we would not expect landlords to lose out from the excess removal in isolation. Indeed for tenants with one-year tenancies whose claim anniversary coincides with the anniversary of their tenancy, this point in time would be the last formal opportunity to change rents before the remaining reforms took effect.

The loss of HB rose to £8.31 per week (7% of pre-reform mean entitlements) nine months later, when the rest of the reforms took effect. But the impact on rents did not change, remaining at below £1 a week (and no longer statistically significant). These point estimates imply that around 90% (£7.58 of £8.31 per week) of the incidence of the cuts was on tenants at this point, with the remaining 10% on landlords, and the 95% confidence interval for the pass-through parameter – the share of the cut incident on landlords – spans from -7% to +25%. The estimated impact on rents remains small over the period covered by our data, changing from just £0.73 a week to £0.79 a week eleven months after the main impact, by which point the vast majority of recipients will have had at least one opportunity to renegotiate rents or move house. We remain unable to reject the hypothesis that the pass-through parameter is zero. We obtain similar results for a wide range of demographic and geographic subgroups (including age, family type and region).¹⁴ As noted in Section 5, the proportion of the incidence that fell on tenants would be even higher than we estimate if tenants adjust to the reforms

¹³ Source: authors’ calculations using the 2011–12 Family Resources Survey.

¹⁴ Brewer et al., 2014, reports full results for a variety of sub-groups.

Table 4
Estimated impact of cuts to Housing Benefit on new claimants, £/wk.

Model specification	(1)	(2)	(3)	(4)
Housing benefit	−9.28 (1.18)	−6.33 (0.46)	−7.86 (0.52)	−8.20 (0.50)
Rent	−1.57 (1.11)	1.64 (0.42)	−0.14 (0.65)	−0.38 (0.64)
Rent net of HB	7.71 (0.34)	7.97 (0.35)	7.72 (0.49)	7.82 (0.49)
Estimated pass-through (proportion)				0.05 (0.08)
Controls for:				
Local area, number of bedrooms and interactions		X	X	X
Neighborhood deprivation		X	X	X
BRMA-level linear time trends			X	X
Family type, age and interactions				X
N	667,276	661,961	661,961	659,095

Source: authors' calculations using SHBE data.

Note: Standard errors in brackets are robust to heteroskedasticity and clustering at the BRMA level. Standard errors on estimated pass-through parameters account for correlation in the errors from the two parent regressions. When controlling for local area, number of bedrooms and interactions we control for LA, BRMA, number of bedrooms in the property (shared accommodation, 1 bedroom, 2 bedrooms, 3 bedrooms, 4 bedrooms, 5 or more bedrooms), and interaction terms that capture all possible combinations of number of bedrooms and BRMA. When controlling for BRMA-level time trends, we allow them to differ before and after the reform. When controlling for family type, age and interactions, we define 40 mutually exclusive combinations of family type and age: families without children are split jointly by family type (single men, single women, couples) and age of claimant (under 25, 25–34, 35–44, 45–54, 55–64, 65 or more); families with dependent children are split jointly by whether lone parents or couple parents, age of claimant (under 25, 25–34, 35–44, 45 or more), and number of children (1 or 2 or more for under 25s, and 1, 2 or 3 or more for other ages). The results here differ slightly from those presented in Beatty et al. (2013), as controls for neighborhood deprivation were not included in that analysis.

by choosing lower quality properties in ways not captured by our controls for local area, number of bedrooms and neighborhood quality.

6.2. Results for new claimants

Table 4 shows the estimates of the impact of the reform on the rents and HB entitlements of new claimants, based on Eq. (2). Again, the different columns build up to our preferred specification. Focussing on our preferred specification (column 4), we estimate that the reforms reduced HB awards for new claimants – conditional on claimants' characteristics and property type and location – by an estimated average of £8.20 per week. Almost all of the incidence of this – an estimated £7.80, or 95% – was on the tenants. Again, we obtain similar results for a wide range of demographic subgroups.¹⁵ These estimates do not rely on the assumption that a contemporaneous control group is unaffected by the treatment (see Section 5), and so the fact that the estimates on incidence are very similar to our main results is evidence against the possibility that our main estimates were biased downwards by a violation of that assumption. It also provides evidence that our main estimates are not driven by short-run rigidities or fixed costs of moving, as these concerns apply to a lesser extent when looking at the impact on new claimants: we estimate that around a third of new claims to HB coincide with new tenancies, rendering short-run rigidities due to renegotiation costs or fixed costs of moving irrelevant. Even if rents are only affected for that one third of new claimants who also have a new tenancy (i.e. scaling up the estimates in Table 4 by a factor of 3), then this would imply that rents fell by £1.38 as a result of the reform for that third, still leaving the vast majority (83%) of the incidence on those tenants.

6.3. Heterogeneity in incidence

Table 5 shows separate estimates of the impact of the HB cuts for three subgroups:

- Single adults without dependent children due to be aged 25–34 at the point that the reforms took effect, who would have been entitled to the 1-bedroom LHA rate in the absence of reform but instead were entitled only to the shared accommodation rate;
- Large families entitled to the 5-bedroom LHA rate before the reform but the 4-bedroom rate after the reform;
- Recipients living in one of the five BRMAs in which the overall nationwide caps on LHA rates now bind (those five BRMAs are all in London, though the majority of London BRMAs were not affected).

For brevity, we report only estimated impacts 11 months after being fully rolled onto the reformed system. For each group, the first row (labelled 'covariate-adjusted') shows estimates with controls analogous to those presented in Tables 5 and 6. The second row (labelled 'unadjusted') shows estimates without controls for contemporaneous property characteristics (but with a control for initial BRMA, based on circumstances in January 2011). The difference between the unadjusted and covariate-adjusted estimates provides a summary of the extent to which tenants ended up living in lower-value types of properties as a result of the reforms.

Single adults aged 25–34 without dependent children lost an average of £13 per week in HB from the reforms (conditional on property size, local area fixed effects and neighborhood quality); we estimate that their quality-adjusted rents fell by about £5 per week, implying that just over one third of the incidence was on their landlords. The 95% confidence interval for the pass-through parameter – the share of the incidence that fell on the landlords – runs from 16% to 66%. The estimates that do not adjust for property characteristics show larger falls in both HB entitlements and rents, suggesting that some of the individuals affected responded by moving to cheaper properties. We corroborate this directly below by looking at additional outcome variables.

Families who were entitled to the 5-bedroom LHA rate in January 2011 lost an average of about £29 per week in HB entitlement from the reforms; we estimate that their rents fell by almost £12 per week, implying that about 40% of the incidence was on their landlords, with a 95% confidence interval running from 16% to 64%. As above, a comparison

¹⁵ Reported in Beatty et al., 2013.

Table 5

Estimated impact of cuts to Housing Benefit on existing recipients likely to be affected by certain elements of the reform package (11 months after main impact), £/wk.

		Housing benefit	Rent	Rent net of HB	Estimated pass-through (proportion)	N
Increased scope of shared accommodation rate	Covariate-adjusted	−13.05	−4.80	8.25	0.36	49,569
	Unadjusted	(1.36) −15.55 (1.59)	(1.31) −7.36 (1.55)	(1.73) 8.18 (1.78)	(0.10)	49,635
Abolition of 5-bedroom LHA rate	Covariate-adjusted	−29.21	−11.69	17.52	0.40	5699
	Unadjusted	(8.49) −31.60 (9.99)	(5.48) −19.04 (9.27)	(5.44) 12.56 (5.36)	(0.12)	5703
National caps on LHA rates	Covariate-adjusted	−41.93	−5.68	36.25	0.14	16,992
	Unadjusted	(9.96) −48.48 (12.59)	(10.19) −17.07 (14.20)	(12.31) 31.41 (12.13)	(0.24)	16,992
All existing recipients	Covariate-adjusted	−6.84	−0.79	6.06	0.11	238,782
	Unadjusted	(0.92) −7.40 (0.91)	(1.09) −1.11 (1.12)	(0.83) 6.28 (0.84)	(0.15)	239,279

Source: authors' calculations using SHBE data.

Note: Standard errors in brackets are robust to heteroskedasticity and clustering at the BRMA level. Standard errors on estimated pass-through parameters account for correlation in the errors from the two parent regressions. "Covariate-adjusted" figures include controls for BRMA, local authority, number of bedrooms in the property, local area deprivation, 'claim-cohort', calendar month, linear time trends in each BRMA, and family type and age. "Unadjusted" figures do not include controls for contemporaneous BRMA, LA, number of bedrooms and local area deprivation, but do include controls for BRMA in January 2011.

Table 6

Estimated impact of cuts to maximum entitlements on property choices of existing recipients likely to be affected by certain elements of the reform (11 months after main impact).

			N
Increased scope of shared accommodation rate	Probability of moving (ppts per month)	1.0	43,655
		(0.5)	
Abolition of 5-bedroom LHA rate	Probability of living in shared accommodation (ppts)	17.0	43,564
		(0.4)	
National caps on LHA rates	Probability of moving (ppts per month)	0.6	5406
		(0.8)	
All existing recipients	Number of bedrooms	−0.14	5700
		(0.16)	
Increased scope of shared accommodation rate	Probability of moving (ppts per month)	0.8	16,163
		(0.6)	
Abolition of 5-bedroom LHA rate	Probability of moving out of capped area (ppts per month)	0.3	16,163
		(0.5)	
National caps on LHA rates	Probability of moving (ppts per month)	0.4	219,592
		(0.2)	

Source: authors' calculations using SHBE data.

Note: Standard errors given in brackets are robust to heteroskedasticity and clustering at the BRMA level. When the outcome variable is probability of moving, results are from a probit regression including controls for BRMA in January 2011, calendar month, claim-cohort, family type and age and rent and claim anniversaries, run on data from April 2011 onwards. When outcome variable is probability of living in shared accommodation, results are from a probit regression with all controls listed above as well as linear time trends in each BRMA, run on data from April 2011 onwards. When outcome variable is number of bedrooms, results are from an OLS regression with same controls as given above.

with the estimates that do not adjust for property characteristics suggests that these subsidy recipients also responded by living in cheaper types of properties. The finding that a large share of the incidence of the reforms falls on the landlords of this group does not support the possibility that the heterogeneity we find is driven by differences in search costs or the costs of moving, since such costs are likely to be relatively large for large families.

Finally, subsidy recipients who, in January 2011, were living in one of the five BRMAs in which the overall national caps on LHA rates bind lost an average of about £42 per week in HB entitlement (conditional on property characteristics) from the reforms. We estimate that their

quality-adjusted rents fell relatively little, though their raw rents fell more, suggesting some possible quality adjustments (but neither the covariate-adjusted nor unadjusted rent changed by a statistically significant amount). The finding that the incidence of the cuts has fallen on this group more than the other subgroups examined can be rationalised if demand for housing is less elastic along the margin of which area to live in than along the margin of property size within a given area.

Table 6 provides direct evidence of whether recipients in these subgroups moved properties in response to the reforms (for dichotomous outcomes, we used a probit specification and report marginal effects at the mean of the covariates). A large proportion of single adults aged

25–34 responded by moving into shared accommodation: the reforms increased the probability of living in shared accommodation by 17 percentage points for this group. This is strong evidence of a relatively high elasticity of demand among this group of recipients. We also find some evidence that, consistent with a relatively high elasticity of demand, some of those previously entitled to the 5-bedroom rate responded by moving to smaller properties, and that both those affected by the extension of the Shared Accommodation Rate and those affected by the 5-bedroom rate are around 1 percentage point more likely to move house each month as a result of the reforms (the average (pre-reform) monthly moving rate among LHA recipients was 2.2%). Finally, although the reforms increased the probability of moving for those affected by the introduction of the national caps, less than half of those additional moves are out of the affected area: this is supporting evidence that, as argued above, demand is less elastic along the margin of which area to live in than along the margin of property size within a given area.

7. Discussion and conclusions

During 2011 and 2012 the UK government reduced the generosity of the rent subsidy it provides to low-income private renters. Using monthly administrative panel data on subsidy recipients, and exploiting the phased roll-out of the reforms, we estimate that, on average, about 90% of the incidence of these cuts in the short-run was on tenants. But we find significant heterogeneity: for two groups singled out by the reforms on the basis that they were previously subsidised to rent properties that were large relative to their needs, we estimate that less than two-thirds of the incidence of the cut fell on them (and we can reject the possibility that all of the incidence was on them).

Why is this? One possible explanation for our results is that these two groups were hit harder than average by the changes, and this simply made them quicker to notice and to make an effort to respond, explaining the large share of the incidence that fell on their landlords. We do not find this explanation plausible. First, it relies on short-run rigidities explaining the absence of larger changes in rents for other groups. For several reasons outlined in detail in [Section 5](#), we do not believe this to be the case. Second, our results show that tenants affected by the national caps lost by far the *most* HB of the subgroups considered, and yet the estimated incidence on them is high (and we cannot reject the hypothesis that it was 100%).

We instead argue that it is variations in the elasticity of housing demand that is likely to explain the heterogeneity in incidence we observe. Theory suggests that the incidence of rent subsidies on landlords should be higher when demand for rented housing is more elastic. The two groups who shifted more of the incidence of the cuts onto their landlords are both groups to whom the previous system was deemed to be particularly generous – subsidising young single people to rent self-contained properties, and large families to rent 5-bedroom properties. The logic was essentially that, at the margin, housing for these groups had the characteristic of a luxury rather than a necessity. One might therefore expect the housing demand of these groups to be relatively elastic at the margin. For example, compared to individuals with dependent children or partners, 25–34 year-old single adults without children may be relatively willing to substitute between self-contained and shared accommodation when their HB entitlements are cut. In fact, we find direct supporting evidence of this, in that a significant number of individuals in this group *did* choose to move into shared accommodation as a result of the subsidy cut. The group affected by the abolition of the 5-bedroom rate are families with large numbers of children who were, in many cases, fully subsidised to rent some of the largest properties in their area. After the subsidy cut this group might not be prepared to pay much for an additional bedroom, rather than having more children sharing a room. Again, we find some evidence that some members of this group moved to smaller accommodation as a result of the reforms, though the estimated reduction in the average number of bedrooms is not statistically significant. Importantly, the between-group heterogeneity

in the elasticity of demand that our results indicate is not necessarily a property of the groups themselves. Rather, it is a consequence of the fact that the demand elasticity is higher at the margin once housing consumption is higher, and the previous system was arguably subsidising some groups to consume an especially large amount of housing (relative to their needs). In other words, while the heterogeneity is visible empirically by comparing different tenants who (due to the pre-reform system) were located at different points in the housing quality distribution relative to their needs, fundamentally the heterogeneity we find suggests there is variation in demand elasticities within-tenant along the housing quality distribution.

Heterogeneity in the elasticity of demand could also explain some of the contrasting findings of different empirical studies on the incidence of rent subsidies. The most natural comparison for this study is [Gibbons and Manning \(2006\)](#), the other paper to look at the incidence of the HB programme. They studied reforms in the mid-1990s that introduced caps on the size of rents eligible for HB, based on average market rents in the local area. As such, the subsidy recipients who were directly affected by the mid-1990s reforms were a relatively high-rent minority who might be able to substitute more easily towards cheaper accommodation. In contrast, the reforms studied here extended these sorts of restrictions much further down the rent distribution – typically to the 30th percentile of local rents, and sometimes lower – and affected the large majority of subsidy recipients. Where we do focus on subgroups to whom the size of the subsidy was more generous relative to their needs prior to the reform, our results on incidence are closer to those of [Gibbons and Manning](#). Meanwhile [Eriksen and Ross \(2015\)](#) – who exploit a similarly broad-based change to the generosity of US housing subsidies to the change we study here – come to conclusions about their overall incidence that are similar to ours.

A natural interpretation of this variation in demand elasticities along the housing quality distribution is that housing is a necessity at basic levels but becomes a luxury at the margin when housing consumption increases. The insight for the likely incidence of alternative rent subsidy regimes is potentially substantial. If, as cuts to subsidies bite further down the distributions of rent and housing quality, the average demand elasticity of affected tenants falls (and hence the effect of subsidies on total demand and rental prices falls), then the incidence of less generous subsidies will tend to fall proportionately more on tenants.

Acknowledgements

Funding: Browne, Emmerson, Hood and Joyce acknowledge funding from the Economic and Social Research Council ([ESRC](#)) through the Centre of the Microeconomic Analysis of Public Policy at IFS (grant number [RES-544-28-5001](#)), and Brewer acknowledges support from the [ESRC](#) through the Research Centre on Micro-Social Change at the University of Essex (grant number [ES/L009153/1](#)).

The authors thank the UK Department for Work and Pensions, which made data and funding available for an assessment of the reforms. They also thank conference participants at the SOLE/EALE 2015 conference in Montreal, the Work and Pensions Economic Group at the University of Sheffield and the 2016 PEUK conference in Oxford, and seminar participants at the Centre for the Analysis of Social Exclusion at the London School of Economics, the Centre for Regional Economic and Social Research at Sheffield Hallam University, the Economic and Social Research Institute in Dublin and the Institute for Fiscal Studies in London for useful comments. The views expressed and any errors are the authors' alone.

Supplementary materials

Supplementary material associated with this article can be found, in the online version, at doi:[10.1016/j.jue.2019.103198](#).

Appendix A: Policy background

In 2015–16, spending on Housing Benefit (HB) in the UK is projected to be £24.5 billion: 12% of all government spending on cash transfers.¹⁶ £9.1 billion of that total is spent on rent subsidies for recipients in the private rented sector (the focus of this paper), with the remainder spent on tenants in public housing (who are additionally subsidised indirectly through having a sub-market rent). Spending on HB for private renters increased by 136% in real terms between 2000–01 and 2010–11, thanks to a 94% increase in the caseload and a 22% increase in average entitlements during a period of rising real rents. Since then, real expenditure has been roughly flat: further growth in the number of subsidy recipients has been offset by the impact of the reforms analysed in this paper, which cut the generosity of entitlements as part of a wider package of measures aimed at reducing government borrowing following the global financial crisis and associated recession.

For subsidy recipients who rent from a private landlord and whose claim began in April 2008 or later, HB entitlement is a function of actual rent and a cap known as the Local Housing Allowance (LHA) rate.¹⁷ For a subsidy recipient with no private income or assets who lives with no more than a partner plus any dependent children,¹⁸ the function under the pre-reform system was:

$$HB = \min(LHA \text{ rate}, \text{rent} + £15 \text{ per week})$$

The LHA rate varies geographically, and by the subsidy recipient's family type. The geographical variation is between areas known as Broad Rental Market Areas, which are deemed to represent self-contained housing markets. There are 192 BRMAs in Great Britain, and a further 8 in Northern Ireland. The variation by family type arises through a set of rules that maps a subsidy recipient's family type to a reasonable accommodation size (ranging from a room in a shared property to a five bedroom property), known as the 'size criteria'. Under the LHA rules, subsidy recipients are allowed one bedroom for each of the following occupiers, each coming only into the first category for which they are eligible:

- A couple each aged 16 or over.
- An individual aged 16 or over.
- Two children under 16 of the same sex.
- Two children under 10.
- A child.

For example, a couple with two children aged 12 of opposite sex are entitled to three bedrooms, but a couple with two children aged 12 of the same sex are entitled to two bedrooms. Before the reforms, the maximum number of bedrooms was capped at 5, but the reforms lowered

¹⁶ HB is an entitlement-based program, rather than a cash-constrained, rationed program like the Housing Choice Voucher Program in the US, whereby any renter with sufficiently low income and financial assets is entitled to it. Figures on UK HB spending from Department for Work and Pensions (2015) "Benefit expenditure and caseload Tables 2015", available at <https://www.gov.uk/government/statistics/benefit-expenditure-and-caseload-tables-2015>.

¹⁷ Claims that began before April 2008 are not assessed under the LHA rules, were not affected by the reforms studied here and are ignored in the rest of the paper.

¹⁸ For subsidy recipients living with an adult other than their partner, 'non-dependent deductions' (NDDs) are subtracted from 'rent' in the formula. In addition all HB claims are subject to a means test. This withdraws entitlement at a rate of 65p for each £1 by which income, after direct tax, exceeds a threshold that varies by family type. The system of NDDs and the rules of the means test were unaffected by the set of reforms studied here, so we abstract from them throughout and focus simply on 'maximum' (pre-means test) entitlements ignoring the impacts of any NDDs.

this to 4. Individuals living in shared (rather than self-contained) accommodation are entitled to the lower 'shared accommodation rate'. Before January 2012, this 'shared accommodation rate' also automatically applied to all single childless individuals aged under 25; from January 2012, that age threshold was raised to 35.

Before the reforms analysed in this paper, LHA rates were set equal to the median of private sector rents (not including those being rented by HB recipients) among properties of a similar size and in the same geographical area. As a result, the LHA rate that applied to a particular subsidy recipient should have been sufficient to cover the full rent of the median property rented by non-subsidy recipients in their area, of the size deemed appropriate for their family circumstances. If they rented a cheaper property than that, then subsidy recipients could effectively keep the first £15 a week of the difference.

The reform package studied in this paper had several elements. One element removed the weekly 'excess' of £15 that subsidy recipients could keep if their rent was less than their applicable LHA rate, so that the function became:

$$HB = \min(LHA \text{ rate}, \text{rent})$$

The other elements of the reform package affected the calculation of subsidy recipients' applicable LHA rates. These changes were:

- Setting LHA rates at the 30th percentile of local private sector rents among non-HB recipients (for the relevant property type) rather than at the median;
- Abolishing the 5-bedroom rates, so that large families previously entitled to this became entitled only to the 4-bedroom rate;
- Capping the rates at £250, £250, £290, £340 and £400 per week for the shared accommodation, 1-bedroom, 2-bedroom, 3-bedroom and 4-bedroom rates respectively (reducing rates below the 30th percentile of local rents in the highest-rent areas, which in practice means parts of inner London);
- Reducing the entitlement of most single adults without dependent children aged 25–34 to the amount for a room in a shared property (known as the Shared Accommodation Rate or SAR), rather than the rate for a 1-bedroom property.

The switch to the 30th percentile and the removal of the excess affected a wide group of subsidy recipients. The other changes affected only small subgroups. In our empirical analysis we look separately at those subgroups.

The removal of the £15 excess applied to new claimants from April 2011, and to existing recipients on their first annual claim anniversary after April 2011 (i.e. at some point between April 2011 and March 2012). The changes to the calculation of LHA rates applied to new claimants from April 2011 (at the same time as the excess removal); typically, they applied to existing recipients nine months after their first annual claim anniversary after April 2011 (i.e. nine months after the excess removal, at some point between January and December 2012). Where there has been a claim 'reassessment', the relevant anniversary is the anniversary of the most recent reassessment rather than the anniversary of the start date of the claim. For ease of exposition, the rest of the discussion abstracts from this and just refers to 'claim anniversaries'. More precisely, the nine-month interval was a period of "transitional protection" from the cuts to LHA rates. This protection could expire before end of those nine months, if a claimant had a change of circumstance which triggered a claim reassessment, such as a change in family type or a move to another area.

The changes that reduced the value of the LHA rate lower the cap on HB payments. They therefore give subsidy recipients an incentive to seek cheaper properties or to pay less for a given property, and the empirical issue that we explore in Section 6 is how much of the incidence then falls on landlords. The removal of the £15 excess has different effects on incentives. The pre-reform system (which allowed subsidy recipients to

keep £15 of any difference between their actual rent and their LHA rate) gave subsidy recipients an incentive to keep rent up to £15 below their LHA rate, either by choosing cheaper accommodation or by negotiating with landlords. Removing the excess means that subsidy recipients no longer have this incentive so, if they change their behaviour in response, we would expect them to choose more expensive types of accommodation or to accept a higher rent for a given property. Hence this change could effectively transfer the excess from tenants to landlords, rather than from either group to the taxpayer. There is no plausible mechanism by which it could lead to lower rents.

Appendix B: Data appendix

SHBE is made up of returns submitted to central government each month by local authorities (LAs). It includes monthly information on contractual levels of rent and characteristics of the subsidy recipients. Much of this information – including rent, location and family type – is required for the HB entitlement calculation and recipients are therefore legally required to keep it up to date.

B.1 Definition of key variables

The derivation of weekly contractual rents in the SHBE data is typically straightforward, using a combination of the rent amount reported and the periodicity that it is reported to cover (weekly, monthly, etc.).

Additional data cleaning was required in some cases where the periodicity was recorded as weekly when in fact it was monthly. This issue was almost exclusively confined to cases recorded by a single software provider (Civica) and for monthly records no later than early 2011. Misrecording is evident from the fact that average weekly rents in affected Local Authorities appeared to fall by approximately 75% in a single month when the issue was resolved. We corrected for this error by identifying subsidy recipients for whom, when comparing one month's record with the next, periodicity changed from weekly to monthly with no change to the reported rent. For such subsidy recipients we assume that the periodicity had always been monthly when reported weekly in prior months, and hence multiplied reported rents in prior months by (12/52) in order to convert them into weekly amounts. For the small number of Civica cases with periodicity recorded as weekly where the claim ended no later than early 2011 (specifically, where the last record of the claim is from a scan submitted before 1st March 2011), we record weekly rents as missing. This is because we know that these periodicities are relatively likely to be incorrect, but some will be correct (i.e. some subsidy recipients genuinely report weekly amounts), and we are unable to distinguish between the two without being able to observe a change in periodicity when the error was corrected.

We set rents to missing in four other circumstances:

- A joint tenancy is recorded and the software provider is Saf-ron/Camino, as there appears to be a tendency for the full rent for the dwelling to be recorded in such cases (rather than just the share of the rent for which the subsidy recipient is liable);
- Rent is recorded as zero;
- Dummy values (beginning 9999) appear to have been used for recorded rents;
- Periodicity is recorded as daily, as implied weekly rents tend to be very high in these cases.

Maximum weekly HB entitlements, ignoring non-dependent deductions, are known functions of rent and the applicable LHA rate. Where the excess 'rule' still applies, we define them as the minimum of the LHA rate and the rent plus £15. Otherwise, we define them simply as the minimum of the LHA rate and rent. We set maximum HB entitlement to missing in rare cases where the LHA rate is recorded as zero.

Analyses that use rent, maximum HB, or rent net of HB as the dependent variable are all conducted on the common sample for which all three of these variables are non-missing.

Data cleaning on other variables was also carried out where necessary. For example, certain local authorities at certain times incorrectly record whether or not subsidy recipients are in shared accommodation. Instances of this are identifiable from the fact that, in certain local authorities in certain months, a clear majority of subsidy recipients are recorded as residing in shared accommodation – with the proportion very close to the proportion of subsidy recipients in self-contained accommodation elsewhere. It seems clear that these cases have simply been recorded the wrong way round, and it is therefore straightforward to correct.

B.2 Sample selection – existing recipients

The basis for our analysis of all existing recipients is a random one-in-three sample of all HB recipients assessed under the LHA rules in January 2011, which is a sample of 283,574 subsidy recipients. Some individuals renting in the private sector were receiving HB assessed under a different set of rules: such recipients are disregarded in our analysis.

43,851 subsidy recipients are dropped from this sample because the point in time at which they would be affected by the reforms analysed cannot be robustly determined, leaving us with a final sample of 239,723 subsidy recipients. In the absence of behavioural response (which we do not incorporate in order to preserve the exogeneity of our treatment), the point at which a subsidy recipient was affected by the reforms was determined by the date of the last claim reassessment or claim anniversary in the year prior to April 2011 (or the date on which the claim began, if it began in the year prior to April 2011 and there had been no reassessment since). For full details on how this date is calculated, see [Brewer et al. \(2014\)](#). In short, there are three reasons why the point at which a subsidy recipient would have been affected can be impossible to determine robustly:

1. Some individuals whose claim began before April 2010 do not appear to have had any claim reassessments or anniversaries between April 2010 and March 2011, because their LHA rate remained constant throughout this period. For most of these individuals, it is therefore impossible to determine the anniversary of their claim. It is possible for a subsidy recipient's LHA rate after a claim reassessment or anniversary genuinely to be the same as their previous one. We can use publicly available LHA rates in different BRMAs over time to identify the subsidy recipients for which this was the case (and those subsidy recipients are not dropped).
2. Some subsidy recipients have large gaps in their records, because local authorities do not always submit scans every month. If a gap of more than 60 days occurs prior to the point at which we identify a subsidy recipient as having had their last claim reassessment or anniversary before April 2011, we are unable to calculate the date on which it occurred with sufficient accuracy.
3. Where an individual's claim has never been visibly reassessed, and they have not been dropped as a result of rule 1 (because their claim began after April 2010 or because a reassessment or anniversary during 2010–11 should not have changed their LHA rate), the point at which they will be affected (in the absence of behavioural response) depends on the start date of their claim. For some of these cases, the start date recorded in the SHBE data extract is not deemed sufficiently reliable, for one of the following reasons:
 - a. The start date recorded is more than three months earlier than the first observation we have for that individual;
 - b. The start date recorded is later than the first observation we have for that individual;
 - c. The start date is in April 2009, and the individual lives in one of a number of local authorities in which all start dates from 2008 to 2009 were reset to April 2009.¹⁹

¹⁹ These local authorities are Stockton-on-Tees, Gateshead, Blackpool, Rochdale, Fylde, Rushcliffe, South Staffordshire, Taunton Deane and Wrexham.

Table A1
Characteristics of example claim-cohorts in January 2011.

<i>Characteristic</i>	February claim-cohort	May claim-cohort	August claim-cohort	November claim-cohort
Household type				
Single man	27%	28%	30%	31%
Single woman	16%	15%	16%	16%
Couples without children	7%	6%	6%	7%
Single parents	33%	34%	32%	31%
Couples with children	17%	16%	16%	16%
Age				
Under 25	17%	14%	17%	18%
25–34	31%	31%	33%	32%
35–44	25%	26%	25%	25%
45–59	19%	20%	18%	19%
60 and above	7%	9%	7%	7%
Employment status				
In-work family	31%	33%	33%	32%
Out-of-work family	69%	67%	67%	68%
Neighborhood deprivation				
Bottom overall quintile	36%	33%	34%	32%
2nd overall quintile	26%	28%	26%	27%
3rd overall quintile	19%	19%	18%	19%
4th overall quintile	12%	13%	13%	12%
Top overall quintile	7%	7%	8%	7%
Mean rent	£130pw	£138pw	£135pw	£130pw
Mean HB entitlement	£122pw	£131pw	£127pw	£122pw
N	15,754	16,822	17,754	30,586

Source: authors' calculations using SHBE data.

When looking at the impact on the probability of moving home, we restrict the estimation sample to observations from April 2011 onwards because an individual's claim-cohort is mechanically related to moves prior to that date. We apply the same sample restriction when looking at the probability of living in shared accommodation, because the sample we select for that analysis excludes anyone living in shared accommodation prior to the start of the roll-out (since those already in shared accommodation were not affected by the extended coverage of the Shared Accommodation Rate).

Appendix Table A1 shows that four example claim-cohorts (February, May, August and November) look almost identical in terms of demographic characteristics, employment status, local-area deprivation, rents and housing benefit entitlements.

B.3 Sample selection – new claimants

For our analysis of new claimants, we ignore any SHBE records for LHA claims that had already started before the period of data used for analysis (i.e. before June 2010). For the records that remain – those of new LHA claims – we look at the circumstances of the claimant the first time that they were recorded. Since local authorities submit scans of their records once per month, this means that we extract the first monthly scan for each claim, and ignore all subsequent monthly scans.

One piece of data cleaning was required in order to ensure that we were defining new claims robustly. Scans from some Local Authorities have a tendency to include claim start dates that have been erroneously reset on a particular date, making the number of new claims appear larger than it really is in that Local Authority on that day and making the start dates of some existing claims appear more recent than they actually are. We were able to detect instances of this by identifying claims which appear to have started soon after (within six months of) a previous active claim by the same claimant, and looking at the proportion of apparent new claims in each Local Authority on each date which have

those characteristics. This proportion is far higher than normal in certain Local Authorities on particular days. Where the proportion exceeds 70% on a day in which at least five apparent new claims were made in a certain Local Authority, we conclude that any apparent new claim in that Local Authority on that day which shortly follows a previous active claim by the same claimant is likely to be erroneous. We therefore exclude such claims.

To guard against using information that did not genuinely apply at the beginning of a claim, we exclude from analysis claims for which the first monthly scan appears more than four months after the recorded start date of the claim. For example, if a claim is recorded as having started in January 2011, but the first scan of the relevant Local Authority's records which included that claim was submitted in or after June 2011, we would exclude this claim from the analysis.

Figs. 1 and 2 in the paper show a large rise in the rents and entitlements of those claiming HB in the run up to the April 2011 reform. Beatty et al. (2013) additionally shows that these spikes in rents and entitlements were accompanied by a large increase in the volume of claims, and these spikes in volumes, entitlements and rents can be explained by the financial incentives created by the way the reforms were rolled out. Because someone making a new claim just after 1st April 2011 would face the reformed, less generous, HB system, but someone starting a new claim just before 1st April would not face this system in full for another 21 months, claimants faced a strong incentive to make their claim before the cut-off date. Furthermore, the difference in the size of entitlements between the unreformed and reformed system was, in general, increasing in the size of (pre-reform) entitlement, so those with higher entitlements (and rents) faced a stronger incentive to claim before the cut-off date; this explains the change in the composition of new claims that lies behind the spike in mean entitlements and rents. For example, the proportion of new claims occurring in London rose by three percentage points between January and March 2011, from 14.3% to 17.3%; the same proportion did not fluctuate by more than one percentage point over any other two-month period in these data. Similarly,

Table A2
Demographic characteristics of new claimants.

Characteristic	June 2010 to November 2010 (% of claimants)	June 2011 to November 2011 (% of claimants)
Family type		
Single man	36%	35%
Single woman	19%	19%
Couples without children	7%	8%
Single parents	24%	25%
Couples with children	14%	15%
Age		
Under 25	23%	23%
25–34	33%	33%
35–44	23%	23%
45–54	13%	13%
55–64	6%	6%
65 and above	3%	3%
Employment status		
In-work family	29%	30%
Out-of-work family	71%	70%
Neighborhood deprivation		
Bottom overall quintile	33%	33%
2nd overall quintile	27%	27%
3rd overall quintile	19%	19%
4th overall quintile	13%	13%
Top overall quintile	8%	8%
N	336,486	334,093

Source: authors' calculations using SHBE data.

the average number of individuals in the household of new claimants rose from 1.86 to 1.95 between January and March 2011, also a larger fluctuation than over any other two-month period in the data.²⁰ The grey line on Fig. 2 plots mean residuals from a regression of rent on a

set of indicators for BRMA and the number of bedrooms, and shows that the spike in raw rents is largely (though not entirely) explained just by these two factors (it also shows that the decline in mean rents over this period was more than explained by changes in the composition of new claimants).

Appendix Table A2 shows, the pre- and post-reform claimants in our estimation sample are extremely similar with respect to observed characteristics.

References

- Beatty, C., Cole, I., Powell, R., Crisp, R., Brewer, M., Browne, J., Emmerson, C., Joyce, R., Kemp, P., Hall, S., Pereira, I., 2013. Monitoring the Impact of Changes to the Local Housing Allowance System of Housing Benefit. Department for Work and Pensions Research Report no. 838.
- 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 Research Report no. 871.
- Collinson, R., Ganong, P., 2018. How do changes in housing voucher design affect rent and neighborhood quality? *Am. Econ. J.* 10 (2), 62–89.
- Eriksen, M.D., Ross, A., 2015. Housing vouchers and the price of rental housing. *Am. Econ. J.* 7 (3), 154–176.
- Fack, G., 2006. Are housing benefits an efficient way to redistribute income? Evidence from a natural experiment in France. *Labour Econ.* 13 (6), 747–771.
- Gibbon, S., Manning, A., 2006. The incidence of UK housing benefit: evidence from the 1990s reforms. *J. Public Econ.* 90 (4–5), 799–822.
- Genesove, D., 2003. The nominal rigidity of apartment rents. *Rev. Econ. Stat.* 85 (4), 844–853.
- Grislain-Letremy, C., Trevien, C., 2014. The Impact of Housing Subsidies on the Rental Sector: The French Example INSEE WP G2014/08.
- Hills, J., 2007. Ends and Means: The Future Roles of Social Housing in England CASEreport 34. ISSN 1465-3001.
- Laferrière, A., Le Blanc, D., 2004. How do housing allowances affect rents? An empirical analysis of the French case. *J. Hous. Econ.* 13 (1), 36–67.
- Olsen, E.O., 2003. Housing programs for low-income households. In: Moffitt, R.A. (Ed.), *Means-Tested Transfer Programs in the United States*. NBER.
- Sayag, D., Zussman, N., 2015. The Distribution of Rental Assistance Between Tenants and Landlords: The Case of Students in Central Jerusalem. Bank of Israel DP 2015.01.
- Susin, S., 2002. Rent vouchers and the price of low-income housing. *J. Public Econ.* 83 (1), 109–152.
- Viren, M., 2013. Is the housing allowance shifted to rental prices? *Empir. Econ.* 44, 1497–1518.

²⁰ Both analyses are available on request.