

DISCUSSION PAPER SERIES

No. 9649

HOUSEHOLD DEBT AND THE DYNAMIC EFFECTS OF INCOME TAX CHANGES

James Cloyne and Paolo Surico

INTERNATIONAL MACROECONOMICS



Centre for Economic Policy Research

www.cepr.org

Available online at:

www.cepr.org/pubs/dps/DP9649.php

HOUSEHOLD DEBT AND THE DYNAMIC EFFECTS OF INCOME TAX CHANGES

James Cloyne, Bank of England, CFM and CESifo
Paolo Surico, London Business School and CEPR

Discussion Paper No. 9649
September 2013

Centre for Economic Policy Research
77 Bastwick Street, London EC1V 3PZ, UK
Tel: (44 20) 7183 8801, Fax: (44 20) 7183 8820
Email: cepr@cepr.org, Website: www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programme in **INTERNATIONAL MACROECONOMICS**. Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: James Cloyne and Paolo Surico

CEPR Discussion Paper No. 9649

September 2013

ABSTRACT

Household Debt and the Dynamic Effects of Income Tax Changes*

Using a long span of expenditure survey data and a new narrative measure of exogenous income tax changes for the United Kingdom, we show that households with mortgage debt exhibit large and persistent consumption responses to tax changes. Home-owners without a mortgage, in contrast, do not appear to react, with responses not statistically different from zero at all horizons. Splitting the sample by age and education yields only limited evidence of heterogeneity as the distributions of these demographics tend to overlap across housing tenure groups. We interpret our findings through the lens of traditional and more recent theories of liquidity constraints, providing a novel interpretation for the aggregate effects of tax changes on the real economy.

JEL Classification: E21, E62 and H31

Keywords: liquidity constraints, mortgage debt and narrative tax changes

James Cloyne
Bank of England
Threadneedle St
London EC2R 8AH

Paolo Surico
London Business School
Regent's Park
London NW1 4SA

Email:

james.cloyne@bankofengland.co.uk

Email: psurico@london.edu

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=177540

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=159564

*We are grateful to James Banks, Joao Cocco, Matteo Iacoviello, Ethan Ilizetzki, Andrew Leicester, Karel Mertens and Federico Ravenna for very useful suggestions and seminar participants at the April 2013 London Macro Workshop, University of Bonn, University of Rome `La Sapienza', LUISS, ERC Workshop on Macroeconomic Dynamics with Heterogeneous Agents, 9th CSEF-IGIER Symposium and the NBER Summer Institute for comments. Evangelos Aspiotis provided valuable research assistance. The views in this paper are those of the authors and do not necessarily reflect the views of the Bank of England, the Monetary Policy Committee or the Financial Policy Committee. Surico gratefully acknowledges financial support from the European Research Council (Starting Independent Grant 263429) and Leverhulme Trust. The graphs in this paper are best viewed in color.

Submitted 09 September 2013

1 Introduction

The persistent rise in mortgage debt across many industrialized economies prior to the recent financial crisis has drawn considerable attention to the role of private indebtedness in the transmission of macroeconomic shocks. On the empirical side, Mian et al. (2012), Dynan (2012) and the IMF (2012) report that high levels of household debt are likely to have amplified and prolonged the Great Recession of 2007–08. On the theoretical side, Kaplan and Violante (2013), Eggertsson and Krugman (2012) and Andres et al. (2012) lay out models with heterogeneous agents in which fiscal policy is more effective the larger the proportion of debt-constrained households in the economy.

A common presumption behind these studies is that debtors are more likely to face liquidity constraints and thus adjust their consumption significantly in response to (macroeconomic) conditions that unexpectedly change their income. An important implication of this line of research is that it is not net wealth per se that determines the consumption reaction to fiscal policy changes: households who made a large durable purchase — such as housing — may well be wealthy and liquidity constrained at the same time, depending on their level of indebtedness as well as their mortgage repayments relative to earnings.

Despite the clear relevance of this transmission channel for both policy and academic research, little is known in the data about whether the dynamic effects of income tax changes on consumption vary with a household debt position and whether the mortgagors' response is larger or smaller than the outright home-owners'. At least three considerations make this task particularly challenging. First, survey data with good expenditure coverage typically lack equally detailed and reliable information on the household finance position over a sufficiently long period of time. Furthermore, whether a household holds mortgage debt or not is the outcome of a selection into group. Last but not least, consumption and income changes are jointly determined, therefore one needs to identify the exogenous component of income changes. As we

will discuss later, the presence of a relatively large number of income tax changes in the United Kingdom motivates our focus on the expenditure data from the UK Family Expenditure Survey (FES).

To elicit individual debt positions, we propose to group households by their housing tenure, which allows us to distinguish between home-owners with a mortgage and home-owners without. The motivation for this novel grouping strategy is threefold. First, mortgages are the most prominent form of household debt, in both incidence and value. Second, the extensive margin of whether a household holds a mortgage or not is likely to be less prone to measurement errors than the intensive margin of its outstanding value (which is only observed in the FES over a shorter period of time characterized by few tax changes). Third, commonly used expenditure survey data are either a repeated cross-section, like the FES, or a panel with a short time series dimension, like the US Consumption Expenditure Survey. Following Browning et al. (1985), some grouping estimator is therefore needed to aggregate individual observations into pseudo-cohorts.

A further advantage of looking at heterogeneity in consumption through housing tenure is that we can investigate the dynamic effects of tax changes on another interesting group of British households, namely those renting from local authorities or housing associations: ‘social renters’. As we will show, a typical household in this group is characterized by little (if any) net wealth, low income and compulsory education. On the one hand, these features seem to fit well the *traditional* stereotype of liquidity constrained households in one-asset models. On the other hand, the share of social renters — which is about 20% of the British population — appears too small to account for the large and persistent effects of tax change on the aggregate economy reported in the empirical macro literature.

A potential drawback of grouping by housing tenure is the selection effect associated with possible transitions from one tenure status to another over time. However, the very gradual rate at which ownership has risen in the U.K. suggests this may be less of a concern. Indeed, our results are robust to using the grouping strat-

egy proposed by Attanasio et al. (2002), which addresses explicitly the possibility of compositional changes and selection effects.

To address the endogeneity of income changes, we exploit exogenous variation in aggregate income taxes identified using the narrative approach pioneered by Romer and Romer (2010) and applied to the U.K. by Cloyne (2013). The U.K. is a natural choice for our purposes because there have been a large number of income tax changes in the last forty years. Furthermore, detailed information from official documents allows us to identify individual tax measures and their motivation. Tax changes that were introduced for reasons unrelated to the business cycle can then be used to identify exogenous variation in household income.

Using a long span of household survey data from the FES and a new narrative measure of fiscal shocks, we report significant aggregate effects of exogenous income tax changes on private consumption. More importantly, our disaggregated approach allows us to identify the expenditure response of different groups of households to an income change. We find that the estimated dynamic effects are highly heterogeneous across housing tenures, in a way that is understated — if not missed — when households are grouped by age and/or education.

We establish a number of specific results. First, mortgagors exhibit the largest and most significant response. In contrast, outright home-owners hardly adjust their expenditure, with effects that are never statistically different from zero. Second, the response of social renters is significant but smaller than (though rarely statistically different from) the mortgagors'. Third, the composition of mortgagors' net wealth is significantly different from those of outright owners and social renters. More specifically, we show that a typical household with mortgage debt holds very little *liquid* net wealth despite owning sizable *illiquid* assets. Since mortgagors account for about half of the British population, this *new* type of liquidity-constrained households can make a significant contribution to the large aggregate effects of tax changes typically found in empirical studies and also documented here.

This paper contributes to a growing empirical literature on consumption hetero-

geneity. Recent studies, including Anderson et al. (2012), De Giorgi and Gambetti (2012a,b), Ercolani and Pavoni (2012), Giavazzi and McMahon (2012) and Misra and Surico (2013) have focused on age, liquidity and income to account for the diverse responses of household expenditure to a fiscal shock.¹ The findings from earlier contributions have been convincingly interpreted as supportive of theories based on precautionary saving, partial insurance and limited participation. The results presented here highlight the role of an additional channel for the transmission of structural shocks: mortgage debt. This is consistent with a framework where the decision to purchase a large durable good through borrowing makes some households ‘wealthy’ hand-to-mouth, as in the consumption model with transaction costs on accessing illiquid wealth by Kaplan and Violante (2013).

Our results are also of interest for the large body of research in empirical macroeconomics on the effects of fiscal policy on real activity. While our estimates for the whole economy are consistent with those reported by Mountford and Uhlig (2009), Romer and Romer (2010), Mertens and Ravn (2012) and Cloyne (2013), among others, our approach allows us to identify which households drive the aggregate result as well as which individual characteristics tend to predict a higher sensitivity of consumption to income changes. These appear important dimensions along which to evaluate the effectiveness of tax interventions both on the aggregate economy and across different groups in society.

The rest of our analysis is structured as follows. Section 2 describes the household survey data and how these are used to construct expenditure measures by housing tenure. We also present our narrative-identified tax changes and the empirical specification we employ. Section 3 verifies that the macroeconomic impact of income tax changes is similar using FES and National Accounts data. Section 4 presents the main finding of heterogeneity in the consumption response to tax changes across housing

¹Campbell and Mankiw (1989), Attanasio and Weber (1993) are earlier contributions in the vast empirical literature on the correlation between consumption and income changes. Recent examples include Jappelli and Pistaferri (2010), Acconcia et al. (2011) and Pistaferri and Saporta (2012).

tenures. Section 5 looks at the credit conditions and financial position of the British households by tenure group. These are used to interpret our empirical findings in the light of traditional and newer theories of liquidity constraints that predict heterogeneous responses of household consumption to income changes. In Section 6, we confirm the robustness of our findings to controlling for other exogenous tax changes, anticipation effects, compositional changes, heterogeneity in the amount of tax wind-fall and using other expenditure categories. The Appendix contains a description of the data and further econometric results.

2 Data and empirical specification

In this section, we describe the different data sets we employ. We first discuss the household survey data and the grouping strategy used to construct time series of consumption for cohorts based on housing tenure status. We then move to the narrative data on UK tax changes and the way we exploit these to construct an exogenous income tax measure. The final part of this section explains the estimation strategy, which brings together the expenditure and tax data.

2.1 Household consumption data

The focus of our analysis is on whether households with mortgage debt respond more to income tax changes than those without. For this purpose, we need both good quality household expenditure data and information on household debt positions.

Household expenditure data are obtained using 32 waves, from 1978 to 2009, of the Living Costs and Food Survey, commonly known as the Family Expenditure Survey. This survey has high quality, detailed information on expenditure and household characteristics.² Each wave of the survey contains around 7,000 households, generating over 200,000 observations in total.

²Although the survey has run annually from 1968, a consistent measure of educational attainment is only available from 1978.

Ideally, we would like to observe individual balance sheet positions *and* expenditure. Unfortunately, there are no micro data sets that jointly record detailed information on consumption and wealth over a sufficiently long time period and the FES is no exception. Nevertheless, the lack of household debt data can be bypassed by looking at housing tenure status. Since loans secured on housing represent the majority of household debt, we propose to examine whether home owners with mortgage debt react more to tax changes than those without. Our primary focus is therefore on the response of mortgagors compared to outright owners.

We also make use of a third tenure category, those living in accommodation rented from local authorities or housing associations. For short, we refer to these households as ‘social renters’. These households tend to be poorer, have compulsory education and — as we will show using less frequent data from the British Household Panel Survey — have little liquid or illiquid wealth.³ The social renters therefore fit the demographic characteristics of those more likely to be traditionally credit constrained and are an interesting additional comparison group.

Since the FES is not a panel dataset, it is common to use a grouping estimator along the lines proposed by Browning et al. (1985). We therefore aggregate households into pseudo-cohorts by housing tenure. Given our focus on mortgage debt, housing tenure appears a natural dimension along which to aggregate households into pseudo-cohorts. This produces a time series for consumption for each of our three tenure categories: mortgagors, outright owners and social renters.⁴

Our analysis focuses on non-durable goods and services consumption, in keeping with earlier contributions such as Attanasio and Weber (1993), Attanasio et al.

³The BHPS only collects information of wealth and assets in three years: 1995, 2000 and 2005. Furthermore, the BHPS does not have rich expenditure microdata. As a result, the consumption literature has made extensive use of the FES.

⁴Mortgagors represent on average about 50% of the full sample whereas social renters (owners outright) cover a share a bit below (above) 20%. Unfortunately, private renters account for about 10% of observations and therefore they are excluded from the analysis below. Preliminary attempts to run a specification based on private renters revealed that the point estimates for this group are similar to the point estimates for mortgagors, but the standard errors are so wide (probably due to the fewer number of observations) that they cannot be discriminated statistically from the point estimates of any other group.

(2002) and Campbell and Cocco (2007). The data are discussed in more detail in the Appendix; in Section 6, we show that the results are robust to using non-housing expenditure. For each group, we gross-up the individual household responses using household weights, divide by the number of people in the household to generate a per capita measure and divide by the retail prices index excluding mortgage repayments (RPIX). We construct a quarterly series by assigning households to quarters based on their interview date. To address seasonality, we use the annual change in quarterly expenditure for each housing tenure group.

Since we examine the response of consumption to tax changes, we only want to include taxpayers in our sample. The FES contains information on income taxes paid and we therefore exclude households who reported they did not pay any tax.⁵ However, there may be some measurement error associated with this reporting. After excluding these households, we also then drop any remaining households whose income was below the threshold for paying income tax.

One issue using a grouping estimator is that the dimension along which the aggregation is performed needs to be constant or fully predictable over time. This explains why birth cohort has proved so popular in the literature. In our case, we do not know whether a household with a particular tenure status had the same tenure status in the previous period, or will still have the same tenure status in the next period. Figure 1 shows that there have been changes in the shares of the tenure groups over time, although these changes have been relatively slow moving: it seems unlikely that individuals fluctuate between groups at a quarterly or even annual frequency. However, to ensure robustness of our findings, we also consider grouping households according to their predicted probabilities of being a mortgagor, following Attanasio et al. (2002). In Section 6, we present evidence based on this approach.

⁵There are a number of issues using the FES tax information directly. First, the FES is not a panel so we cannot compute the tax changes over time for a particular household. Second, even if we grossed-up the tax data to the housing tenure group level this measure would be endogenous. Third, these data refer to taxes already paid and therefore already factored into the household's consumption decision. In contrast, the tax shock series we construct refers to the change in future liabilities as a result of changes in tax policy.

It is useful to examine the demographic properties of the three housing tenures. For each group, Figure 2 shows kernel density estimates of age and weekly household real income per adult, as well as the share of households with different education levels and non-mortgage debt. Mortgagors are on average younger, more educated and relatively richer but it is worth noting that the distributions of the three tenure groups overlap significantly.⁶ In particular, the estimated densities for both groups of home owners are characterised by a long right tail. This means that average income of outright owners is relatively closer to the average income of mortgagors rather than social renters.⁷ Finally, the mortgagor group has the largest share of households with non-mortgage debt whereas outright owners have the lowest. This is based on the shorter sample 1986-2009 for which information on non-mortgage debt is available in the FES. We will return to credit market access by housing tenure in Section 5 where we use other surveys on the financial position of the British households to interpret our econometric results.

2.2 UK income tax changes and narrative identification

Endowed with household expenditure data by housing tenure, we now need to construct a suitable measure of tax changes. The identification challenge we face is that tax changes may affect consumption and other macroeconomic variables but common measures of taxes, such as total tax revenues, are also affected by the state of the economy (either automatically or because of endogenous discretionary policy actions). Our household tenure groups are large shares of the population. The simultaneity between fiscal policy and consumption therefore prevents consistent estimation of the dynamic effects of tax changes by regressing consumption on, for example, tax

⁶A similar picture emerges considering household disposable income.

⁷The sub-sample evidence in the Appendix reveals that the distributions of education and age have not changed much over time. In contrast, income has become more unequal, although this has largely been a feature of the top of the income distribution. Consequently, it has affected the average income of the owners and mortgagors but less so the average income of social renters.

revenues.⁸

To address this identification problem, we employ a narrative approach following Romer and Romer (2010) for the United States and Cloyne (2013) for the United Kingdom. We use detailed documentation from historical sources to identify ‘exogenous’ legislated changes from the motivations given by lawmakers at the time of the policy intervention.⁹

Unfortunately, the narrative measures of *aggregate* tax changes used in earlier contributions contain changes to a variety of taxes such as income, consumption and capital taxes, each of which may affect household groups differently (e.g. Stamp Duty, Vehicle Excise Duty etc). Ideally, we seek tax changes that affect all housing tenure groups. One advantage of the narrative approach is that, by going through detailed official budget documents, we can isolate changes in specific income taxes that affect all income taxpayers.

In the UK, income tax accounts for around 25–30 per cent of total government revenue and consists of a set of allowances, bands of income and marginal tax rates that apply to each income band. Each individual has a personal allowance which is deducted from their income to calculate their ‘taxable income’. An income taxpayer is therefore someone who earns more than their allowance (although there is a system of tax credits that lowers the tax liability of the poorest taxpayers). In the tax year 2012-2013, for instance, the first 34,370 pounds above the allowance was subject to a 20 per cent rate, any further earnings up to 150,000 pounds were subject to a 40 per cent rate and then 50 per cent was charged on all earnings over 150,000. These bands and allowances are increased each year in line with inflation unless the UK Parliament decides otherwise. We do not treat automatic inflationary increases as tax shocks in our data set.

⁸Or using taxes paid data from the FES. Even the construction of cyclically adjusted revenue can be problematic as this still requires an estimate of the output elasticity of revenue. See, for example, Caldara and Kamps (2012).

⁹The idea has also been used to identify government spending shocks (Ramey and Shapiro (1998); Ramey (2011)), monetary policy shocks (Romer and Romer (1989, 2004)) and fiscal consolidations (Guajardo et al. (2011)).

Income tax is payable on a wide range of income including earnings from employment, property, interest, retirement pensions and some social security benefits. To focus on changes that affect all income taxpayers we collect a data set of changes affecting the lowest bracket of income tax. Specifically we consider changes in the tax free allowances (that determine the level of income above which income tax starts to be paid), the basic rate of income tax (currently 20 per cent) and the income bands defining the basic rate. We refer to this group of tax changes as the allowance and basic rate of income tax.¹⁰ This set of tax changes will therefore be paid by all households in our sample.

We first collect data on all the discretionary income tax changes from narrative sources such as UK Budget documents. We isolate around 150 changes in the allowances and basic rate of income tax. For the quantitative magnitude, we follow Romer and Romer (2010) and use the revenue forecasts from the Budget documents. However, the focus is on the change in tax liabilities rather than any short-run revenue effect due to the timing of revenues reaching the Exchequer. We therefore use the ‘full year’ revenue estimate, which is the projected on-going annualised revenue effect. This value is then assigned to the implementation date of the policy change. As explained by Mertens and Ravn (2012), it is possible that this implementation date is anticipated if it is announced some time earlier. We address this possibility in Section 6 and show our findings are robust to considering only those tax changes that were more likely to be unanticipated.

The source for the policy changes and revenue estimates is the Financial Statement and Budget Report (FSBR) which is published alongside the Budget speech. For actions between Budgets (not already covered in the FSBR), we use the estimates given by the Chancellor of the Exchequer to Parliament. The source is the official parliamentary record, Hansard. The tax changes are then categorised by their given

¹⁰Mertens and Ravn (2013) split the Romer and Romer dataset into corporate and personal tax liabilities and study the macroeconomic effects of these tax changes. As we examine sub-groups of the population, we need to construct a more specific measure of income tax changes.

motivation (using a variety of UK government, parliamentary and historical documents and speeches), isolating decisions that were taken for reasons uncorrelated with macroeconomic fluctuations. Following Romer and Romer (2010), we refer to these changes as ‘exogenous’.¹¹ The classification of tax changes keeps as close as possible to the motivation given in the narrative documents (although the history literature was also used to frame and cross reference the interpretation of the given statements). For illustrative purposes, the Appendix provides an example of tax changes that we regard as ‘exogenous’ and ‘endogenous’.

Individual exogenous income tax changes are assigned to quarters and aggregated. Figure 3 shows, as solid line, our newly constructed tax series scaled by nominal GDP, together with the aggregate tax change series in Cloyne (2013) as dashed line. There have been a sizable number of income tax changes and many of these have been quite large, providing good variation in our narrative tax series over time. The large majority of these legislated changes were supply-side reforms designed to encourage long-run economic performance, sharpening incentives and lowering the burden of taxation. Reassuringly, the correlation between all other exogenous *income* tax changes and our measure of allowance and the basic rate is low, at 0.12. Similarly, the correlation with all other exogenous tax changes is 0.05. This suggests that changes in our measure were not contemporaneously offset by changes in, for example, higher rates of income tax. In Section 6, we show that our findings are not sensitive to controlling for other tax changes.

While income tax changes appear persistent in the UK, we do not regard these as changes in permanent income. First, households may have expected the changes to be reserved or offset elsewhere. Second, taxes were changed frequently over our sample (for example, income taxes were changed in 52 of the 128 quarters in our sample) so it seems unlikely households believed any particular change was permanent. Third, over

¹¹The authors refer to these as tax changes “not taken to offset factors pushing growth away from normal” (Romer and Romer (2010), 770). Detailed information on the classification of individual measures and supporting evidence for each classification can be found in Cloyne (2012).

time there were increases in other taxes and, during the 1980s and 1990s, government consumption fell as a share of GDP. The debt to GDP ratio was also relatively stable over this period (at least until 2008). Furthermore, we have confirmed that our results are robust to controlling for other fiscal variables. We present evidence on this in Section 4.3.

Having constructed our new tax series, it is important to verify that it is not predictable using past information at either the aggregate or cohort-specific level. Specifically, we conduct Granger causality tests based on a VAR which contains the change in consumption per capita for each household group (our dependent variables in the next section), the change in real GDP per capita and the central bank's policy rate. Reassuringly, we could not reject the hypothesis that these variables do not Granger cause our income tax series: the p-values using various lag lengths were high, over 0.4, for 4, 6 and 8 lags.¹²

2.3 Empirical specification

In Romer and Romer (2010) and Cloyne (2013), the tax measure can be thought of as the change in an aggregate average tax rate. For our purpose, it makes less sense to divide income tax liabilities by aggregate GDP as this would not reflect an average tax rate. Instead we transform our nominal tax liabilities series into a (real) income tax change per taxpayer. We divide the (narrative) projected change in nominal liabilities by RPIX and the total number of individual income taxpayers. Over a three year horizon, this amounts to an average windfall of about 500 pounds per household at 2009 prices.

Two features of this specific transformation make it particularly useful for our analysis. First, the size of the estimated coefficients on tax changes are readily interpretable as the pound change in consumption following a *one* pound change in taxes. Second, the imputed windfall does not vary systematically with household income, guarding against the concern that household consumption and income are

¹²Similar results are obtained using the change in real gross income per capita for each group.

jointly determined within the same time period. In Section 6, we use an alternative transformation which explicitly accounts for heterogeneity in the tax windfall. We discuss the drawbacks of this but show that our results are not overturned.

We estimate the following regression:

$$\Delta C_{i,t} = \alpha_i + \sum_{j=0}^{12} \beta_j^i \tau_{t-j} + \gamma^i \Delta Z_{i,t} + u_{i,t} \quad \text{with } i = M, S, O. \quad (1)$$

where $\Delta C_{i,t}$ is the pound change in real per capital non-durable consumption, τ_t is our proposed measure of income tax change, which is the prospective annual change in real tax liabilities per taxpayer, and $\Delta Z_{i,t}$ is a vector containing the change in the mean of the demographic variables age and education as well as the proportion of employed households and retired households per group. The index i refers to the cohorts: mortgagors (M), social renters (S) and outright owners (O).

By construction, τ_t are uncorrelated with other shocks to consumption $u_{i,t}$. For a £1 change in taxes, the sum of the β coefficients in equation (1) gives the total four quarter consumption relative to trend at a point in time. For example, after four quarters the sum of the first four β coefficients gives the total additional consumption (in pounds) over the first year after the tax reform. In the next section, we show that the aggregate effect of a change in exogenous income taxes is the same when using expenditure data from the FES and the UK National Statistics. For this exercise, the dimension of i reduces to one and the vector Z_t is excluded.

Given the relatively long time dimension but small cross-section dimension of our data, the system of equations (1) is estimated using the method of Seemingly Unrelated Regressions. The rationale for this choice (as opposed to equation-by-equation OLS) is twofold. First, each of the three tenure groups cover a large share of aggregate consumption and we cannot exclude a priori that the errors terms might be correlated across the equations. Omitted correlation in the covariance matrix of the disturbances does not pose a challenge for consistency but it would reduce efficiency. Second, the SUR method provides a natural framework in which to test formally for heterogenous responses across groups.

3 Aggregate results

Before presenting evidence on the response of different housing tenure groups, it is important to verify that the aggregate effects found using the household survey data are consistent with results found using National Accounts data.

Figure 4 reports the dynamic effects of an exogenous income tax liability cut on two measures of aggregate expenditure on non-durable goods and services. The first measure comes from the Office for National Statistics (ONS) and it represents the official aggregate series on non-durable consumption for the UK economy. The second measure comes from the FES and it is constructed by aggregating household-level expenditure in each quarter using household-specific weights that ensure representativeness of the FES sample in the UK population.

The point estimates for the impulse responses based on aggregate FES data are shown as the continuous black line with circles and results using the ONS data are shown in the blue dotted line. The shaded area represents the 68% confidence intervals based on 20,000 bootstrap repetitions; the grey lines show the 95% intervals. Under a normal distribution, the 68% band corresponds to adding one standard deviation either side of the point estimates and it is commonly reported in the empirical literature on fiscal policy (see, for example, Blanchard and Perotti (2002) and Romer and Romer (2010) among many others).

Three findings are worth noting in Figure 4. First, the point estimates of the impulse response based on FES data are very close to those based on ONS data, with the latter always well within the 68% confidence bands of the former. The two impulse responses virtually overlap during the first year and, after that, the discrepancy is smaller than 5% on average. Second, both sets of impulse responses are associated with real effects of fiscal policy that peak just below 1.5 after three years. Third, these real effects are statistically larger than zero at the 5% significance level for all horizons after the second quarter. The results in this section lead us to conclude that the FES appears a reliable source for conducting a disaggregated analysis of the effect

of fiscal changes on different groups in society.

The aggregate effects estimated using a narrative approach are typically larger than found using specific episodes of fiscal stimulus, such as the 2001 and 2008 income tax rebates in the US.¹³ However, it is important to note the effects reported in that literature refer to *temporary* payments aimed at stimulating aggregate demand. In this paper, as in the vast majority of empirical macro studies on fiscal policy, the focus is on *persistent* tax changes, which are likely to have both a longer-term supply-side impact and larger general equilibrium effects.

Finally, it is worth emphasizing that the size of the estimated responses in Figure 4 is consistent with the findings reported in the macro literature, where a one percentage change in aggregate taxes to GDP is shown to trigger large and persistent increases in total consumption and GDP, with peak effects above 2 and 3 percent respectively (see Romer and Romer (2010) and Cloyne (2013)). Using aggregated FES data, we also find evidence of large and persistent effects of the narrative tax changes on household *pre-tax* income, which is consistent with a supply side interpretation of the magnitude of the consumption responses reported in this paper. Indeed, this latter result holds for the *pre-tax* income of each housing tenure group.

4 The heterogeneous effects of tax changes

In this section, we document significant heterogeneity in the dynamic effects of fiscal policy on household consumption across housing tenures. We begin with the estimates for the three groups: mortgagors, outright owners and social renters. We then contrast our main results with those from more ‘traditional’ group splits based on birth year and education, and find that the heterogeneity is far more pronounced (and significant) when the sample is divided according to housing tenure status. Finally, we show that our results are not sensitive to using a VAR in group-specific consumption, aggregate output, inflation, the short-term rate and government spending.

¹³See, for example, Parker et al. (2013), Misra and Surico (2013) and the references therein.

4.1 Housing tenure

The dynamic effects of a one-pound exogenous tax liability change on the non-durable expenditure of each household group are displayed in Figure 5. Recall that these charts show the four-quarter response. For example, after four quarters this is the cumulative effect on consumption over the first year.

The first row shows that the consumption of mortgagors (M) responds significantly at the 5% level beyond the first quarter and it reaches a peak above 2 pounds after three years. In contrast, the response of owners without mortgage (O) in the last row is never statistically different from zero at the 32% significance level and peaks below 0.8 or 80 pence five quarters after the shock. Finally, the point estimates in the second row suggest that social renters (S) adjust their non-durable expenditure by less than one pound with responses that are significant at the 32% level from quarter 1 and at 5% between quarters 2 and 5.¹⁴

The estimates in Figure 5 map into cumulative multipliers for mortgagors and outright owners of about 1.9 and 0.1 pounds, following a one-pound tax windfall per year, with the differences across groups being more marked during the second and third year. The multiplier for social renters is about 1.1.¹⁵ As on average owners without a mortgage have a significantly higher gross income than social renters, it seems unlikely that the heterogeneity in Figure 5 is driven by possible heterogeneity in the tax windfall. We return to this issue in the Appendix.

In the first column of Table 1, we formally test the extent of heterogeneity shown in Figure 5. For each tenure group, we report the effect of the tax cut after one, two and

¹⁴The finding of heterogeneity and the ranking across groups are robust to using the log difference of consumption, as opposed to the consumption change. Under this specification, however, the size of the responses are difficult to interpret as we only observe the tax liability *changes* projected by HM Treasury.

¹⁵In Section 4, we present evidence suggesting that social renters are likely to have a worse credit access, if any, than mortgagors. To the extent that households with access to unsecured lending are subject to a loan-to-income ratio constraint, the income windfall generated by a tax change would represent a relaxation of credit conditions. This simple mechanism may offer a rationale for the larger-than-one cumulated impact on mortgagors' consumption, even without appealing to general equilibrium effects. Furthermore, it may also account for the finding that the point estimates of the mortgagors' response is larger than (though not statistically different from) the social renters'.

three years. For each yearly horizon f , we report p-values for three null-hypotheses: (i) homogeneity among the responses of all three groups, (ii) homogeneity between the responses of mortgagors and owners without mortgage, and (iii) homogeneity between the responses of mortgagors and social renters.

In the Appendix, we provide further details on the specifications behind the estimates in the rest of the table, which assess robustness to controlling for all other exogenous tax changes (second column with heading ‘including all other tax changes’) and to focussing only on the personal allowance and basic rate of income tax changes implemented within 90 days from announcement (third column with heading ‘unanticipated ABR’).

A number of salient points emerge from Table 1. The consumption responses of all groups to the tax change seem stable across the three specifications. In particular, mortgagors and social renters always adjust their expenditure on non-durable goods and services by a large and significant amount. Outright home-owners, in contrast, respond with a modest and insignificant change.

As for the formal tests of homogenous responses, the inference one can draw from tests of the joint null hypothesis across the three groups is mixed: the p-values in the first column are above 0.1 whereas those in the other columns tend to be below 0.1. The bilateral tests, however, reveal that the results of the joint hypothesis are driven by the behaviour of mortgagors and social renters, whose estimates are never significantly different at the 10% level. On the other hand, the null hypothesis that the outright owners adjustment is statistically similar to the mortgagors’ is rejected at several horizons across all specifications. Furthermore, the p-values associated with the null $\beta_{yf}^M = \beta_{yf}^O$ are always far lower than those associated with $\beta_{yf}^M = \beta_{yf}^S$. We conclude that owners with a mortgage and owners without are likely to react differently to the tax change. On the other hand, social renters and mortgagors are likely to behave similarly (though the difference in the point estimates for the latter pair does not seem negligible).

In summary, our estimates suggest that housing tenure is highly correlated with

the unobserved characteristics driving the heterogeneous response of household consumption to an exogenous income tax change. Specifically, whether a household has mortgage debt seems an important dimension through which tax changes affect consumption. In Section 5, we will show that the composition of the mortgagors' asset portfolio — in particular the lack of significant liquid net wealth — is likely to make this group more responsive to income variation.

4.2 ‘Traditional’ sample splits

In this section, we examine the degree of heterogeneity across two alternative groupings that have been used extensively in the micro literature: age and education.¹⁶ As these variables are either exogenous or fully predictable over time, they are ideal dimensions along which to construct pseudo-cohorts.

We re-estimate equation (1) over these different groupings. The first column of Figure 6 refers to a specification where households are grouped depending on whether the head is born before 1920 (first row), between 1920 and 1939 (second row) and between 1940 and 1959 (third row). The second and third columns then split these groups further by education, depending on whether the household head has only compulsory or also post-compulsory education.

In line with the descriptive estimates in Figure 2, heterogeneity across groups is more muted (and far less precise) along these ‘traditional’ dimensions. Taken at face value, the point estimates in Figure 6 suggest that, if anything, younger and more educated households tend to adjust their non-durable consumption by a relatively larger amount following a tax change.

The evidence of heterogeneity based on the traditional characteristics appears far less compelling than the evidence based on housing tenure. While Figure 2 shows that mortgagors tend to be younger and more educated, it is important to note that,

¹⁶A further advantage of grouping households by housing tenure, relative to using age, liquidity, leverage or income is that we do not need to take a stand — prior to estimation — on the specific (and somewhat arbitrary) threshold levels below which a household is considered to be, for example, younger, poorer or more levered.

crucially, not all young and educated households are mortgagors.¹⁷ To the extent that a sizable share of households with mortgage are liquidity-constrained (as suggested by the evidence in Section 5), then the housing tenure grouping has a considerable advantage for the purpose of identifying heterogeneity in the consumption responses to income changes.

4.3 Controlling for business cycle conditions

Household consumption is likely to be influenced by macroeconomic conditions as well as group-specific dynamics. The narrative identification procedure, as well as the formal Granger causality tests in Section 2, suggest that our tax changes are uncorrelated with macroeconomic fluctuations or group-specific consumption changes. However, to control for the possibility of residual correlation between the tax measures, other fiscal variables and the dynamics of the business cycle, we control for lagged values of GDP, inflation, the central bank's policy rate and government spending. Furthermore, to control for group-specific dynamics, we add the lagged values of non-durable consumption.

For each tenure group, we run a Vector Autoregression (VAR) using four lags of the variables above but the results are robust to using shorter and longer lag lengths. In keeping with the SUR analysis, we use twelve lags of the narrative measure of tax changes and include this as exogenous variable. The VAR results are shown in Figure 7: the responses of the three groups are largely unaffected and, if anything, the confidence bands for the outright owners appear smaller. More importantly, the point estimates for the outright owners are still outside the 95% confidence bands around the mortgagors' response at horizons beyond five quarters.

¹⁷This is consistent with the finding in Bottazzi et al. (2010) of considerable variation in the rate at which different birth cohorts transit to home ownership.

5 Interpretations

In this section, we interpret the findings above through the lens of traditional and newer theories of consumption heterogeneity based on liquidity constraints. We first report evidence on credit conditions for each housing tenure group. Then, we analyze the net wealth position of British households and emphasize the distinction between liquid and illiquid asset accumulation by housing tenure. Finally, we discuss the relative merits of different theoretical frameworks.

5.1 Credit conditions

In the bottom right panel of Figure 2, based on a shorter FES sample, we showed that the proportion of mortgagors with non-mortgage debt is far higher than the proportion of outright owners and, to a lesser extent, social renters. In this section, we explore further the extent of access to credit markets by looking at the only three sets of waves from the British Household Panel Survey (BHPS) that collect information on the household financial position.

Figure 8 shows the ratio of outstanding non-mortgage debt to disposable income by housing tenure for 2005 (top panel), 2000 (middle panel) and 1995 (bottom panel). In each year, the ranking of the ratios for the three groups match the ranking of the dynamic effects of an income tax change reported in the previous section. Mortgagors, who tend to adjust consumption by the largest amount, are also the households with the highest debt to income ratio. The converse is true for outright owners. Furthermore, the mortgagors' ratios are about three times larger than those of social renters, with the exception of 2000 where the proportion between the ratios of these two groups is 1.5 to 1.

While this evidence suggests that households with mortgage debt may have better access than social renters to both secured and unsecured loans, it says little about the extent to which these two groups may be credit constrained relative to their desired level of consumption. Unfortunately, the BHPS does not contain direct ques-

tions on credit availability which we can draw upon. However, the NMG Consulting Survey, which has been running from 2006 to 2012 for the Bank of England, does. In particular, interviewees are asked whether they have been “*put off spending by credit availability concerns*”. Figure 9 reports the proportion of households with an affirmative answer by housing tenure.

The evidence from the NMG survey reveals that, despite being the most responsive group to a tax change, mortgagors do not appear the most credit constrained. In fact, social renters appear the most likely to be credit constrained with about 32% sufficiently concerned about their access to credit so as to pause their spending. The shares of mortgagors and outright owners are only 22% and 8%, respectively.

Altogether, Figures 8 and 9 raise an important question: if mortgagors have, at least, some access to credit and seem less likely to be put off spending by credit availability concerns, why does their consumption change significantly following a tax change? We will show below that net wealth, and in particular its composition between liquid and illiquid assets, may provide an answer to this question.

5.2 Net wealth: liquid versus illiquid assets

Traditional explanations for a significant consumption response to unexpected income changes emphasize net wealth as an important driver of heterogeneous behaviour. In short, wealthier households are less likely to be liquidity constrained. As this argument is typically made in the context of one-asset models, the academic and policy discussion seems to have implicitly abstracted from the distinction between liquid and illiquid assets. To the extent that most household wealth is held in the form of housing and thus it is not immediately accessible, looking at *liquid* net wealth (as opposed to *total*) may, however, shed light on the heterogeneous consumption responses across the three tenure groups.

Table 2 reports summary statistics for the distributions of financial and housing wealth by tenure group, using the three years over which the BHPS has asked questions on the household’s financial position. Following Crossley and O’Dea (2010), net

financial wealth is defined as the value of saving and investment net of non-mortgage debt and is meant to provide a measure of the stock of *liquid assets*.¹⁸ Net housing wealth is the difference between the property value estimated by the household and the value of any outstanding mortgage.

Three important findings emerge from Table 2. First, social renters — who account for about 20% of the sample — are characterized by little *liquid* financial net wealth and no housing wealth. Together with the evidence on credit access from the previous section, and the fact that they tend to be younger and lesser educated than the other groups, social renters appear to fit well the traditional stereotype of liquidity constrained households. Second, outright owners — who make up just over 20% of the population — score high in both financial and housing wealth and seem unlikely to face significant credit constraints. Third, mortgagors — approximately half of the British population — seem in between the other two groups as they display low *liquid* net wealth but high housing wealth. Indeed, in each of the three years, there are more than 50% of mortgagors who hold either non-positive financial net wealth or only a small positive amount. As the vast majority of mortgagors have at least some equity in their house, their *total* net wealth tends to be high, although it might not be immediately accessible.

The importance of the distinction between liquid and illiquid wealth is formalized in an elegant model by Kaplan and Violante (2013), who show that the interaction between liquidity constraints and transaction costs of accessing illiquid assets make some consumers ‘wealthy’ hand-to-mouth whenever the size of the income change is small relative to the transaction costs. These are households holding low net liquid wealth despite owning sizable illiquid assets and they seem to describe well the financial position of the mortgagors tenure group.

¹⁸‘Saving’ includes: Savings or Deposit Accounts, National Savings Bank Accounts and Cash ISAs (or TESSAs). ‘Investment’ comprises: National Savings Certificates, Premium Bonds, Unit trusts/Investment trusts, Stocks and shares ISAs (or PEPs), Shares, National Savings Bonds (capital, income or deposit) and Other investments (gilts, government or company securities). ‘Non-mortgage debt’ refer to: Hire purchase agreements, Personal Loans, Credit and store cards, Catalogue or mail order purchase agreements, DWP Social Fund loans, Overdrafts and Student Loans.

To verify whether mortgagors are indeed ‘wealthy’ hand-to-mouth consumers, Figure 10 shows the response of non-durable consumption for this group to ‘smaller’ and ‘larger’ tax changes. The ‘smaller’ tax changes refer to changes in personal allowance only, which on average amount to about 250 pounds per household at 2009 prices over a period of three years. The ‘larger’ tax changes, on the other hand, are associated with changes in the basic and higher rates of income tax and, once adjusted for the share of tax-payers in the mortgagors group subject to these rates, average around 750 pounds. Using this categorization, we identify 25 smaller tax changes and 41 larger tax changes, with a correlation between the two tax series of -0.03 over the quarters in which at least one of them is different from zero.¹⁹ It is worth noting that the classification into ‘smaller’ and ‘larger’ changes refers to the size of the windfall received by the household. What we deem as ‘smaller’ at the individual level, however, may well be large in the aggregate if most households are subject to that tax change. Analogously, the individual changes we refer to as ‘larger’ may, in fact, be small relative to the size of the economy if they are concentrated among a small groups of households.

A key finding from Figure 10 is that mortgagors adjust their non-durable consumption significantly only in response to ‘smaller’ tax cuts. The response to ‘larger’ changes, in contrast, is never statistically positive. Furthermore, the difference between the effects of these tax changes is significant at most horizons, suggesting that the size of the windfall triggers very different behaviours. We also repeated this exercise for outright owners and social renters, finding little evidence for a significantly different response to the two types of tax changes for each housing tenure group.

5.3 Discussion

At this stage it is useful to discuss the extent to which different — but not necessarily mutually exclusive — theories may be consistent with the evidence presented in this

¹⁹In keeping with the analysis in the rest of the paper, we have restricted our attention here to tax cuts. Results are robust to including all changes in the basic and higher rates.

paper. As the behaviour of social renters seems to accord well with the predictions of traditional liquidity constraint models, we will focus here on the diverse responses of mortgagors and outright owners.

Impatience and risk aversion may play a role in explaining why mortgagors adjust their consumption significantly following a tax change. However, these explanations explain less well why their response is concentrated on smaller income changes. Rational inattention, in contrast, may offer a rationale for the latter finding: to the extent that the ‘smaller’ changes are small relative to the household income, mortgagors may rationally choose to allocate little attention to the newly available resources and not re-optimize their consumption-saving plan. But, as we have already mentioned, the response of outright home-owners does not differ significantly with the size of the tax change and this group has an average income only slightly lower than mortgagors’ average income. On the other hand, the interaction between liquidity constraints and transactions costs on accessing illiquid wealth, along the lines of the model recently developed by Kaplan and Violante (2013), does appear consistent with both the heterogeneous responses across home-owner groups and the significantly higher response of mortgagors’ consumption to ‘smaller’ income tax changes.

To assess further the empirical merits of this latter mechanism, it would be useful to quantify the magnitude of the transaction costs of accessing housing wealth. This task, however, is very challenging as the combination of many lenders and a wide range of products per lender means that there are a large number of mortgage deals available in the UK market. This vast amount of information can nevertheless be surveyed using specialized web search engines, such as ‘comparethemarket’, ‘money.co.uk’ and ‘moneysavingexpert’. These are designed to scan and compare the numerous offers available on- and off-line. All searches we have conducted on typical mortgage deals reveal that arrangement fees as well as charges on over-payments, early exit and equity withdrawal are a diffuse practice.

According to the Miles (2004) report, average arrangement fees are about 900 pounds, consistent with the estimate of the British Consumers Association ‘Which?’

for 2009.²⁰ On the other hand, our internet searches suggest that early exit fees and early repayment charges tend to average around 1.25% of the repaid amount, whenever this goes above 10% of the monthly installment. According to Halifax, the price of a typically transacted house for which a mortgage has been approved in 2009 was about 160,000 pounds.²¹ Considering an average loan-to-value ratio of 75% and assuming an early repayment of about one third of the borrowed amount, implies transaction costs in the neighbourhood of 500 pounds. Finally, Kaplan and Violante (2013) show that transaction costs of 1000 US dollars (or 650 British pounds) allow their model to match the fraction of wealthy hand-to-mouth consumers in the U.S. Survey of Consumer Finance, which they estimate at around 30%. In summary, there seems to be suggestive evidence that the average transaction costs to access housing wealth in the U.K. may lie between 500 and 900 pounds at 2009 prices.

The discussion above suggests that households with sizable illiquid assets but little liquid wealth may change consumption significantly after a (relatively smaller) income tax change. In the final part of this section, we ask whether the number of UK consumers with those financial characteristics is large enough to be quantitatively important for the aggregate economy. To this end, we follow Kaplan and Violante (2013) who define ‘wealthy’ hand-to-mouth those households with sizable illiquid assets but liquid wealth balances below half of their monthly earnings. For mortgagors in our sample, this averages around 1,350 pounds at 2009 prices and, in each set of BHPS financial waves, there are about 60% of mortgagors whose net liquid wealth is below this number. On average mortgagors account for around half of population, which implies that some 30% of British households may be ‘wealthy’ hand-to-mouth.²² Together with another 20% represented by the social renters, these two groups of

²⁰<http://conversation.which.co.uk/money/mortgages-fees-comparison-confusio/>

²¹http://www.lloydsbankinggroup.com/media/excel/2013/060913_historicdata.xls

²²Ideally, one would like to split the mortgagors group further on the basis of the number of years since origination: more recent mortgagors would seem less likely to hold positive net financial wealth. Information on the mortgage characteristics, however, are only available in the FES for short periods of time that are typically characterized by a small number of tax changes. Unfortunately, the outstanding value of the mortgage does not get reported in the FES over the full sample.

(very different) liquidity-constrained consumers seem large enough to account for the sizable aggregate effect of tax changes on the real economy typically documented in empirical macro studies and also reported in this paper.

6 Sensitivity analysis

We now examine the robustness of our findings to a range of exercises. First, we control for other exogenous changes in tax liabilities by also including the distributed lags of the remaining exogenous tax changes. Second, we consider only the unanticipated component of our baseline measure of exogenous income tax changes. Third, we show that our results carry over to a broader measure of household expenditure. Fourth, we examine whether the heterogeneity we find might be due to heterogeneity in the windfall. Finally, we address the possibility of changing composition and selection effects in the housing tenure groups, applying the method in Attanasio et al. (2002).

6.1 Other tax changes and anticipation effects

In this subsection, we confirm the robustness of our baseline results to two possibilities: that the allowance and basic rate income tax changes were offset by other tax changes and that our results reflect anticipation effects.

In considering a subset of tax changes, one possible concern is that these are correlated with other tax changes. For example, increases in the higher income rate may be correlated with a reduction in the lower rates, inducing spurious heterogeneity if these other tax changes are omitted. To address this possibility, the third column of Table 1 builds on the empirical specification in the second column (discussed in Section 4) by adding the significant coefficients on contemporaneous and (up to twelve quarters) lagged values of any *other* exogenous tax changes. This is our new series subtracted from the aggregate exogenous tax changes in Cloyne (2013).

The estimates reveal that our baseline results are not altered by controlling for these other tax changes. The null hypothesis of homogeneous responses among the

three tenure groups is overwhelmingly rejected at any of the yearly horizon with more than 95% confidence. Mortgagors again respond the most, with responses that are always significant at the 5% level. The converse is true for the outright owners. Social renters adjust consumption by less than mortgagors but the difference is not statistically significant.

In the last column of Table 1, we perform a similar exercise except that we only use the ‘unanticipated’ component of our exogenous income tax changes. We follow Mertens and Ravn (2012) by defining an unanticipated change as one that was implemented within 90 days of announcement. The estimates and hypothesis tests broadly confirm our earlier findings.

6.2 Other expenditure categories

So far our analysis has focused on non-durable consumption. We now explore the extent to which the heterogeneous responses documented in Section 4 depend on the composition of expenditure and, more specifically, on the purchase of durable goods and services. To this end, we consider three alternative categories on (i) ‘food’, (ii) ‘strictly non-durable’ and (iii) ‘non-housing’ expenditure. Following Lusardi (1996), strictly non-durable consumption excludes semi-durable categories such as ‘apparel’, ‘health’ and ‘reading’ from non-durable goods and services. Following Attanasio et al. (2011), non-housing consumption excludes items such as ‘rents’, ‘mortgage payments’ and ‘water bills’ from total consumption and therefore includes durable categories such as ‘vehicles’ and ‘electrical appliances’.²³

Figure 11 reports the point estimates of the dynamic effects of a one pound income tax cut on food (dashed line), strictly non-durable (solid line) and non-housing expenditure (dotted line) for each housing tenure group. Following the tax shock, the expenditure on food is characterised by small adjustments and virtually no heterogeneity. The responses of the strictly non-durable categories are more pronounced

²³The results in this section are based on the method of Seemingly Unrelated Regressions but are robust to using the VAR model employed earlier.

as well as more diverse across housing tenures. But it is only when we focus on non-housing expenditure that the dynamic effects for mortgagors and social renters become large and (as shown in the Appendix) mostly significant. This contrasts with the behaviour of owners whose response in the third row tends to be smaller than those for the other groups and it is never (rarely) statistically different from zero at the 5% (32%) significance level.

Compared with the non-durable results in earlier sections, the response of non-housing expenditure for each tenure group tends to be larger and the response of strictly non-durable expenditure tends to be smaller. Furthermore, the differences across expenditure categories appear largest for mortgagors and smallest for outright owners, at most horizons. The degree of heterogeneity in the response of household consumption to income tax changes therefore seems to increase with the degree of ‘durability’ of the goods and services purchased.

6.3 Accounting for heterogeneity in the amount of windfall

We now examine whether the heterogeneous responses in Figure 5 might reflect (omitted) heterogeneity in the tax windfalls. To this end, we use cohort-specific average income levels to construct a measure of exogenous tax liability changes that varies across housing tenure groups according to group-level income. Specifically, we multiply the change in projected liabilities scaled by GDP — proportional to the average income tax rate — by the cohort-specific average income in the previous period. We further rescale the revenue to GDP figure by the inverse of the average share of household income in GDP and the average share of revenue raised by the basic rate of income tax in total income tax revenue. This should then proxy the average tax rate for income subject to the basic rate.

The drawback of using such a measure is that it makes the tax variable depend on income, raising the concern that the variation in the amount of the windfall across groups becomes endogenous. Specifically, to the extent that the amount of tax change increases with income and that higher tax rates reduce consumption, using

a cohort-specific tax series is likely to bias our estimates downward. Furthermore, the bias would be more pronounced for groups that react the most. Bearing these caveats in mind, Figure 12 presents the results, which broadly confirm our previous findings. The responses of social renters and owners are largely unaffected. Furthermore, mortgagors still respond the most, although the peak effect is reduced to around 1.5 pounds (consistent with the sign of the bias discussed before). It therefore remains the case that households with debt respond the most and those without — the outright owners — hardly respond at all.

6.4 Compositional changes and selection effects

Whilst the shares of social renters, mortgagors and outright owners have varied slowly over time (see Figure 1), both compositional changes and selection into group might bias our estimates. To explore this possibility, we adapt the methodology proposed by Attanasio et al. (2002) to generate individual predicted probabilities of owning mortgage debt. Specifically, we run a probit regression over the full sample to generate individual predicted probabilities of having a mortgage based on a high order polynomial in age, education, a time trend and their interactions. For households observed in period t , we compute the probability that they had a mortgage four quarters earlier. For these two periods, we classify households as ‘likely’ or ‘unlikely’ mortgagors’ depending on whether the probability in the first of the two periods is larger or smaller than the share of mortgagors in that period. We then take the difference in consumption across these two periods for each group.²⁴

The results for this sample split are shown in Figure 13. Consistent with the evidence based on actual housing tenure, the response of households who are unlikely to hold a mortgage is largely never statistical significant at the 32% level. In contrast, the dynamic effects of the exogenous tax change on the non-durable consumption of the ‘likely’ mortgagors is always significant at the 5% level and, in line with the

²⁴Running separate probit regressions for each year or using a fixed threshold produces very similar results. The estimates of the probit regression for the whole sample are reported in Appendix B.

estimates in Figure 5, peak at values above 2 pounds after three years. It is still the case that the 95% confidence bands for the likely mortgagors do not include the point estimates for the unlikely mortgagors at horizons beyond six quarters. We conclude that the potential bias associated with compositional changes and selection effects is likely to be negligible when grouping households by their housing tenure status.

7 Concluding remarks

Recent years have witnessed a renewed interest in the role of household debt in the transmission of macroeconomic shocks. Theoretical studies have formalised the idea that some agents may become liquidity constrained by making a large durable purchase. A main implication of this ongoing research effort is that, following an exogenous change in taxes, households with mortgage debt could increase their consumption by more than those without.

This paper finds empirical support for this hypothesis. Using a long span of household survey data and a new narrative measure of exogenous income tax changes for the U.K., we find that the consumption response of mortgagors is significantly larger than the response of outright owners. As the demographics of these groups are similar in at least some dimensions, we argue that traditional sample splits (based for instance on age and education) may understate the full extent of heterogeneity in the data. In contrast, mortgage debt appears to correlate better with the unobserved characteristics driving the diverse consumption behaviours.

Liquidity constraints for lower income/less educated households is an often cited explanation for the sizable aggregate effects of tax changes on GDP and consumption that are typically reported in the empirical macro literature. On one hand, we show that social renters — who tend to have no wealth, low income and compulsory education — do respond significantly to tax changes. This type of liquidity constrained households, however, accounts for only around 20% of our sample and therefore appears too small a share to drive the aggregate results. On the other hand,

mortgagors — who tend to hold little liquid wealth despite owning sizable housing equity — account for about half of the British population, thereby providing a novel interpretation for the sizable impact of tax changes on the aggregate economy.




Figure 1: Shares of social renters, mortgagors and outright home owners.

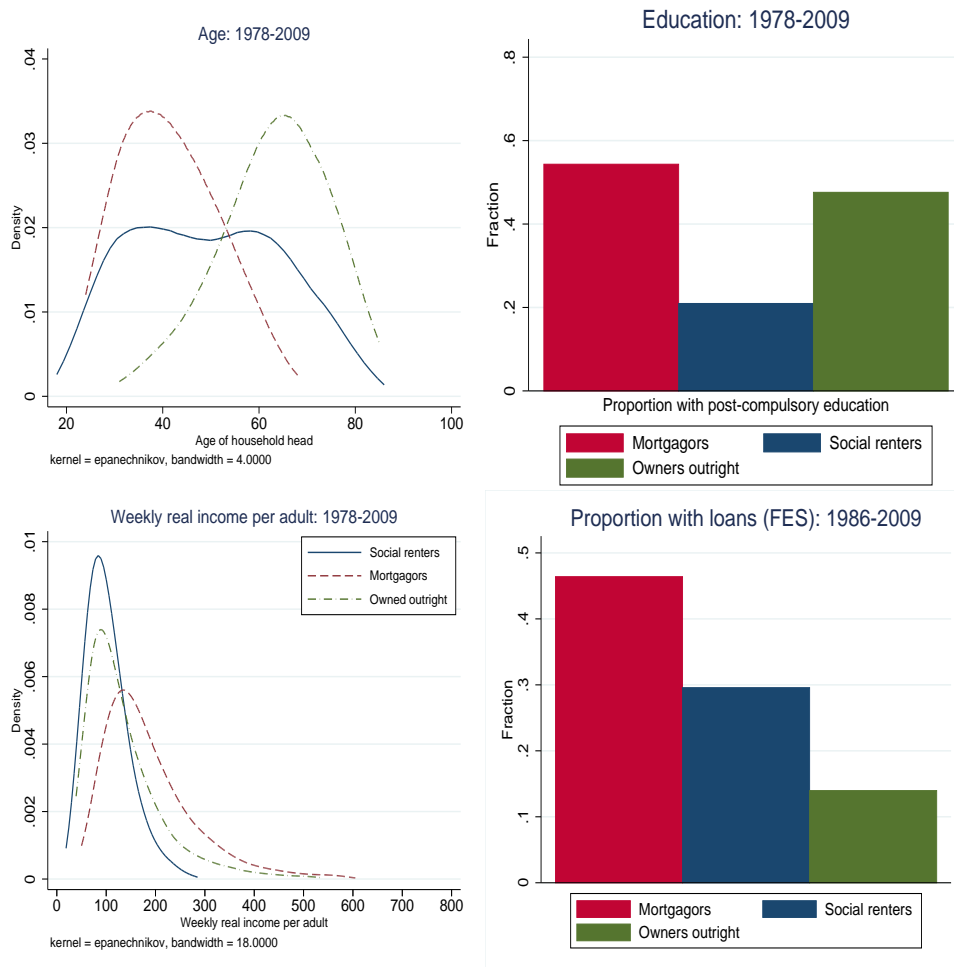


Figure 2: Distributions of age and income and proportions of households with post-compulsory education and non-mortgage loans across housing tenures.




Figure 3: Tax liability changes over GDP: income tax measure (red) vs. all exogenous tax changes (black dashed)




Figure 4: Dynamic effects of a per-taxpayer liability change in the allowance and basic rate of income tax on the change in per-capita non-durable goods and services consumption: aggregated Family Expenditure Survey (FES) vs. Office for National Statistics (ONS). Point estimates based on FES (ONS) data are reported as solid black line with circles (dashed blue line). Shaded areas (grey lines) represent 68% (95%) confidence bands.

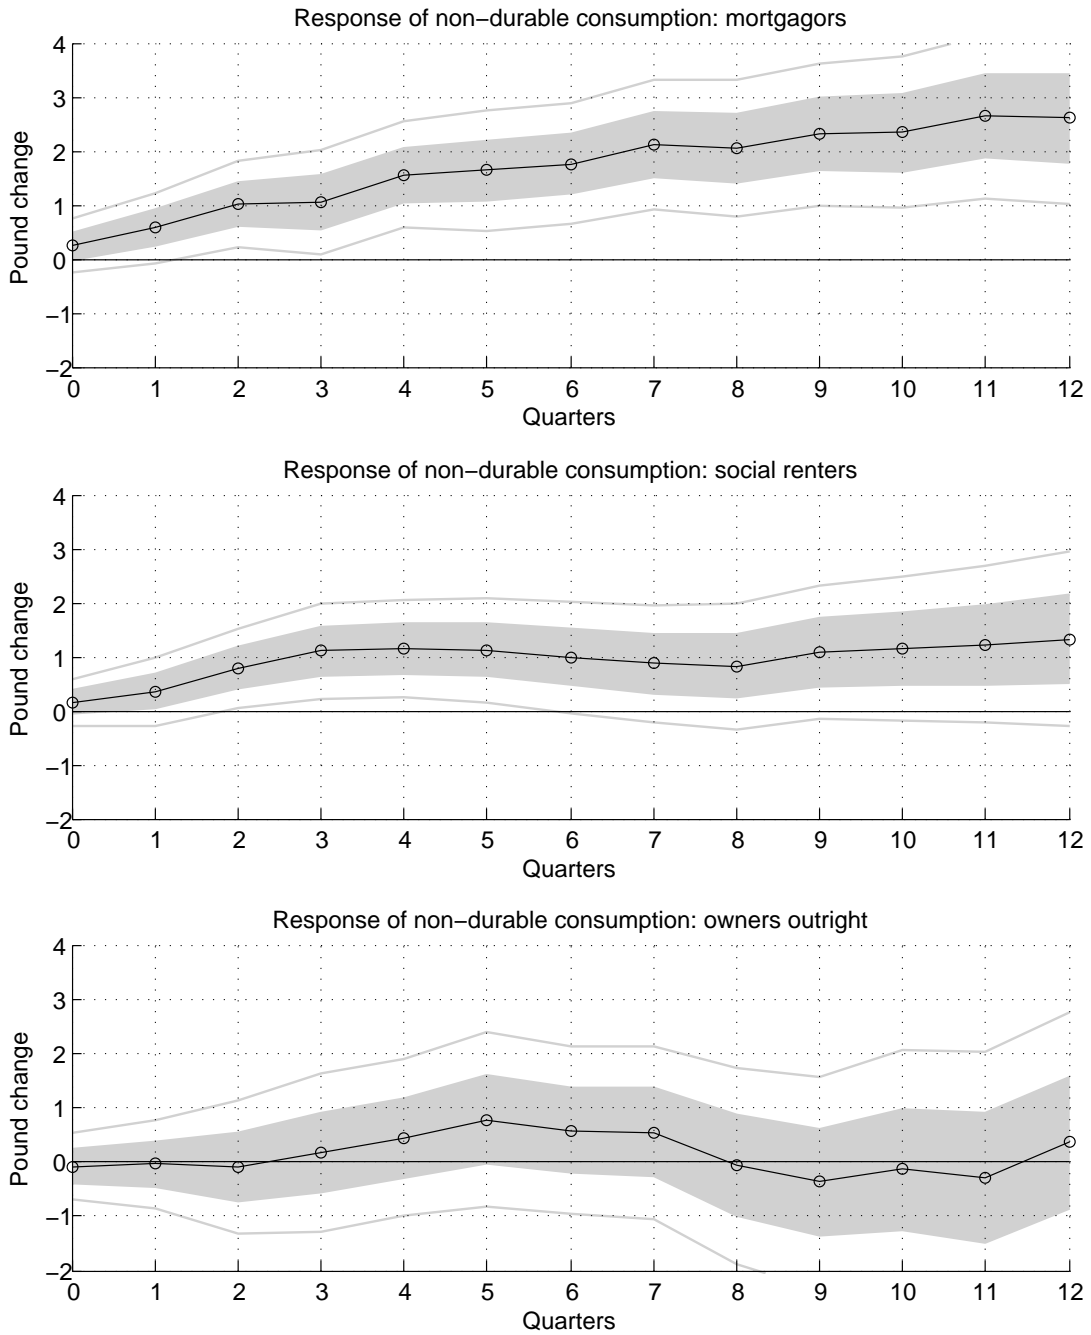


Figure 5: Dynamic effects of a per-taxpayer liability change in the allowance and basic rate of income tax on the change in per-capita non-durable goods and services consumption across housing tenures. Shaded areas (grey lines) represent 68% (95%) confidence bands.

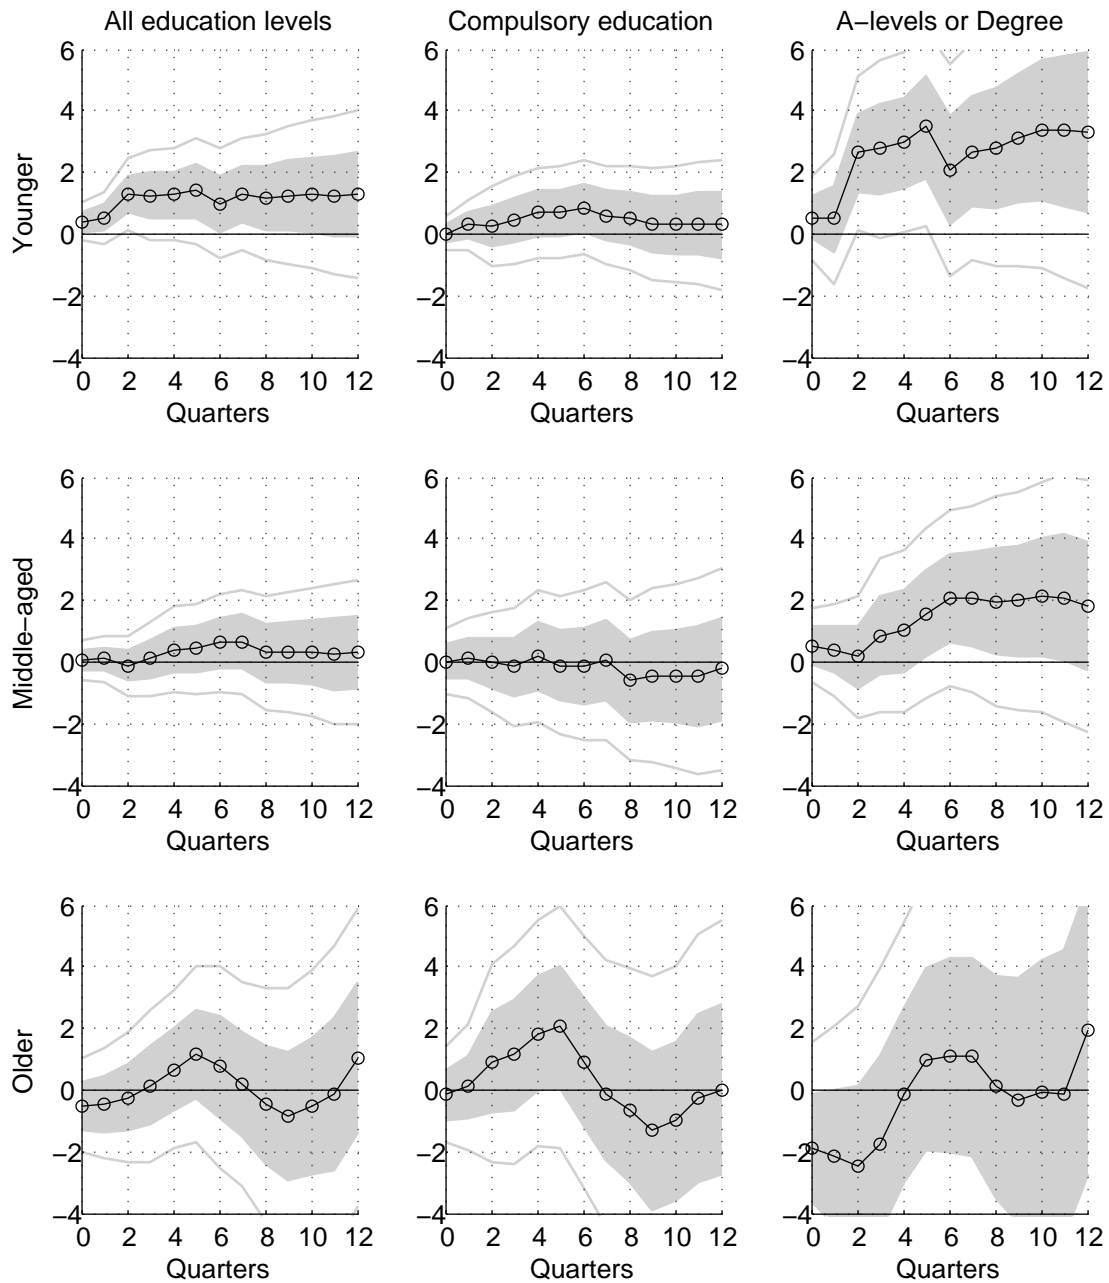


Figure 6: Dynamic effects of a per-taxpayer liability change in the allowance and basic rate of income tax on the change in per-capita non-durable goods and services consumption across birth cohorts and education. Shaded areas (grey lines) represent 68% (95%) confidence bands.

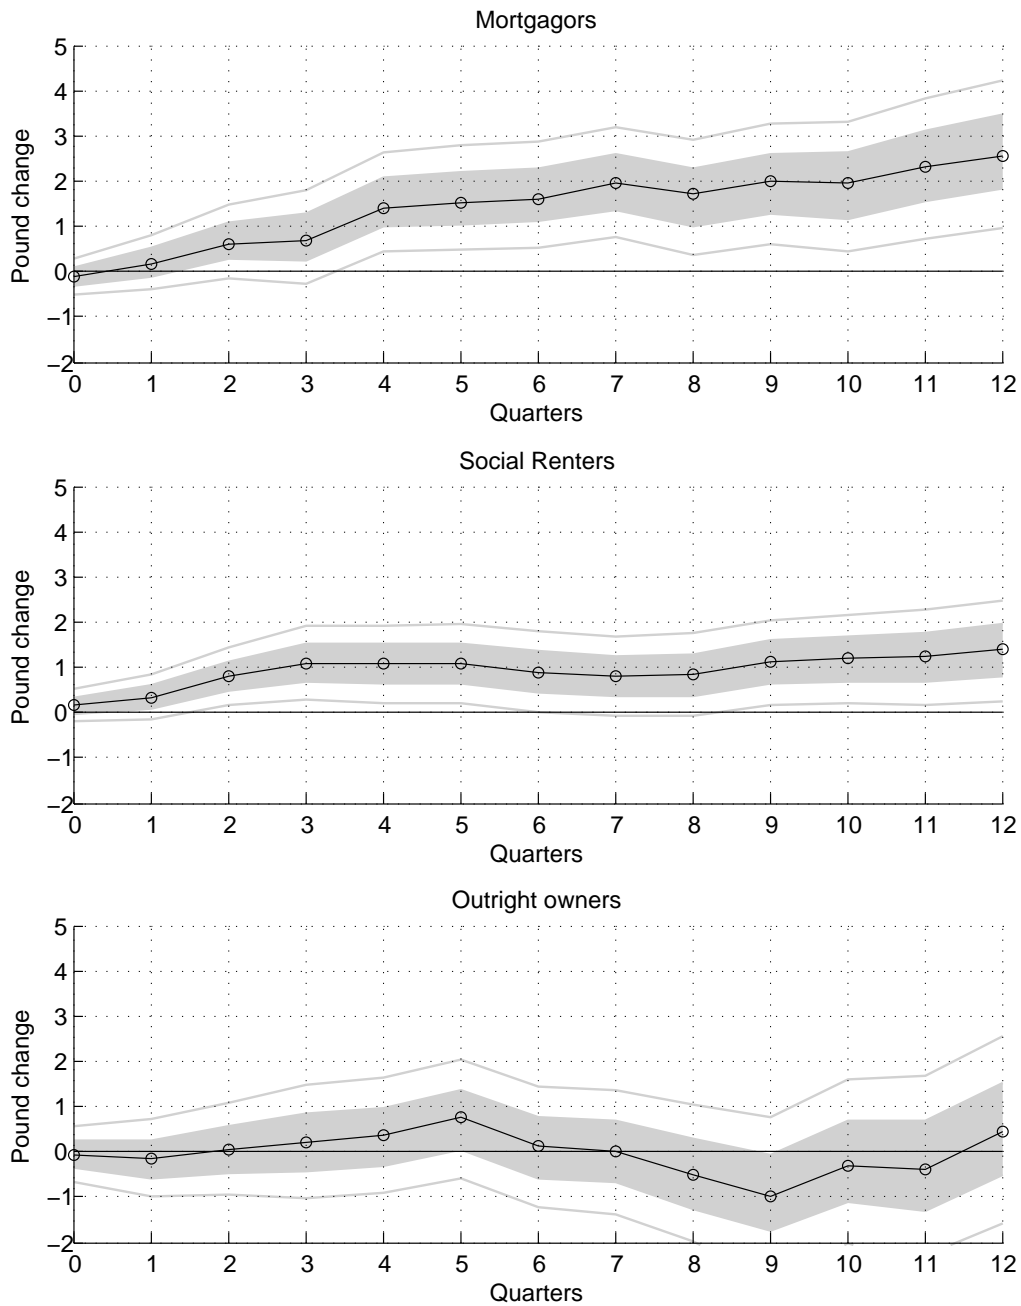


Figure 7: Dynamic effects of a per-taxpayer liability change in the allowance and basic rate of income tax on the change in per-capita non-durable goods and services consumption across housing tenures using a VAR(4) in non-durable consumption per-capita change, real GDP per-capita change, real government spending per capita change, inflation and the Bank rate. Shaded areas (grey lines) represent 68% (95%) confidence bands.

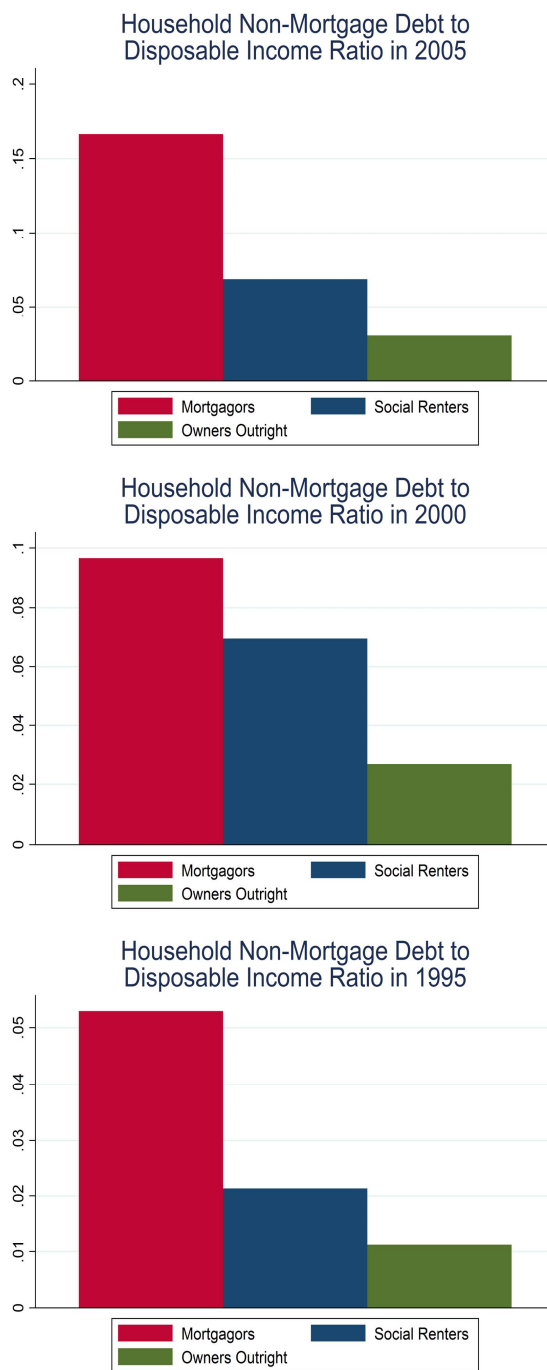


Figure 8: Non-mortgage debt includes hire purchase agreements, personal loans, credit and store cards, catalogue or mail order purchase agreements, DWP social fund loans, loans from private individuals, overdrafts and student loans. Disposable income refers to net annual income where the appropriate adjustments are made for income taxes, national insurance, pension contributions and local taxes (following Levy and Jenkins (2012)). Source: British Household Panel Survey.

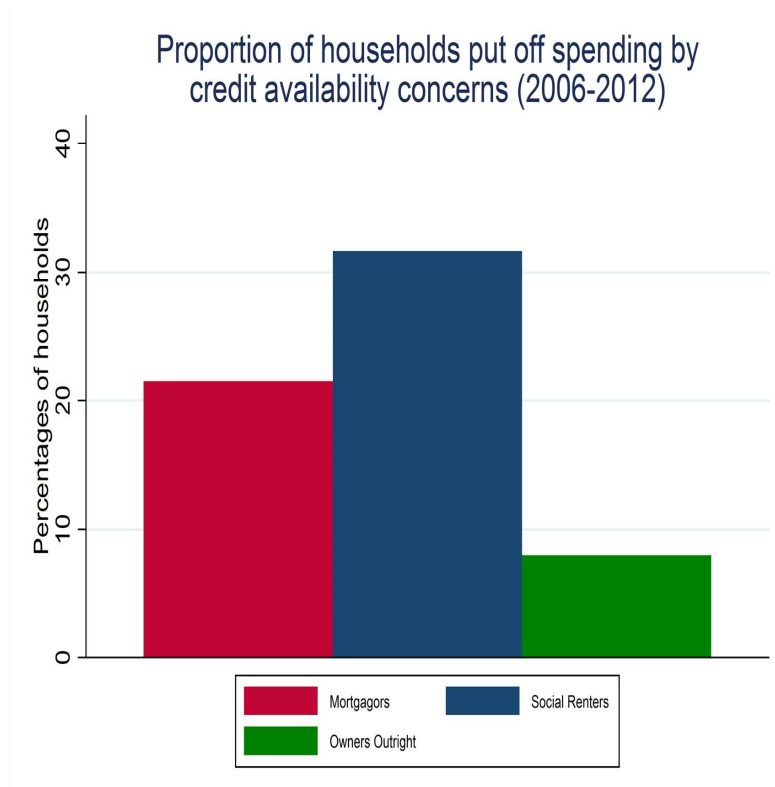


Figure 9: Proportion of households put off spending by credit availability concerns by housing tenure. Sample: 2006-2012. Source: NMG Consulting survey carried out for the Bank of England.

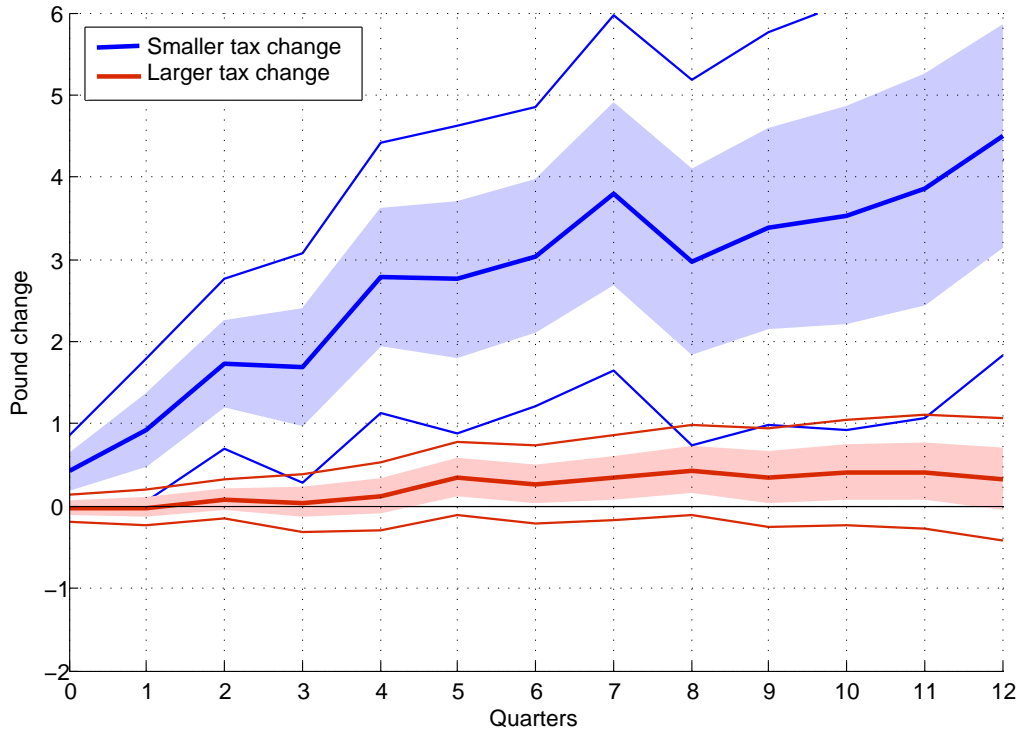


Figure 10: Dynamic effects of a per-taxpayer liability change on the change in per-capita non-durable goods and services consumption for mortgagors. ‘Smaller change’ refers to an average liability change in the allowance corresponding to about 250 pounds per household over three years at 2009 prices. ‘Larger change’ refers to an average liability change in the basic and higher rates of income tax corresponding to about 750 pounds per household over three years at 2009 prices. Shaded areas (solid lines) represent 68% (95%) confidence bands.

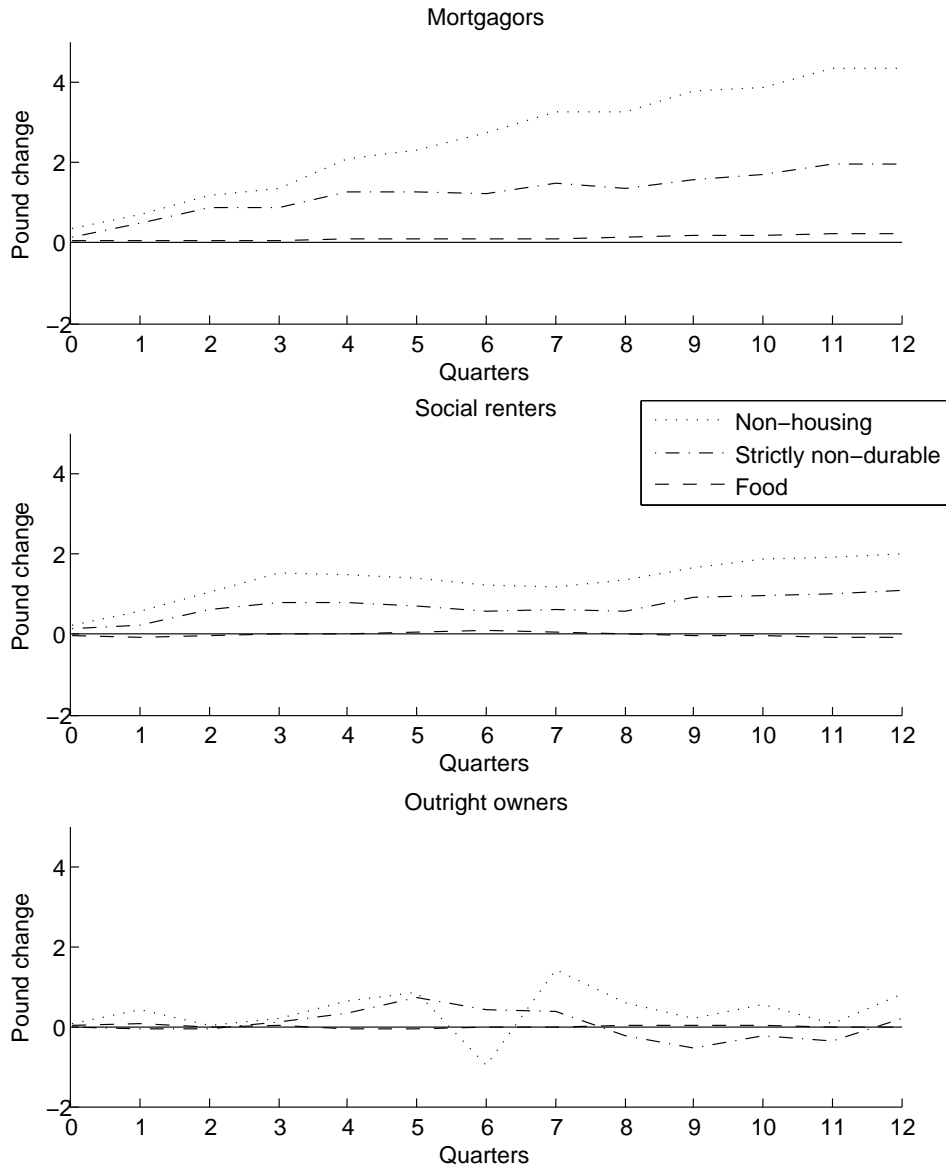


Figure 11: Point estimates of the dynamic effects of a per-taxpayer liability change in the allowance and basic rate of income tax on the change in per-capita expenditure on food, strictly non-durable and non-housing goods and services across housing tenures. Shaded areas (grey lines) represent 68% (95%) confidence bands.

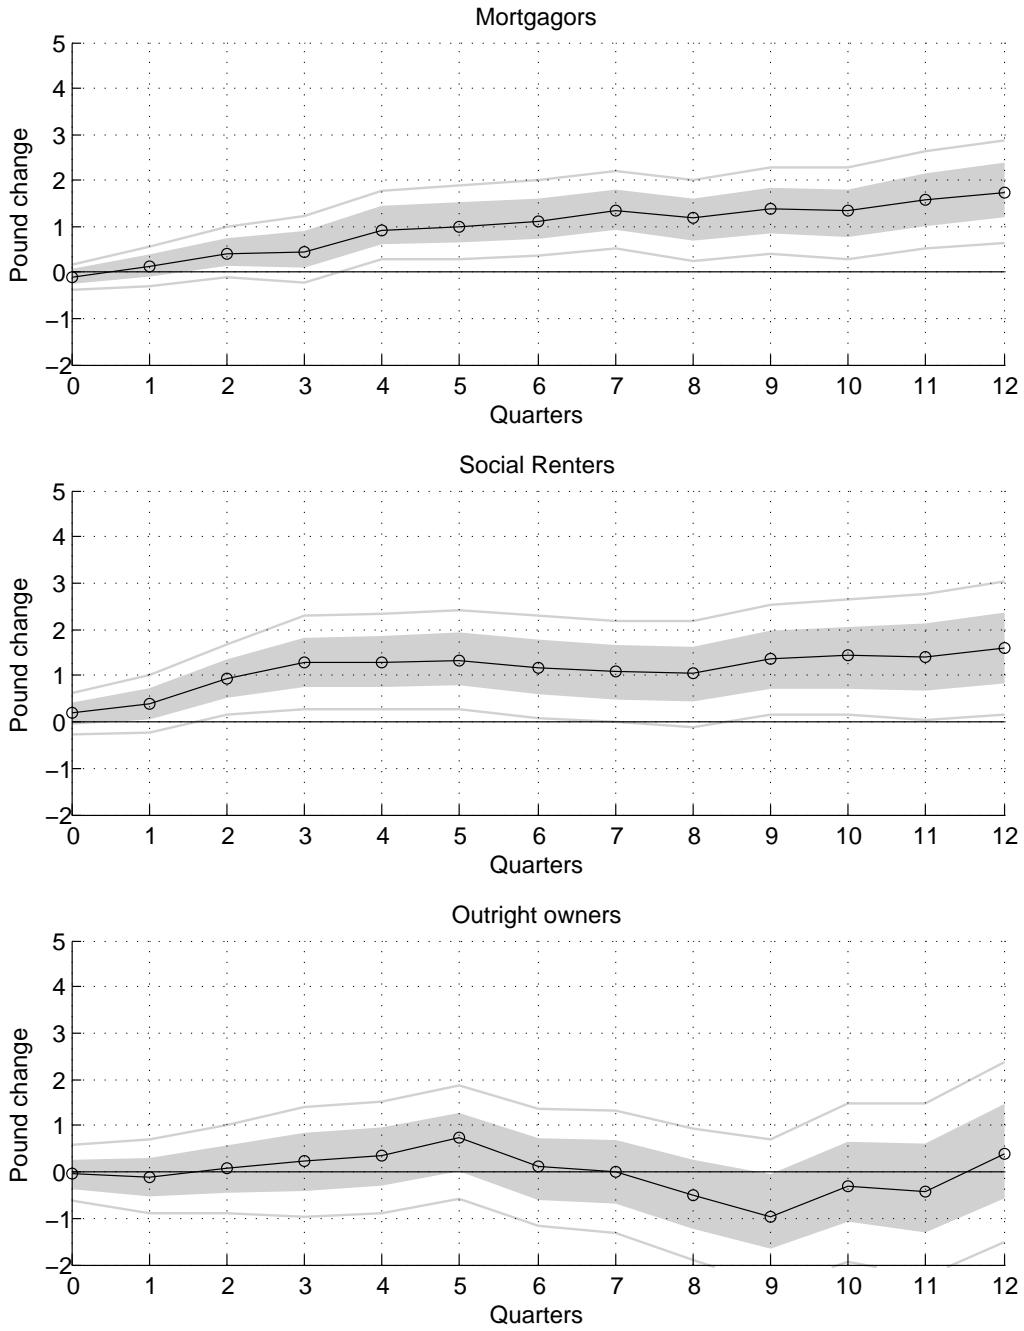


Figure 12: Dynamic effects of a per-taxpayer liability change in the allowance and basic rate of income tax on the change in per-capita non-durable goods and services consumption across housing tenure controlling for the amount of the tax windfall. Shaded areas (grey lines) represent 68% (95%) confidence bands.

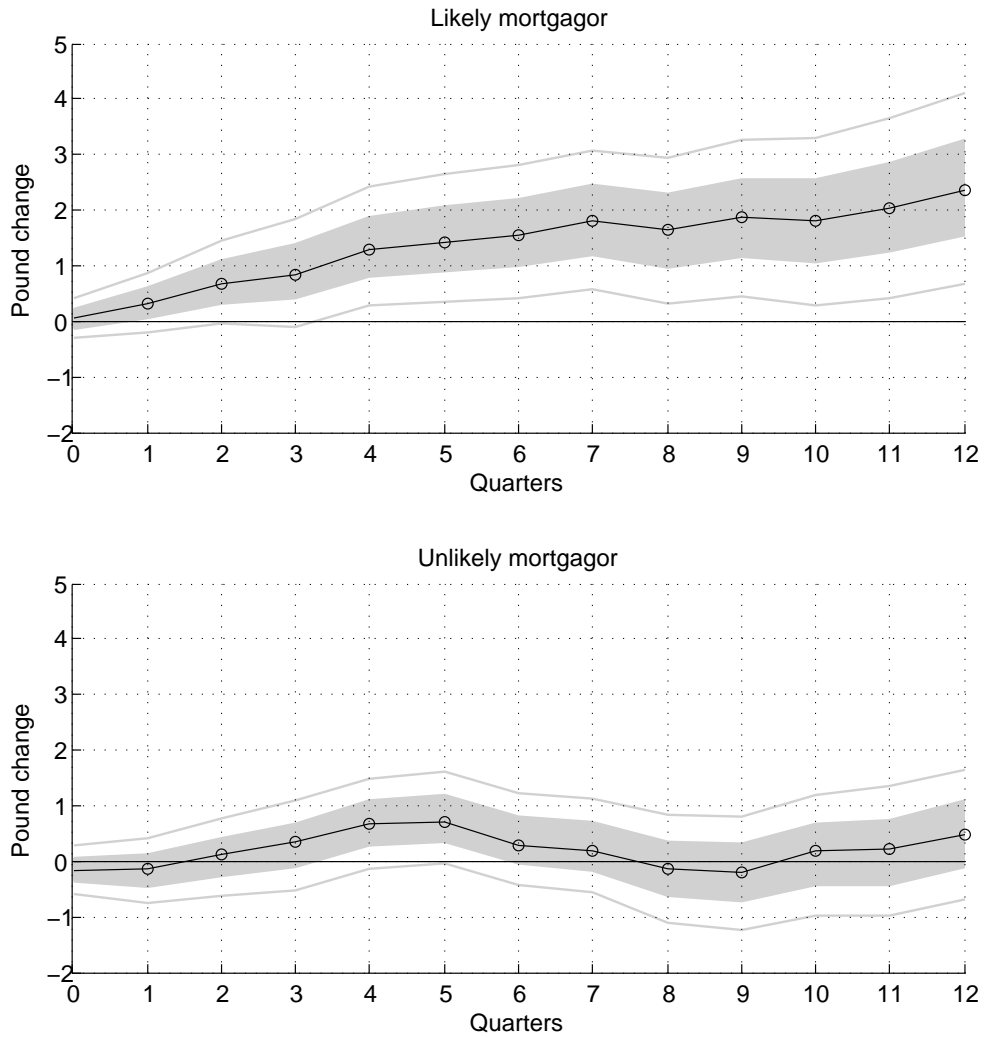


Figure 13: Dynamic effects of a per-taxpayer liability change in the allowance and basic rate of income tax on the change in per-capita non-durable goods and services consumption across housing tenure likelihoods predicted by probit estimation. Shaded areas (grey lines) represent 68% (95%) confidence bands.

TABLE 1: TESTING FOR HETEROGENEOUS CONSUMPTION RESPONSES

	allowance & basic: ABR	incl. all other tax changes	unantici- pated ABR
response by tenure	<i>effect at the one year horizon (β_{y1}^i)</i>		
<i>mortgagors (M)</i>	1.56** (0.50)	1.72** (0.42)	1.77** (0.59)
<i>social renters (S)</i>	1.15** (0.47)	1.52** (0.52)	1.10* (0.57)
<i>owners outright (O)</i>	0.44 (0.74)	0.31* (0.25)	0.11 (0.29)
<i>p-values for H0:</i>			
$\beta_{y1}^M = \beta_{y1}^S = \beta_{y1}^O$	0.38	0.01	0.03
$\beta_{y1}^M = \beta_{y1}^O; \beta_{y1}^M = \beta_{y1}^S$.16;.52	.01;.75	.01;.37
	<i>effect at the two year horizon (β_{y2}^i)</i>		
<i>mortgagors (M)</i>	2.05** (0.64)	2.17** (0.50)	2.22** (0.68)
<i>social renters (S)</i>	0.83* (0.61)	1.15** (0.52)	0.91* (0.61)
<i>owners outright (O)</i>	-0.07 (0.93)	-0.28 (0.52)	-0.44 (0.65)
<i>p-values for H0:</i>			
$\beta_{y2}^M = \beta_{y2}^S = \beta_{y2}^O$	0.11	0.01	0.02
$\beta_{y2}^M = \beta_{y2}^O; \beta_{y2}^M = \beta_{y2}^S$.05;.15	.01;.13	.01;.10
	<i>effect at the three year horizon (β_{y3}^i)</i>		
<i>mortgagors (M)</i>	2.61** (0.81)	2.48** (0.60)	2.45** (0.77)
<i>social renters (S)</i>	1.33* (0.84)	1.41** (0.55)	1.20* (0.64)
<i>owners outright (O)</i>	0.35 (1.22)	0.43 (0.47)	0.29 (0.63)
<i>p-values for H0:</i>			
$\beta_{y3}^M = \beta_{y3}^S = \beta_{y3}^O$	0.24	0.03	0.11
$\beta_{y3}^M = \beta_{y3}^O; \beta_{y3}^M = \beta_{y3}^S$.10;.26	.01;.17	.04;.18

Note: based on $\Delta c_t^i = \alpha^i + \sum_{j=0}^{12} \beta_j^i \Delta T_{t-j} + \gamma^i Z_t^i + u_t^i$ with $i=S, M$ and O where i stands for social renters (S), owners with (M) or without a mortgage (O). $\beta_{yf}^i = \sum_{j=0}^{f*4} \beta_j^i$ with $f=1, \dots, 3$. c_t is per-capita non-durable expenditure on goods and services. Z_t^i include Δage_t^i , Δemp_t^i , $\Delta educ_t^i$ and $\Delta retire_t^i$ where age refers to the household head, the dummy emp ($retire$) is one if the household head is employed (retired). $educ$ is 0, 1 or 2 for compulsory education, A level and degree level, respectively. In columns 1 and 2, ΔT_t refers to the change in our Allowance and Basic Rate measure of income tax liability divided by RPIX and total numbers of tax payers (ABR). In columns 2-3, we add all per-payer liability changes other than ABR as further control and drop the coefficients on ΔT_{t-j} that are not statistically significant at 32% level. In column 3, $\Delta T_{t-j's}$ become the unanticipated part of ABR. *(**) denotes 32(10)% significance level. Standard errors in parenthesis.

TABLE 2: NET WEALTH POSITION BY HOUSING TENURE

PANEL A: 2005 BRITISH HOUSEHOLD PANEL SURVEY AT 2005 POUNDS					
	<i>p25</i>	<i>median</i>	<i>p75</i>	<i>mean</i>	<i>obs</i>
NET FINANCIAL WEALTH					
<i>social renters</i>	-400	0	150	470	1,337
<i>mortgagors</i>	-3,250	0	4,600	3,014	3,179
<i>owners outright</i>	0	3,000	21,540	18,293	2,385
NET HOUSING WEALTH					
<i>social renters</i>	0	0	0	0	1,337
<i>mortgagors</i>	55,000	95,000	150,000	113,257	3,179
<i>owners outright</i>	100,000	150,000	230,000	189,385	2,385
PANEL B: 2000 BRITISH HOUSEHOLD PANEL SURVEY AT 2000 POUNDS					
	<i>p25</i>	<i>median</i>	<i>p75</i>	<i>mean</i>	<i>obs</i>
NET FINANCIAL WEALTH					
<i>social renters</i>	-300	0	250	959	1,959
<i>mortgagors</i>	-2,000	82	5,500	4,197	3,234
<i>owners outright</i>	0	3,500	20,000	15,889	2,009
NET HOUSING WEALTH					
<i>social renters</i>	0	0	0	0	1,959
<i>mortgagors</i>	16,000	38,000	71,400	54,572	3,234
<i>owners outright</i>	50,000	75,000	120,000	97,284	2,009
PANEL C: 1995 BRITISH HOUSEHOLD PANEL SURVEY AT 1995 POUNDS					
	<i>p25</i>	<i>median</i>	<i>p75</i>	<i>mean</i>	<i>obs</i>
NET FINANCIAL WEALTH					
<i>social renters</i>	-50	0	400	1,185	944
<i>mortgagors</i>	-750	193	3,690	4,209	2,024
<i>owners outright</i>	25	3,000	16,540	15,711	1,054
NET HOUSING WEALTH					
<i>social renters</i>	0	0	0	0	944
<i>mortgagors</i>	29,000	10,000	55,000	37,162	2,024
<i>owners outright</i>	44,000	60,000	85,000	72,069	1,054

Note: Net financial wealth is the value of savings and investments less outstanding non-mortgage debt. Housing wealth is the household's estimate of the property value net of any outstanding mortgage. *p25* is 25th percentile and *p75* is 75th percentile.

A Data description

A.1 Aggregate data

- *Non-housing, non-durable goods and services expenditure*: UK Office for National Statistics categories Non-durable goods plus Services, Semi-Durable Goods (such as clothing and household maintenance) minus housing and water (codes: utiq+utii+utim-adft-adfu-adfw).
- *Price index*: Retail Prices Index excluding mortgage interest payments (RPIX), ONS code chmk.
- We scale our tax measure by the total number of taxpayers, available from Her Majesty's Revenue and Customs' website (www.hmrc.gov.uk)

A.2 Household micro data

We use the Family Expenditure Survey (later called the Expenditure and Food Survey and, recently, the Living Costs and Food Survey) from 1978 to 2009. These data are available from the UK Data Archive.

Household consumption expenditure

- *Non-durable goods and services expenditure*: includes food, alcohol, tobacco, fuel, light and power, clothing and footwear, personal goods and services, fares, leisure services, household services, non-durable household goods, motoring expenditures and leisure goods.
- *Durable expenditure*: durable household goods, motor vehicles and durable leisure goods. This includes expenditure such as furniture and furnishings, electrical appliances and audio-visual equipment.
- *Total non-housing expenditure*: total expenditure minus housing expenditures (including rents, rates and water charges).

- We gross-up using household weights and divide by the number of household members to construct a per capita measure.

Housing tenure and non-mortgage loans

The FES records the tenure status of households. Social renters are defined as those living in local authority housing or accommodation provided by housing associations. Mortgagors and owners are taken directly from the FES. The renter category is dropped due to the limited number of observations. The FES also asks a household for any loans they have. We use this variable to define non-mortgage borrowers and non-borrowers.

Demographics

The demographic variables are taken from the FES: age, education and employment status of the head of household.

B Demographics over time

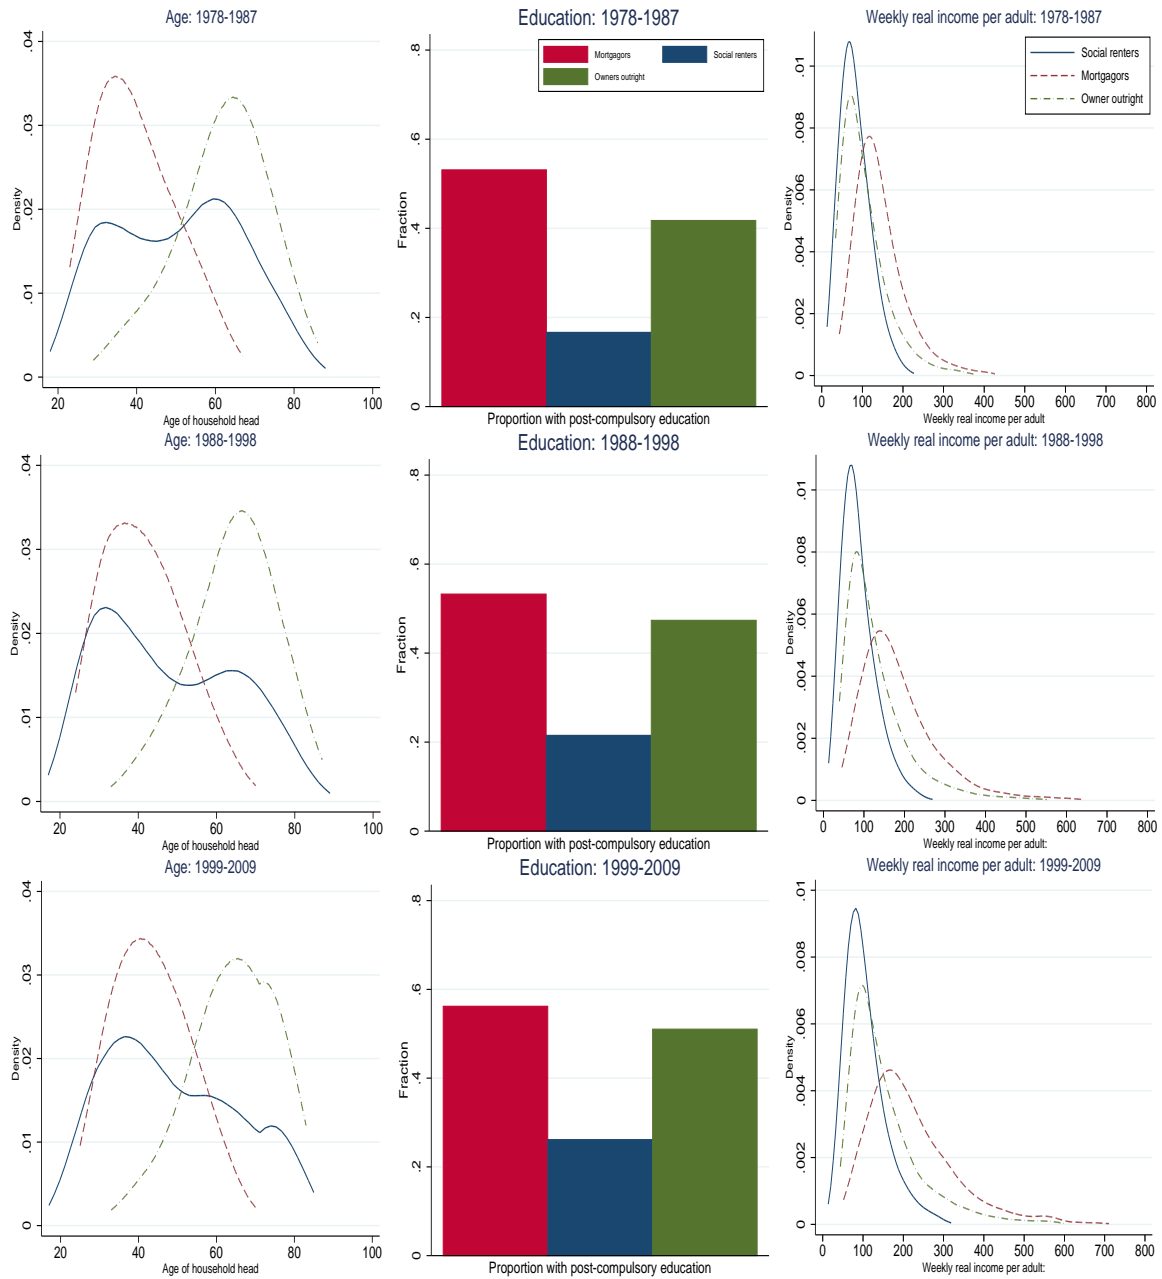


Figure 14: Evolution of age, education and income distributions across housing tenures.

C Probit regressions results

Table A.1: Probit estimation results for the likelihood of being a mortgagor

Variable	Coefficient	Std. Err.
Age of head	0.087	(0.002)
(Age of head) ²	-0.001	(0.000)
Head has A levels	1.061	(0.034)
College education	1.821	(0.050)
Age × A levels	-0.010	(0.001)
Age × College	-0.018	(0.001)
Trend	-0.045	(0.004)
Trend ²	0.001	(0.000)
Trend ³	0.000	(0.000)
Trend ⁴	0.000	(0.000)
Trend ⁵	0.000	(0.000)
Age × Trend ²	0.000	(0.000)
Age × Trend	0.000	(0.000)
A levels × Trend	-0.003	(0.000)
College × Trend	-0.005	(0.000)
Constant	-0.653	(0.064)

D Tax change examples

In this section we illustrate how the exogenous tax change series is constructed. As noted above, the Financial Statement and Budget Report (FSBR) categorises all the individual tax changes in each budget, together with their projected revenue effects and implementation dates.

Table A.2 illustrates two examples of the identification approach. The first row considers one of the larger changes in income tax announced in 1986. The FSBR records a “Reduction of 1p in the basic rate”, which the Treasury projected to cost 830 million each year.²⁵ The measure came into effect from April 6th 1986 and is assigned to 1986 quarter 2.

The second step is to determine whether the policy change was reacting to other macroeconomic fluctuations. The table shows the motivation given by Chancellor Nigel Lawson in the budget speech. There is no mention of this tax reform reflecting business cycle movements and was a purely supply-side measure designed to improve long-run economic performance. Providing a one sentence illustration is limiting but gives the essential idea behind the classification. In practice, we use a richer variety of statements, commentaries and contexts to construct our specific income tax change series, building on the classification strategy in Cloyne (2012).

The second row illustrates a tax change that would be regarded as endogenous — reacting to other macroeconomic fluctuations. Consider the stimulus package announced in 2008 as a response to the deepening recession. Chancellor Alistair Darling announced in November 2008 that an earlier temporary increase in the personal allowance would be made permanent. As the quote shows, this was part of a package of measures designed to stimulate the economy. This income tax change (and the other tax changes announced) is endogenous and excluded from our series.

²⁵For more information on how the UK Treasury calculates these revenue projections see, for example, H.M. Treasury (2005) page 205.

Table A.2: Example tax changes

Policy change	Announced	Implemented	Treasury projected revenue effect (£m)	Motivation given	Classification
1pp reduction in basic rate of tax	18/3/1986	6/4/1986	-830	“It is no accident that the two most successful economies in the world, both overall and specifically in terms of job creation those of the United States and Japan - have the lowest level of tax as a proportion of GDP. Reductions in taxation motivate new businesses and improve incentives at work. They are a principal engine of the enterprise culture, on which our future prosperity and employment opportunities depend”. ^a	Exogenous
Increase in personal allowance	24/11/2008	6/4/2008	-3370	“To prevent the recession from deepening, we also need to take action to put money into the economy immediately”. ^b “The Government has already taken action to help people through the current global economic difficulties. Building on this, the Government announces further packages of targeted support, providing additional help to those who need it most now”. ^c	Endogenous

Sources: (a) Hansard, HC Deb 18 March 1986 vol 94 c182. (b) Hansard, HC Deb 24 November 2008, c495. (c) HM Treasury (2008): Pre-Budget Report.

E Other expenditure categories: confidence bands

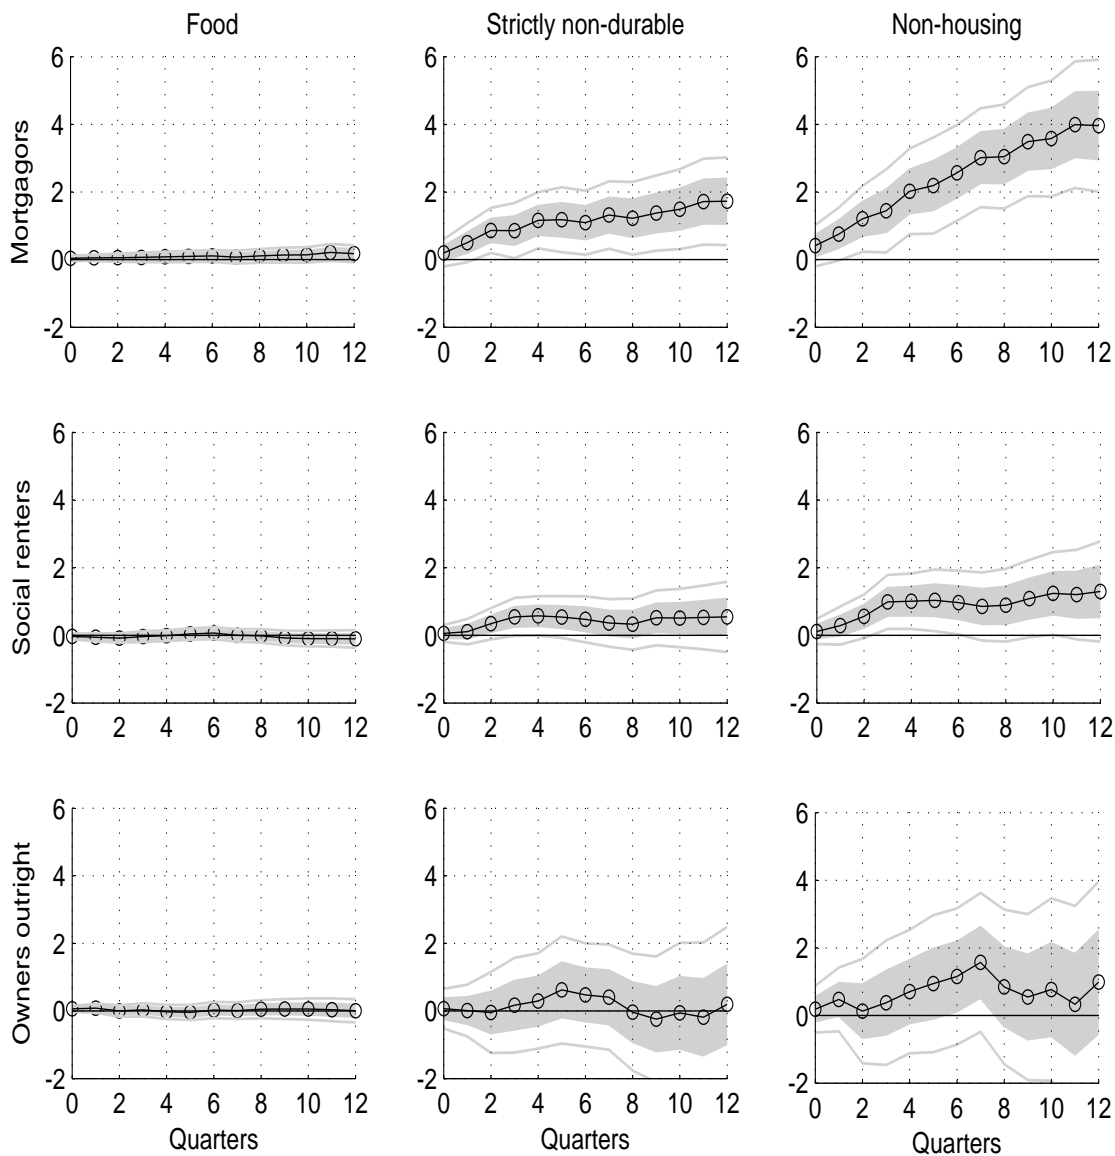


Figure 15: Dynamic effects of a per-taxpayer liability change in the allowance and basic rate of income tax on the change in per-capita expenditure on food, strictly non-durable and non-housing goods and services across housing tenures. Shaded areas (grey lines) represent 68% (95%) confidence bands.

References

- Acconcia, A., G. Corsetti, and S. Simonelli (2011). Mafia and public spending: Evidence on the fiscal multiplier from a quasi-experiment. *CSEF WP No. 281*.
- Anderson, E., A. Inoue, and B. Rossi (2012). Heterogeneous consumers and fiscal policy shocks. *mimeo, Duke University, NCSU and UPF*.
- Andres, J. A., J. E. Bosca, and J. Ferri (2012). Household leverage and fiscal multipliers. *Banco de Espana Documentos de Trabajo No. 1215*.
- Attanasio, O., J. Banks, and S. Tanner (2002). Asset holding and consumption volatility. *Journal of Political Economy* 110, 771–792.
- Attanasio, O., A. Leicester, and M. Wakefield (2011). Do house prices drive consumption growth? The coincident cycles of house prices and consumption in the UK. *Journal of the European Economic Association* 9, 399–435.
- Attanasio, O. and G. Weber (1993). Consumption growth, the interest rate and aggregation. *The Review of Economic Studies* 60, 631–649.
- Blanchard, O. and R. Perotti (2002). An empirical characterization of the dynamic effects of changes in government spending and taxes on output. *Quarterly Journal of Economics* 117, 1329–1368.
- Bottazzi, R., T. Crossley, and M. Wakefield (2010). Late starters or excluded generations? a cohort analysis of catch up in home ownership in england. *Institute for Fiscal Studies WP 12/10*.
- Browning, M., A. Deaton, and M. Irish (1985). A profitable approach to labour supply and commodity demands over the life-cycle. *Econometrica* 53(3), 503–44.
- Caldara, D. and C. Kamps (2012). The analytics of SVARs: A unified framework to measure fiscal multipliers. *FEDS Working Paper 2012-20, Board of Governors of the Federal Reserve System*.

- Campbell, J. and J. Cocco (2007). How do house prices affect consumption? Evidence from micro data. *Journal of Monetary Economics* 54, 591–621.
- Campbell, J. and G. Mankiw (1989). Consumption, income and interest rates: Reinterpreting the time series evidence. *NBER Macroeconomics Annual* 4, 185–216.
- Cloyne, J. S. (2012). Discretionary tax shocks in the United Kingdom 1945-2009: a new narrative account and dataset. <http://ideas.repec.org/e/pcl85.html>.
- Cloyne, J. S. (2013). Discretionary tax changes and the macroeconomy: new narrative evidence from the United Kingdom. *American Economic Review* 103(4), 1507–28.
- Crossley, T. and C. O’Dea (2010). The wealth and saving of uk families on the eve of the crisis. *Institute for Fiscal Studies*.
- De Giorgi, G. and L. Gambetti (2012a). Consumption heterogeneity and the effects of aggregate shocks. *mimeo, Stanford and Autonoma de Barcelona*.
- De Giorgi, G. and L. Gambetti (2012b). The effects of government spending on the distribution of consumption. *mimeo, Stanford and Autonoma de Barcelona*.
- Dynan, K. (2012). Is a household debt overhang holding back consumption? *Brooking Papers on Economic Activity Spring*, 299–362.
- Eggertsson, G. and P. Krugman (2012). Debt, deleveraging, and the liquidity trap: A Fisher-Minsky-Koo approach. *Quarterly Journal of Economics* 127, 1469–513.
- Ercolani, V. and N. Pavoni (2012). The precautionary saving effect of government consumption. *Mimeo, Bank of Portugal and Bocconi University*.
- Giavazzi, F. and M. McMahon (2012). The household effects of government spending. *in Fiscal Policy after the Financial Crisis, Eds. Alesina, A. and Giavazzi, F., University of Chicago Press*.

- Guajardo, J., D. Leigh, and A. Pescatori (2011). Expansionary austerity: New international evidence. *IMF WP/11/158*.
- H.M. Treasury (2005). *Budget 2005*. The Stationary Office.
- IMF (2012). Dealing with household debt. *Global Economic Outlook Ch. 3*, 89–124.
- Jappelli, T. and L. Pistaferri (2010). The consumption response to income changes. *Annual Review of Economics 2*, 479–506.
- Kaplan, G. and G. L. Violante (2013). A model of the consumption response to fiscal stimulus payments. *Manuscript, Princeton University and NYU* <http://scholar.princeton.edu/gkaplan/files/kaplan-violante-wp-2013.pdf>.
- Levy, H. and S. Jenkins (2012). Documentation for derived current and annual net household income variables, bhps waves 1-18. *UK Data Archive Study Number 3909*.
- Lusardi, A. (1996). Permanent income, current income, and consumption: Evidence from two panel data sets. *Journal of Business and Economic Statistics 14*, 8190.
- Mertens, K. and M. O. Ravn (2012). Empirical evidence on the aggregate effects of anticipated and unanticipated U.S. tax policy shocks. *American Economic Journal: Economic Policy 4(2)*, 145–181.
- Mertens, K. and M. O. Ravn (2013). The dynamic effects of personal and corporate income tax changes in the United States. *American Economic Review 103(4)*, 1212–47.
- Mian, A., K. Rao, and A. Sufi (2012). Household balance sheets, consumption, and the economic slump. *mimeo, Princeton University and University of Chicago*.
- Miles, D. (2004). The uk mortgage market: Taking a longer-term view. *Report prepared for HM Treasury*.

- Misra, K. and P. Surico (2013). Consumption, income changes and heterogeneity: Evidence from two fiscal stimulus programmes. *CEPR Discussion Paper No. 9530*.
- Mountford, A. and H. Uhlig (2009). What are the effects of fiscal policy shocks. *Journal of Applied Econometrics* 24(6), 960–992.
- Parker, J., N. Souleles, D. Johnson, and R. McClelland (2013). Consumer spending and the economic stimulus payments of 2008. *American Economic Review Forthcoming*.
- Pistaferri, L. and E. I. Saporta (2012). Changes in the income distribution and aggregate consumption. *mimeo, Stanford University*.
- Ramey, V. and M. Shapiro (1998). Costly capital reallocation and the effects of government spending. *Carnegie Rochester Series on Public Policy* 48, 145–194.
- Ramey, V. A. (2011). Identifying government spending shocks: It’s all in the timing. *Quarterly Journal of Economics* 126(1), 1–50.
- Romer, C. D. and D. H. Romer (1989). Does monetary policy matter? A new test in the spirit of Friedman and Schwartz. *NBER Macroeconomic Annual* 1989.
- Romer, C. D. and D. H. Romer (2004). A new measure of monetary shocks: Derivation and implications. *American Economic Review* 94(4), 1055–84.
- Romer, C. D. and D. H. Romer (2010). The macroeconomic effects of tax changes: Estimates based on a new measure of fiscal shocks. *American Economic Review* 100(3), 763–801.