The curious incidence of rent subsidies: evidence from administrative data[1]

MIKE BREWER*, JAMES BROWNE†, CARL EMMERSON†, ANDREW HOOD† AND ROBERT JOYCE†

**University of Essex and Institute for Fiscal Studies*

†Institute for Fiscal Studies

Abstract

This paper provides new evidence on the incidence of rent subsidies, making use of a package of reforms that cut the value of subsidies for about a million households in the UK. Using administrative panel data on the universe of subsidy recipients, and exploiting the phased roll-out of the reforms, about 90% of the incidence of the reforms is found to be on tenants - a considerably higher estimate than is found in most of the existing, small, literature. But we also find significant heterogeneity in the degree of incidence, with the incidence on tenants being substantially lower for groups previously subsidised to rent relatively expensive properties. We argue that this may reflect differences in the elasticity of housing demand, with potentially important implications for the optimal design of rent subsidies.

JEL classification: H22, H53, I38

Keywords: housing, incidence, subsidy

[1] Correspondence to robert_j@ifs.org.uk. 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 and the Work and Pensions Economic Group at the University of Sheffield, 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.

1. Introduction

Targeted demand-side subsidies for rented housing are a major and growing element of modern welfare states. Governments might therefore be concerned to know that, according to most of a small empirical literature, these subsidies result in higher rents such that a large portion of them are incident on landlords.

In this paper, we use administrative monthly panel data on the universe of housing benefit claimants in Great Britain, and exploit a natural experiment provided by a substantial package of cuts to housing benefit that was rolled out during 2011 and 2012, to estimate the incidence of those cuts. Overall, we find that a large majority of the incidence was on tenants, but we also uncover significant heterogeneity in the balance of incidence between tenants and their landlords. This heterogeneity may hold more general lessons for the determinants of the incidence of rent subsidies and hence their optimal design.

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 spends about 0.01% of GDP ($18 billion in 2014) subsidising the rents of 2.2 million families through Housing Choice Vouchers, the largest such federal program. Demand-side subsidies are substantially more significant 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 housing benefit.[2] These demand-side subsidies have increasingly taken the place of intervention through public housing projects, and rising rent levels have further increased their cost.[3] To the extent that rent subsidies raise the price of rented accommodation, governments are transferring resources to landlords rather than the intended recipients.

We estimate the incidence of a package of cuts to housing benefit in Great Britain, which were phased in over a period of nearly two years. The precise date at which existing claimants were affected depended, for the most part, on the calendar month in which their

[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 housing benefit in the private rented sector, the UK government subsidises tenants in public housing through a combination of sub-market rents and housing benefit. 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).

[3] The evolution of the US system is described in Susin (2002). Hills (2007) provides a detailed account of the UK case.

claim had started, while the reformed system applied to all new claims after a particular date.[4] We use a difference-in-differences design to estimate effects on rents for existing claimants, effectively using those not yet rolled onto the reformed system as a control group at each point in time. Separately, by looking at new claims we are able to assess (and rule out) the potential for general equilibrium effects to confound our analysis.

Our central estimates indicate that the rents of subsidy recipients, conditional on property characteristics, were little changed on average by the reforms, meaning that about 90% of the cuts to subsidies were incident on the tenants rather than their landlords. But we also uncover significant heterogeneity in the balance of incidence between tenants and their landlords. This leads us to the most plausible explanation for why our findings seem at odds with most previous work, and has potentially important implications for the design of rent subsidies. We find that the extent to which the cuts were incident on tenants was much lower among those tenants who were previously subsidised to rent some of the highest-rent properties in their area – less than two-thirds of the cut was incident on them. This subgroup is likely to be more similar to the group affected by previous UK reforms analysed in Gibbons and Manning (2006), the most similar of the previous papers to ours, and which found between one-third to 40 percent to be incident on tenants.

We argue that a likely explanation for the significant heterogeneity in incidence is differences in the elasticity of demand for rented housing. The pre-reform housing benefit system meant that certain groups were subsidised to consume a quantity of housing that was arguably high relative to their needs. This meant that their demand was relatively responsive when the subsidy was cut back: whilst basic housing is a necessity, at the margin it is likely to be a luxury good if housing consumption is already relatively high. As a result, more of the incidence of the subsidy cut was shifted to their landlords. This would have important policy implications: for a given overall cost to the government, a system with lower subsidies for a larger number of recipients would be more incident on tenants (i.e. lead to a smaller aggregate rent rise) than a system with higher subsidies for a smaller number of recipients.

[4] The precise month in which subsidy recipients were rolled onto the new system was also determined by the last time they moved house, or saw a change in their family composition, but is closely correlated with the calendar month in which claims began. This is discussed in more detail in Appendix B.

This is in contrast, for example, to the current system in the US, where less than a quarter of those with incomes low enough to be eligible receive any subsidy for their rent.[5]

As mentioned, we add to evidence about the incidence of rent subsidies provided by a number of previous studies. For the United States, Susin (2002) effectively compared rent trends between areas where housing voucher supply has been expanded by different amounts: he estimated that the existence of the voucher system had increased the rents of unsubsidised low-income households by 16%. As discussed by Olsen (2003), it is possible that these results were confounded by unobserved differences between areas, correlated with housing voucher supply, which caused the variation in rent trends. However, Collinson and Ganong (2015) exploited a plausibly exogenous change in county-level price ceilings for housing vouchers in 2005 and similarly concluded that the primary impact of raising the ceiling was to increase rents (this time using data on the rents of voucher recipients, rather than unsubsidised households), estimating that 89% of the incidence was on landlords. Eriksen and Ross (2015), on the other hand, look at the impact of area-varying voucher expansion in the early 2000s on the rents of individual housing units, controlling for unit fixed effects, and find no significant impact of the supply of vouchers on rents on average. They did find impacts on the distribution of rents that are consistent with "voucher recipients increasing their demand for higher-quality units after receiving the subsidy, but decreasing their demand for lower-quality units they would have occupied without the study".

A small body of evidence from outside the United States has tended to support the Susin finding that the incidence of rent subsidies is largely on landlords. For the UK, using data on voucher recipients only, Gibbons and Manning (2006) looked at a cut to housing benefit which applied to new claimants in the mid-1990s. They effectively compared the rent levels of new claimants and existing claimants after controlling for tenure length effects. They estimated that most (60% to two-thirds) of the cut was incident on landlords via reduced rents. For France, Laferrère and Le Blanc (2004), Fack (2006) and Grislain-Letremy and Trevien (2014) using different reforms and slightly different identification techniques all find that higher rent subsidies lead to significantly higher rents. Viren (2013) used Finnish panel data on rented properties and, having assumed that variation in trends in maximum allowable

[5] Collinson et. al (2015).

subsidies across areas is exogenous, estimated that one-third to one-half of a Finnish rent subsidy is incident on landlords. Sayag and Zussman (2015) study the introduction of a voucher scheme for students in central Jerusalem, finding that this increased the rents of both recipients and non-recipients, such that landlords captured four-fifths of the value of the grants.

The rest of the paper proceeds as follows. Section 2 outlines the reforms to housing benefit that form the basis of our estimates, describes our empirical strategy and the data we use. Section 3 presents some descriptive analysis of the data, and Section 4 presents our preferred estimates of the incidence of rent subsidies. Section 5 discusses our results and concludes.

2. Empirical approach and data

Our main estimate of the incidence of rent subsidies derives from an assessment of how a set of reforms to Housing Benefit in the UK in 2011 and 2012 affected the entitlements to HB and the rent paid by existing subsidy recipients. We first give brief details of the reform (with more details in Appendix A), and then discuss our empirical specification and the data used for estimation.

The reform

For subsidy recipients whose claim began in April 2008 or later, entitlement to housing benefit (before the means-test is applied) for those who rent from a private landlord was a function of actual rent and a cap known as the Local Housing Allowance (LHA) rate, given by $HB = \min(LHA\,rate, rent + £15\,per\,week)$. The applicable LHA rate varies geographically according to Broad Rental Market Areas (BRMAs)[6] and by the subsidy recipient's family type (which, as explained in Appendix A, determines the size of property that a recipient is deemed to be entitled to). If subsidy recipients rented a property whose rent was below their LHA rate, then they could effectively keep the first £15 per week of the difference.

The reform package we exploit to estimate the incidence of the rent subsidy had several elements. One element removed the weekly 'excess' of up to £15 that subsidy recipients could keep if their rent was below the LHA rate, so that HB entitlements were calculated

[6] There are 192 BRMAs in Great Britain.

simply as $\min(LHA\,rate, rent)$. The other elements of the reform package affected the calculation of the applicable LHA rate: all changes led to lower LHA rates, but the reductions were proportionately larger for large households (those previously deemed entitled to rent a 5-bedroom property), some of those in central London (who were the only ones affected by new nationwide ceilings), and single adults aged 25-34 without dependent children.

The removal of the £15 excess applied to new claimants from April 2011, and to existing claimants 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 claimants nine months after their first annual claim anniversary after April 2011 (i.e. at some point between January and December 2012).

The changes that reduced the value of the LHA rate lowered the cap on housing benefit 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 Sections III and IV is how much of the incidence then falls on their 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. If the latter, 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.

Empirical specification

For our main estimates of the extent to which entitlements to HB and rents changed as a result of the reforms, we exploit the between-cohort variation in when the reforms took effect: the nature of the roll-out means that otherwise-identical individuals observed at the same point in time but whose claim anniversary fell in a different calendar month faced different levels of subsidy. We implement this with a difference-in-differences (DiD) style, timing-of-events linear regression specified as:

$$y_{iact} = f_a(t) + \pi_c + z_\tau \beta + X'_{iact}\alpha + \varepsilon_{iact} \ , \tag{1}$$

where i indexes subsidy recipients who are observed at multiple points in time t, live in area a, and are members of cohort c, defined by the calendar month in which their annual claim anniversary falls.

The main outcome variables are rent, housing benefit entitlement and the difference between the two (in £s per week).[7] The key variables of interest, z_τ, are a 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. These dummies run from twelve months before the point of transition onto the reformed system (i.e. three months before the loss of excess) to eleven months after. In this way, we allow the reform's impact to evolve as subsidy recipients approach the point of transition through anticipation effects, and through gradual adjustment as they spend longer under the reformed system. The variables z_τ vary at the cohort-time level, and so we also include a set of twelve cohort (i.e. month-of-claim-anniversary) fixed effects, π_c, and control for time in $f_a(t)$ with a full set of monthly dummy variables plus a linear trend in each BRMA (buffeted by the national-level monthly shocks).

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). This is so that we do not confound changes in property choices with changes in the unit-price of accommodation. These are, of course, imperfect measures of housing quality; if subsidy recipients adjust along other (unobserved) margins of property characteristics, then we will pick this up as a price change rather than a quality change, and

[7] 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 (Beatty et al, 2014). 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 rather than the contractual rent.

this would lead us to over-estimate the incidence on the subsidy on landlords and under-estimate the incidence on tenants.

We also control for family type and age, as these may change over time for reasons unrelated to the reform in ways that are not adequately captured by our time trends; as shown in the next section, these demographic controls make a negligible difference to our estimates. We estimate equation (1) by OLS. Estimated standard errors allow for heteroskedasticity and for errors clustered at the BRMA level.[8]

Threats to identification

The usual 'common trends' assumption in a DiD design is that, in the absence of the reforms, trends in the outcome variables would have been unrelated to treatment status. In our context, this means that there should be no systematic, unobserved, time-varying differences between subsidy recipients in different cohorts (i.e. whose claim began in different months of the year). As always this is untestable, but the inclusion of cohort fixed effects has negligible impacts on our estimates[9], which increases confidence that the cohorts are very similar.

A further assumption required for identification in a DiD design is that the treatment does not affect the outcomes of the control group. In our context, this assumption might not hold if there were general equilibrium effects. Our estimates for existing claimants effectively use subsidy recipients not yet rolled onto the reformed system as a control group at each point in time. Because this group is defined only by the calendar month in which their housing benefit claim began, it is likely to be highly comparable to the group already being treated – a useful feature in that common trends is likely to hold, but potentially problematic in that different cohorts are surely part of the same market. In a rental market with no frictions and with perfect competition, there would be a single (quality-adjusted) rental price at all times. In such a world, the rents paid by all cohorts would change at the same time, regardless of whether they had been affected by the cuts in subsidy, and estimation of equation (1) would

[8] When looking at additional outcomes (see Table 7) we use a slightly different set of control variables. We remove controls for housing quality (as the quality adjustment is not necessary for other outcomes) and replace them with initial BRMA fixed effects. When the outcome is whether the claimant moved house, we include controls for estimated rental contract and claim anniversaries, as one might expect these to have a (potentially confounding) effect on the probability of moving. In addition, some of the additional outcomes examined in Table 7 are binary: for these we use probit estimation and report marginal effects at the mean.

[9] Results isolating the impact of adding the cohort fixed effects are available on request.

therefore find that the reform had no impact of rents, regardless of the true impact. More generally, general equilibrium effects of a less extreme nature would attenuate our estimates of the extent to which the reforms were incidence on the landlords of existing claimants. This issue was recognised by Gibbons and Manning (2006). They also used a control group of existing claimants when looking at how cuts to rent subsidies affected rents of new claimants, and noted: "our methods would seriously underestimate the impact of the reforms if the rental housing market is perfectly competitive and frictionless, when any reduction in benefit payments to one claimant group drives down rents throughout the rental sector".

In reality, though, the rental market is not characterised by spot prices: tenants' rents typically change at no more than annual frequency (and sometimes cannot be changed at all within-year). Of course, given this infrequency of changes, it is possible that subsidy recipients or their landlords might have reacted in advance of the actual date on which a recipient's subsidy was cut, but our specification will explicitly capture any adjustments that happen in the 12 months before tenants are rolled onto the fully-reformed system. Furthermore, Beatty et al, (2013) presents evidence of limited awareness of the reforms in question on the part of both tenants and landlords before entitlements were reduced, which suggests that adjustments were more likely to take place once claims were actually being assessed under the new system than before.

However, we can relax this assumption of no general equilibrium results by looking at the impact of the reforms on the flow of new recipients of housing benefit (HB), who were affected by the reformed, less generous housing benefit system in full from April 2011. To implement this, we extract information on the circumstances of subsidy recipients that applied when a claim for HB began. The equation we estimate for the new subsidy recipients is effectively a before-after design, or a regression discontinuity design with time as the running variable. However, the exclusion of a window of data around the discontinuity (which we justify below) means that in effect we do not compare recipients who are arbitrarily close to the discontinuity but located on different sides, as is usual in RDDs. It is specified as follows:

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

As before, i indexes new subsidy recipients observed to make a claim at time t, who live in area a.[10] In the terminology of an RDD, the running variable is the time that the recipient claimed HB, t, and the coefficient of interest is β, as the treated group are those who claimed HB after April 2011. Based on initial descriptive analysis (shown in Appendix Figures C1 and C2) we specify $f_a(t)$ as a BRMA-specific linear trend that is allowed to vary between the pre- and post- reform periods. As before, 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. Equation (2) is estimated using OLS, and standard errors are robust to heteroskedasticity and clustered at the BRMA level.

Data and sample selection

We use administrative monthly panel data on the universe of live housing benefit claims in Great Britain, from the Single Housing Benefit Extract (SHBE). SHBE is made up of returns submitted to central government each month by local authorities (LAs).[11] It includes monthly information on contractual levels of rent and characteristics of the subsidy recipients. Details of the construction of key variables, including data cleaning, are given in Appendix B.

Our main estimates are based on the sample of individuals receiving HB assessed under the LHA rules in January 2011, shortly before the reforms were implemented.[12] We use monthly observations for these subsidy recipients between January 2010 and November 2013 inclusive.[13] As described in Section 2, these claimants 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.[14] More precisely, the nine-month interval was a period of "transitional protection" from the reforms (other than the removal of the excess). This protection expired immediately if a claimant had a change of

[10] Our main analysis measures time in months. However, although each LA submits a scan of its records only on one day each month, the exact start date of an HB claim is part of each record, and so we can observe the characteristics of new claims for LHA in continuous time

[11] Our analysis therefore excludes Northern Ireland, although it was affected by the reforms in question.

[12] Some individuals renting in the private sector were receiving HB assessed under a different set of rules: such recipients are disregarded in our analysis.

[13] We restrict our sample to observations from April 2011 onwards when looking at some of the additional outcomes in Table 7, for reasons discussed in Section 4.

[14] Where there had 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'.

circumstance which triggered a claim reassessment, such as a change in family type or a move to another area. To avoid endogeneity, we therefore allocate recipients to cohorts on the basis of when they would have been affected by the roll-out had there been no changes in circumstances from April 2011.

For computational reasons, a random one-in-three subset of subsidy recipients was taken for most of the regression analysis in Section 4. After dropping a further 15% of the sample because they are missing important information, we are left with an estimation sample of 239,723 subsidy recipients, observed on average for 28 months. Table 1 shows basic demographic characteristics of subsidy recipients in January 2011, both for the universe and for our final estimation sample: the two groups are almost identical in terms of these key characteristics.

TABLE 1

Demographic characteristics of existing claimants in January 2011

Characteristic	*Full SHBE sample (% of claimants)*	*Estimation sample (% of claimants)*
Household type		
Single man	28.9	29.2
Single woman	15.6	15.6
Couples without children	6.4	6.3
Single parents	32.7	32.4
Couples with children	16.4	16.3
Age		
Under 25	16.1	16.5
25-34	31.6	31.8
35-44	25.2	25.2
45-59	19.0	18.9
60 and above	8.2	7.6
N	850,249	239,723

For the robustness check using the impact of the reforms on new claimants, we use the new claims between June and November 2010 (giving us a sample of 336,486 observations from before the reforms took place) and between June and November 2011 (giving us a sample of 334,093 observations from after the reform took place). Table 2 shows the similarity of these two samples with respect to basic demographic characteristics. We exclude all new claims between December 2010 and May 2011 to avoid the contamination of our estimates by behavioural responses in the timing of claims around the introduction of the reforms in April 2011 (see Appendix C for more details).

TABLE 2

Demographic characteristics of new claimants

Characteristic	*June 2010 to November 2010 (% of claimants)*	*June 2011 to November 2011 (% of claimants)*
Family type		
Single man	35.7	34.5
Single woman	18.8	18.7
Couples without children	7.4	7.6
Single parents	24.3	24.7
Couples with children	13.7	14.7
Age		
Under 25	23.2	22.6
25-34	32.8	33.1
35-44	22.8	22.8
45-54	13.0	13.4
55-64	5.6	5.6
65 and above	2.6	2.7
N	336,486	334,093

3. Descriptive analysis

Figure 1 illustrates the nature of the phased roll-out by cohort by showing the (covariate-adjusted) average weekly housing benefit entitlements of existing subsidy recipients by calendar month.[15] The "All" line shows entitlements gradually declining between April 2011 and December 2012, the period over which recipients were migrating to the new system.[16] This gradual decline when averaged across all existing recipients is the result of successive cohorts being sequentially rolled onto the new system and seeing sudden reductions. Two example cohorts are shown in the figure. The August cohort (i.e. those whose claim anniversary fell in the month of August) saw entitlements fall in August 2011, when they lost any excess. They then saw a further fall nine months later, in May 2012, as their transitional protection expired and they were affected by the rest of the reform package. The same pattern holds for the November cohort but with the drops in entitlements occurring three months later.

FIGURE 1

Average maximum entitlement of existing claimants by month

(Residual from regression of entitlements on BRMA and number of bedrooms, £pw)

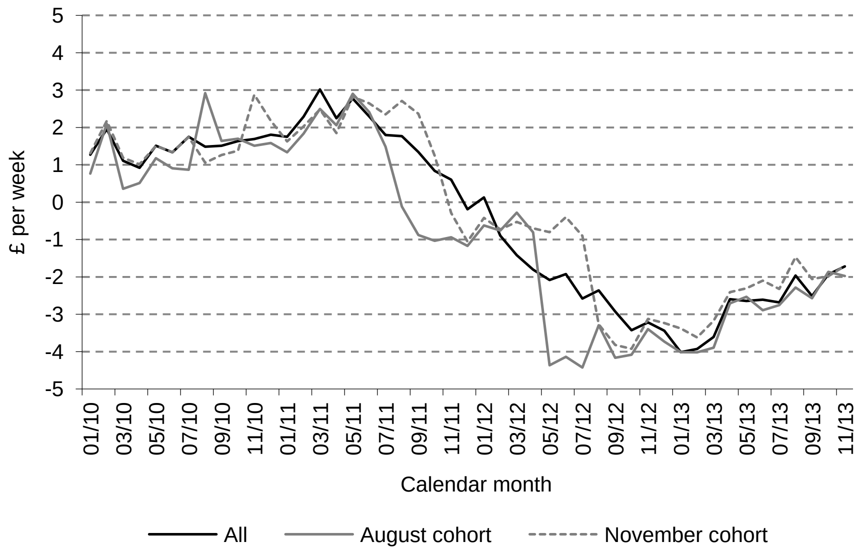

[15] Specifically, we plot mean residuals from a regression of LHA entitlements on all interactions between BRMA and number of bedrooms.

[16] Existing claimants began to lose their excess from April 2011, but if there had been no changes in circumstance none of the other reforms would have affected existing claimants until January 2012.

Figure 2 pools all cohorts but plots how entitlements changed not by calendar month, but by the number of months since being rolled onto the new system. We define month "0" as that in which subsidy recipients are rolled fully onto the new system i.e. the month in which transitional protection expires; month "-9" is therefore the point at which any excesses are removed. It also shows the equivalent for rents (again, having partialed out interactions between BRMA and property size). The Figure shows both that large changes in entitlement occurred at month -9 and month 0, and also that these reductions in housing benefit entitlements seems to have had little impact on rents being paid, i.e. the incidence seems to have been largely on tenants. The next section confirms this formally with our regression results.

FIGURE 2

Average maximum entitlements and rents of existing claimants by months since main impact

(Residual from regression on BRMA and number of bedrooms, £pw)

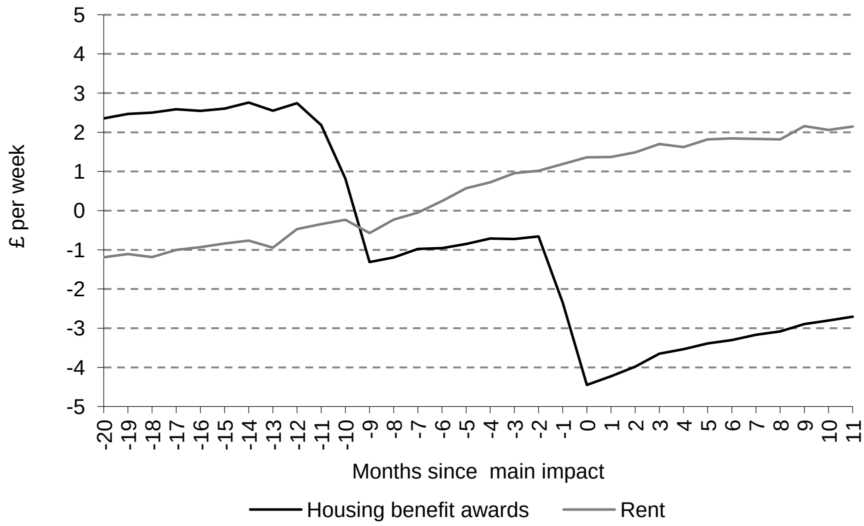

4. Regression results

Main results

Table 3 presents estimates based on equation (1) for the effect of the reforms on rents, entitlement to housing benefit and the difference between them. As one moves from left to right, we add regressors to build up to our preferred specification. For each specification we show three sets of coefficients, capturing impacts at three different stages: the point at which subsidy recipients reach their first annual claim anniversary after April 2011, at which any excess they had was removed; the point nine months later when they were fully rolled onto the new system and subject to all other elements of the reform package; and eleven months after that, which is the latest point in our data at which we observe all cohorts. Standard errors, which are robust to clustering at the BRMA level and heteroskedasticity, are in parentheses.

Column 1 shows the results from regressions of housing benefit, rent, and the difference between the two, on just a 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 time trends, rents increase as the reform is rolled out, while housing benefit entitlements fall. Columns 2 and 3 add controls for housing quality (local area, number of bedrooms and their interaction, along with a measure of local-area deprivation), showing that the rise in rents over time is smaller once we control for a rise in quality. Column 4 adds the time trends and cohort fixed effects: once these are included, we estimate the impact of the cut in subsidy to be a small fall in quality-adjusted rents. Finally, our preferred specification (column 5) adds controls for family type and age, on the basis that they might be changing over time for reasons unrelated to the reforms. This makes little difference to the estimates.

Focussing on our preferred specification, claimants lost an average of about £5 per week as a result of the removal of excesses, with around 40% of claimants affected. At this point their rents were also reduced slightly by about £0.80 per week, which is statistically significant. As discussed in Section 2, we would not expect landlords to lose out from the excess removal in isolation: there is no reason why this should result in lower rents, and it might even increase them. Instead, we attribute this fall in rents to an anticipatory adjustment to the remaining, impending reforms (given, as discussed in Section 2, the likely rigidity of rents mid-contract and the fact that tenancy agreements tend to last for at least one year). Indeed for tenants whose claim anniversary coincides with the anniversary of their tenancy, this point in time would be the last opportunity to change rents before the reforms took effect without renegotiating mid-tenancy.

TABLE 3

Estimated impact of cuts to housing benefit on existing claimants, £/wk

		Model				
		(1)	(2)	(3)	(4)	(5)
Loss of excess	Housing benefit	-0.72	-2.09***	-3.40***	-4.90***	-4.98***
		(0.73)	(0.31)	(0.41)	(0.44)	(0.42)
	Rent	3.73***	2.44***	0.93***	-0.74**	-0.81***
		(0.74)	(0.27)	(0.31)	(0.29)	(0.27)
	Rent net of HB	4.46***	4.53***	4.33***	4.17***	4.17***
		(0.23)	(0.24)	(0.27)	(0.34)	(0.34)
Point of main impact	Housing benefit	-4.12**	-4.65***	-6.67***	-8.31***	-8.31***
		(1.93)	(0.84)	(0.85)	(1.09)	(1.01)
	Rent	5.17***	4.81***	2.78***	-0.65	-0.73
		(1.43)	(0.47)	(0.40)	(0.78)	(0.68)
	Rent net of HB	9.29***	9.46***	9.45***	7.66***	7.58***
		(0.71)	(0.79)	(0.79)	(0.84)	(0.85)
11 months after main impact	Housing benefit	-1.35	-1.64*	-5.04***	-7.08***	-6.84***
		(2.02)	(0.99)	(0.82)	(1.07)	(0.92)
	Rent	6.97***	6.83***	3.43***	-0.79	-0.79
		(1.84)	(0.83)	(0.56)	(1.19)	(1.09)
	Rent net of HB	8.32***	8.47***	8.47***	6.29***	6.06***
		(0.39)	(0.45)	(0.48)	(0.80)	(0.83)
	N	239,576	239,576	239,094	239,094	238,782

Note: *** Statistically significant at 1% level, ** Statistically significant at 5% level, * Statistically significant at 10% level. Standard errors in brackets are robust to heteroskedasticity and clustering at the BRMA level. Figures given in UK pounds per week. Model (1) contains only a series of dummy variables corresponding to the number of months since being rolled onto the new system; model (2) adds controls for BRMA and LA; model (3) adds controls for the number of bedrooms in the property (shared accommodation, 1 bedroom, 2 bedrooms, 3 bedrooms, 4 bedrooms, 5 or more bedrooms); interaction terms that capture all possible combinations of number of bedrooms and BRMA; and local-area deprivation; model (4) adds linear time trends for each BRMA, a full set of month dummy variables and 'cohort' fixed effects; and model (5) adds joint controls for family type and age. 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).

The loss of housing benefit rose to £8.30 per week nine months later at the point when the rest of the reforms took effect. There was essentially no further impact on rents, meaning that about £7.60 per week (90%) of the cuts were incident on tenants at this point. One might expect this to be partly due to short-run rent rigidity. However, the proportion of the cut estimated to be incident on tenants remained very stable, at around 90%, eleven months later. In fact, the estimated effect on rent (and hence the estimated incidence on landlords) is not statistically significantly different from zero. We obtain similar results for a wide range of demographic and geographic subgroups (including age, family type and region).[17]

It is important to recall that these estimates control for local area and number of bedrooms. As noted in Section 2, the proportion of the cuts to subsidies incident on tenants would be even higher than our results suggest if tenants adjust to the reforms partly by choosing lower quality properties in ways not captured by our controls.

Robustness to general equilibrium effects: estimates for new claimants

As Section 2 discussed, estimates of the incidence of housing benefit that are based on new claimants are robust to the existence of general equilibrium effects, and therefore provide a first robustness check on our preferred estimates. Table 4 shows the estimates for new claimants, based on equation (2). Again, the different columns build up to our preferred specification by gradually adding regressors.

Focussing on our preferred specification (column 5), we estimate that the reforms reduced housing benefit 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 this – an estimated £7.80, or 95% – was incident on the tenants. Again, we obtain similar results for a wide of demographic subgroups.[18] The fact that our analysis of new claimants provides a very similar estimate of the incidence of the reforms to our main results suggests the small estimated impact on rents is not simply the result of general equilibrium effects.

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

[18] These are reported in full in the interim impact evaluation of these reforms carried out by some of the present authors (Beatty et. al., 2013).

TABLE 4

Estimated impact of cuts to housing benefit on new claimants, £/wk

	Model				
	(1)	(2)	(3)	(4)	(5)
Housing benefit	-9.28***	-7.73***	-6.36***	-7.87***	-8.21***
	(1.18)	(0.63)	(0.48)	(0.52)	(0.50)
Rent	-1.57	0.12	1.62***	-0.21	-0.46
	(1.11)	(0.52)	(0.43)	(0.66)	(0.64)
Rent net of HB	-7.71***	7.85***	7.97***	7.66***	7.76***
	(0.34)	(0.35)	(0.35)	(0.49)	(0.49)
N	667,278	667,278	662,764	662,764	659,892

Note: *** Statistically significant at 1% level, ** Statistically significant at 5% level, * Statistically significant at 10% level. Standard errors in brackets are robust to heteroskedasticity and clustering at the BRMA level. Figures given in UK pounds per week. Model (1) contains only a post-reform dummy variable; model (2) adds controls for BRMA and LA; model (3) adds controls for the 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; model (4) adds linear time trends for each BRMA, which are allowed to differ before and after the reform; and model (5) adds joint controls for family type and age. 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).

Heterogeneity in incidence

In this subsection we look at subsidy recipients with pre-reform characteristics that meant they would be affected by specific elements of the reform package as long as their circumstances remained the same. Estimating the model in equation (1) for different groups of claimants, defined according to pre-reform circumstances, we find substantial heterogeneity in the results between these subgroups and the population as a whole. We argue that this may be informative about why the results presented in the previous section differ from those elsewhere in the literature, with implications for the optimal design of rent subsidies.

Table 5 shows our estimates for the incidence of the reform package as a whole on those who stood to lose the full £15 excess (we define this group as those whom, in January 2011, had a weekly rent level at least £15 lower than their LHA rate). As discussed in Section 2, there is no reason to expect the removal of the £15 excess to reduce rents, and it might increase them: all else equal, this would suggest that tenants in this group should bear a greater share of the incidence of the reform package, relative to their landlords, than other tenants. However, as this is a group paying relatively low rents, one might also expect these recipients to be selected on important unobserved characteristics, such as their preferences between housing and other consumption, and their willingness and ability to negotiate. The results in Table 5 show that, at the point when subsidy recipients lost their excess, the falls in both housing benefit and rents were slightly larger for this group than the population of subsidy recipients as a whole, but that the share of the cuts to housing benefit incident on the tenants was similar for this group as for the population as a whole. Estimates for later periods, when the other elements of the reform package applied, show that, if anything, a smaller share (but still a majority) of the cuts were incident on the tenants than in the subsidy recipient population as a whole, with the reform leading to falls in (quality-adjusted) rents that are statistically significant from zero. This may indicate that this selected group of individuals were more responsive to the cuts than others, perhaps because their housing demand is more elastic, or because they are more pro-active negotiators.

TABLE 5

Estimated impact of cuts to housing benefit on existing claimants with a £15 excess in January 2011, £/wk

	Housing benefit	Rent	Rent net of HB
Loss of excess	-12.44***	-1.20***	11.24***
	(0.52)	(0.24)	(0.37)
Point of main impact	-14.80***	-3.13***	11.68***
	(1.09)	(0.56)	(0.91)
11 months after main impact	-14.31***	-3.59***	10.72***
	(0.80)	(0.79)	(0.79)

Note: *** Statistically significant at 1% level, ** Statistically significant at 5% level, * Statistically significant at 10% level. Standard errors in brackets are robust to heteroskedasticity and clustering at the BRMA level. Figures given in UK pounds per week. Includes controls for BRMA, local authority, number of bedrooms in the property, local area deprivation, 'cohort', calendar month, linear time trends in each BRMA, and joint controls for 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). N = 134,246.

Table 6 shows separate estimates for three additional subgroups, defined according to their January 2011 characteristics: single adults without dependent children due to be aged 25-34 at the point that the LHA reforms took effect; large families entitled to the 5-bedroom LHA rate; and recipients living in one of five London BRMAs in which the overall nationwide caps on LHA rates bind. For brevity, we report just estimated impacts 11 months after being fully rolled onto the reformed system. For each group, the first row (labelled 'quality-adjusted') shows estimates with controls for contemporaneous property characteristics analogous to those presented in Tables 3 and 4. The second row (labelled 'unadjusted') shows estimates without those controls (but with a control for initial BRMA, based on circumstances in January 2011).

TABLE 6

Estimated impact of cuts to housing benefit on existing claimants likely to be affected by certain elements of the reform package (11 months after main impact) , £/wk

		Housing benefit	Rent	Rent net of HB	N
Increased scope of shared accommodation rate	Quality-adjusted	-13.05*** (1.36)	-4.80*** (1.31)	8.25*** (1.73)	49,569
	Unadjusted	-15.55*** (1.59)	-7.36*** (1.55)	8.18*** (1.78)	49,635
Abolition of 5-bedroom LHA rate	Quality-adjusted	-29.21*** (8.49)	-11.69** (5.48)	17.52*** (5.44)	5,699
	Unadjusted	-31.60*** (9.99)	-19.04** (9.27)	12.56** (5.36)	5,703
National caps on LHA rates	Quality-adjusted	-41.93*** (9.96)	-5.68 (10.19)	36.25*** (12.31)	16,992
	Unadjusted	-48.48*** (12.59)	-17.07 (14.20)	31.41*** (12.13)	16,992
All existing claimants	Quality-adjusted	-6.84*** (0.92)	-0.79 (1.09)	6.06*** (0.83)	238,782
	Unadjusted	-7.40*** (0.91)	-1.11 (1.12)	6.28*** (0.84)	239,279

Note: *** Statistically significant at 1% level, ** Statistically significant at 5% level, * Statistically significant at 10% level. Standard errors in brackets are robust to heteroskedasticity and clustering at the BRMA level. Figures given in UK pounds per week. "Adjusted" figures include controls for BRMA, local authority, number of bedrooms in the property, local area deprivation, '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.

Under the pre-reform system, single adults without dependent children aged 25-34 would have been able to claim the 1-bedroom rate of LHA, but after the reforms they could claim only a rate deemed sufficient for them to rent a room in shared accommodation. Conditional on property characteristics, they lost an average of about £13 per week in housing benefit from the reforms; but we estimate that their rents fell by about £4.80 per week, implying that just over one third of the incidence was on their landlords. The estimates that do not adjust for property characteristics show larger falls in both housing benefit entitlements and rents, suggesting that some of the individuals affected responded by moving to cheaper properties – something we show directly below.

Families who were entitled to the 5-bedroom LHA rate in January 2011 lost an average of about £29 per week in housing benefit entitlement from the reforms (as, post-reform, they were deemed to require only a 4-bedroom property); but we estimate that their rents fell by almost £12 per week, implying that about 40% of the incidence was on their landlords. Again, comparison with the estimates that do not adjust for property characteristics suggests that these subsidy recipients may also have responded by living in cheaper types of properties than they would otherwise have done.

Finally, subsidy recipients who, in January 2011, were living in one of five London BRMAs in which the overall national caps on LHA rates bind lost an average of about £42 per week in housing benefit 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 quality adjustments (but neither the adjusted nor unadjusted rent changed by a statistically significant amount).

Table 7 explores some of the behavioural responses of these subgroups, looking at whether they moved to a different property in response to the reforms, and, if so, to what type of accommodation. For binary outcomes, a probit specification is used and marginal effects at the mean are reported. When looking at the impact on the probability of moving home, we restrict the estimation sample to observations from April 2011 onwards. This is because an individual's cohort is mechanically related to moves prior to that date.[19] We apply the same sample restriction when looking at the probability of living in shared accommodation, because we exclude those not living in shared accommodation prior to the reforms (since they were not affected by the extended coverage of the Shared Accommodation Rate).

[19] For example, individuals in the April cohort cannot have moved house between April 2010 and March 2011 (as moving house would trigger a reassessment of their claim, changing their cohort).

TABLE 7

Estimated impact of cuts to maximum entitlements on property choices of existing claimants 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*** (0.5)	43,655
	Probability of living in shared accommodation (ppts)	17.0*** (0.4)	43,564
Abolition of 5-bedroom LHA rate	Probability of moving (ppts per month)	0.6 (0.8)	5,406
	Number of bedrooms	-0.14 (0.16)	5,700
National caps on LHA rates	Probability of moving (ppts per month)	0.8* (0.6)	16,163
	Probability of moving out of capped area (ppts per month)	0.3 (0.5)	16,163
All existing claimants	Probability of moving (ppts per month)	0.4*** (0.2)	219,592

Note: *** Statistically significant at 1% level, ** Statistically significant at 5% level, * Statistically significant at 10% level. Standard errors given in brackets are robust to heteroskedasticity and clustering at the BRMA level. Figures given in UK pounds per week. When the outcome variable is probability of moving, results are from a probit regression including controls for BRMA in January 2011, calendar month, 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.

The estimates suggest that, after 11 months under the new system, each of these subgroups was about 1 percentage point more likely to move house in that month as a result of the reforms.[20] The group affected by the extended coverage of the Shared Accommodation Rate were 13 percentage points more likely to be living in shared accommodation. There is also some evidence that those affected by the abolition of the 5-bedroom rate responded by renting smaller properties, and that those affected by the national caps responded by moving out of the capped areas, although these effects are not statistically significant.

5. 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 previously unexploited administrative monthly

[20] The impact is stable, at about 1 percentage point a month, from five months after being rolled onto the reformed system.

panel data on the universe of subsidy recipients, and exploiting the phased roll-out of the reforms, we estimate that, on average, about 90% of the incidence of these cuts to UK housing benefit was on tenants. But there was significant heterogeneity and, for two groups affected much more than average by the cuts, we estimate that less than two-thirds of the cut was incident on them (and we can reject the possibility that it was all incident on them).

Why is this? Theory suggests that the incidence of rent subsidies on landlords should be higher when demand is more elastic, and we consider that variation in elasticities may explain our findings. 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 housing benefit entitlements are cut. We have found 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, and hence highest-rent, 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.

An alternative explanation for our results might be that the groups for which a smaller proportion of the cut was incident on them (and a larger proportion on their landlords) were hit harder than average by the changes, and this simply made them quicker to notice and to make an effort to respond. If this were the key explanation, it would suggest that the small effects on rents for other subsidy recipients simply reflect short run rigidity. However, there are good reasons to doubt this. Our results show that tenants affected by the national caps lost by far the *most* housing benefit of the subgroups considered, and yet the estimated incidence on them is high (and we cannot reject that it was 100%). There seems no obvious reason why they should adjust more slowly as a group. In addition, previous studies of changes to rent subsidies, which have tended to find higher incidence on tenants than this paper, have also either typically looked either at near-contemporaneous responses of rents to subsidies (e.g.

Viren, 2013) or have shown that effects on rents occurred quickly (Gibbons and Manning, 2006).

These findings may explain why our main result – that the incidence of changes to rent subsidies were so heavily incident on tenants – is at odds with the findings of most previous empirical estimates. The only other paper to look at the incidence of UK housing benefit is Gibbons and Manning (2006). They studied reforms in the mid-1990s that introduced caps on the size of rents eligible for housing benefit, based on average market rents in the local area. As such, the subsidy recipients who were directly affected by the mid-1990s reforms were paying rents above the local average; 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, rather than a relatively high-rent minority who might be able to substitute more easily towards cheaper accommodation. Where we do focus on a subgroup whose ability to live fully subsidised in some of their area's highest-rent properties was removed or restricted, our results on incidence are closer to those of Gibbons and Manning.

If the heterogeneity in the incidence of the cuts that we find is indeed explained by heterogeneity in demand elasticities, it would have important implications for the likely impacts of other reforms to rent subsidy regimes, and for the optimal design of such regimes. If, as cuts to subsidies bite further down the distributions of rent and housing quantity and/or quality, the average demand elasticity of affected tenants falls, then less generous subsidies will tend to be proportionately more incident on tenants. This would have important implications for the optimal design of rent subsidies in general: given a fixed level of total expenditure, a system with lower subsidies for a larger number of recipients would lead to a larger share of the subsidy being incident on tenants than a system with higher subsidies for a smaller number of recipients. For example, our findings suggest a potential source of inefficiency in the current system of rent subsidies in the US: given that currently less than a quarter of households with income low enough to qualify receive a rent subsidy[21], the share of that subsidy lost to landlords could be cut by spreading it more widely across low-income households.

[21] Collinson et. al. (2015).

References

Beatty, C., Cole, I., Powell, R., Crisp, R., Brewer, M., Browne, J., Emmerson, C., Joyce, R., Kemp, P., Hall, S. and 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. and 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., Ellen, I.G. and Ludwig, J. (2015) "Low-income housing policy", NBER Working Paper 21071, available at http://www.nber.org/papers/w21071.

Collinson, R. and Ganong, P. (2015) "The Incidence of Housing Voucher Generosity". http://dx.doi.org/10.2139/ssrn.2255799.

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.

Browne, J. and Hood, A. (2015) "Options for further cuts to social security" in Emmerson, C., Johnson, P. and Joyce, R. eds. (2015) *The IFS Green Budget*, available at http://www.ifs.org.uk/uploads/gb/gb2015/ch9_gb2015.pdf.

Eriksen, M.D., and Ross. A. 2015. "Housing Vouchers and the Price of Rental Housing." *American Economic Journal: Economic Policy*, 7(3): 154-76.

Fack, G. (2006) "Are Housing Benefits an efficient way to redistribute income? Evidence from a natural experiment in France". *Labour Economics,* 13 (6), pp. 747-771.

Gibbon, S. and Manning, A. (2006) "The Incidence of UK Housing Benefit: Evidence from the 1990s Reforms". *Journal of Public Economics,* 90 (4-5), pp. 799-822.

Grislain-Letremy, C. and 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.

Laferrère, A. and Le Blanc, D. (2004), "How do housing allowances affect rents? An empirical analysis of the French case". *Journal of Housing Economics*, 13 (1) pp.36-67.

Olsen, E. O. (2003) "Housing Programs for Low-Income Households" in Moffit, R. A (ed.) *Means-Tested Transfer Programs in the United States*, NBER.

Sayag, D. and 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". *Journal of Public Economics.* 83 (1), pp. 109-152.

Viren, M. (2013) "Is the housing allowance shifted to rental prices?" *Empirical Economics* 44, pp. 1497-1518.

Appendix A: Policy background

In 2015–16, spending on housing benefit in the UK is projected to be £24.5 billion: 12% of all government spending on cash transfers.[22] £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).[23] Spending on housing benefit 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.[24] 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.

For subsidy recipients who rent from a private landlord and whose claim began in April 2008 or later, housing benefit entitlement is a function of actual rent and a cap known as the Local Housing Allowance (LHA) rate.[25] For a subsidy recipient with no private income or assets who lives with no more than a partner plus any dependent children,[26] the function under the pre-reform system was:

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

The LHA rate varies geographically, and by the subsidy recipient's family type. The geographical variation is between areas are 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'.

[22] 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.

[23] Department for Work and Pensions (2015).

[24] Department for Work and Pensions (2015).

[25] 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.

[26] 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 housing benefit 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.

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 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.[27] 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:

$$HB = \min\,(LHA\ rate, rent)$$

[27] These cuts were part of a wider post-recession fiscal consolidation implemented by the UK government in an attempt to reduce a large structural budget deficit. They accounted for about 12% of the cuts to social security and 1.7% of the whole consolidation package between 2010–11 and 2015–16 (Browne and Hood (2015)).

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 claimants 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 claimants 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).

The changes that reduced the value of the LHA rate lower the cap on housing benefit 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 Sections III and IV 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

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 Saffron/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 housing benefit 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 housing benefit entitlement to missing in rare cases where the LHA rate is recorded as zero.

Analyses that use rent, maximum housing benefit, or rent net of housing benefit 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.

Sample selection – existing claimants

The basis for our analysis of all existing claimants is a random one-in-three sample of all LHA claimants in January 2011. We take a one-in-three sample purely for computational reasons. (In our analysis of the removal of the five bedroom rate, the extension of the Shared Accommodation Rate, and the introduction of national LHA rate caps, we use data on all those particularly likely to be affected by these reforms, rather than a one-in-three subset.) Focusing on LHA claimants in the one-in-three subset yields a sample of 283,574 subsidy recipients.

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 LHA 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 2008–09 were reset to April 2009.[28]

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

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

which included that claim was submitted in or after June 2011, we would exclude this claim from the analysis.

Appendix C: The exclusion of observations around the introduction of reforms for analysis of new claimants

Appendix Figure C1 shows a seven-day moving average of housing benefit entitlements for new claims made between June 2010 and November 2011 inclusive. The average claim was for about £110 per week. Appendix Figure C2 shows the same trend for the rents of new claimants, which averaged just above £120 per week. Perhaps the most striking feature of both figures is the large spike in mean entitlements and mean rents just before the reformed system starts to apply to new claimants on 1st April 2011. [29] This is consistent with the financial incentives created by way the reforms were rolled out: someone making a new claim just after 1st April would immediately face the reformed, less generous, housing benefit system, whereas someone starting a new claim just before 1st April would not face this system in full for another 21 months. Furthermore, the incentive to make a claim before rather than after 1 April, as measured by the difference in the size of LHA entitlements between the unreformed and reforms system is, in general, increasing in the size of (pre-reform) LHA entitlement. This is reflected in the data: 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, 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.[30] The grey line on Appendix Figure C2 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.

This spike in the volume and nature of new claims for HB is highly likely to be due to individuals bringing forward a claim. To avoid these shifts in the timing of new claims from affecting our analysis that is based on new claimants, we exclude a window of data around the reform time when estimating equation (2). There is clearly a trade-off here between purging these short-run timing responses and losing sample size and having to extrapolate further the pre-reform time trend. However, our sample size is very large and the time trends

[29] Figure 3.3 of Beatty et al (2013) shows there is also a large rise in the number of claims per week in this period.

[30] Both analyses are available on request.

in our outcomes of interest look uncomplicated, we so take a conservative approach and exclude all new claims made between 1st December 2010 and 31st May 2011, marked with vertical lines on Appendix Figures C1 and C2.[31]

APPENDIX FIGURE C1

Average housing benefit entitlement of new claimants by date of claim (seven-day moving average)

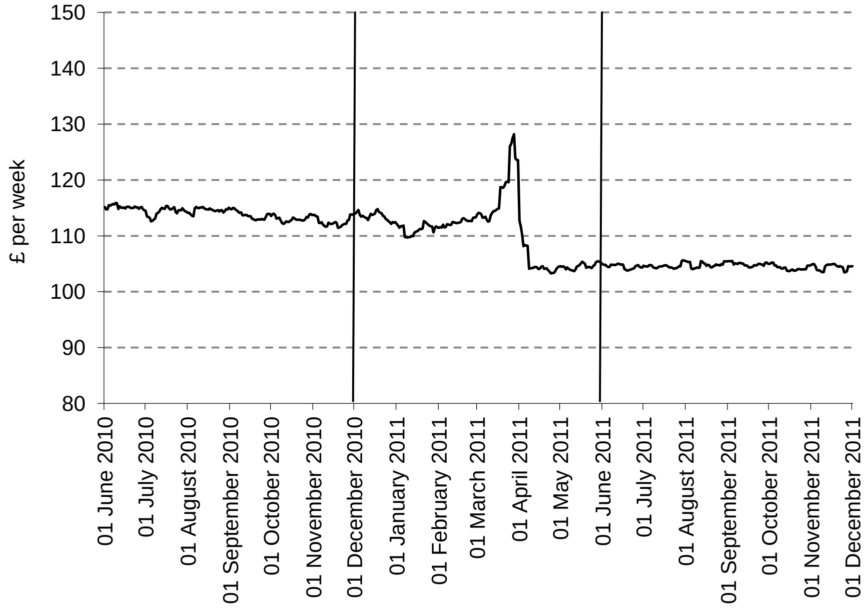

[31] We have conducted sensitivity analysis and, as the figure suggests, our estimates are robust to small shifts in the window of data excluded.

APPENDIX FIGURE C2

Average rent of new claimants by date of claim

(seven-day moving average)