

Etheridge, Ben

Working Paper

House prices and consumption inequality

IFS Working Papers, No. W19/24

Provided in Cooperation with:

The Institute for Fiscal Studies (IFS), London

Suggested Citation: Etheridge, Ben (2019) : House prices and consumption inequality, IFS Working Papers, No. W19/24, Institute for Fiscal Studies (IFS), London, <https://doi.org/10.1920/wp.ifs.2019.1924>

This Version is available at:

<https://hdl.handle.net/10419/223268>

Standard-Nutzungsbedingungen:

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

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

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

Terms of use:

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

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

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

House prices and consumption inequality

IFS Working Paper W19/24

Ben Etheridge

House Prices and Consumption Inequality*

Ben Etheridge[†]
University of Essex

16th January 2019

Abstract

I characterize how house price shocks affect consumption inequality using a life-cycle model of housing and non-housing consumption with incomplete markets. I derive analytical expressions for the dynamics of inequalities and use these to analyze large house price swings seen in the UK. I show that movements in consumption inequality were large, that they correspond with the theoretical predictions qualitatively, and that the model explains a large fraction of the movements quantitatively. I demonstrate the accuracy of this analysis using an extended model's full non-linear solution. Finally, accounting for house price shocks alters estimates of labour-income risks using cross-sectional data.

JEL Classification: D31, D91, E21

Keywords: Consumption inequality, Wealth effect, Collateral effect, Housing

*Thanks to Richard Blundell, Tom Crossley and Hamish Low for invaluable advice and encouragement. Sincere thanks also to the editor, Dirk Krueger, and to three anonymous referees, whose comments helped substantially improve the paper. Further thanks to Michael Graber and to Cormac O'Dea for comments. I also thank numerous individuals and seminar audiences for further comments. Financial assistance was provided by the ESRC through the Institute for Fiscal Studies. All errors are my own.

[†]E-mail: bsethe@essex.ac.uk, www: <https://sites.google.com/site/benetheridge/>

Introduction

Across much of the world, house prices have fluctuated substantially in the past 30 years, seemingly in connection with the wider macroeconomy. Accordingly, the effect of house price shocks on aggregate consumption has been extensively researched.¹ Yet given that housing wealth and home ownership are unevenly distributed, it seems important to ask not only how house price shocks affect mean consumption, but also how they affect consumption inequality. I contribute to the literature on house prices and consumption by addressing this question.

Of course, an extensive literature has examined how house prices affect consumption across different groups. These groups include age and housing tenure (for example, [Li and Yao, 2007](#)) and, perhaps most saliently, region (for example, [Mian, Rao, and Sufi, 2013](#)). Nevertheless, in this paper I find that the effect of house prices shocks on inequality is both starker and richer. I show that even a common house price shock affects inequality within groups defined by age, education and region. This within-group measure of inequality is important because it explains the bulk of total inequality.² Specifically I find that a positive common house price shock generally causes such consumption inequality to increase.

This paper therefore builds on two literatures. First, on that assessing the effect of house prices on consumption. Second, on that explaining the evolution of consumption inequality. This second literature has typically focused on the relationship between consumption inequality and income inequality, and particularly the role of shocks from the labour market.³ I contribute to this second literature by exploring the role of shocks to asset prices: in this case house prices.

I organize the analysis around a relatively standard partial equilibrium model of household choices over (non-housing) consumption, housing tenure and housing size in an incomplete-market setting. This benchmark model includes exogenous aggregate house price shocks, idiosyncratic income shocks, transactions costs and borrowing restrictions determined by collateral constraints.

I use this framework to make three specific contributions. First I provide an analytic characterization of how house price shocks affect consumption inequality. To do this, I simplify the benchmark model by abstracting from transactions costs and collateral constraints. I also use inequality measures which decompose appropriately: these are the variance-of-the-logarithm measure, and associated covariances, as used since [Deaton and Paxson \(1994\)](#). By simplifying the benchmark model, and then by approximating the shock to consumption in terms of the house-price and labour-market shocks, I track how house price shocks, in particular, drive the inequality measures. More precisely, I decompose the consumption response to a house price shock into three factors: an

¹For an early literature using similar data to that used here, see [Campbell and Cocco \(2007\)](#) and [Attanasio, Blow, Hamilton, and Leicester \(2009\)](#) for the UK and [Carroll, Otsuka, and Slacalek \(2011\)](#) for the US. See also the literature discussed below in the dedicated literature review.

²On within-group inequality in labour earnings, the literature is vast. It has focussed in particular on the inequality ‘boom’ in the 1980s. Perhaps most relevant in the present context is [Moffitt and Gottschalk \(2002\)](#). For consumption inequality, see for example [Cutler and Katz \(1992\)](#), [Krueger and Perri \(2006\)](#), [Blundell, Pistaferri, and Preston \(2008\)](#) and [Aguiar and Bils \(2015\)](#). [Meyer and Sullivan \(2017\)](#) provide a wide ranging overview of inequality trends in both income and consumption in the U.S since the 1960s. They informally discuss the role of asset prices in the joint evolution of these inequality measures.

³In addition to those papers focusing more explicitly on the inequality boom in the 1980s, see also [Storesletten, Telmer, and Yaron \(2004\)](#), [Krueger and Perri \(2005\)](#), [Blundell, Low, and Preston \(2013\)](#), [Heathcote, Storesletten, and Violante \(2014\)](#) and [Blundell, Pistaferri, and Saporta-Eksten \(2016\)](#).

income effect, a substitution effect, and an endowment/wealth effect.⁴ As such, my decomposition relates to that in [Berger, Guerrieri, Lorenzoni, and Vavra \(2018\)](#), who derive an interpretable sufficient statistics formula for the approximate effect of house price shocks on consumption. My formula is similar to theirs, but is more easily used to study heterogeneity and inequality movements, because each component of my formula is observable at the household level.

In terms of the three factors, I argue that consumption inequality is driven mainly by heterogeneity in the endowment effect. This heterogeneity arises not only because of the division between owners and renters, but also, importantly, because of dispersion in leverage. Ultimately, I express the effect of house price shocks on inequality in terms of key statistics, which can be computed from cross-sectional micro-data. At the level of the household, the key statistic is an endowment, or housing-wealth share: the proportion of *gross* housing wealth in total life-time wealth. In terms of inequality movements, the key statistic is how this housing-wealth share covaries with the consumption and income distributions. It is important to remember that, while inequality is driven by the house price shocks, it is also, of course, driven by idiosyncratic shocks from the labour market. This is the first paper, therefore, to characterize how these shocks affect consumption inequality in combination.

By simplifying the model I am able to derive this analytic characterization. However, this characterization comes at the cost of abstracting from important features of the housing and mortgage markets, such as collateral constraints. I therefore test the accuracy of my approximations by solving the full benchmark model and disciplining it using UK data. I show that the approximations driven by a housing wealth effect are close to the true house price elasticity, and that the approximations pick up variation in the true elasticity well in several dimensions.

My second contribution is to apply the framework quantitatively to observed movements in inequality. I examine the UK from the late 1980s until the financial crisis around 2008. This is a suitable period for the application. House prices displayed large swings: a large decline in the early 1990s followed by a boom of unprecedented length. Over this period the labour market was also comparatively stable.⁵ More practically, during this period, detailed data are available on household wealth in 1995, 2000 and 2005 from the British Household Panel Survey. This dataset also contains data on incomes and on food consumption. My headline inequality measures come from the Family Expenditure Survey (FES), which has detailed cross-sectional data on incomes and on broad expenditures.⁶

I first show that the approximated model is highly consistent with broad movements in the data. Specifically, I analyse fixed groups of households given by decade-of-birth cohorts. As discussed, the inequality movements depend most importantly on how the housing-wealth share is distributed. The data show that, for each group, this share is positively correlated with consumption, and increasingly so over time. Because of this positive

⁴Following a house price increase, the income effect captures the result of the reduction in real life-time resources owing to the increase in cost of housing services; the substitution effect captures the shift away from housing services holding fixed life-time utility; the endowment effect captures the pure wealth effect from the increase in value of housing assets.

⁵The UK experienced a mild recession in the early 1990s during which the unemployment rate rose by 3 percentage points to a little over 10%. Thereafter unemployment declined and stayed below 6% from 1999 onwards. On the other hand, inequality was also influenced by redistributive reforms beginning in 1997. These reforms do not affect the present analysis. See [Etheridge \(2013\)](#), an earlier working paper version of the current study, for a fuller discussion.

⁶As is relatively standard in the literature, I drop the distinction between consumption and expenditure in this paper.

correlation, the model implies that house price increases should cause the gap between consumption inequality and other measures to grow. Similarly, house price declines should cause the gap to fall. For both the groups I study, this gap fell in the 1990s when house prices declined and grew strongly thereafter, consistently with the model.

I then quantify these movements, first using the wealth share approximation. Focussing on the peak years of the house price boom over 1997-2004, the gap between consumption inequality and the other measures grew by 4.1 log points for the 1950s cohort and by 5.8 log points for the 1960s.⁷ These are large changes compared with previous inequality movements. By computing the statistics described above, I show the simplified model explains around 50% of the observed gap for the 1950s cohort, although for the 1960s cohort the model explains less. A key strength of my approach, here, is that I can impose the cross-sectional distributions of the relevant variables into the calculations directly.

I then extend this analysis by examining the role of several additional features, most importantly collateral effects.⁸ To this end I use the full benchmark model's non-linear solution, and measure the strength of collateral constraints at the household level using housing leverage. I use the model to compute a simulation-based consumption elasticity which I can then take to the data. Qualitatively and empirically, I find that leverage covaries positively with the income and consumption distributions. Therefore the housing wealth and collateral channels work in the same direction: Through both channels, a positive house price shock causes consumption inequality to increase. Quantitatively, collateral effects are less important in driving inequality because households with high leverage do not tend to lie at the extremes of the consumption distribution.

My third and final contribution is to assess how accounting for house price shocks affects estimates of income risks. As discussed above, researchers have used the cross-sectional distributions of income and consumption to quantify risks and risk sharing at least since since [Deaton and Paxson \(1994\)](#) and [Blundell and Preston \(1998\)](#). In my model, ignoring house price increases causes estimates of the variance of permanent income shocks to be biased upwards. This is because the researcher mis-attributes growth in consumption inequality to income shocks rather than to shocks from asset prices. Likewise, I find that accounting for the house price declines in the first part of the sample period increases the estimated variance of permanent shocks during the recession in the early-1990s.

This paper proceeds as follows. I first provide a dedicated discussion of the literature on housing, which has recently developed rapidly. Section 1 presents the benchmark model. I discuss the simplified model's approximated solution and its solution for the effect of house price shocks on inequality measures in section 2. I discuss the various sources of data in Section 3, where I also discuss the method used to estimate income risks. Section 4 presents the primary results based on housing wealth shares. I examine the accuracy of the wealth-share approximations in section 5, where I also extend the quantitative analysis to include collateral effects, and two additional features: a time-varying rent-price ratio and an elasticity of substitution between consumption and housing that differs from 1. Section 6 includes estimates of income risk, while section 7

⁷In other words, the house price increases caused the variance-of-log measure to increase by 0.058 for the 1960s cohort. I use the term 'log points' similarly throughout the paper.

⁸The collateral effect captures the result of relaxing or tightening borrowing constraints, holding house prices fixed.

concludes. A lengthy online appendix contains detailed derivations, discussion of the model’s non-linear solution, and additional results. In particular, appendix C extends the theoretical framework to allow for the presence of shocks to aggregate income, and provides a general equilibrium interpretation.

Literature Review

In terms of house prices and consumption my paper is closest to Berger et al. (2018). As discussed, they derive an interpretable sufficient statistics formula for the approximate effect of house price shocks on consumption, which can be taken to micro-data. My formula is different from theirs but contains important similarities, which I discuss in detail in section 2. Briefly, although my formula is a less close approximation to the consumption elasticity than theirs, it can be used in the present context more easily. Specifically, implementing their formula requires calculating marginal propensities to consume out of transitory income, at the household level. Calculating household-specific *MPCs* is infeasible with survey data. My formulae, on the other hand, can be implemented easily with data on income, consumption and wealth. Moreover, when using the simulation approach, I am able to build on my baseline approximations to account for some of the additional features that they include, in particular collateral effects.

Aside from Berger et al., the literature on house prices and consumption has expanded rapidly. Progress has been made both on empirical, and on more theoretical and quantitative grounds.

On the empirical side, much of the earlier literature on house price elasticities focused on UK data, in part because of the importance of housing wealth in household portfolios in the UK. Examples include Attanasio et al. (2009), Disney, Gathergood, and Henley (2010) and Campbell and Cocco (2007), who find the largest elasticities. Using US data both Cooper (2013) and Aladangady (2017) find sizeable elasticities, and argue that these effects are driven by those with high leverage, implying an important role for collateral effects. Relative to these empirical estimates it is worth emphasizing that my results do not require the consumption elasticity to be large on average. My focus, instead, is on how this elasticity varies across the population. Nevertheless my quantitative analysis does imply a larger importance of wealth over collateral effects.

More recently the literature has turned to administrative data on broader measures indicative of household consumption, such as automobile sales, credit card expenditure, and mortgage deals. Much of this work is surveyed by Mian and Sufi (2016), alongside other lines of research.⁹ In influential contributions among this strand, Mian and Sufi (2011) argue that US households boosted their consumption before the financial crisis by extracting home equity. Mian et al. (2013) then argue that the distribution of housing wealth and leverage explain differential declines in consumption across US locations during the great recession. The literature has since proceeded to explore in detail how the relationship between housing and consumption is mediated by various features of mortgage contracts: Bhutta and Keys (2016) compare the effects of interest rate changes versus house price changes on equity extraction, while Ganong and Noel (2016) examine the effects of mortgage debt forgiveness. In this paper I abstract from detailed features of mortgage contracts, but, in section 5 in

⁹Also see Piazzesi and Schneider (2016) for a survey based on an asset-pricing perspective.

particular, I embed the discussion within some of these recent results. Also relevant to the present study on inequality, recent papers have emphasized that the expansion of mortgage availability in the boom in the US, and subsequent delinquency during the bust, was seen evenly across the distribution of income and credit scores. Papers in this strand are [Adelino, Schoar, and Severino \(2016\)](#), [Foote, Loewenstein, and Willen \(2016\)](#) and [Albanesi, De Giorgi, and Nosal \(2017\)](#).¹⁰ These papers are consistent with my results because I find that leverage is, if anything, positively correlated with the income and consumption distributions in the UK.

On the quantitative side, the present paper is most similar to those using partial equilibrium models, such as [Li and Yao \(2007\)](#), [Li, Liu, Yang, and Yao \(2016\)](#) and [Díaz and Luengo-Prado \(2010\)](#), who focus on the distribution of wealth, in an infinite horizon setting. I use a model in this class to assess the life-cycle evolution of consumption inequality. More recently the literature has seen the development of several general equilibrium models. In many of these models house prices are driven by changes to credit conditions ([Garriga, Manuelli, and Peralta-Alva, 2017](#); [Midrigan and Philippon, 2011](#); [Huo and Ríos-Rull, 2016](#); [Favilukis, Ludvigson, and Van Nieuwerburgh, 2017](#)). Typically in these models, consumption is then driven by a combination of the changes to credit conditions directly, and indirectly through the subsequent changes to house prices. A crucial component to the quantitative relevance of these models is correctly capturing the dispersion in various asset and debt positions. Focussing on this heterogeneity more explicitly, but without modeling housing, [Krueger, Mitman, and Perri \(2016\)](#) review evidence from the PSID on the role of the distribution of income, consumption, assets and debt in aggregate fluctuations. Perhaps most relevantly to the current paper, [Kaplan, Mitman, and Violante \(2017\)](#) establish a model which allows for shocks not only to credit conditions and to aggregate productivity, but also to beliefs about future housing needs. In part because of their modeling of a rental market they find that changes to credit conditions have little effect on house prices and on aggregate consumption. In their model house prices and consumption are driven more by changes to beliefs. Given that their belief shock has little direct effect on consumption, but works through the resulting effect on house prices, this shock works similarly to the exogenous house price shock used here.¹¹ Additionally, they find that house prices affect aggregate consumption mainly through the wealth channel, rather than through relaxing or tightening constraints.

1 Theoretical Framework

In this section I specify a fairly standard model of non-housing consumption and housing choices with an exogenous income process and house price risk.¹² This benchmark model is rich enough to feature choices over

¹⁰Among other relevant studies using administrative data, [Browning, Gørtz, and Leth-Petersen \(2013\)](#) examine house price effects directly using administrative panel data from Denmark.

¹¹In related work [Landvoigt \(2017\)](#) uses a life-cycle model, and observed choices, to infer expectations about future house price movements.

¹²For the rest of the paper I use the terms ‘consumption’ and ‘non-housing consumption’ interchangeably. The shorter version is usually preferred for brevity.

ownership and renting and to feature transactions costs in housing and collateral constraints.

1.1 Choices and Preferences

Households are born at time $t = 0$, work until $t = T_w$ and die at time $t = T$. At time t , household i maximizes expected lifetime utility by choosing a state-contingent plan for housing size, H_{ik} , and another composite consumption good, C_{ik} , in each future period $k \geq t$:

$$\max_{\{C_{ik}, H_{ik}\}_{k=t}^T} \mathbb{E}_t \left(\sum_{k=t}^T \beta^{k-t} u(C_{ik}, H_{ik}) + \beta^{T-t+1} B(\tilde{W}_{iT+1}) \right) \quad (1)$$

where β is a subjective rate of time preference, assumed to be common across households, $u()$ is the within-period felicity function during life, $B()$ is a warm-glow bequest function and \tilde{W}_{iT+1} is the stock of wealth bequeathed. I omit formalizing the contingency of C_{ik} , H_{ik} and \tilde{W}_{iT+1} on the history of shocks here for brevity of notation. In addition to the size or quality of housing, households can choose between renting and owning. This tenure decision is addressed below.

I restrict preferences to be separable and homothetic across time. Within each period, preferences can be fully general. For my main results, however, I require $u()$ to be homothetic across consumption and housing. This intratemporal homotheticity restriction is common when modeling the demand for housing.¹³¹⁴ A suitable general class of homothetic preferences is those with the constant elasticity of substitution (CES) form:

$$u(C, H) = \frac{1}{1-\gamma} \left(\left(aC^{\frac{\tau-1}{\tau}} + (1-a)H^{\frac{\tau-1}{\tau}} \right)^{\frac{\tau}{\tau-1}} \right)^{1-\gamma}$$

where τ is the elasticity of substitution between consumption and housing, a captures the relative preference for non-housing consumption and the elasticity of intertemporal substitution is constant and given by $\frac{1}{\gamma}$. In the special case when $\tau = 1$, preferences are Cobb-Douglas with the form $u(C, H) = \frac{1}{1-\gamma} (C^a H^{1-a})^{1-\gamma}$. These Cobb-Douglas preferences give the cleanest final results.

I ignore deterministic changes to consumption (due, for example, to the presence of children) for simplicity of exposition. These could be re-introduced and would affect the anticipated gradient on consumption growth. In an empirical application, this gradient can be netted out by the econometrician, using observations on household size and other characteristics.

The bequest function is given by

$$B(\tilde{W}) = \Phi \frac{1}{1-\gamma} \tilde{W}^{1-\gamma}$$

where Φ measures the strength of preferences for bequests versus within-life consumption.

¹³As is standard, I distinguish between *inter*-temporal and *intra*-temporal hometheticity. Preferences are *inter*-temporally homothetic if both rich and poor households allocate total expenditures in the same proportions across time. Preferences are *intra*-temporally homethetic if both rich and poor households allocate expenditures across goods in the same proportion within time periods.

¹⁴Davis and Ortalo-Magné (2010), for example, argue that preferences are Cobb-Douglas, using evidence on wages and house prices across time and across US metropolitan statistical areas.

1.2 House Price Process

House prices evolve as follows:

$$p_t^H = p_{t-1}^H e^{\mu + \zeta_t^H}$$

where ζ_t^H is the shock to (the log of) house prices, p_t^H , with $\mathbb{E}_{t-1} e^{\zeta_t^H} = 1$, and μ is a constant reflecting trend house price increases.

This specification is simple, and arguably standard among related studies. Specifically, most studies treat house prices as a random walk with drift.¹⁵ On the other hand, house prices have historically shown long swings both up and down. Arguably, therefore, house prices could be modeled with richer dynamics.¹⁶ I justify the simple specification, as usual, by noting that the random walk model is a suitable benchmark for asset prices. Nevertheless, I empirically assess how house prices evolve in section 3.

My specification is perhaps less standard by imposing that house prices are common across all households, and by neglecting more local changes. Many related studies, in contrast, use regional price movements for identification.¹⁷ Of course, if house price shocks are heterogeneous then they are more likely to increase consumption dispersion.¹⁸ My specification is therefore on the conservative side. Again, in section 3, I empirically assess national versus more local price movements.

1.3 Constraints

The flow budget constraint for those renting at time t is:

$$A_{it} + C_{it} + \chi p_t^H H_{it} = RA_{it-1} + Y_{it} + (1 - \kappa)(1 - \delta) p_t^H d_{it-1} H_{it-1} \quad (2)$$

The left hand side of equation 2 gives expenditures at time t . These expenditures comprise financial savings, A_{it} , consumption, C_{it} , and housing rent. Rent is determined by the time-invariant rent-price ratio, χ , current house price, p_t^H , and house size, H_{it} . The right hand side of equation 2 shows time- t resources. These resources comprise liquid savings from the previous period, A_{it-1} , returned at the risk free interest rate, R , alongside net income, Y_{it} and housing wealth, $p_t^H d_{it-1} H_{it-1}$, net of depreciation δ , and sales costs κ . d_{it-1} is an indicator of ownership in time $t - 1$. Housing wealth, if it exists, is inherited from choices made at time $t - 1$ and re-valued at time t prices.

The flow budget constraint for those owning at time t is:

$$A_{it} + C_{it} + p_t^H H_{it} = RA_{it-1} + Y_{it} + (1 - \mathbf{1}_{d_{it-1} H_{it-1} \neq H_{it}}(\kappa))(1 - \delta) p_t^H d_{it-1} H_{it-1} \quad (3)$$

¹⁵See, for example, [Campbell and Cocco \(2007\)](#) in the context of the UK.

¹⁶See, for example, [Ortalo-Magné and Rady \(2006\)](#), and references therein. More recently [Guren \(2018\)](#) provides a model of house price momentum.

¹⁷Again, [Campbell and Cocco \(2007\)](#) and [Attanasio et al. \(2009\)](#) are examples in the UK setting.

¹⁸Of course, the truth of this statement depends on the correlation of house price shocks with other features of the distribution. Heterogeneous house price shocks would increase consumption inequality if, for example, they are unforecastable from current information and hence uncorrelated with income and wealth.

Similarly to the constraint for renters, the left hand of equation 3 gives (gross) expenditures at time t , the right hand side gives (gross) resources. The household does not pay transactions costs, κ , if it stays in a house of the same size, in which case the net expenditures on housing are given by depreciation, $\delta p_t^H H_{it}$, only.

Finally, in the benchmark model, household borrowing is restricted by a collateral constraint:

$$A_{it} \geq (1 - \Xi) \frac{1 - \delta}{R} p_t^H d_{it} H_{it} \quad (4)$$

A household can only borrow by owning a house and putting it up for collateral. Financial net borrowing is limited by a downpayment or equity requirement Ξ . Of course, most houses are purchased with mortgage financing. As discussed in the introduction, much of the recent literature has focussed on the intricacies of mortgage contracts, and their role in mediating the link between house price shocks and consumption. Many of these intricacies are apparent in the UK context. These include: counter-cyclical mortgage pricing, a counter-cyclical rent-to-price ratio, long-term mortgage contracts, time-varying rules on mortgage initiation, and costly re-financing. The model presented here abstracts from all these detailed features, but nevertheless generates the key mechanism which can drive consumption: house price rises loosen borrowing restrictions while house price declines tighten them. I return to some of these further features of mortgage contracts throughout the paper, and especially in section 5.

Finally, bequeathed assets, \tilde{W}_{iT+1} are determined as follows:

$$\tilde{W}_{iT+1} = RA_{iT} + Y_{iT} + (1 - \kappa) (1 - \delta) p_{T+1}^H d_{iT} H_{iT}$$

1.4 Income Process

During working life, latent income evolves according to the following standard permanent-transitory process:

$$\ln Y_{it} = X_{it} \varphi_t + \ln P_{it} + \epsilon_{it} \quad (5)$$

$$\ln P_{it} = \ln P_{it-1} + \nu_{it} \quad (6)$$

where X_{it} captures cohort, age, education, and household size and $X_{it} \varphi_t$ captures the deterministic forecastable component of income depending on these factors. $\ln P_{it}$ is (log) permanent income, ν_{it} is the mean-zero permanent shock and ϵ_{it} is the mean-zero transitory process. The innovations are independent of one another and independent across time. These innovations are allowed, however, to have variances that change over time. These variances are estimated using cross-sectional moments, as discussed in section 3. I have the usual interpretation that permanent shocks represent long-term productivity changes such as promotions or changes in health status within the household, while transitory shocks represent bonuses, temporary lay-offs or other short-term changes in hours of work. I could allow for more general and richer processes: For example the transitory shocks could be modeled as an MA(1) process, as in Meghir and Pistaferri (2004), but identification of these lagged effects is weak with cross-sectional data. Similarly I could allow a more complex autoregressive process on the permanent/persistent process. These additions would change the analysis little.

1.5 Other Model Features

The retirement age is fixed. After retirement, households have access to a pension, which is riskless, and the value of which is determined by final salary. Households die at a fixed age and face no mortality risk beforehand. The recursive problem for the household is discussed in fuller detail in appendix D.

2 The Role of Wealth Effects in Consumption Inequality

The benchmark model just specified contains four key channels driving the relationship between house prices and consumption: these are not only income and substitution effects, but, more importantly, wealth effects and collateral effects. In this section I abstract from collateral effects and assess the role in consumption inequality of the key remaining channel, the wealth effect. I do this by simplifying the benchmark model and approximating its solution for the response of consumption to income and house price shocks. These approximations offer two main advantages over solving the dynamic program fully. First I can show what the model implies for inequality profiles transparently. Second, to quantify effects, I can impose the observed and relevant distributions of consumption, wealth and income into the calculations directly. The main weakness of the approach is that it abstracts from the potentially important non-linearities in consumption. I discuss this trade-off below and further assess the role of collateral effects in particular in section 5.

To proceed, I consider the household at generic time t and modify the benchmark model described above. Most importantly I drop the collateral constraint given by equation 4. Without the collateral constraint, households are limited only by their terminal condition on borrowing, and therefore have a natural borrowing constraint. I also set transactions costs, κ , to zero. Moreover I abstract from the household tenure decision. Accordingly a household can arrive at time t either as a renter or an owner, but as I argue below, the choice thereafter does not affect consumption responses: for concreteness we could imagine that the household sells its house and rents thereafter. I solve the approximated model fully in appendix A. Here I give an overview of the key ideas.

2.1 Expected Lifetime Budget Constraint

A key step in approximating the solution is examining the budget constraint in expectation over the whole lifetime. This lifetime budget constraint also provides an intuitive guide to how heterogeneity in endowments drives heterogeneity in consumption responses. For this reason I focus on deriving this constraint first.

I begin by generalizing constraints 2 and 3 into the following constraint defined for time $k = t, \dots, T$:

$$A_{ik} + C_{ik} + \Upsilon p_k^H H_{ik} = RA_{ik-1} + Y_{ik} + p_k^H \tilde{H}_{ik} \quad (7)$$

where \tilde{H}_{ik} is the housing endowment carried forward from time $k - 1$ to time k and is specified as follows:

$$\begin{aligned} \tilde{H}_{ik} &= \Theta H_{ik-1} \text{ for } k > t \\ \tilde{H}_{it} &\text{ given} \end{aligned} \quad (8)$$

Here Θ and \mathcal{Y} are fixed constants, and κ has been set to 0. In the model of renting, then $\Theta = 0$ and $\mathcal{Y} = \chi \in [0, 1]$, while the model of homeownership uses $\mathcal{Y} = 1$ and $\Theta = 1 - \delta$. In appendix A I solve the model explicitly in the case of homeownership, but much of the analysis follows for any choice of \mathcal{Y} and Θ .¹⁹

Let ρ_{iT} be a sequence of shocks from time t to time T , with a subsequence of shocks from time t until time k given by ρ_{ik} (ρ_{iT}). Conditioning on any ρ_{iT} , we can sum discounted versions of equation 7 over time to obtain the following lifetime constraint:

$$\sum_{k=t}^T q^{k-t} C_{ik} + \underbrace{\sum_{k=t}^T q^{k-t} r_k^H H_{ik}}_{\text{Cost of housing services}} + q^{T-t} A_{iT} = \sum_{k=t}^T q^{k-t} Y_{ik} + R A_{it-1} + \underbrace{p_t^H \tilde{H}_{it}}_{\text{Housing endowment}} \quad (9)$$

where, in addition to the notation defined previously $q = \frac{1}{R}$ is the price of a unit of the risk-free asset, and $r_k^H = \varsigma_k p_k^H$ is a generalised rental, or user, price of housing, with $\mathbb{E}_t \varsigma_k$ constant for all $k \geq t$. Note that, precisely speaking, stochastic variables should be written as, for example, $C_{ik}(\rho_{ik}(\rho_{iT}))$. I remove the explicit dependence in the equation to save space.

The price on housing services is therefore defined as follows:

$$r_k^H = \begin{cases} (\mathcal{Y} p_k^H - q \Theta p_{k+1}^H) & \text{if } k = t, \dots, T-1 \\ \mathcal{Y} p_k^H & \text{if } k = T \end{cases}$$

Notice that for any choice of \mathcal{Y} and Θ , then

$$\frac{\mathbb{E}_t(r_k^H)}{\mathbb{E}_{t-1}(r_k^H)} = \frac{\mathbb{E}_t(p_k^H)}{\mathbb{E}_{t-1}(p_k^H)} = e^{\zeta_t^H} \quad (10)$$

for $k = t, \dots, T$, where $\frac{\mathbb{E}_t(p_k^H)}{\mathbb{E}_{t-1}(p_k^H)}$ captures the time- t innovation to expected future house prices. These innovations to prices are the crucial aspect for the solution for consumption given below. That the prices of housing services display the same innovations for any choice of \mathcal{Y} and Θ provides the rationale for leaving tenure decisions implicit.²⁰

Finally, to obtain an expected lifetime budget constraint I take expectations with respect to information set \mathcal{I} . The information set \mathcal{I} can be specified quite generally, and could equal, for example, the set of all information

¹⁹Of course, the approximate solution may be more or less accurate given extreme choices of these parameters.

²⁰In the case of compulsory homeownership, when $\mathcal{Y} = 1$ and $\Theta = 1 - \delta$, then the prices imply that the household must purchase a house in the final period at present cost $q^{T-t} p_T^H H_{iT}$. In this case, households would prefer to rent in the final period before death, leaving open more resources for non-housing consumption. I could therefore reduce and rationalize the difference between the renting and owning versions of the model further by assuming that homeowners are also able to rent in the final period.

available at $t - 1$, or the set of all information available at t . The expected constraint used is then as follows:

$$\mathbb{E}_{\mathcal{I}} \left[\sum_{k=t}^T q^{k-t} C_{ik} + \underbrace{\sum_{k=t}^T q^{k-t} r_k^H H_{ik}}_{\text{Cost of housing services}} \right] = \mathbb{E}_{\mathcal{I}} \left[\sum_{k=t}^T q^{k-t} Y_{ik} + RA_{it-1} + \underbrace{p_t^H \tilde{H}_{it}}_{\text{Housing endowment}} \right] \quad (11)$$

The left hand side of equation 11 captures lifetime expenditures, the right hand side captures lifetime resources.

The key point of this subsection is that, as given by the left hand side of equation 11, homeowners and renters face a similar pattern of lifetime expenditures. The main difference between homeowners and renters is in resources, as given by the right hand side. Homeowners have a valuable endowment of housing, whereas the renters do not. In notation, the inherited housing endowment is $p_t^H (1 - \delta) H_{it-1}$ for single-home owner-occupiers and 0 for renters. More generally, if a household owns positive housing assets which differ from the housing services they consume then, again, the housing endowment need not correspond exactly to H_{it-1} . This would be the case, for example, if the household owns a second home.

Some further discussion of the refined model, and the lifetime budget constraints presented here. First, it is worth emphasizing that all housing decisions here are frictionless. I re-introduce transactions costs in section 5. Second, by encapsulating housing prices into a single r_t^H I assume implicitly that the cost of housing is common across tenure types. For example, I assume there is no cost premium on renting. Ensuring that expected housing costs are common across types is equivalent to placing further restrictions on \mathcal{T} and Θ . However, in the standard case of Cobb-Douglas preferences, homogenous across the population, the introduction of heterogeneous user costs of housing would not affect the main result. I discuss this issue further in the next subsection.

Third, and on the same note, recall that the benchmark model imposes a constant rent-price ratio, a strong assumption. In fact, both in the US and the UK, the rent-price ratio is counter-cyclical, to the extent that assuming a constant rental price (and varying rent-price ratio) is likely as close to the data. (See [Belfield, Chandler, and Joyce \(2015\)](#) for the UK.) For this reason I numerically explore an extended model with constant rent prices in section 5, and show that the conclusions from the main analysis hold. Finally, notice I have imposed that $A_{iT} = 0$. This is the solution to the household's problem in the absence of bequests, and when $A_{iT} \geq 0$ is a terminal condition on assets. Given the preferences described in equation 1, this is without loss of generality. Preferences for bequests could be re-introduced without affecting the final result.²¹

2.2 Consumption Dynamics

I now approximate the solution to the refined model, adapting the approach used in, for example [Blundell et al. \(2013\)](#). Specifically, the household maximizes lifetime utility subject to the sequence of constraints in equation

²¹In the current model household preferences between consumption and bequests are homothetic, therefore even in the presence of a strong bequest motive the income elasticity on current consumption is 1. The results in this section would only be affected if bequests are modelled as either a necessary, or, more realistically, a luxury good.

7, and law of motion in equation 8. I simplify notation by using lower case letters c_{it} , h_{it} and y_{it} to capture the log of consumption, housing and income net of effects from observable household characteristics. In appendix A I show that approximate solutions for the growth of consumption and housing are:

$$\begin{aligned}\Delta c_{it} &\approx (\eta_{C,p^H} - \eta_{C,p^C}) \left(\Gamma_{it}^\lambda + v_{it}^\lambda \right) + \eta_{C,p^H} (\mu + \zeta_t^H) \\ \Delta h_{it} &\approx (\eta_{H,p^C} - \eta_{H,p^H}) \left(\Gamma_{it}^\lambda + v_{it}^\lambda \right) - \eta_{H,p^H} (\mu + \zeta_t^H)\end{aligned}\quad (12)$$

where λ_{it} is the marginal utility of wealth, v_{it}^λ is its innovation and Γ_{it}^λ is its predictable change, forecastable at time $t - 1$. μ is the predictable change to house prices, and ζ_t^H is the innovation, as above. The other parameters are the Frisch (or λ -constant) elasticities. η_{x,p^y} is the elasticity of good x with respect to price p^y of good y .²² These Frisch elasticities could, in general, vary by household, but here I omit the i subscript for brevity of notation.

Expression 12 implies that house price shocks affect consumption both by shocking marginal utility (through v_{it}^λ) and by inducing pure substitution (Frisch) effects. Holding marginal utility fixed, a positive house price shock, for example, induces a Frisch response towards non-housing consumption of $\eta_{C,p^H} \zeta_t^H$. On the other hand, holding fixed this pure substitution effect, the resulting downwards shock to marginal utility causes consumption to grow by $\eta_{C,p^H} - \eta_{C,p^C}$ and housing to grow by $\eta_{H,p^C} - \eta_{H,p^H}$. These rates depend on the income elasticities on preferences, and in general will differ. However, when preferences are CES, then these coefficients on the marginal utility of wealth terms are identical and equal to $-\frac{1}{\gamma}$. In this case, consumption and housing move in parallel. The coefficient of relative risk aversion, γ , then determines, for example, how responsive consumption and housing are to interest rate changes, and to what extent consumption risk lowers life-time utility.

A problem with expression 12 is that changes to marginal utility are unobservable. I therefore improve on it by expressing marginal utility shocks in terms of the real shocks that households receive: shocks to house prices and to income. I do this by approximating innovations to the income side of the expected lifetime budget constraint given in equation 11. In appendix A I therefore show that, when preferences are intratemporally homothetic, and focussing on changes to consumption:

$$\Delta c_{it} \approx \Gamma_{it}^C + \pi_{it}(\nu_{it} + \alpha_{it}\epsilon_{it}) + \left(\underbrace{\psi_{it-1}^H}_{\text{Endowment effect}} + \underbrace{s_{it}^H (\eta_{C,p^H} + \eta_{H,p^H})}_{\text{Substitution effect}} - \underbrace{s_{it}^H}_{\text{Income effect}} \right) \zeta_t^H \quad (13)$$

where Γ_{it}^C is a predictable component reflecting long-term shifts due to trend house-price increases, and saving due to time preferences, interest rates and the precautionary motive. ζ_t^H is the realized common shock to house prices, ν_{it} is the idiosyncratic permanent shock to income and ϵ_{it} the transitory shock to income. The remaining terms capture the transmission of these shocks through to consumption: α_{it} is an annuitization factor capturing the contribution of time- t income to discounted expected future (or ‘lifetime’) wealth, and is typically small; π_{it} is the share of discounted expected future labour income (‘human capital wealth’) in lifetime

²²The Frisch elasticity is defined as $\eta_{x,p^y} \equiv -\frac{u_{\tilde{x}\tilde{y}}u_y}{u_{CC}u_{HH}-u_{CH}^2} \frac{1}{x} > 0$ for $x, y \in \{C, H\}$. \tilde{x} refers to the good other than x : for example, if $x = C$, then $\tilde{x} = H$.

wealth, and ψ_{it-1}^H is the share of gross housing wealth in lifetime wealth at time t , given the housing stock inherited from time $t - 1$. In notation, ψ_{it-1}^H is approximately equal to $\frac{1}{\Lambda_{it}} \mathbb{E}_{t-1} p_t^H \tilde{H}_{it} \approx \frac{1}{\Lambda_{it}} p_{t-1}^H \tilde{H}_{it}$, where lifetime wealth, Λ_{it} , corresponds roughly to the right hand side of equation 11: Λ_{it} is approximately the sum of current financial wealth, RA_{it-1} (which includes mortgage debt as a liability), expected gross housing wealth, $\mathbb{E}_{t-1} p_t^H \tilde{H}_{it}$, plus expected discounted life-time labour income, $\mathbb{E}_{t-1} \sum_{k=t}^T q^{k-t} Y_{ik}$.²³ s_{it}^H is the budget share of housing in discounted expected future (‘lifetime’) expenditure, equalling $1 - a$ in the Cobb-Douglas benchmark. η_{C,p^H} and η_{H,p^H} are the Frisch elasticities described above. See appendix A for precise details of these terms. For example the appendix shows that the approximations are made using *geometric* means for the expectations rather than arithmetic. I wait until section 3 before discussing the empirical implementation.

Expression 13 warrants further interpretation. I focus first on the effect of the house price shock, ζ_t^H . In this homothetic case, the shock’s effect can be decomposed simply into endowment, income and substitution components. These components bear similarities to a Slutsky-type decomposition in elasticity form. The endowment effect is captured by ψ_{it-1}^H , the share of gross housing wealth in lifetime wealth. To interpret this effect, suppose that the substitution and income effects cancel. Then a house price shock, ζ_t^H , causes consumption to grow proportionately to how much housing is owned. At one extreme, for example, suppose that consumption and housing are financed purely by housing assets, with no labour income and no other financial wealth. Then a positive 1% house price shock raises consumption by 1%. More generally, when labour income and financial wealth are present, then the pure endowment effect on consumption is muted.

Of course, a positive house price shock also makes housing services more expensive, inducing a negative income effect. Intuitively, this negative income effect is equivalent to the result of reducing all of labour income, financial wealth and housing wealth in identical proportions. When preferences are homothetic, a 1% negative income effect, for example, then reduces consumption by 1% also. In the present case, when the house price shock is ζ_t^H then housing costs, ignoring behavioural changes, rise by $s_{it}^H \zeta_t^H$. Therefore the house price increase is equivalent to reducing lifetime income by $s_{it}^H \zeta_t^H$ and consumption drops accordingly.

Finally, a positive house price shock makes households substitute away from housing services. This substitution effect is the product of the expected budget share of housing expenditures, s_{it} , with the sum of the house-price Frisch elasticities, $\eta_{C,p^H} + \eta_{H,p^H}$. When preferences are CES, with elasticity of substitution τ , then $\eta_{C,p^H} + \eta_{H,p^H} = \tau$. Therefore, when $\tau > 1$ the substitution effect dominates the income effect, while when $\tau < 1$ the reverse is true. When preferences are Cobb-Douglas, with $\tau = 1$, then the substitution and income effects cancel. In this case, changes to consumption are driven ultimately by the endowment effect.

Aside from house price shocks, consumption changes are also determined by the shocks to income, ν_{it} and ϵ_{it} . The effect of these shocks depends on the elasticities π_{it} and $\alpha_{it} \pi_{it}$. This result is standard and discussed in much further detail in, for example, [Blundell et al. \(2013\)](#), so I do not pursue it further here. As a final point, in appendix A I derive an approximation for consumption growth for general $u(\cdot)$, without imposing homotheticity. However, in that case the approximation is not so easily interpretable, and it cannot be used in empirical applications without knowledge of the income elasticities on consumption and housing.

²³Notice that the human capital wealth share, π_t , has a time t subscript, while the housing wealth share ψ_{t-1}^H has been given a time- $t - 1$ subscript. In fact all wealth shares are computed using asset and debt positions at time $t - 1$ and income summed from time t onwards.

The decomposition in expression 13 bears similarities to that discussed by Berger et al.. They break down the effect of house price gains into endowment, substitution and income effects plus collateral effects. Collateral effects arise when a house price increase (decline) relaxes (tightens) borrowing restrictions. As discussed, I cannot capture these effects in this Euler approach framework. The relevant question in my application is how this collateral channel affects the inequality measures. I return to this issue in section 5.

2.3 The Effect of Income and House Price Shocks on Consumption and Income Moments

I now use expression 13 together with the income process specification to calculate how income and house price shocks affect inequality measures. Broadly, and focussing on house prices, the effect of an aggregate shock depends on the dispersion of each of ψ_{it-1}^H , s_{it}^H and $\eta_{C,p^H} + \eta_{H,p^H}$. Clearly ψ_{it-1}^H , the endowment share, is highly unevenly distributed. This is not only because both renting and homeownership are common, but also because of heterogeneity in extra housing assets and, more importantly, in leverage.

Before exploring this heterogeneity in endowments in detail, I consider the role of heterogeneity in the expected budget share, s_{it}^H , and the Frisch elasticities, $\eta_{C,p^H} + \eta_{H,p^H}$. First, recall that when preferences are CES, then the Frisch elasticities combine to equal τ , the elasticity of substitution. Therefore, as previously discussed, when preferences are Cobb-Douglas, the terms in s_{it}^H and the Frisch elasticities disappear entirely. More generally, with regard to the elasticity of substitution, the literature has produced various estimates of its mean that differ in either direction from 1, but has little to say about heterogeneity. With regard to budget shares, s_{it}^H is equal across households if, for example, preferences (including the consumption weight, a) are homogeneous, and the price of housing services is common. Of course, s_{it}^H may be unevenly distributed if these conditions do not apply, as well as for other reasons, such as housing adjustment costs.

However, overall, even when preferences are not Cobb-Douglas, then, as long as τ is close to 1, the budget share, s_{it}^H , likely has a much lesser role on the heterogeneous impact of house price shocks than the endowment share ψ_{it-1}^H . For this reason, and because s_{it}^H is often difficult to measure empirically, I drop it from the main analysis of inequality profiles from now on. However, I return to some of these issues in section 5, where I specifically explore the effects of housing adjustment costs combined with $\tau \neq 1$.²⁴

Concentrating on endowment heterogeneity, therefore, appendix B shows that the processes given by expressions

²⁴Alongside that analysis, it is worth briefly considering another relevant channel: renters typically face a cost premium on their housing services and have lower-than-average consumption. In this case, if consumption and housing are complements, with $\tau < 1$, then such heterogeneity in s_{it}^H would lead to consumption inequality growing following a house price increase. Of course, if $\tau > 1$ then the effect is reversed.

6 and 13 imply the following formulae for growth in inequality of income and consumption:

$$\Delta \text{Var}(y_t) \approx \text{Var}(\nu_t) + \Delta \text{Var}(\epsilon_t) \quad (14)$$

$$\Delta \text{Cov}(c_t, y_t) \approx \underbrace{\bar{\pi}_t \text{Var}(\nu_t) + \Delta [\bar{\pi}_t \alpha_t \text{Var}(\epsilon_t)]}_{\text{Contribution from income}} + \underbrace{\zeta_t^H \text{Cov}(\psi_{t-1}^H, y_{t-1}^P)}_{\text{Contribution from housing}} \quad (15)$$

$$\Delta \text{Var}(c_t) \approx \underbrace{(\bar{\pi}_t^2 + \text{Var}(\pi_t))(\text{Var}(\nu_t) + \alpha_t^2 \text{Var}(\epsilon_t))}_{\text{Contribution from income}} + \underbrace{(\zeta_t^H)^2 \text{Var}(\psi_{t-1}^H) + 2\zeta_t^H \text{Cov}(\psi_{t-1}^H, c_{t-1})}_{\text{Contribution from housing}} \quad (16)$$

where I use the following notation: $\text{Var}(x_t)$ denotes the cross-sectional variance of variable x_{it} ; $\text{Cov}(x_t, z_t)$ similarly denotes the cross-sectional covariance of variables x_{it} and z_{it} , and \bar{x}_t denotes the cross-sectional mean of variable x_{it} . Finally I use $y_{it}^P \equiv \ln P_{it}$ for the log of permanent income.

Clearly, changes to income inequality as given by expression 14 are determined solely by the income process. But both the covariance of income and consumption and the variance of consumption given in expressions 15 and 16 are determined by income fluctuations and by house price shocks jointly. My main focus is on the contribution of house price shocks, which I now discuss in detail. How house price shocks contribute to the covariance of income and consumption in expression 15 is simplest to describe: these shocks induce changes in the covariance only if the housing wealth share is correlated with the distribution of permanent incomes. This would happen if, for example, those with higher permanent income have more housing wealth. In my set-up there is no reason why the housing wealth share and the distribution of permanent incomes should be correlated. Of course, the wealth share is likely positively correlated with the income distribution for other unmodelled reasons, such as heterogeneous mortgage availability or uneven preferences for homeownership. Ultimately, I assess the size of this correlation empirically.

More importantly, house price shocks also contribute to the consumption variance, as given in equation 16. This contribution consists of two parts. First a positive house price shock can increase the consumption variance purely because the housing wealth share is unequally distributed. This effect is given by $\text{Var}(\psi_{t-1}^H)$. Second, a positive house price shock can cause an increase or decrease in the consumption variance depending on whether or not the wealth share is positively correlated with the consumption distribution. To see how these effects differ, imagine two distributions. First, if housing wealth is spread randomly across the consumption distribution and roughly in proportion to households' life-time wealth, then the wealth share is uncorrelated with consumption and the house price shock induces just an orthogonal shock to the consumption distribution of size $(\zeta_t^H)^2 \text{Var}(\psi_{t-1}^H)$. As a second, more realistic case, suppose the wealth share and household consumption are correlated positively. This might be justified theoretically if households who receive good transitory shocks accumulate sufficient funds to provide a downpayment on a home and can also afford higher consumption. Then household consumption covaries positively with the wealth share too. In this case the positive shock increases inequality by an additional factor $2\zeta_t^H \text{Cov}(\psi_{t-1}^H, c_{t-1})$.²⁵

²⁵In contrast, if the share of a generic type of wealth ψ_t^j covaries negatively with the consumption distribution, then a positive shock to its value, ζ_t^j reduces inequality if $-2\zeta_t^j \text{Cov}(\psi_t^j, c_{t-1}) > (\zeta_t^j)^2 \text{Var}(\psi_t^j)$. Such is the case with, for example, social security (state pension) wealth. The share of life-time in pensions varies negatively with total life-time wealth, because of the redistributive nature of the social security system. Therefore, a positive shock to such wealth, due to, say, an unexpected increase in generosity would reduce consumption inequality.

The contribution of income risks to movements in the inequality measures is discussed by [Blundell et al. \(2013\)](#). Briefly, in their framework the consumption variance and the covariance measures are both largely driven by the variances of permanent shocks ($\text{Var}(\nu_t)$), so should move roughly in parallel. They move perfectly in parallel when households have no assets ($\pi_t = 1$). More realistically, the first term on the right hand side of expression [16](#), $\bar{\pi}_t^2 + \text{Var}(\pi_t)$, is less than the first term on the right hand side of expression [15](#), $\bar{\pi}_t$, and so permanent income shocks induce consumption inequality to grow more slowly than the covariance.²⁶ The contribution from transitory shocks to both the consumption-based measures given in equations [15](#) and [16](#) is generally small.

2.4 Further Discussion and Relationship to [Berger et al. \(2018\)](#)

I obtained the formula in expression [13](#) by using an Euler equation approach. To the best of my knowledge, this is the only way to obtain an expression for inequality changes that can be taken directly to standard micro-data. Notice, however, that expression [13](#) takes a very simple form, which can also be viewed more generally. It states that the change to log consumption is the sum of a predictable component and changes coming from income shocks and from house price shocks, and that these latter components are the product of the raw shock times the relevant elasticity.

Viewed more generally, therefore, we can use the formula:

$$\Delta c_{it} \approx \Gamma_{it}^C + \omega_{it}^P \nu_{it} + \omega_{it}^T \epsilon_{it} + \omega_{it}^H \zeta_t^H$$

where ω_{it}^P is the elasticity of consumption with respect to permanent income shocks, ω_{it}^T is the elasticity with respect to transitory income shocks, and ω_{it}^H is the elasticity with respect to house price shocks. This house price elasticity has been explored recently by [Berger et al. \(2018\)](#), within a general class of housing models, encompassing the benchmark model I describe in section [1](#). In their simplest model, with no rental markets and frictionless adjustment, and with Cobb-Douglas preferences, they show that the elasticity of consumption ω_{it}^H can be given by

$$\omega_{it}^H = MPC_{it} \times \frac{p_{t-1}^H \tilde{H}_{it}}{C_{it}}$$

where \tilde{H}_{it} is the gross housing endowment for household i entering time t , defined previously, and MPC_{it} is the marginal propensity to consume out of transitory income.

[Berger et al.](#) then use simulations of more complex models involving transactions costs, rental markets, more realistic mortgage contracts and richer house price expectations. They show that, even though the formula no longer holds exactly in these more complex models, it generally approximates the true elasticity well.

How does [Berger et al.](#)'s formula compare to mine? In fact, it bares striking similarities. Remember, in the Cobb-Douglas case, I am approximating the elasticity ω_{it}^H as

$$\omega_{it}^H \approx \psi_{it-1}^H = \frac{p_{t-1}^H \tilde{H}_{it}}{\Lambda_{it}}$$

²⁶If π is distributed as a Bernoulli variable taking values 0 with frequency $1-p$ and 1 with frequency p , then $\bar{\pi}_t^2 + \text{Var}(\pi_t) = p = \bar{\pi}_t$. Whenever π is distributed more uniformly between 0 and 1, then $\bar{\pi}_t^2 + \text{Var}(\pi_t) < \bar{\pi}_t$.

where Λ_{it} is discounted expected life-time wealth at the start of time t . Both approximations indicate that the elasticity is proportional to *gross* housing wealth, and not *net*. The formulae differ in that whereas my baseline formula has discounted lifetime wealth in the denominator, [Berger et al.](#) have $\frac{C_{it}}{MPC_{it}}$. In general, of course, these expressions are different; they line up precisely, however, in the permanent income model with perfect certainty. Their similarity is explained by [Berger et al.](#)'s further analysis. [Berger et al.](#) decompose the effect of shocks into four components: the income effect, substitution effect, endowment effect, as well as a collateral effect. They show numerically that the income, substitution and collateral effects cancel out, leaving the endowment effect to drive the final elasticity. My formula is also driven ultimately by the endowment effect.

The main attraction of my formula is that it is easier to take to the data. The formula used by [Berger et al.](#) requires computing *MPCs* at the household level. This computation is not realistically possible, not least because of measurement error in consumption and income.²⁷ [Berger et al.](#) therefore compute average *MPCs* across the population, splitting by wealth level. On the other hand, with data on wealth and incomes it is possible to compute discounted expected lifetime wealth and wealth shares, even if this computation requires some further modelling choices. One main challenge in computing a household's expected lifetime wealth is pinning down the permanent component of income. Intuitively, this challenge is more easily met using income panel data. In contrast, if only cross-sectional data is available, this computation has to be performed using observed current income, which may be a poor proxy. In appendix [F.3](#) I explore in further detail the data requirements for accurately implementing the formulae derived here. I find that implementation is accurate as long as one has access to a short panel of the type available in the present analysis. Briefly, with a short panel, a one-year lead of current income can be used as a sort of instrument for permanent income when computing key statistics such as $\text{Cov}(\psi_{t-1}^H, c_{t-1})$. For these reasons, I pursue the Euler equation approach in the empirical application. I describe in detail how I estimate discounted expected lifetime wealth and the wealth shares in practice in section [3.5](#).

Despite the similarities to [Berger et al.](#)'s formula, the main problem with my approach, however, is that it does not capture collateral effects. However, I assess the importance of collateral effects when I explore the non-linear solution to the full benchmark model in section [5](#). I use the analysis there to assess further the tradeoff between [Berger et al.](#)'s approach and mine.

3 Data and Estimation Method

3.1 Overview and Nature of the Empirical Exercise

In my main empirical exercise I use the approximated model solution discussed in section [2](#). I take it to data by computing the empirical counterparts of the theoretical moments expressed in equations [14](#), [15](#) and [16](#). This implementation requires measures of the cross-sectional variances of income and consumption and their covariance. It also requires measures of wealth shares: their means and their covariance with income and with

²⁷More fundamentally, there is also the identification problem of splitting permanent shocks from transitory. Note that self-reported MPCs are collected by [Japelli and Pistaferri \(2014\)](#), but I do not investigate these data here.

consumption. Finally it requires estimates of the ex-post house price shock in each year. I use the model and these inputs to estimate the raw drivers of inequality, and, in particular, the role of the housing market.

It is perhaps worth emphasizing that the only estimands in equations 14, 15 and 16 are the unobserved variances of the income components. According to the model used, on the other hand, the contribution from the housing market is directly observable. In brief, therefore, the point of the empirical exercise is to examine whether accounting for housing market factors improves model fit. This exercise takes two specific forms: first I examine to what extent accounting for the housing market explains movements in consumption inequality. Second I examine to what extent accounting for the housing market changes estimates of income risk.

I use two primary data sources for the household: the Family Expenditure Survey (FES) and the British Household Panel Survey (BHPS).²⁸ The FES provides cross-sectional data on household income and consumption over 1989-2008. The BHPS provides yearly panel data on income and food consumption over 1991-2008, data on housing wealth and mortgages over 1993-1995 and 1997-2008, and data on other financial wealth for 1995, 2000 and 2005. Both these datasets have been used extensively, and are described in detail in, for example, [Blundell and Etheridge \(2010\)](#). Here I include some description of the datasets themselves as well as detailed descriptions of my treatment. Finally, I use national house price data over 1969-2008 from the UK Office for the Deputy Prime Minister (ODPM).

The main empirical exercise is cross-sectional in nature, and relies on computing cross-sectional variances and co-variances. I focus on cross-sectional properties of the data mainly because the quality of the consumption data in the BHPS, as in most household panel surveys across the world, makes a detailed dynamic analysis difficult. Nevertheless, and as extra supporting evidence, I show dynamic regressions of consumption and of equity withdrawal in appendix F.4, table A4. These regressions indicate a strong response to house price shocks.

Finally, it is worth emphasizing that the analysis is partial equilibrium in nature, and so takes shocks to prices as exogenous. As discussed elsewhere, the recent literature has explored house price effects in general equilibrium and has explored a variety of underlying drivers, such as credit conditions or beliefs. In this light, and to extend the main analysis presented here, in appendix C I consider equilibrium effects on inequality quantitatively in a simple setting. Specifically, I consider the effect of aggregate income shocks, using stylized assumptions on housing supply.

3.2 FES Cross-Sectional Data

The FES is an annual survey conducted chiefly for determining the basket of goods used to construct the retail price index. The FES has been running since 1957, although it has only collected data in its present form on a consistent basis since the 1970s. In a typical year the FES contains information on around 6500 households. In general the households form a representative sample, but excluded are those not living in private houses, such as residents of residential homes or students. For households participating in the FES, each member over 16

²⁸The FES has since been renamed the Expenditure and Food Survey (EFS) and then the Living Costs and Food Survey (LCFS). In this paper I use its historic name.

is asked to complete a diary detailing all their spending, both home and abroad, over a two week period. In addition to this diary, household members perform an interview in which they are asked questions about their demographic background, and asked to recall expenditures on large infrequently-purchased items (such as cars).

The main sample period is 1989-2008. I focus on the pre-crisis period for reasons I discuss later. The baseline sample contains 133,804 households. Each household is one data point. To each household I allocate a head, for example the male in a household consisting of a married couple with children. The head then determines the household's assigned characteristics. My measure of education is the head's age on leaving school or university. The low education group comprises those leaving school at the compulsory schooling age (15 for those born before 1957 and 16 thereafter). I use as population all households with heads aged 25-60. The sample is formed as follows: I drop households for which the head is outside the age range, or where food consumption or disposable income is negative, leaving 85,119 households. For robustness of the results I trim the top and bottom 1% of observations of each distribution. This leaves 84,569 household consumption observations.

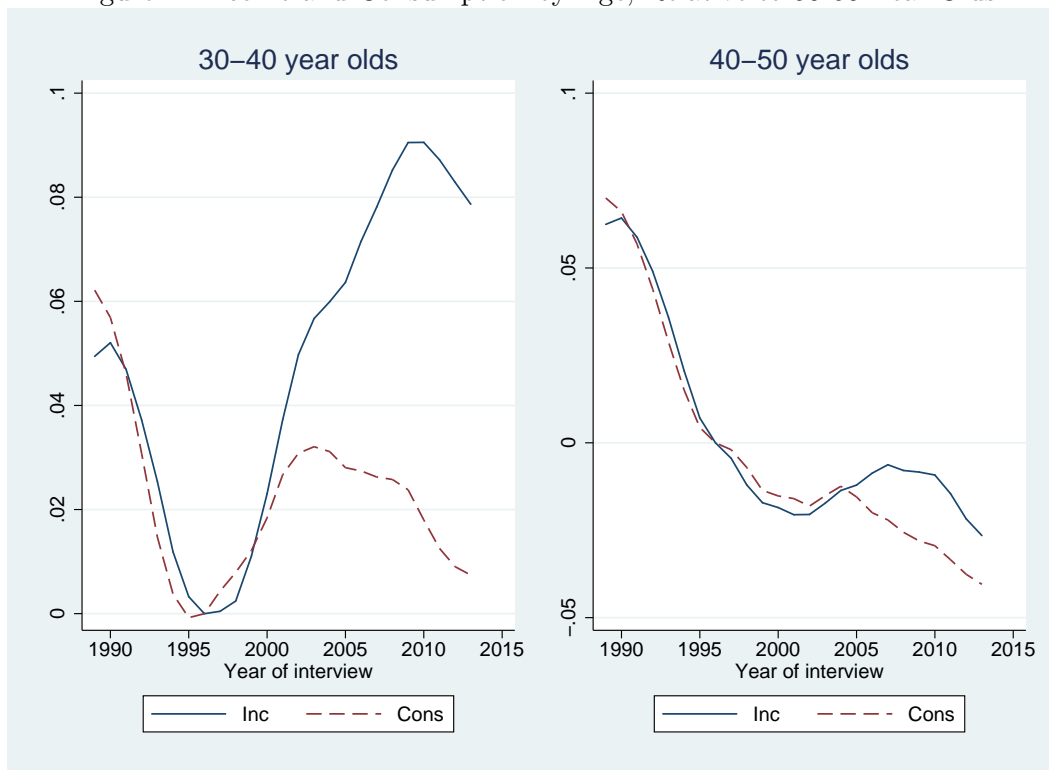
The income measure is total disposable income: labour earnings net of taxes, plus benefits and private transfers, plus asset returns excluding the drawing down of capital or capital gains. Income therefore includes interest and dividends from wealth but excludes imputed rent from a housing asset. The consumption measure is expenditure on all items aside from housing and durables. It therefore consists mainly of spending on services and semi-durables. In more detail, it includes spending on food, catering, alcohol, tobacco, 'housing services' (such as water bills), energy, household goods, household services, clothing, personal goods, motoring expenses, travel expenses, leisure goods and leisure services. I deflate both income and consumption using the RPI index. I compute residual measures of income and consumption after regressions on demographics, education, region and a quartic polynomial in age.

To motivate the analysis I show patterns of mean income and consumption for various age groups. Here I analyse rotating groups of those aged 30 – 40, those aged 40 – 50 and those aged 50 – 60. Figure 1 shows mean income and consumption for the younger two groups relative to the means for the older group, and normalized to 0 in 1996, the bottom of the housing market. The figure also shows data to 2013 to show some of the post-crisis period. Notice that income for each group relative to the older group shows cyclical behaviour. Most striking, however, are the patterns for the youngest group after the trough in the housing market in the mid-1990s. Here, consumption for the youngest group lags far behind relative growth in income, implying large extra saving. This extra saving indicates that the older age groups benefitted strongly from the house price boom. Notice that the implied saving for the middle age group is very similar to that for the oldest group, deviating only in the post-crisis period.

3.3 BHPS Panel Data

The BHPS is a comprehensive longitudinal study for the UK for general use in the social sciences, covering 1991-2008. The survey typically records around 5000 household interviews per year. The BHPS has detailed information on earnings, hours worked and other income, and information on housing and durables.

Figure 1: Income and Consumption by Age, Relative to 50-60 Year Olds



Source: FES.

Notes: Figure shows relative mean income and consumption of each age group in each year relative to 50-60 year olds. Profiles are re-based to 1996=0. Profiles are smoothed using a 3-year moving average..

Despite being the main UK household panel survey, the BHPS has limited information on consumption. The survey only contains questions about food consumed within the home (food ‘in’) and about small durables purchases such as TVs and kitchen appliances, as well as about energy use. Within this group, only for food does consumption plausibly equal expenditure. I therefore focus on these responses. Food consumption is measured precisely in 1991 but then banded into twelve intervals for all years thereafter. The top interval is unbounded above and the bottom is bounded by 0; the log of food consumption is therefore unbounded both above and below. For all intermediate intervals I impute as my observation the interval midpoint. For the top interval (over £160 per week) I impute £180 spending, for the bottom (less than £10 per week) I impute £5 spending. This procedure is similar to that reported in [Etheridge \(2015\)](#), who also provides a detailed validation exercise. He shows that this simple assignment gives reliable results.

As for wealth, the BHPS has comprehensive information on housing values for most years: Homeowners are asked to provide a self-assessed valuation of their property. Information on the main residence is available for all years, and for other properties for all years except for 1991, 1992 and 1996. The extant value of mortgages is given in all years except for 1991 and 1992. However, comprehensive information on financial wealth is available for 1995, 2000 and 2005 only. These data include information on bank savings accounts, tax-free savings accounts (such as Individual Savings Accounts), financial investments, unsecured debts and private loans. The data for 1995 do not account for student loans and credit card debts, but I ignore this consideration and treat the data as comparable across waves. While the value of the first house and the value of all mortgages are reported exactly, the value of second homes and many items of financial wealth are reported in bands only. For these banded data I follow, for example, [Banks, Smith, and Wakefield \(2002\)](#) in using imputations on the value of each type of asset. The data on financial wealth come from the derived datasets reported in [Banks et al. \(2002\)](#) and [Crossley and O’Dea \(2010\)](#).

The measure of income is identical to that in the FES. The income data come from [Levy and Jenkins \(2008\)](#), who calculate total net household income based on the full reports of detailed income sources, together with yearly details of the UK tax system. I use a current, rather than annual income measure: usual monthly income at the time of interview. More precisely, net disposable income is defined as the sum of earned income, asset income and transfers (public and private) minus state taxes (income tax and national insurance contributions). Capital gains, or the drawing down of capital, are excluded. Pension income, which often comes from drawing down capital, is included in the definition, but because I focus on heads of working age, its contribution is small. I form residual measures of food consumption and income just as for the FES data.

The sample selection proceeds as follows: I use only the core BHPS sample and ignore the low-income booster sample. I follow similar sample formation and selection procedures for the BHPS as followed for the FES. The raw BHPS sample over 1993-2008 contains 69,770. In the three years with financial wealth data, 1995, 2000 and 2005, it contains 13,105 households. I remove households for whom the head is outside the age range, and those for whom educational status is missing.²⁹ For the cohorts I focus on, those born in the 1950s and 1960s, this leaves a baseline sample comprising 26,543 observations over 1993-2008 and 4,985 observations in

²⁹The measure of education in the BHPS is qualification level. I define low education as those with no qualification higher than an O-level, the national exam at age 16 in the UK. I remove those whose educational status is missing because education is an important regressor when forming residual food and income measures.

the years with financial wealth data. I then trim the top and bottom 1% of the distribution of disposable income and remove observations. By taking logs I also remove any negative income observations, comprising 0.6% of the initial sample themselves. I do not trim the food consumption distribution: because expenditures are assigned to 12 bands there is not the same chance of reporting implausibly high or low observations through mis-coding, or omission of a component. Unlike the FES, where each questionnaire is completed in entirety, the BHPS questionnaires are often incomplete, so the quoted statistics are computed using fewer observations. For example, the number of income observations remaining in the three years with financial wealth data is 4,258 and the number of food observations is 4,465. Over 1993-2008 the equivalent sample sizes are 22,420 and 24,214.

3.4 ODPM House Price Data

The national house price data come from the UK Office for the Deputy Prime Minister (ODPM). The ODPM published a quarterly house price index over 1968-2011 based on data from the UK Survey of Mortgage Lenders.³⁰ For most of its history, the survey has involved a variety of mortgage lenders supplying a 5% sample of their completions from the preceding month. The main disadvantage of the ODPM data is that it excludes cash purchases, around 25% of all deals. Ultimately, however, the time series provided is very similar to that from other sources, such as the Land Registry. I use the annual time-series over 1969-2008 and deflate by the UK retail price index.

3.5 Estimating the Wealth Share Measures

The empirical application requires measures of the shares of the different components of household wealth in total expected life-time wealth. The important components, which appear in equations 15 and 16, are housing wealth and discounted labour (human capital) wealth. The other components in total life-time wealth are: mortgage debt (as a liability) and other financial wealth, such as holdings in savings accounts. The housing wealth share ψ_{t-1}^H , for example, is therefore operationalized as:

$$\frac{\text{Gross housing wealth}_{t-1}}{\text{Gross housing wealth}_{t-1} + \text{Expected discounted labour income}_t + \text{Other current wealth}_{t-1}}$$

To compute this measure I perform similar steps to that reported in Etheridge (2015). I sum jointly-held assets with those held individually by the head, spouse and other household members. By including all household members I am assuming complete resource pooling within the household. This assumption is consistent with using income and food consumption at the household level. I compute discounted human capital wealth by the following procedure. First, I estimate permanent income by averaging income at time t , $t + 1$ and $t + 2$, to smooth measurement error and transitory shocks. I take the twice-forward income rather than $t - 1$ income, because later I calculate the covariance with time $t - 1$ variables, and so remove spurious correlation from, say, measurement error. I discuss the effect of using the lead of current income in these calculations in more detail

³⁰This house price index was taken on by the Department of Communities and Local Government before being discontinued in 2011. As with the FES, I use the historic title.

in appendix F.3. I then assume a ‘net’ discount rate on future income of 4%pa. This discount rate consists of growth in income net of pure discounting. This discount rate can be rationalized, for example, as consisting of income growth of -1% pa and a market interest rate of 3% pa, which is consistent with the data and calibration for middle-aged households used in section 5. Asset and debt positions are taken using time $t - 1$ values as discussed in section 2. I then remove outliers of the wealth share measures by trimming the top and bottom 1% of their distributions. Finally, because complete wealth data are only available for 1995, 2000 and 2005, I interpolate wealth share moments in other years. Briefly, I do this as follows: recall that the BHPS contains data on housing wealth in all years and complete data on housing wealth (including second homes and mortgage debt) for all years except for 1991, 1992 and 1996. I therefore compute approximate wealth share moments after 1992 (and excluding 1996) using these incomplete data and compute a final profile for the wealth share moments using a combination of the precise and approximate measures. This procedure allows for estimates of the housing wealth share (involving ψ_{it-1}^H) for 1994-2007.³¹ Details are provided in appendix F.1. When estimating the income process I also require wealth share measures before 1994. My procedure for computing these is discussed below. I perform robustness checks against all these assumptions. Most importantly, and as discussed below, varying the discount rate for future income changes the results little.

The main omission from my measure of life-time resources is wealth from pensions. I ignore pensions because they are very illiquid and households likely discount this far-distant wealth strongly (Carroll, 1997). As a practical matter, pensions wealth data are available for those who appear in the 2000 wave over 1991-2001, but not up to 2005. The data for 1991-2001 come from Emmerson, Disney, Wakefield, and Tetlow (2008), who impute values based on pension scheme participation, earnings histories observed within the BHPS time frame, and imputed earnings in previous periods from retrospective questions on employment histories. I discuss how including pension wealth affects results for 2000 in section 4.

Finally, my empirical application requires measures of the covariances of the wealth shares with total non-housing consumption. Because the BHPS only contains data on consumption of food, I infer broad non-housing consumption from this narrower variable. To do this I use results from Etheridge (2015), who in turn follows Blundell et al. (2008) by estimating a food demand equation using the FES data. In this paper I use the central estimate of the income elasticity and scale up covariances and variances accordingly. For example, the covariance $\text{Cov}(\psi_{t-1}^H, c_{t-1})$ in equation 16 is computed as $\frac{1}{\chi_f} \text{Cov}(\psi_{t-1}^H, f_{t-1})$ where f_{t-1} is log residual food consumption in time $t - 1$, and the income elasticity is χ_f . In this paper, I use a value of 0.38 for χ_f in keeping with established estimates more widely.

3.6 Estimating the House Price Process

As discussed in section 2, I assume real house prices follow a random walk with drift. Using this model, I estimate an average real return on housing of 0.034 with a standard deviation of shocks of 0.089. These are estimated over 1969-2008 from the ODPM data. In the empirical application, I take the residuals from this estimation as the shocks that households face.

³¹I obtain wealth share measures in 2007 by omitting this time $t + 2$ income and using just time $t + 1$ values. I do not compute wealth share measures for 2008.

To validate this approach, I conduct a simple investigation of the house price time series. I run the following regression:

$$\Delta \ln p_t^H = b_0 + b_1 t + b_H \ln p_{t-1}^H + \zeta_t^H$$

where $\ln p_t^H$ is the log real house price and ζ_t^H is the innovation. I estimate $\hat{b}_H = -0.135$ with a t-value of -1.55 (39 observations). This raw estimate implies some autoregression in prices. However, a Dickey-Fuller test fails to reject the presence of a unit root at any reasonable level of significance.

As also discussed in section 2, I focus on national house price changes and ignore more local price changes. To explore house prices in more detail, figure 2 shows a price index for both the UK as a whole and across selected sub-regions. I show the series back to the early 1970s, to the left of the first vertical line, to show the boom of the early 2000s in historic context. I also extend the data to 2016 to show the the post-crisis period. The data in this figure come from a mortgage lender, Nationwide, because the ODPM data series do not extend back by region so far. For this figure, I have selected 5 of the 11 sub-regions which are the most atypical. The other sub-regions of the British Isles more closely follow the UK average.³² For the UK as a whole, the figure shows that the rise in house prices leading up to the financial crisis was unprecedented in length and scale. From the trough in 1995 to the peak in early 2008, this rise was over 50 percentage points larger than the previous boom in the late 1980s. Additionally, the preceding decline, from the late 1980s to the mid 1990s was far larger than any previous decline.

At the regional level, the implications of figure 2 are nuanced. For example, since the onset of the financial crisis, regional house price dispersion has increased substantially: Since 2008, prices in London have grown faster than in the North, the cheapest region, by over 5%pa. Even before the financial crisis, regions have differed in the precise timing of booms and busts. Consequently, many studies (such as [Attanasio et al., 2009](#) and [Campbell and Cocco, 2007](#)) have used regional and local house price variation for identification. However, it is clear that the long boom in house prices, and the previous decline, were nation-wide phenomena. For example, in contrast to the later large rise in regional house price dispersion, prices in London grew faster than those in the North only by around 1%pa over 1973-2008.

In terms of my empirical application, note I regress consumption on regional dummies. Therefore, to the extent that regional variation contributes to consumption dispersion, I control for effects *between* regions and measure the average effects on inequality *within* regions. However, one would be concerned that a large increase in regional price variation reflects a stretching of the house price distribution over and above pure regional effects. For this reason, I concentrate on the period before the financial crisis, during with the common price change is most pronounced.

As a final point, house price shocks can also, of course, be local or even idiosyncratic at the house-specific level. Remember, by not considering these shocks, my analysis of how house prices drive consumption inequality

³²The only UK region for which house prices follow a noticeably different path is Northern Ireland (not shown). However, Northern Ireland is a tiny component of the UK and is not included in the FES data which are used in this paper. Accordingly, and strictly speaking, my analysis concerns Great Britain. However, I refer to United Kingdom (which includes Northern Ireland), because this is the more recognized country title.

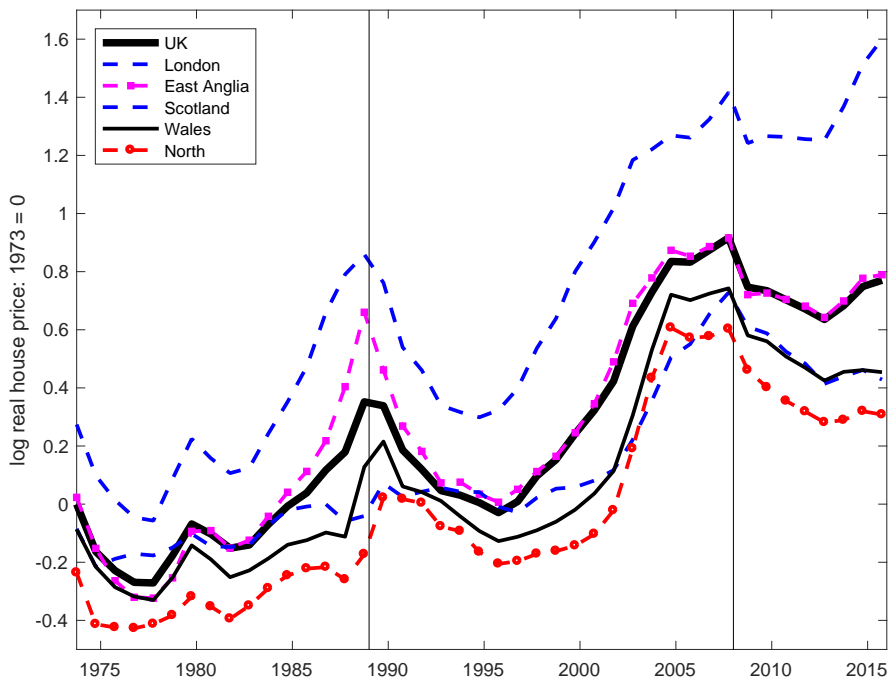


Figure 2: House Prices Across the UK.

Source: Nationwide.

Notes: Selected regions shown only. Omitted regions are: Yorkshire, North West, East Midlands, West Midlands, South East and Northern Ireland. Profiles are log real prices re-based to UK average = 0 in 1973. Profiles are smoothed using a 3-period moving average.

growth is conservative. However, even if I could incorporate idiosyncratic shocks, the analysis here becomes more involved. Idiosyncratic house price changes are often caused by investments in the house such as developments and refurbishment or, in the other direction, by excessive depreciation caused by dereliction. These idiosyncratic changes would not be new information to the household, and need not represent a net change in the household's lifetime wealth.

3.7 Estimating the Income Process Taking into Account House Price Shocks

One motivation for this study is to see how allowing for asset price shocks affects estimates of income risks. I therefore estimate the income process specified by equations 5 and 6, using the moment conditions given by formulae 14-16. As discussed, these cross-sectional moments have been used to estimate this dynamic income model since [Blundell and Preston \(1998\)](#). [Blundell et al. \(2013\)](#) provide a recent analysis. The estimation requires using fixed groups of households across years. I therefore estimate the process for two decade-of-birth cohorts; those born in the 1950s, and the 1960s. I require that all heads be between 25 and 60 in any period. Therefore the sample period for those born in the 1960s begins in 1994. For the estimation procedure, I use a two-step generalized minimum distance procedure. This estimation procedure is slightly more involved than the classical minimum distance procedure used in, say, [Blundell and Preston \(1998\)](#) because the empirical moments are not additively separable from the estimable parameters. The procedure is discussed in detail in appendix E, which also discusses computation of (asymptotic) standard errors. I give an overview here.

In the first step I compute the $3T$ moment conditions which comprise the left hand sides of formulae 14-16. These are the changes to the variances of log residual income and log consumption, and their covariance for each period. I then compute observable components of the right hand side. The first set of components comprises the wealth shares and covariances with the income and consumption distributions from the BHPS. These are available over 1994 – 2007 using the procedure described above. I obtain measures for 1989 – 1994 for the 1950s cohort by using the growth rates at the equivalent ages for the 1960s cohort, observed over 1999 – 2004. This imputation is likely conservative because the housing wealth share grew strongly over 1999 – 2004, whereas it likely declined over 1989 – 1994. I am therefore likely understating the wealth share in the early 1990s. Second, I use the estimates of the house price shocks, ζ_t^H , from the ODPM. Finally, I impose that α_t , the annuity value of transitory shocks, is zero. Estimation allowing for a small positive value of α is discussed by [Blundell et al. \(2013\)](#). In the second step I estimate components of the income process by the minimum distance procedure. The components are $\text{Var}(\nu_t)$, the variance of permanent shocks, and $\Delta\text{Var}(u_t)$, the growth in variance of transitory shocks. I therefore ultimately estimate 2 parameters in each of T years. Using the $3T$ moment conditions results in T degrees of freedom. To emphasize, I extend the standard framework used in the literature by incorporating the extra information on housing wealth and prices. When estimating I use the identity weighting matrix to eliminate well-known bias from using the optimal weighting matrix (see [Altonji and Segal, 1996](#)). Appendix E also gives details of a multi-step estimator when using the (estimated) optimal weighting matrix.

4 Main Results

Before discussing the main results, I present the inequality measures for the two cohorts I study. The income and consumption measures are shown together with their covariance in Figure 3. This figure shows both the raw measures, year by year, alongside smoothed time series computed using a 3-period moving average. Notice that, overall, and on all measures, inequality grew faster for the 1950s cohort than for the 1960s cohort. To reconcile the inequality movements completely requires estimating income processes for each cohort. First, though, I focus on one core feature which is common across cohorts: consumption inequality grew more slowly than the other moments until the mid-1990s, then more quickly thereafter.

I illustrate the main point of the paper in figure 4. It shows the log real house price, alongside the difference between the consumption measure of inequality and the covariance measure for both cohorts. For the rest of the paper, I call changes to this difference the ‘excess’ growth in consumption inequality. Again the figure shows the raw points for the inequality measures alongside a smoothed time-series, now computed using a 4th-order polynomial. Recall from the discussion in section 2 that both the consumption and covariance inequality measures capture dispersion in permanent income. According to that discussion, if changes to permanent differences across households are driven by shocks to flow income, such as from labour earnings, then these inequality measures should grow at a similar rate.³³ Figure 4, however, shows two points clearly. First,

³³It is reasonable that the level of consumption inequality is higher than the covariance, because of measurement error in expenditures. I focus on growth in the inequality measures, not levels, implicitly assuming that the magnitude of measurement error remains constant over time.

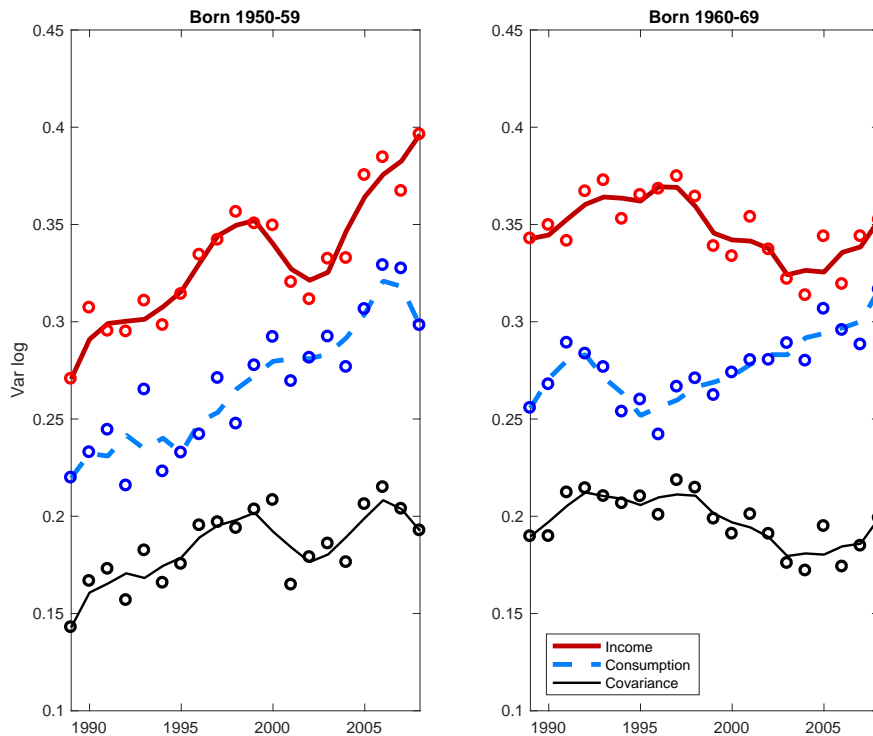


Figure 3: Inequality Profiles by Cohort.

Notes: Solid lines show profiles smoothed using a 3-period moving average.

consumption inequality grew much faster than the covariance measure from the mid-1990s until the mid-2000s for both cohorts. This excess growth for the 1960s cohort from trough to peak is around 6 log points (growth in the log measure by 0.06), according to the smoothed time series. This is a large increase, compared to say, the growth in the 1980s inequality boom.³⁴ Second, this excess growth in inequality is clearly strongly correlated with the housing market. The troughs and the peaks coincide almost perfectly with the house price index for both cohorts.

As discussed in section 2, the relationship between consumption inequality and asset prices depends on how assets are distributed across households. To explore this relationship, I show some features of the distribution of the housing wealth share in figure 5. It shows distributions for 1995 and 2005 wealth values for the two cohorts, using kernel density plots. It is worth emphasizing that this figure uses the *gross* housing wealth share: that is, it doesn't subtract mortgages, but includes the value of all first and second homes. This gross housing wealth share is the relevant measure in the calculations. The density estimates are always bi-modal, with peaks for non-home-owners and then for households with positive housing wealth. The figure further shows that, in both years, the 1950s cohort on average had a higher housing wealth share than the 1960s cohort. It also shows the effect of the housing boom: the 1960s cohort had a much higher share of housing wealth in 2005 than the 1950s cohort did in 1995, when they were the same age. Finally, the figure shows that the housing wealth share is particularly dispersed in the later years and particularly for the younger cohorts. A high housing wealth share

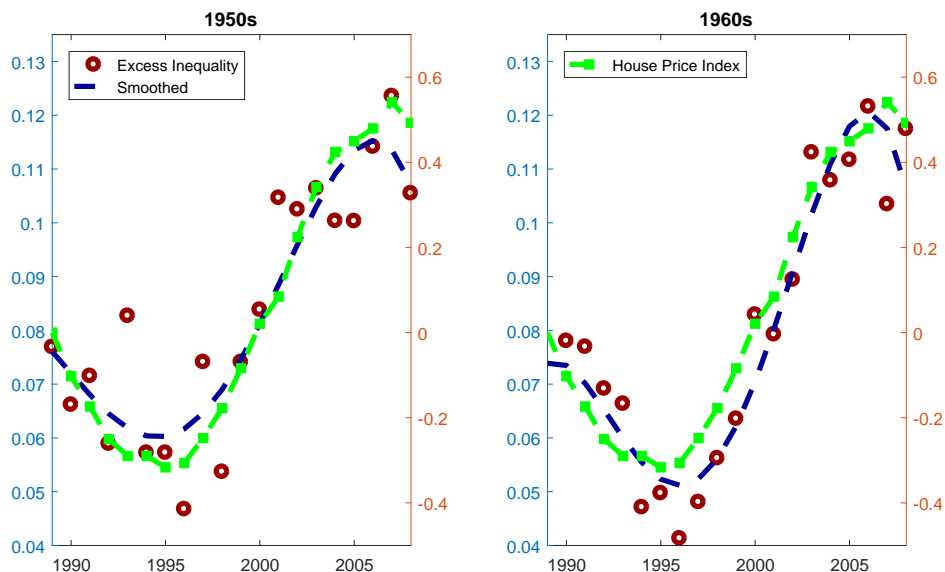


Figure 4: House Prices and Excess Growth in Consumption Inequality.

Notes: Left-hand axes show the difference between var log consumption and the covariance of consumption and income in log points. This is the ‘excess’ growth in consumption inequality. Axes are the same across both graphs. Right-hand axes show detrended real house price index. Difference in inequality measures smoothed using a quartic polynomial.

typically indicates that the household is highly leveraged.

Most importantly, how do these housing wealth shares covary with the consumption and income distributions? Before discussing the estimates that are used in the estimation I first show roughly how food consumption and income correlate with different measures of engagement with the housing market. These correlations are shown in figure 6, for the 1950s cohort, and figure 7, for the 1960s cohort. The top left plot shows correlations of residual income and residual food expenditure with a binary indicator of homeownership. The top right shows correlations with house value and is therefore restricted to homeowners. The bottom left shows correlations with the key statistic used in the analysis, the gross housing wealth share. The bottom right panel introduces a leverage variable. This leverage variable is defined as zero for renters and defined for homeowners as the ratio of net financial *debt* (mortgages and loans minus financial assets) to gross housing wealth, truncated at zero. The overall impression we obtain from these panels is that all measures of engagement are positively correlated with food expenditure and with income. In particular, in the bottom panels, the left-hand graph points to a housing wealth effect, while the bottom right-hand panel points to effects coming from the collateral channel. Both these graphs therefore indicate that a positive house price should cause inequality to grow, while a negative shock should cause inequality to decline. I analyse the housing wealth channel in the present section. The collateral channel is left until section 5.

Figures 6 and 7 show the rough correlations of key housing measures with consumption and with income. However, they cannot be related directly to the model. For this we need not correlations, but covariances. Detailed estimates of these wealth share moments are given in table 1. Here I show the wealth measures using the complete wealth information in 1995, 2000 and 2005. As discussed in section 3, I use partial wealth

³⁴See Blundell and Etheridge (2010) for an analysis of historical movements in inequality measures in the UK.

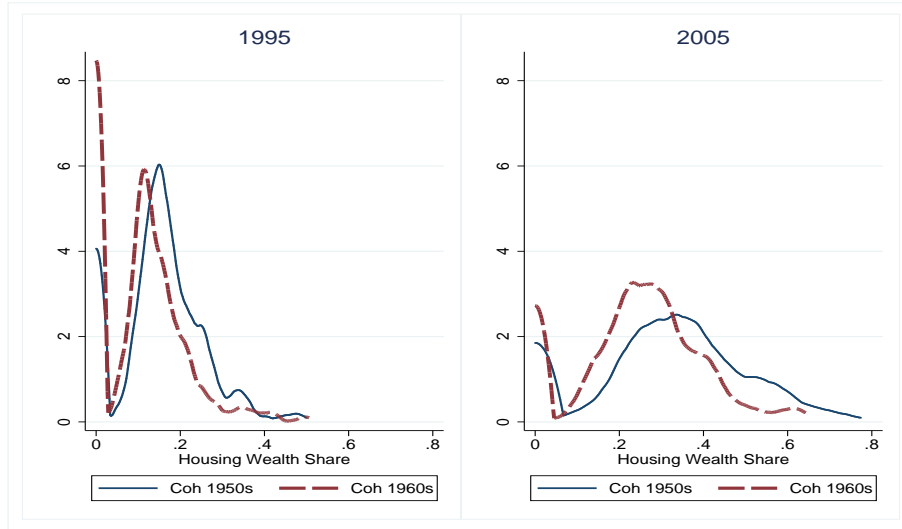


Figure 5: Kernel Density Estimates of the Gross Housing Share by Cohort and by Year.

Notes: Housing wealth share computed as described in text. Kernel densities generated using an epanechnikov kernel with the following bandwidths: 0.015 for 1950s cohort and 0.0125 for 1960s cohort in 1995; 0.03 for 1950s cohort and 0.02 for 1960s cohort in 2005.

information to impute moments in other years. For the years shown, the most important moments are given in the first two columns. They relate to the inputs used in formulae 14-16 in section 2. In particular, the first column gives $\text{Cov}(\psi_t^H, c_t)$, the covariance of the housing wealth share with log household consumption, net of predictable components. Recall that we do not observe total household consumption but rather expenditure on food. This measure therefore uses the covariance with food, scaled up by the inverse of the income elasticity of demand, taken to be 0.38 (see Etheridge, 2015 for a detailed discussion and estimation). Notice that the covariance is positive and statistically significant in most years. This feature implies that positive house price shocks should indeed cause the consumption variance to increase. To provide context to the covariances, and to link to the information shown in figures 6 and 7, I show raw correlations of the housing wealth share measure with *food* expenditure in square brackets.

Similarly, the second column shows the covariance of the housing wealth share with log household income, net of predictable components. This moment relates to the theoretical predictions from expression 15. However, in this case, comparing theory with data is complicated by the fact that the relevant input is not the covariance with current income, but rather with log *permanent* income, y_t^P , which is not observed. Current income might not give a good proxy in this case, because, in theory, it should be more positively correlated with wealth than is permanent income. This is because households save, rather than consume, the bulk of transitory fluctuations. If $\text{Cov}(\psi_t^H, y_t)$ deviates from $\text{Cov}(\psi_t^H, y_t^P)$ for this reason, then estimating permanent income by averaging y_t and y_{t-1} or by instrumenting y_t with y_{t-1} will not help, because transitory income in period $t-1$ is also mostly saved, and so induces a similar bias.³⁵ Therefore I conclude that the estimates in the second column are upper bounds on the relevant input. Most importantly, they are quantitatively smaller than the covariance with consumption. Therefore, for the baseline results I ignore $\text{Cov}(\psi_t^H, y_t^P)$ and treat it as zero, but I also discuss calculations incorporating the covariance with current income.

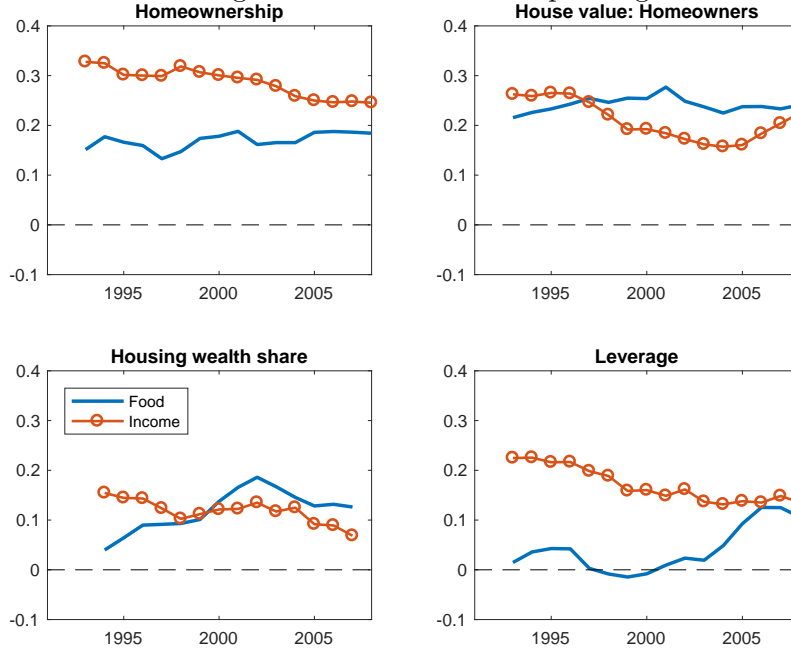
³⁵Likewise, y_t^P should not be instrumented with y_{t+1} because y_{t+1} is included in the calculation of ψ_t^H .

Table 1: WEALTH SHARE ESTIMATES BY COHORT

year	$\text{Cov}(\psi_t^H, c_t)$	$\text{Cov}(\psi_t^H, y_t)$	Mean ψ_t^H	Mean ψ_t^O	Var ψ_t^H	Var ψ_t^O	Var(f_t)	$\text{Cov}(\psi_t^{*H}, c_t)$
1950s Cohort								
1995	0.0065 (0.0061) [0.065]	0.0079 (0.0027) [0.164]	0.1501 (0.0043)	-0.0485 (0.0035)	0.0098 (0.0007)	0.0069 (0.0006)	0.1473 (0.0149)	...
2000	0.0156 (0.0067) [0.114]	0.0079 (0.0034) [0.131]	0.1698 (0.0058)	-0.0404 (0.0035)	0.0168 (0.0013)	0.0059 (0.0004)	0.1626 (0.0150)	0.0093 (0.0055)
2005	0.0271 (0.0103) [0.143]	0.0092 (0.0046) [0.106]	0.3071 (0.0092)	-0.0298 (0.0052)	0.0340 (0.0021)	0.0117 (0.0012)	0.1526 (0.0155)	...
1960s Cohort								
1995	0.0116 (0.0036) [0.132]	0.0050 (0.0023) [0.108]	0.1063 (0.0040)	-0.0627 (0.0028)	0.0085 (0.0006)	0.0045 (0.0004)	0.1291 (0.0096)	...
2000	0.0106 (0.0041) [0.109]	0.0061 (0.0020) [0.129]	0.1282 (0.0038)	-0.0649 (0.0026)	0.0096 (0.0006)	0.0041 (0.0003)	0.1446 (0.0111)	0.0098 (0.0033)
2005	0.0141 (0.0074) [0.100]	0.0023 (0.0031) [0.034]	0.2431 (0.0063)	-0.0744 (0.0037)	0.0214 (0.0013)	0.0071 (0.0005)	0.1335 (0.0110)	...

Notes: Standard errors shown in parentheses and obtained using a bootstrap with 1000 draws. Numbers in square brackets are correlation coefficients. The values shown in the first column are the correlation coefficient of the wealth share with food expenditure. ψ^H is the share of housing in lifetime wealth, ψ^O the share of other financial wealth (including mortgages). y_t is log residual current income, f_t is log residual food. ψ^{*H} is the share of housing in lifetime wealth, including pension wealth. $\text{Cov}(x_t, c_t)$ is computed as $\text{Cov}(x_t, f_t)/\chi_f$ where χ_f is the elasticity of food consumption with respect to total consumption, taken to be 0.38. See text for more details.

Figure 6: Correlations of Housing Indicators with Food Spending and Income: 1950s Cohort

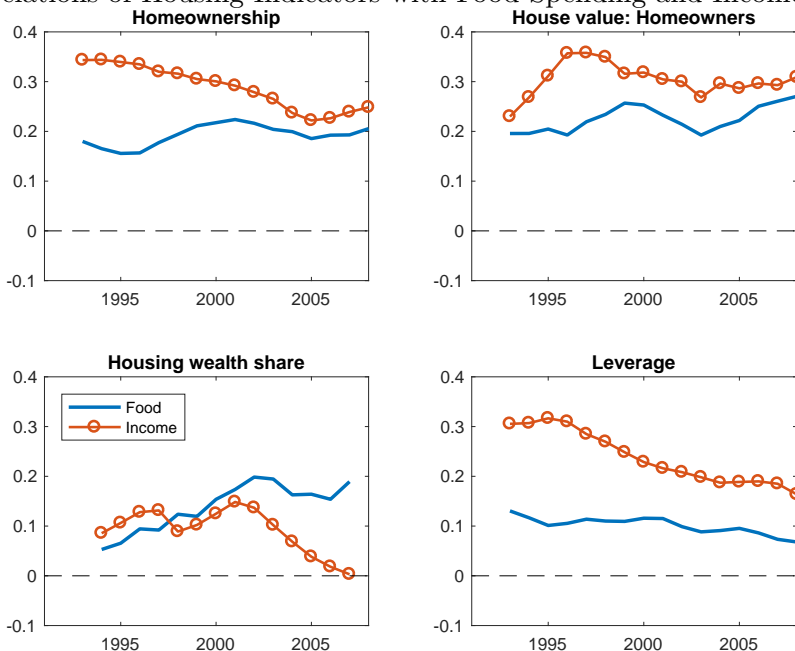


Notes: Profiles smoothed using a 3-period moving average. Leverage is defined as zero for renters and for homeowners with positive current financial wealth, i.e. for whom savings are greater than mortgage debt. Otherwise it is calculated as minus the ratio of current financial wealth (principally savings minus mortgage debt) to gross housing wealth.

It is worth discussing briefly the other empirical moments in table 1, presented in the columns further to the right. The third column corroborates the message from figure 5. It shows that the share of life-time wealth allocated to housing is greater for the 1950s cohort in each time period, but is higher for the 1960s cohort when comparing age by age. The fourth column shows strikingly that other wealth is negative in all years. This feature arises because average mortgage debt exceeds other financial wealth in all these cells. I omit from table 1 the mean share of human capital wealth, $\bar{\pi}_t$; this share can be inferred simply by subtracting the shares of housing wealth and other financial wealth from 1. Moving on, the fifth and sixth columns show that the variance of the housing wealth share is typically higher than the variance of the other wealth share: the variance of the housing wealth share has grown substantially over time, reflecting the effect of the house price boom on stretching this aspect of inequality. The seventh column shows dispersion in residual log food.

Finally, for robustness, the eighth column shows an alternative estimate of the covariance of the housing wealth share with consumption, the most important moment driving inequality. This alternative estimate is obtained by including pension wealth in total life-time wealth. The pension wealth data come from calculations performed by Emmerson et al. (2008) and discussed in section 3. I only show the covariance for 2000 because the data are not available for 2005 and, for 1995, they are only available for those who also appear in the 2000 wave, which gives a selected sample. This final column shows that including pension wealth does bring the covariance down somewhat, particularly for the older cohort, implying a lesser effect of house price shocks on consumption inequality. However, pension wealth, in the estimates shown, is likely overweighted. This is because the pension incomes used to compute this wealth measure have been discounted quite lightly at 2%pa. Pension wealth is, however, likely discounted highly by consumers during working life (Carroll, 1997), and is, of course, highly illiquid. Presumably, the estimate of the covariance of the wealth share with consumption when using a higher

Figure 7: Correlations of Housing Indicators with Food Spending and Income: 1960s Cohort



Notes: Profiles smoothed using a 3-period moving average. For notes on variable definitions see figure 6.

discount rate would lie somewhere between the two estimates shown.

I can plug in the main estimates from table 1 to compute simply the effect of the house price boom on excess growth in inequality, according to expressions 14 to 16. The results are given in table 2. Before discussing these, note that I define the house price boom as between 1997 and 2004, because it is only between these years that house prices increased above trend. Over 1989 to 1995, in contrast, house prices decreased in real terms, while in 1996 and 2005-06, house prices increases were positive but below trend. Over the boom years of '97 to '04, however, the log of average house prices increased above trend by 0.46 (58%).

Table 2: CONTRIBUTION OF HOUSE PRICE SHOCKS TO 'EXCESS' GROWTH IN CONSUMPTION INEQUALITY

Cohort	Accumulated shock growth to house prices $\sum_t \zeta_t$	Average contribution of covariance $\frac{1}{T} \sum_t \text{Cov}(\psi_{t-1}^H, c_{t-1})$	Total contribution to (log) consumption variance	% contribution of observed total excess growth
1950s	0.46 (58%)	0.019	0.019	47.1
1960s	0.46 (58%)	0.017	0.017	29.9

Notes: The observed excess growth in consumption inequality over 1997-2004 was 4.1 log points for the 1950s cohort and 5.8 log points for the 1960s cohort. 'Excess' growth in consumption inequality is defined as growth in the variance of log consumption over and above the covariance of log consumption with log income.

The first column of table 2 shows the size of the house price boom. The remaining columns of table 2 give further details on the calculation of excess inequality growth. Recall from expression 16 that the variance of log consumption is boosted by $2\zeta_t^H \text{Cov}(\psi_{t-1}^H, c_{t-1})$ where ζ_t^H is the house price shock. For example, in 2000,

the shock to prices above trend was 0.071 log points (7.4%) and the wealth covariance for the 1950s cohort was 0.019. Therefore the contribution to the inequality measure was $2 \times 0.071 \times 0.016 = 0.0023$, or 0.23 log points. Accordingly the second column of table 2 shows the average value of $\text{Cov}(\psi_{t-1}^H, c_{t-1})$ over '97-'04. The sum total of all the yearly contributions is given in the third column. For the 1950s cohort the effect of house price increases is calculated to be around 1.9 log points, just under 50% of the observed excess growth in consumption inequality. For the 1960s cohort, the contribution was slightly smaller, and given that the excess growth in inequality was higher, it forms an even smaller proportion of the total.

The remainder of the excess inequality growth can be explained in a number of ways. First, remember that the estimated trend increase in house prices, at 3.4%pa, is quite high. At this level, positive shocks to house prices are perhaps understated. If households have lower expectations for house price increases then the true shocks would be higher, as would be the implied contribution to excess inequality growth. For example, if expected growth in house prices were only 1%pa, then the estimated positive shock over the boom years would be raised by 17 log points. Accordingly the contribution to excess inequality would increase by 17 percentage points for the 1950s cohort and by 11 percentage points for the 1960s cohort. Second, I have assumed a house price process that is common across the distribution. If, as discussed by [Ortalo-Magné and Rady \(2006\)](#), house prices become more dispersed during booms and more compressed during slumps, then this heterogeneity would also contribute to the observed pattern of consumption inequality. Recall I have deliberately used a sample period in this analysis characterized by aggregate rather than heterogeneous shocks, because heterogeneous shocks are harder to assess. Nevertheless heterogeneous price movements are undoubtedly important and likely explain some portion of growth in the inequality that remains.

As a final point I discuss how I can refine these calculations, and check robustness against alternative modelling assumptions. Recall from formulae 14-16 in section 2 that other moments of the wealth data are relevant for explaining the profile of inequality: specifically the variance of the housing wealth share, and the covariance with income. The contribution from the variance of the housing share, given by its product with the square of the housing shock, is negligible. For example in 2000 it was $0.071^2 \times 0.016 < 0.0001$. The covariance of the housing share with permanent income is potentially more important. When we take it into account by proxying permanent income with current income, then the estimated total contribution of housing shocks goes down to 1.6 log points for the 1950s cohort and 1.5 log points for the 1960s cohort. But as discussed above, if, as is likely, current income is an imperfect proxy for permanent income, then these results provide a lower bound on the contribution. In terms of robustness, the most important assumption used is the choice of discount rate used on future income. When the discount rate is lower, for example, housing wealth is less important in total life-time wealth and so the effect of shocks is reduced. Therefore to check how sensitive results are to this choice, I re-compute the statistics using a low net discount rate of 1%pa and a high discount rate of 7%pa. Results are robust to these more extreme choices: for the 1950s cohort a high discount rate raises the estimated contribution by about 10 percentage points, while for the 1960s cohort a high discount rate raises the estimated contribution by around 8 percentage points. The results with the lower discount rate are roughly symmetric. I show the wealth share moments under these alternative discount rates in appendix F.2, table A3.

5 Exploring Non-Linearities in the Model

In this section I solve the full model specified in section 1 and use it to explore the effect on inequality of potentially important consumption non-linearities. After discussing how I parametrize and calibrate the model, I use it for three specific purposes. First I compare the wealth share measure to the true elasticity from the full benchmark model to compare the accuracy of the main approximation. I find the wealth share approximation works well in a number of key dimensions. Second, given that my main approximation is imperfect, and most importantly omits effects from collateral constraints, I then use the benchmark model to compute the contribution to the true elasticity arising from collateral effects. I then use this simulation-based measure to quantify the effect of this collateral channel on the empirical inequality profiles.

Third, and finally, I extend the benchmark model yet further to include additional model elements. Here, I allow for time-varying rent-to-price ratios, arising from a constant rental price. I also explore preferences over housing and consumption that differ from the Cobb-Douglas benchmark, using elasticities of substitution both below and above one. In this third group of exercises I again compare the wealth share measure to the true elasticity. The conclusion from these exercises is that, after exploring a rich variety of models built on the same class, the main results are highly robust. I leave details of the numerical solution of all the models to appendix D.

5.1 Calibration

A model period is a year. Households enter the model aged 25 and retire at 65. They die with certainty at age 82, matching average life expectancy for those aged 25 in the UK.

As in section 2, households' intratemporal preferences are Cobb-Douglas, implying a unitary elasticity of substitution between consumption and housing. I set the mean elasticity of intertemporal substitution equal to 0.5, implying a coefficient of relative risk aversion of 2.

I compute the deterministic component of post-tax income from the BHPS. As with the other life-cycle profiles described later, I proceed by first taking means for each year and age-group cell. I then compute a life-cycle component by regressing these cells on a fourth-order polynomial in age, and time dummies.³⁶ For the stochastic component, household income has transitory and fully permanent components, characterized by a unit root. The variance of permanent shocks is set to 0.005, corresponding to an annual standard deviation of shocks of 7%. This value roughly matches the growth in all inequality measures for both cohorts shown in figure 3. This value further implies that each cohort experiences a growth in the variance of log earnings between ages 25 and 65 of 20 log points. The variance of the initial distribution of permanent differences is set to 0.2 and the variance of transitory shocks is set to 0.05. To interpret these values, it is useful to consider, for example, the initial variance of log income for the 1960's cohort, observed in the FES, and shown in figure 3. According to the parameters chosen, the bulk of the observed initial variance arises from differences in the permanent component.

³⁶Mean life-time income is likely determined more by cohort rather than time effects. However, it is important to control for time effects when estimating life-cycle housing wealth. For consistency, therefore, I use the same procedure when estimating all life-cycle inputs.

The remaining share can then be thought of as being comprised roughly equally of genuine transitory differences and pure measurement error, which is left unmodelled.

I determine households' pension income using a formula based on last-period post-tax income. Specifically, I use a two-part linear formula to match two features of observed life-cycle income: the mean, and the change in income dispersion for those aged 70 – 75 compared to those aged 55 – 60. Accordingly, households in the model receive a pension composed of 90% of any final-period income below the median, and 50% of any final-period income above the median. The full profile for income over the life-cycle is shown in appendix figure A1. All relevant life-cycle variables (income, housing wealth, financial wealth) are equivalized using the modified OECD scale and normalized by setting the median income for incoming households at age 25 to 1.

House prices grow at mean rate 1% per annum, with a standard deviation of innovations of 9% pa. The variance of house price shocks matches the yearly innovations to prices in the data. The house price growth I use in the model is somewhat below the trend growth observed in the data since the late 1960s and used in section 4, which is around 3.4% pa. From the viewpoint of the model, therefore, house prices in the UK have grown more than expected. When simulating the model forward, however, house prices grow on trend, and therefore I do not try to recreate dynamics of the wealth distribution determined by the boom. Depreciation is set at 2% per annum, and is fully offset by maintenance. The yearly interest rate is set to 3%. I set the transaction cost to 5%, in line with established estimates.

I set the collateral requirement to be 10%. This value captures the median loan-to-value on mortgages to first-time buyers in the UK until the financial crisis.³⁷ The collateral constraint here has to be met in each period. The model therefore does not fully account for the long-term nature of standard mortgage contracts. This feature of mortgage contracts is included, for example, in the model of Kaplan et al. (2017). They emphasize its role in their result that a collateral channel is less important than the wealth channel in explaining consumption behaviour. My model therefore likely allows for a greater collateral effect than theirs. However, although mortgage contracts in the UK are long-term, they show important short-term features. Specifically, as Best, Cloyne, Ilzetzki, and Kleven (2018) discuss, almost all households in the UK refinance their mortgages 2 years after origination, because mortgage contracts are sold with low initial rates which thereafter revert to a rate which is much higher. To be able to refinance in this way, households need to meet similar lending conditions as at origination.

I set the the initial (joint) distribution of homeownership, of house values, of (liquid) financial wealth and income as follows: I pool data for those aged between 23 and 27 over all years. I simulate individuals at age 25 to match the means and variances of all these variables. The important associations are a strong positive correlation between house value (for owners) and financial wealth, and a slightly weaker positive correlation between house value and income. All other correlations, including those with homeownership, are weaker and consequently set to zero.

³⁷Data are available from the Council of Mortgage Lenders. See for example <https://data.gov.uk/dataset/median-loan-to-value-ratio-for-first-time-buyers> for data from 2005. Even during the slump in the early 1990s median loan-to-value ratios were at least 90%.

I calibrate the model by matching to three life-cycle profiles. These are: the homeownership rate; mean housing wealth, and the fraction of homeowners who are highly levered. This ‘constrained’ fraction I compute in both the model and the data as the fraction of households with a ratio of net financial debt to house value greater than 60%. Results are robust to using different thresholds.³⁸ As discussed above, I compute these life-cycle profiles from the data by allowing for age and for time effects. I base the profiles on 2000 values, when complete information on wealth is available, and after the market had picked up from the mid 1990s, but long before the market had reached a peak. The choice of moments in the calibration is similar to that in [Berger et al. \(2018\)](#) (henceforth in this section BGLV), except I focus on the constrained fraction rather than on mean liquid financial wealth. I do this to match closely the likely importance of collateral effects across the population.

The parameters used in the calibrated model are shown in table 3. The value for housing services in the utility function is higher than in calibrations using US data, because the ratio of housing wealth to income is higher. To match life-cycle saving I calibrate the rate of time preference β and the bequest motive Φ . The value of the discount rate here is particularly important for matching the share of households that are highly constrained. It therefore gives a good indication of the ultimate strength of the various channels from house prices through to household consumption. In studies which find a strong affect of collateral constraints on consumption (and house prices themselves), such as BGLV or [Favilukis et al. \(2017\)](#), then β is typically low (strong discounting). In contrast a wealth effect is typically more important when β is higher such as in [Kaplan et al. \(2017\)](#). The value of β found here is between these studies, though closer to that in [Kaplan et al.](#).

Table 3: CALIBRATED HOUSING MODEL: PARAMETER VALUES

Based on external evidence:	γ	r	μ	σ_{pH}	δ	Ξ	κ	σ_{yP}^2	σ_{yT}^2
	2	3%	1%	9%	2%	10%	5%	0.005	0.05
Chosen to hit lifecycle	a	β	Φ	χ_{25}	χ_{65}	χ_{ret}			
	0.83	0.953	16.0	0.0405	0.0385	0.0395			

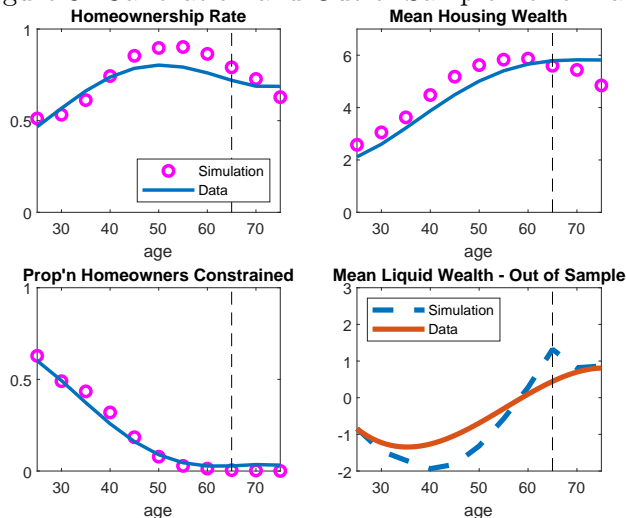
Notes: σ_{pH} is the standard deviation of house price innovations. $r \equiv R - 1$ is the risk-free interest rate. All other variables are defined in section 1, and discussed in the text.

Finally, I choose the rental cost to match life-cycle home ownership. The rental cost parameter here can be thought of as a reduced form which captures various aspects of the true reasons to rent or own. Most importantly, these include the moral hazard premium on renting, and a pure utility taste for renting versus owning. Here I flexibly choose an age-varying rental cost to match homeownership over the lifecycle. I allow for a gradient on the rental cost between ages 25 and 65 and a discrete change to a different rental cost after retirement.

The data and model moments used in the calibration are shown in figure 8 in the top two and the bottom left panels. These panels show I match the targeted statistics well. In particular, apart from young ages, the fraction of households who are constrained is low both in the model and in the data. I also show mean liquid wealth as in BGLV in the bottom right panel for completeness, and as a measure of out-of-sample performance. This panel shows I match mean liquid wealth well on average, although this mean is slightly too low at middle ages, and slightly too high at retirement.

³⁸A noticeable difference arises between the simulations and the data only in the number of households with a leverage greater than 80%. For example, in the data there are a non-negligible fraction of households with leverage greater than 95%. This reflects

Figure 8: Calibration and Out-of-Sample Performance



Notes: ‘Constrained’ is defined as having liquid financial debt (or ‘other current’ debt) greater than 60% of gross housing wealth.

5.2 Accuracy of the Wealth-Share Approximation

I now show the accuracy of the housing wealth share, ψ^H , as an approximation of the true elasticity. I do this by computing ψ^H as the share of gross housing wealth in total discounted expected life-time wealth in the model simulations, as in section 3. In particular, I discount future income at the market interest rate (3%pa). The raw relationship between the two measures is shown in the left-hand plot of figure 9, which is similar to a computation shown by Kaplan et al. (2017).³⁹ To produce this graph I pool simulated households aged 30 – 65, who are representative of those studied in this paper.⁴⁰⁴¹ I then group the simulated households into renters and, for homeowners, into ventiles of the distribution of ψ^H , the housing wealth share. For each ventile I then plot the average of the true elasticity, together with a 45-degree line. The overall impression is that the housing wealth share matches the true elasticity closely. Looking more closely, the relationship is slightly curved, rather than linear. The relationship is flat for those with low housing wealth share because these households are typically young, and the housing wealth share increasingly misses out collateral effects at lower values. For higher values of the housing wealth share, the housing wealth share approximation is slightly above the true elasticity on average, because the true elasticity is brought down by transactions costs. At values of the housing wealth share that are higher still, the housing wealth share matches the true elasticity more closely again. This is because those with a large housing wealth share are generally closer to the threshold of the adjustment decision, and consequently have a higher elasticity.

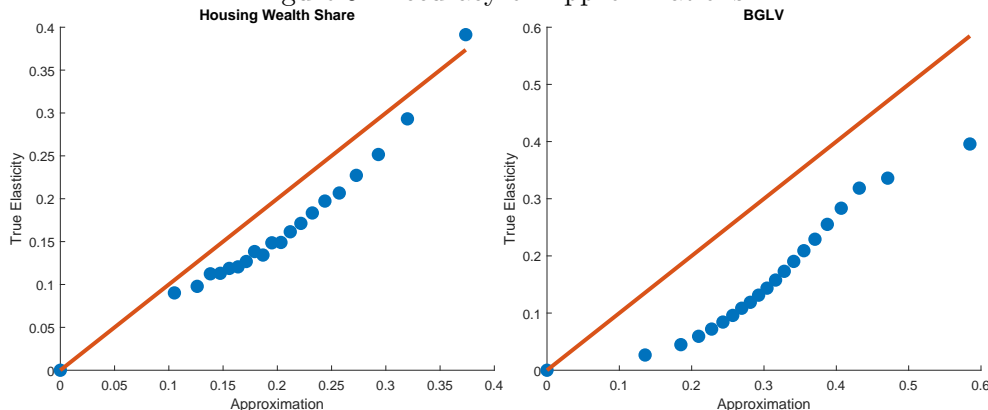
the fact that deposit requirements were generally lower than the 10% imposed in the model, and the availability of credit cards gives households extra debt capacity, which again is not reflected in the model.

³⁹To remove numerical error, these plots use smoothed measures of these variables. Specifically I use predicted values from regressions of the elasticity and the approximation on an interaction of age with cubic polynomial interactions of all the state variables.

⁴⁰I remove those aged between 25 and 30 to allow a ‘burn-in’ for the simulations. In particular, even though I match the initial distribution of homeownership, housing wealth and liquid wealth quite closely, most simulated households are not at their optimal house value, and there is quite a lot of turnover in the housing market in the first couple of years.

⁴¹In fact, if I introduce retirees then the approximation matches the true elasticity much more closely because these households face no income risk and so behave almost like PIH consumers for whom the approximation is exact.

Figure 9: Accuracy of Approximations



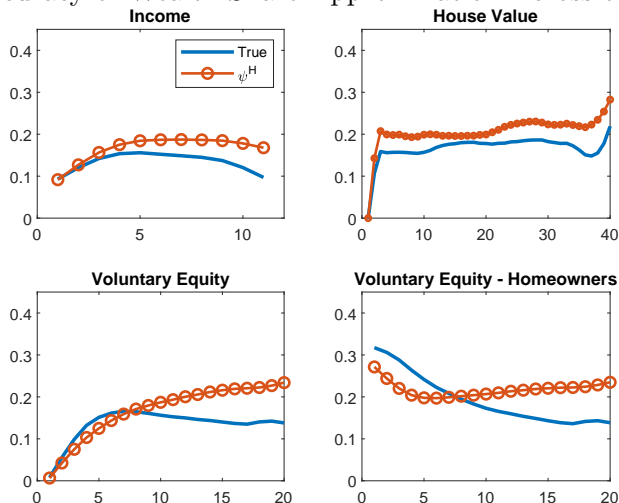
Notes: Each plot is produced by pooling the sample of 30-65 year olds and splitting homeowners into ventiles according to each approximation. I then compute the average value of the approximation and the average value of the true elasticity. Beforehand, the approximations and the true elasticity are smoothed to remove numerical error. See text for more details.

For comparison, on the right hand side I show exactly the same computation, but using BGLV’s approximation to replace the housing wealth share measure. In general, and in terms of pure correlation, BGLV’s approximation is much closer to the true elasticity than the housing wealth share. A raw correlation of the true elasticity with the BGLV approximation for the entire population is 0.94, while the correlation with the housing wealth share is 0.90.⁴² However, BGLV’s approximation is further away from the true elasticity in terms of raw magnitudes. Here the effect of transactions costs drives a wedge between the BGLV and the true elasticity. In fact while the coefficient from a regression of the true elasticity on the housing wealth share approximation is 0.94, the coefficient from the regression of the true elasticity on BGLV’s approximation is 0.80.

I explore the accuracy of ψ^H as an approximation further in figure 10. The figure shows the approximation against the true elasticity across various aspects of the state space. The top left grid shows both measures across the income grid. The top right shows both measures across the grid of house value. Here, the value at zero gives the elasticity for renters. The bottom two plots split the simulated population into ventiles of ‘voluntary equity’. As is standard, voluntary equity is a measure of cash-on-hand and equals the sum of financial wealth and mortgage capacity. The bottom right also shows voluntary equity, but now only for homeowners. The bottom two plots differ mainly in the mix of households who have low equity. Those with low equity in the left-hand plot combine renters and homeowners who are constrained, while the equivalent group on the right hand side omits renters and consequently has a far higher average elasticity. Overall, this figure shows that the wealth share measure captures the true elasticity reasonably well across most aspects of the state space. Overall, I conclude from this exercise that the wealth-share approximation does a good job of matching the true elasticity across much of the population.

⁴²For the prime aged sample studied, these correlations are slightly lower at 0.87 for BGLV’s approximation, and 0.78 for the wealth share approximation.

Figure 10: Accuracy of Wealth Share Approximation Across the State Space



Notes: Plots show pooled simulations of 30-65 year olds. Income plot shows the 11 income grid points. House value plot shows a split of simulated homeowners into 40 bins. Voluntary equity plots show a split of simulated households into 20 bins. In terms of model variables voluntary equity is defined as $A + (1 - \Xi) \frac{1-\delta}{R} p^H H$: financial wealth plus mortgage capacity.

5.3 Quantifying the Role of Collateral Constraints

Although the wealth share elasticity matches the true elasticity fairly well, I can use the model simulations to capture discrepancies in the approximation and to extend the analysis in section 4. My approach is to use a model-based elasticity which can be taken back to the data. This approach therefore complements the main approximation developed in section 2 which is derived from theory.

To keep with the spirit of the paper, I do this in a simple, linear way. I use the following estimating equation on homeowners:

$$\omega_i^{H,TRUE} = \beta_0 + \psi_i^H + \beta_1 lev_i + e_i^\omega$$

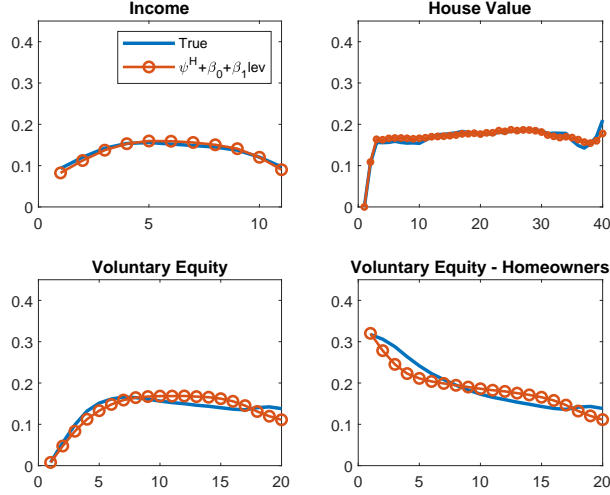
where $\omega_i^{H,TRUE}$ is the true simulated elasticity, ψ_i^H is the housing wealth share, and lev_i is the ratio of net financial debt to gross housing wealth, again truncated at 0. β_0 and β_1 are coefficients, and e_i^ω is a residual. Notice two features of this equation: first, and to maintain consistency with the original approximations, the coefficient on the housing wealth share is restricted to be 1. Second I estimate a common coefficient on leverage for all working-age households. Of course, I could make the revised approximation more accurate by allowing time-varying coefficients or, more generally, by allowing for fully flexible functional forms for all state variables. However, such an approach would be out of keeping with the spirit of the paper. The model based elasticity used here is transparent and informative.⁴³

In terms of results, $\hat{\beta}_1$ is estimated to be 0.24.⁴⁴ Before taking these estimates to the data I examine their effect on helping the approximation match the true elasticity across the state space. I therefore show the revised approximation alongside the true elasticity across various aspects of the state space in figure 11. It shows

⁴³I also used an indicator of being close to constraints, as in that used in the calibration exercise above. The results are very similar. I favour using the linear leverage indicator for ease of interpretation.

⁴⁴ $\hat{\beta}_0$ is estimated to -0.12. This value mainly reflects the effect of transactions costs on the true elasticity.

Figure 11: Accuracy of Extended Approximation



Notes: See notes to figure 10.

broadly that the revised elasticity matches the true elasticity better across all dimensions of the space shown, and particularly for homeowners with low voluntary equity.

Finally, I use the estimated effect of leverage to return to the computations in section 4. Specifically I revise the moments used in equations 15 and 16. These become:

$$\begin{aligned} \Delta \text{Cov}(c_t, y_t) &\approx \text{Contribution from Income} + \underbrace{\zeta_t^H \text{Cov}(\psi_{t-1}^H + \hat{\beta}_1 \text{lev}_{t-1}, y_{t-1}^P)}_{\text{Contribution from Housing}} \\ \Delta \text{Var}(c_t) &\approx \text{Contribution from Income} + \underbrace{(\zeta_t^H)^2 \text{Var}(\psi_{t-1}^H + \hat{\beta}_1 \text{lev}_{t-1}) + 2\zeta_t^H \text{Cov}(\psi_{t-1}^H + \hat{\beta}_1 \text{lev}_{t-1}, c_{t-1})}_{\text{Contribution from Housing}} \end{aligned}$$

where I gloss over the contributions from income for ease of exposition. Put simply, I revise the approximated elasticity from ψ_t^H to $\psi_t^H + \hat{\beta}_1 \text{lev}_t$. Similarly to the exercise in section 4 I plug in the relevant statistics and compute the contribution to growth in excess consumption inequality. The numbers are shown in table 4. Again, the first column shows the estimated above-trend growth in consumption inequality in the peak years of the boom. The second column now shows the total contribution arising from inequality in housing wealth share, taken from table 2 in section 4. The third column shows the contribution arising from dispersion in leverage across the population. To calculate this contribution I compute not only the covariance of leverage with consumption but also the covariance with current income, as a proxy for permanent income. I do this because figures 6 and 7 indicate that the covariance of income with leverage is more quantitatively important than with the housing wealth share. In fact, for the 1960's cohort, the contribution from the covariance of leverage with consumption is a little over 0.7 log points while the (negative) contribution from the covariance of leverage with income is under 0.4 log points. The third column of table 4 therefore shows that the overall contribution from leverage is positive, but small. The final column shows the revised estimates of the fraction of total excess inequality growth explained by housing market factors. These numbers are slightly larger, but very similar, to those shown in table 2 in section 4.

Table 4: CONTRIBUTION OF HOUSE PRICE SHOCKS TO CONSUMPTION INEQUALITY: INCLUDING COLLATERAL EFFECTS

Cohort	Accumulated shock growth to house prices $\sum_t \zeta_t$	Total contribution to excess growth, log points: wealth share	leverage	% contribution of observed total excess growth
1950s	0.46 (58%)	1.9	0.1	48.6
1960s	0.46 (58%)	1.7	0.4	36.1

Notes: The observed excess growth in consumption inequality over 1997-2004 was 4.1 log points for the 1950s cohort and 5.8 log points for the 1960s cohort. ‘Excess’ growth in consumption inequality is defined as growth in the variance of log consumption over and above the covariance of log consumption with log income.

5.4 Allowing for Additional Model Features

As a final exercise, I introduce additional features into the benchmark model discussed in section 5, to test the robustness of the approximation. I explore extensions in two directions. First, I explore model properties when dropping the assumption of a constant rent-price ratio. Second, I drop the assumption of intra-temporal homotheticity, and allow the elasticity of substitution between consumption and housing to deviate from 1. Compared to the models I analyse here, therefore, my approximations are now based on a version of the model that is far simpler. Nevertheless, here I show that my main results hold. In fact, if anything, the analysis here strengthens my main conclusions.

The first extension is motivated by the discussion in section 2: the available evidence implies that the rent-price ratio in the UK has been far from stable. The rent-price ratio is in fact markedly counter-cyclical; housing rents stay fairly constant as house prices rise, and consequently the rent-price ratio falls. I therefore change the model by assuming that the rental price stays on trend despite any house price change. In fact, rental prices themselves rise when house prices rise, so compared to the benchmark model, the modelling here provides an extreme in the opposite direction. Introducing a constant rental price makes the computation more difficult because the house price must be re-introduced as a state variable. To simplify the computation I set growth in house prices in all the extensions in this subsection to 0%.⁴⁵ Accordingly, the model is recalibrated, because even holding house price growth constant, a constant rental price makes renting more attractive relative to the benchmark. Hence, using the same parametrization as in section 5 would imply a counter-factually low homeownership rate. The values of the re-calibrated parameters are shown in appendix table A2.

I show the result of this and the other extensions in figure 12, which builds on figure 9 by showing the relationship between my housing wealth share approximation and the true elasticity for the simulated population. To frame the results, the top left panel of figure 12 replicates the result from figure 9, showing the relationship between the housing wealth share and the true elasticity using the main model from section 5. The top right panel of figure 12 then shows the relationship between these two measures in the model with constant rental prices.

⁴⁵Details of the computation of the model extensions are discussed alongside the computation of the benchmark model in appendix D.

Again, the figure splits the simulated population into renters, and, for homeowners, into ventiles of the housing wealth share distribution. The striking result from the figure is that the house price elasticity is pushed up for everyone. This result arises because a house price increase shifts more households towards renting, which becomes comparatively cheaper. For existing homeowners in particular a shift to renting releases a substantial amount of liquid assets, which boosts consumption. A similar affect arises for renters: when house prices rise they lose the incentive to save for a home, and so increase immediate consumption. For these reasons, these results should be taken with a note of caution: this simple extension does not capture the fact that homeownership itself is procyclical. Therefore to capture a *rise* in homeownership when house prices rise likely requires model features which are beyond the scope of the current paper. However, to the extent that this simple model extension captures cyclicalities in rent-price ratios, we see both that the housing wealth share remains a good approximation of the true elasticity, and that there is little effect on the implications for inequality.

The bottom row of figure 12 shows results when I change the specification for preferences. Both these figures use the calibration used for the current version of the benchmark displayed in the top left corner. In terms of preferences, the literature has found varied estimates for the elasticity of substitution between consumption and housing, and there is as yet no agreement as to whether these two goods are gross substitutes or gross complements.⁴⁶ The bottom left hand of figure 12 uses $\tau = 0.8$, which makes consumption and housing mild complements. With these preferences, holding fixed the wealth effect, a house price increase causes households to increase their expenditure on housing, and to decrease their non-housing consumption. Accordingly, it shows that the total elasticity is now lower than the housing wealth share approximation. However, most importantly, the gradient of the true elasticity with respect to the housing wealth share is now even closer to 1. This implies that using the housing wealth share accurately captures the effect on consumption *inequality*, even if the effect on the consumption *aggregate* is more muted.

The bottom right hand panel uses $\tau = 1.2$, which makes consumption and housing mild substitutes. In contrast to the bottom left hand panel, the consumption response for renters to a positive house price shock is now positive. The true consumption elasticity for all households is now raised compared to the benchmark, although the gradient of the true elasticity with respect to the housing wealth share is now a little below 1.

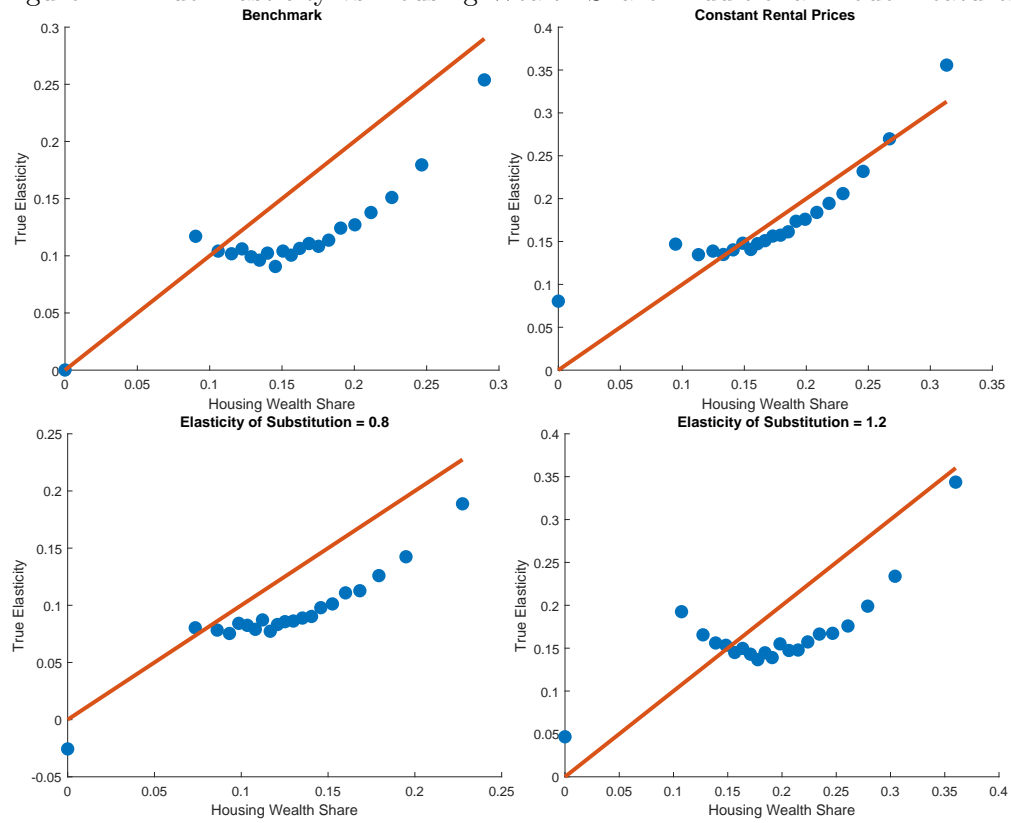
To conclude, it seems therefore, in terms of the relationship between house prices and consumption inequality, the effect of house price shocks is largest when housing and consumption are complements, i.e. when τ is less than 1. Moreover the effect of house prices on consumption inequality is little changed when rental prices are smoother. In fact, overall, it seems that the same pattern emerges; the housing wealth share gives a good indication of the true elasticity across its distribution, and across a variety of models.

6 The Effect of House Price Shocks on Estimates of Income Risk

A main motivation for this study is to see how allowing for asset price shocks affects estimates of income risks. I therefore estimate the model specified by formulae 14-16 both allowing for the effect of house price shocks

⁴⁶As two recent examples, Li, Liu, Yang, and Yao (2016) find an elasticity of substitution of around 0.5, while Piazzesi, Schneider, and Tuzel (2007) find a value of 1.25.

Figure 12: True Elasticity vs Housing Wealth Share: Additional Model Features



Notes: Each plot is produced by pooling the sample of 30-65 year olds and splitting homeowners into ventiles according to each approximation. I then compute the average value of the approximation and the average value of the true elasticity. Beforehand, the approximations and the true elasticity are smoothed to remove numerical error. See text for more details.

and then re-estimating, imposing all effects to be zero. Here, I just take into account the impact on inequality arising from wealth effects; the results in section 5 imply that taking into account collateral effects would yield quantitatively very similar results. I use the minimum distance procedure discussed in section 3, and, in more detail, in appendix E.

Figure 13 gives estimates of the variances of permanent income shocks under both strategies. The estimates are given in roughly 5-year averages to show trends at medium frequencies. I have also chosen the bands to match different periods in the housing market. Looking broadly at the series, the figure shows that the variance of permanent shocks is higher for the 1950s cohort than for the 1960s cohort on average. This result reflects the fact that inequality grew quicker for the older cohort on all measures over the sample period. It is also noticeable that the estimates show a dip in the variance around the turn of the millennium for the 1950s cohort followed by a small spike after 2005. In fact, these estimates are likely biased by the effects of a prominent sequence of tax and benefit changes arriving in the late 1990s. As figure 3 above shows, inequality measures were compressed after the turn of the millennium for both cohorts and then rebounded. This profile arose because of redistributive reforms introduced from 1998 onwards by a new centre-left government, the effect of which were eroded after around 2003.⁴⁷ I could therefore improve the estimates of income risks by taking these reforms, and their erosion, into account. Ignoring these factors, however, does not detract from the point of this section, which is to assess the role of house-price shocks on the income risk estimates.

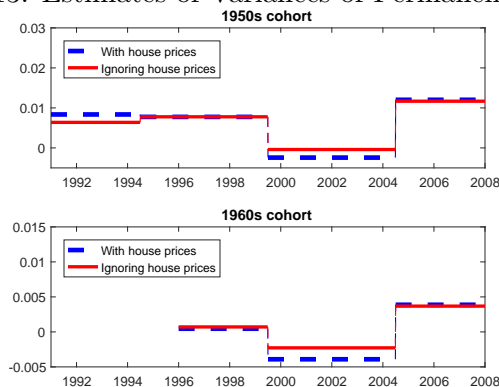
I now turn to the difference between the estimates when allowing and when not allowing for housing wealth shocks. Broadly speaking, when housing shocks are included, the implied variances of permanent income shocks are pushed up during the housing market slump in the early 1990s and down during the boom around 2000. In the boom, for example, when the house price increases are excluded, then growth in consumption inequality is mis-attributed to a high variance of permanent shocks. Conversely, the new estimates are higher at the beginning of the period when house prices were declining. During these early years, house price declines had a downwards effect on consumption inequality. Therefore, the level of permanent shocks must have been quite high to induce the actual growth in consumption inequality that is observed.

In terms of comparing the estimates with and without house price shocks, the magnitudes are not large in any particular year, but accumulate over a prolonged house price boom or slump. For example, for the recession of the early 1990s, including housing wealth effects increases the estimates of yearly permanent shocks for the 1950s cohorts by around 0.25 log points. This is around 30% of the estimated average yearly variance of permanent shocks for the 1950s cohort, or 40% of the average variance in the early 1990s only. Alternatively, we can quantify the differences by considering how they accumulate over time. For example when examining the income process over the peak years of the house price boom, over 1997-2004, then ignoring the effect of house price shocks leads us to *overstate* the effect of permanent income shocks by a cumulative 1.2 log points over the period for the 1950s cohort. In other words we would erroneously attribute inequality movements to a set of permanent income shocks with a standard deviation of 11% of income.

The effect of including house price shocks on estimates of the growth of the variance of transitory shocks is not shown here, but is roughly equal and opposite to the effect on the permanent shock estimates. This is

⁴⁷An earlier working paper of the current study, Etheridge (2013), examines the effect of these reforms in detail.

Figure 13: Estimates of Variances of Permanent Shocks



because the growth in the variance of transitory shocks is largely explained by ‘extra’ growth in the variance of log income: that is, growth that is not generated by shocks to permanent income. Therefore, given any path for the variance of log income, when the estimated variance of shocks to permanent income is higher, then the corresponding estimates of the growth of variance of transitory shocks must be lower.

Before concluding, it is worth relating these results to standard estimates in the literature. [Blundell et al. \(2013\)](#) use the same dataset also to study the 1950s cohort alongside those born in the 1940s for the period up to the mid 1990s. They show that the 1950s cohort, in particular, experienced a spike in the variance of permanent shocks during recessions, such as that in the early 1990s. The spike they observe in the early 1990s is smoothed out in the profiles presented here. However, one result of the analysis here is that, taking into account housing wealth effects *increases* the spike in the estimated variance of permanent shocks in this recession even further.

7 Conclusions

Across much of the world, house prices have fluctuated wildly in the past 30 years, seemingly in connection with the wider macroeconomy, and in particular, aggregate consumption. Given that housing wealth and home ownership are unevenly distributed, it seems important to ask not only how house price shocks affect mean consumption, but also how they affect consumption inequality. This paper is the first to examine the effect of large house price swings on inequality in living standards.

In this paper I have used a standard but detailed model of choices over housing and consumption. I have developed a theoretical framework for capturing several channels through which house price shocks affect various inequality measures. A strength of my main analysis is that it accounts not only for wealth effects but also income and substitution effects. A further strength of my approach is that this framework is developed for general classes of preferences. While I focus in this study on Cobb-Douglas preferences, my application could be extended more widely. Finally I show how to extend this theoretical framework to account for effects arising from collateral constraints, using a simulation-based approach.

The empirical application shows that inequality movements were large over the period I study, and that they coincide strongly with house price movements exactly as is predicted by the theory. Moreover, my approximated model can capture a large fraction of the inequality movements quantitatively. I show that collateral effects work in the same direction as wealth effects, but are quantitatively smaller. Finally, my framework is used to develop better estimates of income risks from cross-sectional data. This is important because, over an entire life, the factor that affects household welfare the most is success in the labour market. However, identifying different types of income risks remains an econometric challenge. I show that accounting for house price movements changes the implied contribution of permanent and transitory risks. In particular, when house prices decline, then ignoring house price movements can lead estimates of permanent income risks to be understated.

The present study addresses inequality *within* observable categories: *within* regions and *within* age groups. Since the financial crisis, in the UK and elsewhere, the most important drivers of inequality have been *across* age groups and *across* regions. These inequalities are influenced in part by the housing market. This role of the housing market would benefit from further research.

References

- ADELINO, M., A. SCHOAR, AND F. SEVERINO (2016): “Loan Originations and Defaults in the Mortgage Crisis: The Role of the Middle Class,” *Review of Financial Studies*, 29, 1635–1670.
- AGUIAR, M. AND M. BILS (2015): “Has Consumption Inequality Mirrored Income Inequality?” *American Economic Review*, 105, 2725–2756.
- ALADANGADY, A. (2017): “Housing Wealth and Consumption: Evidence from Geographically-Linked Microdata,” *American Economic Review*, 107, 3415–3446.
- ALBANESI, S., G. DE GIORGI, AND J. NOSAL (2017): “Credit growth and the financial crisis: A new narrative,” *NBER Working Paper*, wp23740.
- ALTONJI, J. AND L. SEGAL (1996): “Small-sample bias in GMM estimation of covariance structures,” *Journal of Business & Economic Statistics*, 14, 353–366.
- ATTANASIO, O., L. BLOW, R. HAMILTON, AND A. LEICESTER (2009): “Booms and Busts: Consumption, house prices and expectations,” *Economica*, 76, 20–50.
- BANKS, J., Z. SMITH, AND M. WAKEFIELD (2002): “The distribution of financial wealth in the UK: evidence from 2000 BHPS data,” *IFS Working Paper*, w02/21.
- BELFIELD, C., D. CHANDLER, AND R. JOYCE (2015): “Housing: Trends in Prices, Costs and Tenure,” Tech. rep.
- BERGER, D., V. GUERRIERI, G. LORENZONI, AND J. VAVRA (2018): “House Prices and Consumer Spending,” *Review of Economic Studies*, 85, 1502–1542.
- BEST, M., J. CLOYNE, E. ILZETZKI, AND H. KLEVEN (2018): “Estimating the Elasticity of Intertemporal Substitution Using Mortgage Notches,” *NBER Working Paper*, wp24948.
- BHUTTA, N. AND B. KEYS (2016): “Interest Rates and Equity Extraction During the Housing Boom,” *American Economic Review*, 106, 1742–1774.
- BLUNDELL, R. AND B. ETHERIDGE (2010): “Consumption, income and earnings inequality in the UK,” *Review of Economic Dynamics*, 13, 76–102.
- BLUNDELL, R., H. LOW, AND I. PRESTON (2013): “Decomposing Changes in Income Risk Using Consumption Data,” *Quantitative Economics*, 4, 1–37.
- BLUNDELL, R., L. PISTAFERRI, AND I. PRESTON (2008): “Consumption inequality and partial insurance,” *American Economic Review*, 98, 1887–1921.
- BLUNDELL, R., L. PISTAFERRI, AND I. SAPORTA-EKSTEN (2016): “Consumption inequality and family labor supply,” *American Economic Review*, 106, 387–435.
- BLUNDELL, R. AND I. PRESTON (1998): “Consumption Inequality and Income Uncertainty,” *Quarterly Journal of Economics*, 113, 603–640.
- BOHRNSTEDT, G. W. AND A. S. GOLDBERGER (1969): “On the exact covariance of products of random variables,” *Journal of the American Statistical Association*, 64, 1439–1442.
- BROWNING, M., M. GØRTZ, AND S. LETH-PETERSEN (2013): “Housing wealth and consumption: a micro panel study,” *The Economic Journal*, 123, 401–428.
- CAMPBELL, J. AND J. COCCO (2007): “How do house prices affect consumption? Evidence from micro data,” *Journal of Monetary Economics*, 54, 591–621.

- CARROLL, C. (1997): “Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis*,” *Quarterly Journal of Economics*, 112, 1–56.
- CARROLL, C., M. OTSUKA, AND J. SLACALEK (2011): “How Large Are Housing and Financial Wealth Effects? A New Approach,” *Journal of Money, Credit and Banking*, 43, 55–79.
- CLAUSEN, A. AND C. STRUB (2013): “Envelope Theorems for Non-Smooth and Non-Concave Optimization,” *Mimeo*.
- COOPER, D. (2013): “House price fluctuations: the role of housing wealth as borrowing collateral,” *Review of Economics and Statistics*, 95, 1183–1197.
- CROSSLEY, T., P. LEVELL, AND H. LOW (2017): “Consumption Spending, Housing Investments and the Role of Leverage,” *mimeo*.
- CROSSLEY, T. AND C. O’DEA (2010): “The wealth and saving of UK families on the eve of the crisis,” *IFS Reports*, R71.
- CUTLER, D. AND L. KATZ (1992): “Rising Inequality? Changes in the Distribution of Income and Consumption in the 1980’s,” *American Economic Review*, 82, 546–551.
- DAVIS, M. AND F. ORTALO-MAGNÉ (2010): “Household expenditures, wages, rents,” *Review of Economic Dynamics*, 14, 248–261.
- DEATON, A. AND C. PAXSON (1994): “Intertemporal Choice and Inequality,” *Journal of Political Economy*, 102, 437–467.
- DÍAZ, A. AND M. J. LUENGO-PRADO (2010): “The Wealth Distribution with Durable Goods,” *International Economic Review*, 51, 143–170.
- DISNEY, R., J. GATHERGOOD, AND A. HENLEY (2010): “House price shocks, negative equity, and household consumption in the united kingdom,” *Journal of the European Economic Association*, 8, 1179–1207.
- EMMERSON, C., R. DISNEY, M. WAKEFIELD, AND G. TETLOW (2008): “British Household Panel Survey Pension Wealth Derived Variable Data, 1991-2001. 2nd Edition.” *UK Data Service Technical Report SN5725*.
- ETHERIDGE, B. (2013): “Increasing Inequality and Improving Insurance : House Price Booms and the Welfare State in the UK,” *Mimeo*.
- (2015): “A Test of the Household Income Process Using Consumption and Wealth Data,” *European Economic Review*, 78, 129–157.
- FAVILUKIS, J., S. LUDVIGSON, AND S. VAN NIEUWERBURGH (2017): “The Macroeconomic Effects of Housing Wealth, Housing Finance and Limited Risk-Sharing in General Equilibrium,” *Journal of Political Economy*, 125, 140–223.
- FELLA, G. (2014): “A generalized endogenous grid method for non-smooth and non-concave problems,” *Review of Economic Dynamics*, 17, 329–344.
- FOOTE, C., L. LOEWENSTEIN, AND P. WILLEN (2016): “Cross-Sectional Patterns of Mortgage Debt During the Housing Boom: Evidence and Implications,” *NBER Working Paper*, wp22985.
- GANONG, P. AND P. NOEL (2016): “The Effect of Debt on Default and Consumption: Evidence from Housing Policy in the Great Recession,” *Mimeo*.
- GARRIGA, C., R. MANUELLI, AND A. PERALTA-ALVA (2017): “A Macroeconomic Model of Price Swings in the Housing Market,” *Federal Reserve Bank of Saint Louis Working Paper*.
- GOODMAN, L. (1960): “On the exact variance of products,” *Journal of the American Statistical Association*, 55, 708–713.
- GUREN, A. (2018): “House Price Momentum and Strategic Complementarity,” *Journal of Political Economy*, 126, 1172–1218.

- HEATHCOTE, J., K. STORESLETTEN, AND G. VIOLANTE (2014): “Consumption and labor supply with partial insurance: An analytical framework,” *American Economic Review*, 104, 2075–2126.
- HUO, Z. AND J. RÍOS-RULL (2016): “Financial frictions, asset prices, and the great recession,” *mimeo*.
- ISKHAKOV, F., T. JORGENSEN, J. RUST, AND B. SCHJERNING (2017): “The Endogenous Grid Method For Discrete-Continuous Choice Dynamic Models With (or Without) Taste Shocks,” *Quantitative Economics*, 8, 317–365.
- JAPELLI, T. AND L. PISTAFERRI (2014): “Fiscal Policy and MPC Heterogeneity,” *American Economic Journal: Macroeconomics*, 6, 107–136.
- KAPLAN, G., K. MITMAN, AND G. VIOLANTE (2017): “The Housing Boom and Bust: Model meets evidence,” *NBER Working Paper*, wp23694.
- KRUEGER, D., K. MITMAN, AND F. PERRI (2016): “Macroeconomics and Household Heterogeneity,” *Handbook of Macroeconomics: Volume 2*.
- KRUEGER, D. AND F. PERRI (2005): “Understanding Consumption Smoothing: Evidence from the U.S. Consumer Expenditure Data,” *Journal of the European Economic Association*, 3, 340–349.
- (2006): “Does Income Inequality Lead to Consumption Inequality? Evidence and Theory,” *The Review of Economic Studies*, 73, 163–193.
- LANDVOIGT, T. (2017): “Housing Demand During the Boom: The Role of Expectations and Credit Constraints,” *Review of Financial Studies*, 30, 1865–1902.
- LEVY, H. AND S. JENKINS (2008): “Documentation for derived current and annual net household income variables, BHPS waves 1-16,” *Institute for Social and Economic Research working papers*.
- LI, W., H. LIU, F. YANG, AND R. YAO (2016): “Housing over time and over the life cycle: a structural estimation,” *International Economic Review*, 57, 1237–1260.
- LI, W. AND R. YAO (2007): “The Life-Cycle Effects of House Price Changes,” *Journal of Money, Credit and Banking*, 39, 1375–1409.
- MEGHIR, C. AND L. PISTAFERRI (2004): “Income variance dynamics and heterogeneity,” *Econometrica*, 72, 1–32.
- MEYER, B. AND J. SULLIVAN (2017): “Consumption and Income Inequality in the U.S. Since the 1960s,” *NBER Working Paper*, WP 23655.
- MIAN, A., K. RAO, AND A. SUFI (2013): “Household Balance Sheets, Consumption, and the Economic Slump,” *Quarterly Journal of Economics*, 128, 1687–1726.
- MIAN, A. AND A. SUFI (2011): “House prices, home equity-based borrowing, and the US household leverage crisis,” *American Economic Review*, 101, 2132–2156.
- (2016): “Who Bears the Cost of Recessions? The Role of House Prices and Household Debt,” in *Handbook of Macroeconomics: Volume 2*, 255–296.
- MIDRIGAN, V. AND T. PHILIPPON (2011): “Household Leverage and the Great Recession,” *NBER Working Paper*, wp16965.
- MOFFITT, R. AND P. GOTTSCHALK (2002): “Trends in the transitory variance of earnings in the United States,” *Economic Journal*, 112, C68–C73.
- ORTALO-MAGNÉ, F. AND S. RADY (2006): “Housing Market Dynamics: On the Contribution of Income Shocks and Credit Constraints,” *The Review of Economic Studies*, 73, 459–485.
- PIAZZESI, M. AND M. SCHNEIDER (2016): “Housing and Macroeconomics,” *Handbook of Macroeconomics: Volume 2*.

- PIAZZESI, M., M. SCHNEIDER, AND S. TUZEL (2007): "Housing, consumption and asset pricing," *Journal of Financial Economics*.
- STORESLETTEN, K., C. TELMER, AND A. YARON (2004): "Consumption and Risk Sharing over the Life Cycle," *Journal of Monetary Economics*, 51, 609–633.
- WOOLDRIDGE, J. (2002): *Econometric analysis of cross section and panel data*, MIT Press.

Online Appendices to ‘House Prices and Consumption Inequality’

A Appendix: Approximating Changes to Consumption and Housing Expenditures

In this appendix I derive an expression for shocks to non-housing consumption in the presence of income and house price risk. The proof builds on those in [Blundell, Low, and Preston \(2013\)](#) (henceforth referred to as BLP) and [Blundell, Pistaferri, and Saporta-Eksten \(2016\)](#).

The proof involves two steps: first characterizing shocks to consumption, housing choices and income, and second, equating these shocks by examining the expenditure account and the income account of the distribution of (the log of) future life-time resources. Throughout the derivation, I generally ignore 2nd- and higher-order terms in the approximations. See BLP for a detailed discussion of these. As a final point, through most of the derivation I suppress i subscripts: when solving the household’s problem in a life-cycle framework, the distinction between idiosyncratic and common variables is not important.

A.1 The Refined Model

Households are born at time $t = 0$, work until $t = T_w$ and die at time $t = T$. At time t the household maximises lifetime utility:

$$\max_{\{C_k, H_k\}_{k=t}^T} \mathbb{E}_t \left(\sum_{k=t}^T \beta^{k-t} u(C_k, H_k) \right)$$

where β is a subjective rate of time preference, assumed to be common across households. H_k is the housing choice in period $k \geq t$, C_k the non-housing consumption choice, and $u()$ is the per-period felicity function. A suitable benchmark for $u()$ is Cobb-Douglas preferences, $u(C, H) = \frac{1}{1-\gamma} (C^a H^{1-a})^{1-\gamma}$, where a is the share of expenditure on non-housing consumption and γ captures preferences for intertemporal substitution. Note that compared to the preferences described in [section 1](#), the household does not value bequests. As discussed in [section 1](#), and given the preferences for bequests described there, this omission is without loss of generality and simplifies the exposition.

The price process for housing is

$$p_t^H = p_{t-1}^H e^{\mu + \zeta_t^H} \tag{17}$$

where ζ_t^H is the shock to (log) house prices with $\mathbb{E}_{t-1} e^{\zeta_t^H} = 1$ and μ is a constant.

Here I examine a model with home-ownership only, and no renting. The model with renting (and without home-ownership) has the same solution and is conceptually simpler. The flow budget constraint in the current

model is

$$A_t + C_t + p_t^H H_t = RA_{t-1} + Y_t + (1 - \delta) p_t^H H_{t-1} \quad (18)$$

where A_t is liquid wealth at time t , Y_t is income, R the return on liquid wealth and δ is depreciation on housing.

Income evolves as in the standard permanent-transitory model:

$$\begin{aligned} \ln Y_t &= X_t \varphi_t + \ln P_t + \epsilon_t \\ \ln P_t &= \ln P_{t-1} + \nu_t \end{aligned} \quad (19)$$

such that P_t is permanent productivity level, X_t captures household observable characteristics and $X_t \varphi_t$ is the deterministic component of income based on these characteristics. The innovations ϵ_t and ν_{t+k} have mean zero and are independent of each another for all k .

Summarizing, it is useful to characterize the optimization problem in terms of a value function $\mathcal{V}_t(\cdot)$ as follows:

$$\mathcal{V}_t(A_{t-1}, H_{t-1}; p_t^H, P_t, X_t) = \max_{A_t, H_t, C_t} u(C_t, H_t) + \beta \mathbb{E}_t(\mathcal{V}_{t+1}(A_t, H_t; p_{t+1}^H, P_{t+1}, X_{t+1}))$$

subject to 17, 18, 19 and some exogenous deterministic process for X_t which is left unspecified. I characterize the value function differently in appendix D, where I discuss computation.

A.2 Approximate Growth Equations for Consumption and for Housing

With intertemporally separable preferences, the standard arguments of log-linearization apply. I now write $c_t \equiv \ln C_t$, and $h_t \equiv \ln H_t$. Letting λ_t be the marginal utility of wealth we have the standard Euler equation

$$\lambda_t = \mathbb{E}_t(\beta R \lambda_{t+1}) \quad (20)$$

which follows from the the first-order condition on A_t and the envelope condition on A_t at time $t + 1$.

Using a first-order approximation (see, for example [Blundell, Pistaferri, and Saporta-Eksten](#)) we obtain that

$$\Delta \ln \lambda_{t+1} \approx \Gamma_{t+1}^\lambda + v_{t+1}^\lambda \quad (21)$$

where Γ_{t+1}^λ is the predictable change to the marginal utility of wealth and v_{t+1}^λ is the unpredictable innovation.

The first-order conditions further imply:

$$u_C(C_t, H_t) = \lambda_t$$

and

$$u_H(C_t, H_t) + \beta \mathbb{E}_t(R \lambda_{t+1} p_{t+1}^H (1 - \delta)) = p_t^H \lambda_t \quad (22)$$

where $u_x()$ is the partial derivative of $u()$ with respect to argument x . We can insert equations 17 and 20 into equation 22 to obtain the following expression:

$$\begin{aligned} u_H(C_t, H_t) &= p_t^H \lambda_t - \lambda_t (p_t^H e^\mu (1 - \delta)) \left(1 + Cov_t \left(e^{v_{t+1}^\lambda}, e^{\zeta_{t+1}^H}\right)\right) \\ \implies u_H(C_t, H_t) &= \lambda_t p_t^H \left(1 - e^\mu (1 - \delta) \left(1 + Cov_t \left(e^{v_{t+1}^\lambda}, e^{\zeta_{t+1}^H}\right)\right)\right) \\ \implies u_H(C_t, H_t) &= \lambda_t \ddot{r}_t^H \end{aligned}$$

where \ddot{r}_t^H is a risk-adjusted, expected user cost of housing, with risk-adjustment factor $Cov_t \left(e^{v_{t+1}^\lambda}, e^{\zeta_{t+1}^H}\right)$.

These FOCs imply that

$$\Delta \ln u_C(C_{t+1}, H_{t+1}) = \Delta \ln \lambda_{t+1} \quad (23)$$

and

$$\begin{aligned} \Delta \ln u_H(C_{t+1}, H_{t+1}) &= \Delta \ln (\ddot{r}_{t+1}^H \lambda_{t+1}) \\ &\approx \mu + \zeta_{t+1}^H + \Delta \ln \lambda_{t+1} \end{aligned}$$

where ζ_{t+1}^H is the shock to the price of housing services, discussed, amongst elsewhere, alongside equations 10 in section 1. The second expression is an approximation of the first because it ignores innovations to the risk adjustment term $Cov_t \left(e^{v_{t+1}^\lambda}, e^{\zeta_{t+1}^H}\right)$.

Substituting expression 21 into expression 23 gives

$$\Delta \ln u_C(C_{t+1}, H_{t+1}) \approx \Gamma_{t+1}^\lambda + v_{t+1}^\lambda \quad (24)$$

and similarly

$$\Delta \ln u_H(C_{t+1}, H_{t+1}) \approx \Gamma_{t+1}^\lambda + v_{t+1}^\lambda + \mu + \zeta_{t+1}^H$$

Expanding $u_C(C_{t+1}, H_{t+1})$ around $u_C(C_t, H_t)$ gives

$$\begin{aligned} u_C(C_{t+1}, H_{t+1}) &\approx u_C(C_t, H_t) + u_{CC}(C_t, H_t) C_t \Delta c_{t+1} + u_{CH}(C_t, H_t) H_t \Delta h_{t+1} \\ \implies \Delta \ln u_C(C_{t+1}, H_{t+1}) &\approx \frac{u_{CC}(C_t, H_t) C_t}{u_C(C_t, H_t)} \Delta c_{t+1} + \frac{u_{CH}(C_t, H_t) H_t}{u_C(C_t, H_t)} \Delta h_{t+1} \\ &= \frac{u_{CC} C_t}{u_C} \Delta c_{t+1} + \frac{u_{CH} H_t}{u_C} \Delta h_{t+1} \end{aligned} \quad (25)$$

where the abbreviated notation on the last line is used to save space.

Similarly

$$\Delta \ln u_H(C_{t+1}, H_{t+1}) \approx \frac{u_{HC} C_t}{u_H} \Delta c_{t+1} + \frac{u_{HH} H_t}{u_H} \Delta h_{t+1}$$

Substituting expression 24 into expression 25 for consumption, and performing the equivalent manipulation for

housing, gives

$$\begin{aligned}\Gamma_{t+1}^\lambda + v_{t+1}^\lambda &\approx \frac{u_{CC}C_t}{u_C} \Delta c_{t+1} + \frac{u_{CH}H_t}{u_C} \Delta h_{t+1} \\ \Gamma_{t+1}^\lambda + v_{t+1}^\lambda + \mu + \zeta_{t+1}^H &\approx \frac{u_{HC}C_t}{u_H} \Delta c_{t+1} + \frac{u_{HH}H_t}{u_H} \Delta h_{t+1}\end{aligned}$$

Solving these simultaneous equations for Δc_t and Δh_t (i.e. going back one period) we obtain

$$\begin{aligned}\Delta c_t &\approx (\eta_{C,p^H} - \eta_{C,p^C}) (\Gamma_t^\lambda + v_t^\lambda) + \eta_{C,p^H} (\mu + \zeta_t^H) \\ \Delta h_t &\approx (\eta_{H,p^C} - \eta_{H,p^H}) (\Gamma_t^\lambda + v_t^\lambda) - \eta_{H,p^H} (\mu + \zeta_t^H)\end{aligned}\tag{26}$$

where p^C is the price of non-housing consumption (= 1 at time t because non-housing consumption is the numeraire), and $\eta_{x,y}$ is the Frisch elasticity of good x w.r.t price y , assumed constant across individuals and across consumption and housing choices. These elasticities are defined as

$$\begin{aligned}\eta_{C,p^C} &\equiv -\frac{u_{HH}u_C}{u_{CC}u_{HH} - u_{CH}^2} \frac{1}{C} > 0 \\ \eta_{C,p^H} &\equiv -\frac{u_{CH}u_H}{u_{CC}u_{HH} - u_{CH}^2} \frac{1}{C} \\ \eta_{H,p^C} &\equiv -\frac{u_{CH}u_C}{u_{CC}u_{HH} - u_{CH}^2} \frac{1}{H} \\ \eta_{H,p^H} &\equiv -\frac{u_{CC}u_H}{u_{CC}u_{HH} - u_{CH}^2} \frac{1}{H} > 0\end{aligned}$$

Notice further from expression 26 that the pure (marginal-utility constant) substitution effect of a shock to house prices on non-housing consumption is $\eta_{C,p^H}\zeta_t^H$. The pure substitution effect away from housing is $-\eta_{H,p^H}\zeta_t^H$.

As a further discussion of the elasticities, notice that when preferences are given by $u(C, H) = \frac{1}{1-\gamma} (C^a H^{1-a})^{1-\gamma}$, then γ is the coefficient of relative risk aversion for the consumption-housing composite and $-\frac{u_{CC}C}{u_C} = 1 + a(\gamma - 1)$ is the coefficient on consumption alone. The elasticity of intertemporal substitution (EIS) on consumption is therefore $-\frac{u_C}{u_{CC}C} = \frac{1}{1+a(\gamma-1)}$.

Next, it is useful to consider the pure innovations to non-housing consumption and to housing, net of predictable components. These are

$$\begin{aligned}\Delta \hat{c}_t &\approx (\eta_{C,p^H} - \eta_{C,p^C}) v_t^\lambda + \eta_{C,p^H} \zeta_t^H \\ \Delta \hat{h}_t &\approx (\eta_{H,p^C} - \eta_{H,p^H}) v_t^\lambda - \eta_{H,p^H} \zeta_t^H\end{aligned}$$

Finally, it is useful to solve expression 26 for Δh_t in terms of Δc_t . Eliminating $(\Gamma_t^\lambda + v_t^\lambda)$ gives

$$\Delta h_t \approx \frac{(\eta_{H,p^C} - \eta_{H,p^H})}{(\eta_{C,p^H} - \eta_{C,p^C})} (\Delta c_t - \eta_{C,p^H} (\mu + \zeta_t^H)) - \eta_{H,p^H} (\mu + \zeta_t^H)$$

When preferences are homothetic then $\frac{(\eta_{H,p^C} - \eta_{H,p^H})}{(\eta_{C,p^H} - \eta_{C,p^C})} = 1$ and so, in this case

$$\Delta h_t \approx \Delta c_t - (\eta_{C,p^H} + \eta_{H,p^H}) (\mu + \zeta_t^H)$$

Similarly, in terms of innovations,

$$\Delta \hat{h}_t \approx \Delta \hat{c}_t - (\eta_{C,p^H} + \eta_{H,p^H}) \zeta_t^H$$

in the homothetic case, and

$$\Delta \hat{h}_t \approx \frac{(\eta_{H,p^C} - \eta_{H,p^H})}{(\eta_{C,p^H} - \eta_{C,p^C})} (\Delta \hat{c}_t - \eta_{C,p^H} \zeta_t^H) - \eta_{H,p^H} \zeta_t^H$$

in the non-homothetic case.

A.3 Approximating Lifetime Resources

As in BLP I define a function $F : \mathbb{R}^{N+1} \rightarrow \mathbb{R}$ by $F(\xi) = \ln \sum_{j=0}^N \exp(\xi_j)$. By exact Taylor expansion around an arbitrary point $\xi^0 \in \mathbb{R}^{N+1}$

$$\begin{aligned} F(\xi) &= \mathcal{K} + \sum_{j=0}^N \frac{\exp \xi_j^0}{\sum_{k=0}^N \exp \xi_k^0} (\xi_j - \xi_j^0) \\ &\quad + \frac{1}{2} \sum_{j=0}^N \sum_{k=0}^N \frac{\partial^2 F(\bar{\xi})}{\partial \xi_j \partial \xi_k} (\xi_j - \xi_j^0) (\xi_k - \xi_k^0) \end{aligned} \quad (27)$$

where $\mathcal{K} = \ln \sum_{j=0}^N \exp \xi_j^0$ is constant, and $\bar{\xi}$ lies between ξ and ξ_0 .

A.3.1 Approximating the expenditure account of lifetime resources

I now examine the lifetime budget constraint, which must hold in addition to the flow budget constraint in equation 18. The lifetime budget constraint is:

$$\sum_{k=t}^T q^{k-t} C_k + \sum_{k=t}^T q^{k-t} r_k^H H_k = \sum_{k=t}^{T_w} q^{k-t} Y_k + R A_{t-1} + p_t^H \tilde{H}_t \quad (28)$$

where A_{t-1} is the level of liquid assets left at period $t-1$, p_t^H is the house price and $\tilde{H}_t = (1-\delta)H_{t-1}$ is the endowment of housing, inherited from the previous period. R is the risk-free interest rate on liquid assets, $q = \frac{1}{R}$, Y_k is household income and $r_k^H = \left(1 - e^{\mu + \zeta_{k+1}^H} (1-\delta)\right) p_k^H = \varsigma_k p_k^H$ is a generalized rental, or user, price of housing for $t \leq k \leq T-1$, with $\mathbb{E}_t \varsigma_k$ constant for all k and $r_T^H = p_T^H$. Notice that r_k^H differs from \tilde{r}_k^H defined above: whereas \tilde{r}_k^H is an expected user cost of housing, r_k^H is stochastic at time k .

I now expand the expenditure account of lifetime resources around $\mathcal{K}_e = \ln \sum_{j=0}^{T-t} \mathbb{E}_{t-1} \left(q^j C_{t+j} + q^j r_{t+j}^H H_{t+j} \right)$, on the left side of equation 28 the logarithm of expected discounted expenditures. Again I write $c_t \equiv \ln C_t$ and $h_t \equiv \ln H_t$. I define:

$$\begin{aligned} \xi_k &= c_{t+k} - kr && \text{for } k = 0, \dots, T-t \\ \xi_k^0 &= \mathbb{E}_{t-1} c_{t+k} - kr && \text{for } k = 0, \dots, T-t \\ \xi_k &= h_{t+\tilde{k}} + \ln r_{t+\tilde{k}}^H - \tilde{k}r && \text{for } k = T-t+1, \dots, 2(T-t)+1 \\ \xi_k^0 &= \mathbb{E}_{t-1} h_{t+\tilde{k}} + \mathbb{E}_{t-1} \ln r_{t+\tilde{k}}^H - \tilde{k}r && \text{for } k = T-t+1, \dots, 2(T-t)+1 \end{aligned}$$

where $\tilde{k} = k - (T-t+1)$ and $r = R - 1 \approx \ln R$. Applying the approximation formula in 27, ignoring second-order terms, and taking expectations with respect to information set \mathcal{I} :

$$\begin{aligned} \mathbb{E}_{\mathcal{I}} \ln \sum_{j=0}^{T-t} (q^j C_{t+j} + q^j r_{t+j}^H H_{t+j}) &\approx \mathcal{K}_e + \sum_{j=0}^{T-t} \theta_{t+j}^C (\mathbb{E}_{\mathcal{I}} c_{t+j} - \mathbb{E}_{t-1} c_{t+j}) \\ &\quad + \sum_{j=0}^{T-t} \theta_{t+j}^H ((\mathbb{E}_{\mathcal{I}} h_{t+j} + \mathbb{E}_{\mathcal{I}} \ln r_{t+j}^H) - (\mathbb{E}_{t-1} h_{t+j} + \mathbb{E}_{t-1} \ln r_{t+j}^H)) \end{aligned}$$

such that:

$$\begin{aligned} \theta_{t+j}^C &= \frac{\exp[\mathbb{E}_{t-1} c_{t+j} - jr]}{\sum_{j=0}^{T-t} \exp[\mathbb{E}_{t-1} c_{t+j} - jr] + \sum_{j=0}^{T-t} \exp[\mathbb{E}_{t-1} (h_{t+j} + \ln r_{t+j}^H) - jr]} \\ \theta_{t+j}^H &= \frac{\exp[\mathbb{E}_{t-1} (h_{t+j} + \ln r_{t+j}^H) - jr]}{\sum_{j=0}^{T-t} \exp[\mathbb{E}_{t-1} c_{t+j} - jr] + \sum_{j=0}^{T-t} \exp[\mathbb{E}_{t-1} (h_{t+j} + \ln r_{t+j}^H) - jr]} \end{aligned}$$

are the shares of discounted expenditure in total lifetime expenditure. I further define

$$\begin{aligned} s_t^C &= \sum_{j=0}^{T-t} \theta_{t+j}^C \\ s_t^H &= \sum_{j=0}^{T-t} \theta_{t+j}^H \end{aligned}$$

with $s_t^C + s_t^H = 1$. s_t^C and s_t^H give the expected life-time budget shares for consumption and for housing.

A.3.2 Approximating the income account of lifetime resources

Similarly to above, we now expand the income account around $\mathcal{K}_y = \ln \sum_{j=0}^{T_w-t} \mathbb{E}_{t-1} q^j Y_{t+j} + RA_{t-1} + \mathbb{E}_{t-1} p_t^H \tilde{H}_t$, the logarithm of expected discounted incomes. I write $y_t \equiv \ln Y_t$ and $\tilde{h}_t \equiv \ln \tilde{H}_t$. Letting $N = T_w - t + 1$, I

define:

$$\begin{aligned}
\xi_j &= y_{t+j} - jr \\
\xi_j^0 &= \mathbb{E}_{t-1} y_{t+j} - jr \\
\xi_N &= \ln A_{t-1} + r \\
\xi_N^0 &= \ln A_{t-1} + r \\
\xi_{N+1} &= \ln p_t^H + \tilde{h}_t \\
\xi_{N+1}^0 &= \mathbb{E}_{t-1} \ln p_t^H + \tilde{h}_t
\end{aligned}$$

Applying the approximation formula in 27, ignoring higher-order terms and taking expectations with respect to information set \mathcal{I} :

$$\begin{aligned}
\mathbb{E}_I \ln \left(\sum_{j=0}^{T_w-t} q^j Y_{t+j} + RA_{t-1} + p_t^H \tilde{H}_t \right) \\
\approx \mathcal{K}_y \\
+ \pi_t \sum_{j=0}^{T_w-t} \alpha_{t+j} (\mathbb{E}_I y_{t+j} - \mathbb{E}_{t-1} y_{t+j}) \\
+ \psi_{t-1}^H \left[\mathbb{E}_I (\ln p_t^H + \tilde{h}_t) - \mathbb{E}_{t-1} (\ln p_t^H + \tilde{h}_t) \right] \\
+ (1 - \pi_t - \psi_{t-1}^H) [\mathbb{E}_I \ln A_{t-1} - \mathbb{E}_{t-1} \ln A_{t-1}]
\end{aligned}$$

where

$$\begin{aligned}
\alpha_{t+j} &\equiv \frac{\exp [\mathbb{E}_{t-1} y_{t+j} - jr]}{\sum_{k=0}^{T_w-t} \exp [\mathbb{E}_{t-1} y_{t+k} - jr]} \\
\pi_t &\equiv \frac{\sum_{j=0}^{T_w-t} \exp [\mathbb{E}_{t-1} y_{t+j} - jr]}{\Lambda_t} \\
\psi_{t-1}^H &\equiv \frac{\exp \mathbb{E}_{t-1} \ln (p_t^H \tilde{H}_t)}{\Lambda_t} \\
\Lambda_t &\equiv \sum_{j=0}^{T_w-t} \exp [\mathbb{E}_{t-1} y_{t+j} - jr] + \exp \mathbb{E}_{t-1} \ln (RA_{t-1}) + \exp \mathbb{E}_{t-1} \ln (p_t^H \tilde{H}_t)
\end{aligned}$$

Intuitively, α_{t+j} is an annuitization factor for income for which $\sum_{j=0}^{T_w-t} \alpha_{t+j} = 1$, π_t is the share of human capital wealth in lifetime wealth, ψ_{t-1}^H is the share of housing wealth in lifetime wealth, and Λ_t is total discounted lifetime wealth itself.

A.4 Equating Innovations to the Expenditure and Income Accounts

Due to the lifetime budget constraint, the distributions of income and expenditure accounts can be equated with respect to any information set, \mathcal{I} . Applying the operator $\mathbb{E}_t - \mathbb{E}_{t-1}$ to the expenditure account:

$$\begin{aligned}
(\mathbb{E}_t - \mathbb{E}_{t-1}) \circ \ln \sum_{j=0}^{T-t} (q^j C_{t+j} + q^j r_{t+j}^H H_{t+j}) \\
&\approx (\mathbb{E}_t - \mathbb{E}_{t-1}) \circ \mathcal{K}_e + \sum_{j=0}^{T-t} \theta_{t+j}^C [(\mathbb{E}_t - \mathbb{E}_{t-1}) \circ c_{t+j}] + \sum_{j=0}^{T-t} \theta_{t+j}^H [(\mathbb{E}_t - \mathbb{E}_{t-1}) \circ (h_{t+j} + \ln r_{t+j}^H)] \\
&= \sum_{j=0}^{T-t} \theta_{t+j}^C [(\mathbb{E}_t - \mathbb{E}_{t-1}) \circ c_{t+j}] + \sum_{j=0}^{T-t} \theta_{t+j}^H [(\mathbb{E}_t - \mathbb{E}_{t-1}) \circ (h_{t+j} + \ln r_{t+j}^H)] \\
&= s_t^C \Delta \hat{c}_t + s_t^H (\Delta \hat{c}_t - (\eta_{C,p^H} + \eta_{H,p^H}) \zeta_t^H) + s_t^H \zeta_t^H \\
&= \Delta \hat{c}_t + s_t^H (1 - \eta_{C,p^H} - \eta_{H,p^H}) \zeta_t^H
\end{aligned} \tag{29}$$

as long as $u(\cdot)$ is homothetic.⁴⁸ Similarly applying the operator $\mathbb{E}_t - \mathbb{E}_{t-1}$ to the income account and rearranging:

$$\begin{aligned}
(\mathbb{E}_t - \mathbb{E}_{t-1}) \circ \ln \left(\sum_{j=0}^{T-t} q^j Y_{t+j} + RA_{t-1} + p_t^H \tilde{H}_t \right) \\
&\approx \pi_t \sum_{j=0}^{T-t} \alpha_{t+j} [(\mathbb{E}_t - \mathbb{E}_{t-1}) \circ y_{t+j}] + \psi_{t-1}^H [(\mathbb{E}_t - \mathbb{E}_{t-1}) \circ (\ln p_t^H + \tilde{h}_t)] \\
&= \pi_t [\nu_t + \alpha_t \epsilon_t] + \psi_{t-1}^H \zeta_t^H
\end{aligned} \tag{30}$$

where $(\mathbb{E}_t - \mathbb{E}_{t-1}) \circ \tilde{h}_t = 0$ because the housing endowment carried forward into time t is known after the realization of shocks at time $t-1$.

Equating expressions 29 and 30, and solving in terms of total changes to log consumption, gives

$$\Delta c_t \approx \Gamma_t^C + \pi_t (\nu_t + \alpha_t \epsilon_t) + \psi_{t-1}^H \zeta_t^H + s_t^H (\eta_{C,p^H} + \eta_{H,p^H} - 1) \zeta_t^H$$

where Γ_t^C is the predictable gradient on consumption, absorbing not only Γ_t^λ above, but also the shift towards non-housing consumption owing to trend house-price increases.

I now decompose the effect of the house price shock into endowment, income and substitution effects. The endowment effect is given by $\psi_{t-1}^H \zeta_t^H$. The income effect, when preferences are homothetic, is $-s_t^H \zeta_t^H$ to first

⁴⁸If $u(\cdot)$ is not homothetic then the right hand side of expression 29 becomes $\left(s_t^C + \frac{(\eta_{H,p^C} - \eta_{H,p^H})}{(\eta_{C,p^H} - \eta_{C,p^C})} s_t^H \right) \Delta \hat{c}_t + s_t^H \left(1 - \frac{(\eta_{H,p^C} - \eta_{H,p^H})}{(\eta_{C,p^H} - \eta_{C,p^C})} \eta_{C,p^H} - \eta_{H,p^H} \right) \zeta_t^H$

order. This implies that the substitution effect must be $s_t^H (\eta_{C,p^H} + \eta_{H,p^H})$. This gives

$$\Delta c_t \approx \Gamma_t^C + \pi_t(\nu_t + \alpha_t \epsilon_t) + \left(\underbrace{\psi_{t-1}^H}_{\text{Endowment effect}} + \underbrace{s_t^H (\eta_{C,p^H} + \eta_{H,p^H})}_{\text{Substitution effect}} - \underbrace{s_t^H}_{\text{Income effect}} \right) \zeta_t^H$$

as in equation 13. In the Cobb-Douglas benchmark, then $\eta_{H,p^H} + \eta_{C,p^H} = 1$, so the substitution effect reduces to $s_t^H \zeta_t^H = (1 - a) \zeta_t^H$. In this case the income effect cancels out the substitution effect and changes to consumption are driven by the endowment effect alone.

Finally, in the non-homothetic case, then we can write

$$\Delta c_t \approx \Gamma_t^C + \frac{(\eta_{C,p^H} - \eta_{C,p^C})}{(\eta_{C,p^H} - \eta_{C,p^C}) s_t^C + (\eta_{H,p^C} - \eta_{H,p^H}) s_t^H} \times \left(\pi_t(\nu_t + \alpha_t \epsilon_t) + \psi_{t-1}^H \zeta_t^H + s_t^H \left(1 - \frac{(\eta_{H,p^C} - \eta_{H,p^H})}{(\eta_{C,p^H} - \eta_{C,p^C})} \eta_{C,p^H} - \eta_{H,p^H} \right) \zeta_t^H \right)$$

B Appendix: Deriving the Approximate Cross-Sectional Covariance Structure of Income and Consumption

Here I derive the formulae for changes to the variance of consumption and covariance of income and consumption. Now we assume that the house price shock is common to all households. Throughout the following derivations, all variances, covariances and expectations are taken cross-sectionally over households, i . Therefore I use the notation $\text{Var}(x_t)$ to denote the variance of x_{it} , and $\text{Cov}(x_t, y_t)$ for the covariance of x_{it} with y_{it} .

Adapting the notation from appendix A, let v_{it}^C be the period- t idiosyncratic shock to consumption; v_{it}^{inc} the shock to income. As before ζ_t^H is the house price shock; ν_{it} the permanent shock to income, and ϵ_{it} the transitory shock to income. Then, when preferences are Cobb-Douglas:

$$v_{it}^C \approx \pi_{it}(\nu_{it} + \alpha_{it} \epsilon_{it}) + \psi_{it-1}^H \zeta_t^H$$

$$v_{it}^{inc} = \nu_{it} + \Delta \epsilon_{it}$$

I frequently use the following formula from Goodman (1960), that for independent variables:

$$\text{Var}(xy) = E^2(x)\text{Var}(y) + E^2(y)\text{Var}(x) + \text{Var}(x)\text{Var}(y)$$

I also frequently use the formula in Bohrnstedt and Goldberger (1969), that for any three variables:

$$\text{Cov}(xy, v) = E(x)\text{Cov}(y, v) + E(y)\text{Cov}(x, v) + E(\tilde{x}\tilde{y}\tilde{v})$$

Where $\tilde{z} = z - Ez$. Specifically, if one of v, x or y has zero mean and is independent of the other two variables, then $\text{Cov}(xy, v) = 0$.⁴⁹

Finally, I also assume that the predictable change to consumption Γ_{it}^C is common across households within a cohort. See BLP for a further discussion.

B.1 Deriving an Expression for $\Delta\text{Var}(c_t)$

I drop i subscripts to avoid clutter. By simple re-arrangement:

$$\begin{aligned}\Delta\text{Var}(c_t) &= \text{Var}(c_{t-1} + \Delta c_t) - \text{Var}(c_{t-1}) \\ &= \text{Var}(\Delta c_t) + 2\text{Cov}(c_{t-1}, \Delta c_t)\end{aligned}$$

Then

$$\begin{aligned}\text{Var}(\Delta c_t) &\approx \text{Var}(\pi_t(\nu_t + \alpha_t \epsilon_t) + \psi_{t-1}^H \zeta_t^H) \\ &= (\bar{\pi}_t^2 + \text{Var}(\pi_t))(\text{Var}(\nu_t) + \alpha_t^2 \text{Var}(\epsilon_t)) + (\zeta_t^H)^2 \text{Var}(\psi_{t-1}^H)\end{aligned}\tag{31}$$

where the last line follows from the first by application of Goodman's formula.

$$\begin{aligned}2\text{Cov}(c_{t-1}, \Delta c_t) &\approx 2\text{Cov}(c_{t-1}, \pi_t(\nu_t + \alpha_t \epsilon_t) + \psi_{t-1}^H \zeta_t^H) \\ &= 2\text{Cov}(c_{t-1}, \pi_t(\nu_t + \alpha_t \epsilon_t)) + 2\text{Cov}(c_{t-1}, \psi_{t-1}^H \zeta_t^H) \\ &= 2\zeta_t^H \text{Cov}(c_{t-1}, \psi_{t-1}^H)\end{aligned}\tag{32}$$

by application of Bohrnstedt's formula. Putting the terms together we have:

$$\Delta\text{Var}(c_t) \approx (\bar{\pi}_t^2 + \text{Var}(\pi_t))(\text{Var}(\nu_t) + \alpha_t^2 \text{Var}(\epsilon_t)) + (\zeta_t^H)^2 \text{Var}(\psi_{t-1}^H) + 2\zeta_t^H \text{Cov}(\psi_{t-1}^H, c_{t-1})\tag{33}$$

as in expression 16 in the main text.

B.2 Deriving an Expression for $\Delta\text{Cov}(c_t, y_t)$

We have:

$$\begin{aligned}\Delta\text{Cov}(c_t, y_t) &= \text{Cov}(c_{t-1} + \Delta c_t, y_{t-1} + \Delta y_t) - \text{Cov}(c_{t-1}, y_{t-1}) \\ &= \text{Cov}(c_{t-1}, \Delta y_t) + \text{Cov}(\Delta c_t, y_{t-1}) + \text{Cov}(\Delta c_t, \Delta y_t)\end{aligned}$$

⁴⁹Suppose, for example, x is independent of y and v and has zero mean. Then $E(x) = \text{Cov}(x, v) = 0$ and $E(\tilde{x}\tilde{y}\tilde{v}) = E(x\tilde{y}\tilde{v}) = E(E(x\tilde{y}\tilde{v}|\tilde{y}\tilde{v})) = E(\tilde{y}\tilde{v}E(x|\tilde{y}\tilde{v})) = 0$. The argument follows similarly if either y or v is independent of the other two variables and has zero mean.

Looking at each term in sequence:

$$\begin{aligned}
\text{Cov}(c_{t-1}, \Delta y_t) &= \text{Cov}(c_{t-1}, \nu_t + \Delta \epsilon_t) \\
&= \text{Cov}(c_{t-1}, -\epsilon_{t-1}) \\
&\approx -\bar{\pi}_{t-1} \alpha_{t-1} \text{Var}(\epsilon_{t-1})
\end{aligned}$$

And for the 2nd term:

$$\begin{aligned}
\text{Cov}(\Delta c_t, y_{t-1}) &\approx \text{Cov}(\pi_t(\nu_t + \alpha_t \epsilon_t) + \psi_{t-1}^H \zeta_t^H, y_{t-1}) \\
&= \text{Cov}(\pi_t(\nu_t + \alpha_t \epsilon_t), y_{t-1}) + \zeta_t^H \text{Cov}(\psi_{t-1}^H, y_{t-1}) \\
&= 0 + \zeta_t^H \text{Cov}(\psi_{t-1}^H, y_{t-1})
\end{aligned} \tag{34}$$

Finally:

$$\begin{aligned}
\text{Cov}(\Delta c_t, \Delta y_t) &\approx \text{Cov}(\pi_t(\nu_t + \alpha_t \epsilon_t) + \psi_{t-1}^H \zeta_t^H, \nu_t + \Delta \epsilon_t) \\
&= \text{Cov}(\pi_t \nu_t, \nu_t) + \text{Cov}(\pi_t \alpha_t \epsilon_t, \Delta \epsilon_t) + \text{Cov}(\psi_{t-1}^H \zeta_t^H, \Delta \epsilon_t) \\
&= \bar{\pi}_t \text{Var}(\nu_t) + \bar{\pi}_t \alpha_t \text{Var}(\epsilon_t) - \zeta_t^H \text{Cov}(\psi_{t-1}^H, \epsilon_{t-1})
\end{aligned} \tag{35}$$

Putting this together we get:

$$\begin{aligned}
\Delta \text{Cov}(c_t, y_t) &\approx \bar{\pi}_t \text{Var}(\nu_t) + \Delta [\bar{\pi}_t \alpha_t \text{Var}(\epsilon_t)] + \zeta_t^H \text{Cov}(\psi_{t-1}^H, y_{t-1}) - \zeta_t^H \text{Cov}(\psi_{t-1}^H, \epsilon_{t-1}) \\
&= \bar{\pi}_t \text{Var}(\nu_t) + \Delta [\bar{\pi}_t \alpha_t \text{Var}(\epsilon_t)] + \zeta_t^H \text{Cov}(\psi_{t-1}^H, y_{t-1}^P)
\end{aligned} \tag{36}$$

where $y_t^P = y_t - \epsilon_t$ is log permanent income, as in expression 15 in the main text.

C Appendix: Aggregate Income Shocks and an Equilibrium Interpretation

As discussed in the introduction, much of the recent literature has focused on the general equilibrium relationship between house prices and consumption. One particular focus of the literature has been on equilibrium effects of changes to credit conditions. Another possibility is that house prices and consumption are correlated because they are both driven by (revisions to) aggregate income expectations. Consequently, many studies, such as [Campbell and Cocco \(2007\)](#), allow for a connection between house prices and aggregate income shocks in their analyses. These aggregate income shocks are the focus of the current appendix. I first assess how aggregate shocks affect the inequality measures jointly with house price shocks, without considering equilibrium factors. Given some assumptions on the supply side of the economy I then provide an equilibrium interpretation.

C.1 The Covariance Structure of Income and Consumption: Extension to Aggregate Income Shocks

I follow on from sections [A](#) and [B](#) by extending the model to allow for aggregate shocks to income. I first consider these in the absence of general equilibrium effects. Specifically, I model aggregate income as a random walk. I also model interest rates as fixed. I therefore modify the process for household income given by equations [5](#) and [6](#) to become

$$\begin{aligned}\ln Y_{it} &= X_{it}\varphi_t + \ln P_{it} + \epsilon_{it} \\ \ln P_{it} &= \ln P_{it-1} + \nu_{it} + \zeta_t^Y\end{aligned}$$

where ζ_t^Y is an aggregate, proportional, permanent shock that affects all households equally, and all other terms are as before. Note that ζ_t^Y is potentially correlated with ζ_t^H across time.

The introduction of aggregate shocks does not change the basic formulae for consumption innovations. We therefore now have:

$$\begin{aligned}v_{it}^C &\approx \pi_{it}(\zeta_t^Y + \nu_{it} + \alpha_t\epsilon_{it}) + \psi_{it-1}^H\zeta_t^H \\ v_{it}^{inc} &= \zeta_t^Y + \nu_{it} + \Delta\epsilon_{it}\end{aligned}$$

where ζ_t^Y is the aggregate shock to income. The growth in consumption given in expression [13](#) is updated to become

$$\Delta c_{it} \approx \Gamma_{it}^C + \pi_{it}(\zeta_t^Y + \nu_{it} + \alpha_{it}\epsilon_{it}) + \psi_{it-1}^H\zeta_t^H$$

Note that the household observes no difference between an idiosyncratic permanent shock ν_{it} and the aggregate shock ζ_t^Y in my (partial) equilibrium framework. Both shocks affect consumption equally.

The updated version of expression [31](#) is now

$$\begin{aligned}\text{Var}(\Delta c_t) &\approx \text{Var}(\pi_t(\zeta_t^Y + \nu_t + \alpha_t\epsilon_t) + \psi_{t-1}^H\zeta_t^H) \\ &= (\bar{\pi}_t^2 + \text{Var}(\pi_t)) (\text{Var}(\nu_t) + \alpha_t^2\text{Var}(\epsilon_t)) + (\zeta_t^H)^2 \text{Var}(\psi_{t-1}^H) + (\zeta_t^Y)^2 \text{Var}(\pi_t) \\ &\quad + 2\text{Cov}(\pi_t(\nu_t + \alpha_t\epsilon_t) + \psi_{t-1}^H\zeta_t^H, \zeta_t^Y\pi_t) \\ &= (\bar{\pi}_t^2 + \text{Var}(\pi_t)) (\text{Var}(\nu_t) + \alpha_t^2\text{Var}(\epsilon_t)) + (\zeta_t^H)^2 \text{Var}(\psi_{t-1}^H) + \underbrace{(\zeta_t^Y)^2 \text{Var}(\pi_t) + 2\zeta_t^H\zeta_t^Y \text{Cov}(\psi_{t-1}^H, \pi_t)}_{\text{new terms}}\end{aligned}$$

Similarly, the updated version of expression [32](#) is

$$\begin{aligned}2\text{Cov}(c_{t-1}, \Delta c_t) &\approx 2\text{Cov}(c_{t-1}, \pi_t(\zeta_t^Y + \nu_t + \alpha_t\epsilon_t) + \psi_{t-1}^H\zeta_t^H) \\ &= 2\text{Cov}(c_{t-1}, \pi_t(\nu_t + \alpha_t\epsilon_t)) + 2\text{Cov}(c_{t-1}, \pi_t\zeta_t^Y) + 2\text{Cov}(c_{t-1}, \psi_{t-1}^H\zeta_t^H) \\ &= \underbrace{2\zeta_t^Y \text{Cov}(c_{t-1}, \pi_t)}_{\text{new term}} + 2\zeta_t^H \text{Cov}(c_{t-1}, \psi_{t-1}^H)\end{aligned}\tag{37}$$

Therefore the updated version of expression 33 is

$$\begin{aligned} \Delta \text{Var}(c_t) \approx & (\bar{\pi}_t^2 + \text{Var}(\pi_t)) (\sigma_\nu^2 + \alpha_t^2 \sigma_\epsilon^2) + (\zeta_t^H)^2 \text{Var}(\psi_{t-1}^H) + 2\zeta_t^H \text{Cov}(\psi_{t-1}^H, c_{t-1}) \\ & + \underbrace{(\zeta_t^Y)^2 \text{Var}(\pi_t) + 2\zeta_t^H \zeta_t^Y \text{Cov}(\psi_{t-1}^H, \pi_t) + 2\zeta_t^Y \text{Cov}(c_{t-1}, \pi_t)}_{\text{new terms}} \end{aligned} \quad (38)$$

where I use σ_η^2 and σ_ϵ^2 to indicate the variances of permanent and transitory innovations, to save space.

The terms for $\Delta \text{Cov}(c_t, y_t)$ change less. In fact the only changes are that the term in $\text{Cov}(\Delta c_t, y_{t-1})$ in expression 34 becomes larger by the term $\zeta_t^Y \text{Cov}(\pi_t, y_{t-1})$, and the updated version of expression 35 is smaller by the term $\zeta_t^Y \text{Cov}(\pi_t, \epsilon_{t-1})$. The sum of these two changes is $\zeta_t^Y \text{Cov}(\pi_t, y_{t-1}^P)$. Therefore the updated version of expression 36 is

$$\Delta \text{Cov}(c_t, y_t) \approx \bar{\pi}_t \text{Var}(\nu_t) + \Delta [\bar{\pi}_t \alpha_t \text{Var}(\epsilon_t)] + \zeta_t^H \text{Cov}(\psi_{t-1}^H, y_{t-1}^P) + \underbrace{\zeta_t^Y \text{Cov}(\pi_t, y_{t-1}^P)}_{\text{new term}}$$

The cross-sectional moments therefore contain four new terms. First, the growth in covariance between (log) consumption and (log) income features the term $\zeta_t^Y \text{Cov}(\pi_t, y_{t-1}^P)$. This term captures the importance of labour income in life-time wealth across the income distribution. As one example, if all households have the same asset-to-income ratio, then an aggregate income shock will affect all households equally and this term will be zero. More generally, this term may be positive or negative. It may be negative if those at the top of the income distribution have disproportionately more assets and those at the bottom of the income distribution rely more on their income. Then, a shock to aggregate income would compress the covariance.

Similarly, the growth in the variance of log consumption features the term $2\zeta_t^Y \text{Cov}(\pi_t, c_{t-1})$. This term captures any difference in the importance of labour income in life-time wealth across the consumption distribution. Intuitively, and similarly to the case for the covariance with income, suppose, for example, that those at the bottom of the consumption distribution finance consumption through labour income and have few assets, while those at the top of the distribution are asset holders. Then an aggregate income shock will cause consumption inequality to fall.

Finally, the growth in the variance of log consumption also features the terms $(\zeta_t^Y)^2 \text{Var}(\pi_t)$ and $2\zeta_t^H \zeta_t^Y \text{Cov}(\psi_{t-1}^H, \pi_t)$. These terms are best interpreted alongside $(\zeta_t^H)^2 \text{Var}(\psi_{t-1}^H)$, also given in expression 38. They come from an expansion of $\text{Var}(\zeta_t^H \psi_{t-1}^H + \zeta_t^Y \pi_{it})$. Intuitively, these terms capture any raw difference in the importance of labour income and housing wealth across households, irrespectively of how they are correlated with the consumption distribution. Since these terms are composed of second-order products of ζ_t^Y and ζ_t^H , we would expect them to be less important quantitatively.

Before returning to the equilibrium interpretation, notice that these formulae do not contain any terms in $\text{Cov}(\zeta_t^Y, \zeta_t^H)$: the relationship between income and house price shocks does not directly enter the formulae. This is because all variances and covariances in these formulae are cross-sectional, and not taken over time. Of course, in any given period the ex-post shocks ζ_t^Y and ζ_t^H are constants. Nevertheless, any covariance over time

between these shocks affects the empirical analysis in the following ways: First, the term $\zeta_t^H \zeta_t^Y$ which appears in equation 38 is more likely to be positive, ex-post, although its importance is likely small. Second, and more importantly, any correlation should affect households' initial precautionary saving and housing choices. These effects will be picked up in the wealth shares observed in the data.

C.2 An Equilibrium Interpretation Under Fixed Housing Supply

We can now consider an equilibrium interpretation of aggregate income and price movements, using some simple approximating assumptions. First, suppose that housing is in fixed supply. Second, suppose the interest rate is fixed because we are analysing a small, open economy. The first assumption seems highly appropriate for the UK, where the rate of house building has remained extremely low. The second assumption is also reasonably appropriate, although lending conditions presumably deviate away from movements in the risk free rate over the housing market cycle. Finally, suppose that preferences are Cobb-Douglas. Under these assumptions, an aggregate income shock, arising from a shock to labour productivity, say, induces an approximately identical proportionate change to consumption and to house prices.

In this case we can interpret the direct and indirect effect of income shocks on the inequality measures. Ignoring second-order terms, the direct effect of an income shock on consumption inequality (i.e. excluding house price effects) is $2\text{Cov}(\pi_t, c_{t-1})$. The indirect effect driven through house price changes is $2\text{Cov}(\psi_{t-1}^H, c_{t-1})$, yielding a total effect of $2\text{Cov}(\pi_t, c_{t-1}) + 2\text{Cov}(\psi_{t-1}^H, c_{t-1}) = 2\text{Cov}(\psi_{t-1}^O, c_{t-1})$, where ψ_{t-1}^O is the share of other (non-housing) wealth in discounted life-time wealth. Similarly, the direct effect of the aggregate income shock on the covariance measure is $\text{Cov}(\pi_t, y_{t-1}^P)$ and the total effect is $\text{Cov}(\psi_{t-1}^O, y_{t-1}^P)$. I quantify these effects and discuss them further in the next subsection. In brief, ignoring second-order terms the effect of aggregate income shocks on the variance of log consumption is as follows:

$$\Delta\text{Var}(c_t) \approx (\bar{\pi}_t^2 + \text{Var}(\pi_t)) (\text{Var}(\nu_t) + \dots) + \underbrace{2\zeta_t^Y \text{Cov}(\pi_t, c_{t-1})}_{\text{Direct effect of aggregate income shock}} + \underbrace{2\zeta_t^Y \text{Cov}(\psi_{t-1}^H, c_{t-1})}_{\text{Indirect effect of income shock through house prices}} + \underbrace{(\zeta_t^Y)^2 \dots}_{\text{Second-order terms}} \quad (39)$$

Before finishing this subsection it is worth qualifying this equilibrium interpretation. In particular, I have assumed that the interest rate is fixed. Of course, as discussed, a number of papers have modelled the relationship between credit conditions and the housing market, but assessing this channel is beyond the scope of the current paper. More generally, it is clear from the volatility of prices that factors other than aggregate productivity must have a more direct role in driving house prices. The puzzling high volatility of house prices is a focus of the survey chapter by [Piazzesi and Schneider \(2016\)](#). Here, I therefore focus on potentially only a small component of the factors driving house prices. Nevertheless my exercise is potentially informative in assessing whether housing-wealth factors exacerbate or attenuate inequality movements arising from prior forces.

C.3 Results

The previous subsections showed how to adjust the inequality formulae when allowing for shocks to aggregate income and coincident, or resulting, shocks to house prices. I now modify the empirical analysis to assess the effect of these shocks on quantities. The reasons for doing this are two-fold. First, I wish to assess the direction and the magnitude of the bias from ignoring aggregate shocks. Second, and more importantly, I wish to assess whether, and to what extent, housing wealth effects exacerbate or attenuate other forces driving inequality.

To perform the analysis I measure aggregate income as real households' disposable income per head from the national accounts.⁵⁰ According to this series, household incomes had yearly trend growth of 2.6% over 1955 to 2007. Moreover, household incomes did indeed grow above trend coincidentally with the house price boom. Compared to house price changes, however, income shocks were small in proportional terms. Income grew by only a cumulative 3.6% above trend over 1997-2004. In fact, growth in disposable incomes peaked over 1995 to 2001, but even then, it cumulated to only 6.8% above trend. This is much less than the 58% above-trend growth in house prices.

I now quantify the relevant wealth-share moments given in formulae 37 and 38, and in formula 39. The empirical counterparts of these moments are shown in table A1. I turn first to expression 39, which decomposes growth in the variance of log consumption, interpreted in terms of general equilibrium effects. The most important new wealth-share moment for this inequality measure is $\text{Cov}(\pi_t, c_{t-1})$. Table A1 gives its empirical counterpart in the second column. The table shows that these quantities are negative and quantitatively large. A positive aggregate income shock therefore reduces consumption inequality, because those at the bottom of the consumption distribution are more reliant on labour income than on other assets, relative to those who are better off. I quantify the effect on the inequality measure, by performing the same calculation as given in table 2. I find that, for the 1950s cohort, allowing for aggregate income shocks reduces the explained increase in the variance of log consumption by 2.3% of the observed increase over 1997-2004, a small amount. Similarly, the corresponding calculation for the 1960s cohort gives 0.3% of the observed increase. I conclude, therefore, that allowing for aggregate income shocks does not substantially alter the main results.

Although the direct effect of the aggregate income shocks is small, we can nevertheless use the quantities in provided in table 2 to quantify the equilibrium interpretation. As discussed above, when aggregate income increases, houses are in fixed supply and preferences are Cobb-Douglas, we should expect an equal proportionate rise in the relative price of housing. Therefore if the direct effect of an aggregate income shock is to decrease consumption inequality, because $2\text{Cov}(\pi_t, c_{t-1})$ is negative, then house price effects more than offset this factor because the relevant wealth share $2\text{Cov}(\psi_{t-1}^H, c_{t-1})$ is positive and, on average, quantitatively slightly larger. Of course, this analysis does not take into account the forces that likely have more importance than aggregate income in driving house prices and, hence, inequality measures.

To complete the analysis, I turn next to growth in the covariance of consumption and income. I have not shown its role in an equilibrium interpretation, but the effect of coincident shocks is shown in expression 37.

⁵⁰This is variable INXZ in the income accounts.

The relevant wealth-share moment in this decomposition is $\text{Cov}(\pi_t, y_{t-1}^P)$, the covariance of the human capital share with permanent income. Again, I assess this using $\text{Cov}(\pi_t, y_{t-1})$ shown in the table in the third column. The estimates are insignificant and small, implying that, for a given aggregate income shock, the effect on the covariance of consumption with income is also small. For completeness, the first and final columns of table A1 give other components relevant to expression 38. As discussed in section C, however, their role in the inequality measures is negligible.

Table A1: ADDITIONAL WEALTH SHARE ESTIMATES BY COHORT

year	Var π	$\text{Cov}(\pi_{t+1}, c_t)$	$\text{Cov}(\pi_{t+1}, y_t)$	$\text{Cov}(\psi_t^H, \pi_{t+1})$	$\text{Cov}(\psi_t^H, c_t)$
1950s Cohort					
1995	0.0103 (0.0007)	-0.0090 (0.0049)	-0.0061 (0.0027)	-0.0063 (0.0006)	0.0065 (0.0061)
2000	0.0152 (0.0012)	-0.0191 (0.0062)	-0.0042 (0.0031)	-0.0117 (0.0010)	0.0156 (0.0067)
2005	0.0348 (0.0020)	-0.0154 (0.0106)	-0.0021 (0.0048)	-0.0277 (0.0020)	0.0271 (0.0103)
1960s Cohort					
1995	0.0046 (0.0005)	-0.0075 (0.0027)	0.0005 (0.0016)	-0.0040 (0.0004)	0.0116 (0.0036)
2000	0.0060 (0.0006)	-0.0105 (0.0028)	-0.0038 (0.0016)	-0.0056 (0.0005)	0.0106 (0.0041)
2005	0.0160 (0.0010)	-0.0133 (0.0056)	-0.0020 (0.0028)	-0.0143 (0.0010)	0.0141 (0.0074)

Notes: ψ^H is the share of housing in lifetime wealth, π the share of human capital. y_t is log residual current income, f_t is log residual food. $\text{Cov}(x_t, c_t)$ is computed as $\text{Cov}(x_t, f_t)/\chi_f$ where χ_f is the elasticity of food consumption with respect to total consumption, taken to be 0.38. See text for more details. Standard errors obtained using a bootstrap with 1000 draws.

D Model Solution, Calibration and Numerical Details

The model used here is very similar to that in Berger et al. (2018). The model is solved using the method of endogenous grid points extended to discrete choices as discussed in Fella (2014), Iskhakov et al. (2017) and Clausen and Strub (2013). This method uses the exact solution of the Euler equation and so is numerically extremely accurate. I first describe computation of the benchmark model described in section 5. Below I describe computation of the extended models, such as that featuring constant rental prices.

Choices of housing lie on a discrete grid, which is sufficiently fine that I interpret it as continuous. Specifically, the grid for housing has 40 points. These points are spaced at approximately 5% intervals in the middle of the state space, and more widely at the extremes. The grid for financial wealth has 100 points. The income grid has 11 points. I first solve a conditional value function problem for the household: the household maximizes

consumption given a choice of housing. I then solve for the household's *unconditional* value function: the household maximizes over housing values taking optimal consumption as given.

D.1 Conditional Value Functions

D.1.1 Buyers/Movers

In the following exposition I drop t subscripts and use a more compact notation. Therefore X_{+1} , for example, denotes the value of X in the subsequent period, and X_{-1} the value in the previous period. Let $\bar{H} = p^H H$ be housing value. Define $U(C, \bar{H}; p^H) \equiv \frac{1}{(p^H)^{(1-a)(1-\gamma)}} u(C, \bar{H}) = u(C, H)$. I write this new utility function as $U(C, \bar{H})$ for compactness. Let $x_{+1} \equiv \frac{p_{+1}^H}{p^H} = e^{\mu + \zeta_{+1}^H}$ be the shock to house prices. Let G be the inherited house value at the beginning of the period, such that $G = x \times \bar{H}_{-1}$. Let $Y(z, \epsilon)$ be income, which depends on permanent characteristics, here denoted by z , and the transitory shock ϵ . Finally let $\mathcal{V}^{OWN}()$ be the unconditional value function for owners, discussed later.

It can be shown that the conditional value function for those buying or moving a house takes the following form:

$$\begin{aligned} V^{buy, \bar{H}}(X) &= \max_{C, A} U(C, \bar{H}) + \beta \mathbb{E} \left[x_{+1}^{-(1-\gamma)(1-a)} \mathcal{V}^{OWN}(X_{+1}) \right] \\ \text{s.t.} \quad & A + C + \bar{H} = RA_{-1} + Y(z, \epsilon) + (1 - \kappa)(1 - \delta)G \\ & A \geq (1 - \Xi) \frac{1 - \delta}{R} \bar{H}, \quad G_{+1} = \bar{H}x_{+1}, \quad X_{+1} = \langle A, G_{+1}, Y_{+1}, j + 1 \rangle \end{aligned}$$

Define $W = A + (1 - \Xi) \frac{1 - \delta}{R} \bar{H}$ as end-of-period voluntary equity: a choice variable. We can also define $Q = RA_{-1} + (1 - \Xi)(1 - \delta)G + Y(z, \epsilon)$ as beginning-of-period voluntary equity (including income), a state variable. This implies we can write the budget constraint as follows:

$$\begin{aligned} W + C &= Q + \left((1 - \Xi) \frac{1 - \delta}{R} - 1 \right) \bar{H} + (\Xi - \kappa)(1 - \delta)G \\ &= Z(Q, G, \bar{H}) \end{aligned}$$

where $Z()$ gives total resources available for consumption and saving, given endowments and choices over housing.

The problem can therefore be written as follows:

$$\begin{aligned} V^{buy, \bar{H}}(X) &= \max_{C, W} U(C, \bar{H}) + \beta \mathbb{E} \left[x_{+1}^{-(1-\gamma)(1-a)} \mathcal{V}^{OWN}(X_{+1}) \right] \\ \text{s.t.} \quad & W + C = Z, \quad W \geq 0 \\ & Z = Q + \left((1 - \Xi) \frac{1 - \delta}{R} - 1 \right) \bar{H} + (\Xi - \kappa)(1 - \delta)G \\ & G_{+1} = \bar{H}x_{+1}, \quad Q_{+1} = RW + (1 - \Xi)(1 - \delta)\bar{H}(x_{+1} - 1) + Y(z_{+1}, \epsilon_{+1}) \\ & X_{+1} = \langle Q_{+1}, G_{+1}, Y_{+1}, j + 1 \rangle \end{aligned}$$

The algorithm works as in [Fella \(2014\)](#). First we take a grid \hat{W} for W . Then, using endogenous grid points, and for each point on the productivity and housing grids, we can obtain solutions $\hat{C}(\hat{W}, \bar{H})$ and $\hat{Z}(\hat{W}, \bar{H}) = \hat{W} + \hat{C}(\hat{W}, \bar{H})$. These are the policy functions *conditional* on housing choice. This gives a correspondence $\hat{C}(Z(\hat{W}), \bar{H})$ which gives values for consumption as a function of housing choice and an endogenously determined grid of available resources Z .

As a final point note that a difference with the model in [Berger et al. \(2018\)](#) arises in the specification of the income process, which here has a permanent-transitory, rather than an AR(1), structure. This alteration leads to a slight difference in the definition of the state variables: here beginning-of-period voluntary equity includes contemporaneous income.

D.1.2 Renters

For renting, define $W = A$ as end-of-period voluntary equity: a choice variable. Similarly define $Q_{+1} = RW + Y(z_{+1}, \epsilon_{+1}) = RA + Y(z_{+1}, \epsilon_{+1})$ as beginning-of-period voluntary equity, a state variable. Then the problem can be written as follows:

$$\begin{aligned}
V^{rent}(X) &= \max_{C, W, \bar{H}} U(C, \bar{H}) + \beta \mathbb{E} \left[x_{+1}^{-(1-\gamma)(1-a)} \mathcal{V}^{RENT}(X_{+1}) \right] \\
\text{s.t} \quad & W + C + \chi \bar{H} = Z, \quad W \geq 0 \\
& Z = Q + (\Xi - \kappa)(1 - \delta)G \\
& G_{+1} = 0, \quad Q_{+1} = RW + Y(z_{+1}, \epsilon_{+1}) \\
& X_{+1} = \langle Q_{+1}, G_{+1}, Y_{+1}, j + 1 \rangle
\end{aligned}$$

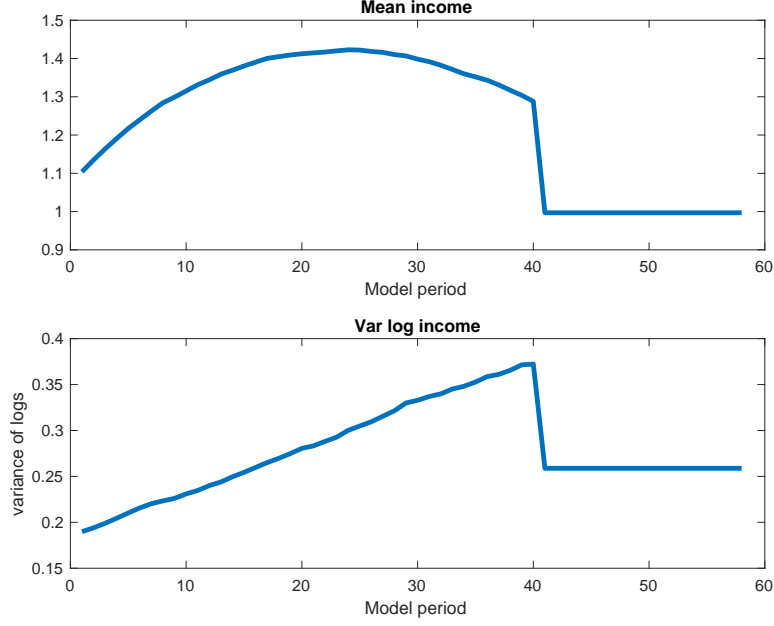
D.1.3 Stayers

Using the same definitions as for buyers we can write the stayer's problem as follows:

$$\begin{aligned}
V^{stay}(X) &= \max_{C, W} U(C, G) + \beta \mathbb{E} \left[x_{+1}^{-(1-\gamma)(1-a)} \mathcal{V}^{OWN}(X_{+1}) \right] \\
\text{s.t} \quad & W + C = Z, \quad W \geq 0 \\
& Z = Q + \frac{G}{R} ((1 - R)(1 - \Xi)(1 - \delta) - R\delta) \\
& G_{+1} = Gx_{+1}, \quad Q_{+1} = RW + Y(z_{+1}, \epsilon_{+1}) + (1 - \Xi)(1 - \delta)G(x_{+1} - 1) \\
& X_{+1} = \langle Q_{+1}, G_{+1}, Y_{+1}, j + 1 \rangle
\end{aligned}$$

When solving the model we do not need to solve explicitly for staying, but can use the solution for buyers. We

Figure A1: Life-cycle Income Profiles



can do this by noticing that

$$Z^{stay} = Z^{buy} (\bar{H} = G) + G\kappa(1 - \delta)$$

D.2 Unconditional Value Functions

We can then define the value functions for in-coming renters and owners as follows:

$$\mathcal{V}^{RENT}(X) = \max \left\{ V^{rent}(X), \max_{\bar{H}} V^{buy, \bar{H}}(X) \right\}$$

$$\mathcal{V}^{OWN}(X) = \max \left\{ V^{stay}(X), V^{rent}(X), \max_{\bar{H}} V^{buy, \bar{H}}(X) \right\}$$

D.3 Extra Calibration Details

Figure A1 shows the mean of income and the variance of log income over the life cycle.

D.4 Solution of the Extended Models: Constant Rental Price and Non-Unitary Elasticity of Substitution

The main change when solving these models is that it becomes impossible to remove the house price as a state variable. Computation is therefore much slower, although the computational procedures are conceptually

simpler. For example the conditional value function for a household buying a house of quality H is:

$$\begin{aligned}
V^{buy,H}(X) &= \max_{C,W} u(C, H) + \beta \mathbb{E} [\mathcal{V}^{OWN}(X_{+1})] \\
\text{s.t.} \quad & W + C = Z, \quad W \geq 0 \\
& Z = Q + \left((1 - \Xi) \frac{1 - \delta}{R} - 1 \right) p^H H + (\Xi - \kappa) (1 - \delta) p^H H_{-1} \\
& d_{-1} \in \{0, 1\} \\
& Q_{+1} = RW + (1 - \Xi) (1 - \delta) H (p_{+1}^H - p^H) + Y(z_{+1}, \epsilon_{+1}) \\
& X_{+1} = \langle Q_{+1}, H, Y_{+1}, p_{+1}^H, j + 1 \rangle
\end{aligned}$$

where I borrow most notation from the benchmark model. I also define the variable $d = 0$ if the individual rents at time t , and $d = 1$ if he buys.

For those choosing to rent, when the rental price is constant (or deterministic) the problem can be written as follows:

$$\begin{aligned}
V^{rent}(X) &= \max_{C,W,H} u(C, H) + \beta \mathbb{E} [\mathcal{V}^{RENT}(X_{+1})] \\
\text{s.t.} \quad & W + C + \chi H = Z, \quad W \geq 0 \\
& Z = Q + (\Xi - \kappa) (1 - \delta) d_{-1} p^H H_{-1} \\
& d_{-1} \in \{0, 1\} \\
& Q_{+1} = RW + Y(z_{+1}, \epsilon_{+1}) \\
& X_{+1} = \langle Q_{+1}, Y_{+1}, p_{+1}^H, j + 1 \rangle
\end{aligned}$$

In terms of numerical choices, I now use sparser grids. The grid for housing quality has 30 points. These points are now spaced at approximately 6.5% intervals in the middle of the state space, and more widely at the extremes. The grid for financial wealth still has 100 points. The income grid now has 9 points. The grid for house prices has 21 points. This grid is such that the house price elasticities are computed using a price change of between 2.5% and 5.5%, depending on age.

Table A2 shows the parameters used for calibrating the extended model version with price as a state variable. The first line ('Benchmark Re-Calibration') shows the parameters chosen to hit the lifecycle in the benchmark model, but with house price growth set to 0. These values are also used in the model extensions where the elasticity of substitution between housing and consumption is 0.8 and when it is 1.2. The second line ('Constant Rental Price') shows the parameters chosen to hit the lifecycle in the model with constant rental prices. See Section 5 for a further discussion.

Table A2: HOUSING MODEL EXTENSIONS: PARAMETERS USED TO HIT LIFECYCLE

	a	β	Φ	χ_{25}	χ_{65}	χ_{ret}
Benchmark Re-Calibration	0.825	0.953	17.0	0.0505	0.0475	0.048
Constant Rental Price	0.825	0.953	17.0	0.0535	0.053	0.051

Notes: ‘Benchmark Re-Calibration’ refers to the re-calibration using the model with house price as a state variable and zero house price growth. This calibration is also used when the CES parameter is chosen to be 1.2 and when it is chosen to be 0.8. ‘Constant Rental Price’ refers to the re-calibration using the model with constant rental prices, and a stochastic price-to-rent ratio. See Section 5 for further discussion.

E Appendix: The Minimum Distance Estimator for the Income Process

In this section I describe the asymptotic properties of the estimator used in section 6. Some of the notation here is recycled from other sections. In particular, k , V , A and Ξ are defined differently than elsewhere.

E.1 Model and Data Structure

Let m be a $P \times 1$ vector of moments from two different data sources (in our application, the FES and BHPS), with sample sizes N and M . Let κ be a $R \times 1$ vector of parameters. Our model specifies:

$$g(m_0, \kappa_0) = 0$$

where $g(\cdot)$ is a potentially non-separable function of raw moments and parameters and forms a $S \times 1$ vector of moment conditions. In our case $S = 3T$ where T is the number of time periods and $g(\cdot)$ consists of all the empirical counterparts of equations 14, 15 and 16. We seek to estimate κ . In our case, κ consists of $\text{Var}(\nu_t)$, the variance of permanent shocks and $\Delta \text{Var}(\epsilon_t)$, the change in the variance of transitory shocks. Furthermore, $R = 2T$. More generally, identification requires that $R \leq S$.

E.2 Estimation

This situation is as in classical minimum distance estimation except that the raw moments cannot be separated from the parameters. To derive the estimator and asymptotic statistics I therefore modify the arguments used in, for example, Wooldridge (2002).

As mentioned, the data in this application come from two sources. I derive asymptotic distributions in the case where both samples have the same sample size, $N (= M)$, letting $N \rightarrow \infty$.

$\hat{\kappa}$ solves the first-order condition:

$$\mathcal{H}(\hat{\kappa})' \text{Ag}(\hat{m}, \hat{\kappa}) = 0$$

such that A is a weighting matrix and $\mathcal{H}(\hat{\kappa}) \equiv \nabla_{\kappa}g(\hat{m}, \kappa)|_{\hat{\kappa}}$ is the gradient of the objective function with respect to estimable parameters.

By a mean-value expansion:

$$\begin{aligned}\sqrt{N}g(\hat{m}, \hat{\kappa}) &= \sqrt{N}g(m_0, \kappa_0) + G\sqrt{N}(\hat{m} - m_0) + \mathcal{H}\sqrt{N}(\hat{\kappa} - \kappa_0) + o_p(1) \\ &= G\sqrt{N}(\hat{m} - m_0) + \mathcal{H}\sqrt{N}(\hat{\kappa} - \kappa_0) + o_p(1)\end{aligned}$$

where $\mathcal{H} \equiv \nabla_{\kappa}g(m_0, \kappa)|_{\kappa_0}$ and $G \equiv \nabla_m g(m, \kappa_0)|_{m_0}$.

Combining the expressions above, and using $\mathcal{H}(\hat{\kappa})' = \mathcal{H}' + o_p(1)$:

$$\sqrt{N}(\hat{\kappa} - \kappa_0) \approx -(\mathcal{H}'A\mathcal{H})^{-1}\mathcal{H}'AG\sqrt{N}(\hat{m} - m_0)$$

If $\sqrt{N}(\hat{m} - m_0) \sim N(0, V)$, then:

$$\sqrt{N}(\hat{\kappa} - \kappa_0) \sim N\left[0, (\mathcal{H}'A\mathcal{H})^{-1}\mathcal{H}'A\Xi A'\mathcal{H}(\mathcal{H}'A\mathcal{H})^{-1}\right]$$

such that $\Xi \equiv GVG'$.

If used, the optimal weighting matrix is therefore $A = \Xi^{-1}$ giving:

$$\sqrt{N}(\hat{\kappa} - \kappa_0) \sim N\left[0, (\mathcal{H}'\Xi^{-1}\mathcal{H})^{-1}\right]$$

In this study, the identity weighting matrix is used. When using this optimal weighting matrix, however, the following multi-step procedure is available:

1. Estimate $\hat{\kappa}$ as the minimizer of $g(\hat{m}, \kappa)Ag(\hat{m}, \kappa)$, using e.g. $A = I_S$, the identity matrix.
2. Form $\hat{\Xi} = \hat{G}\hat{V}\hat{G}'$. \hat{V} is obtained from the raw data via a bootstrap. \hat{G} can be obtained from $\nabla_m g(m, \hat{\kappa})|_{\hat{m}}$.
3. Estimate $\hat{\hat{\kappa}}$ as the minimizer of $g(\hat{m}, k)(\hat{\Xi})^{-1}g(\hat{m}, \kappa)$.

F Appendix: Data and Extra Results

F.1 Computing Wealth Share Statistics In Years With Incomplete Wealth Information

Wealth shares are available using complete information on financial wealth for 1995, 2000 and 2005. Accordingly wealth share statistics are computed for these years simply as described in section 3.5. I compute wealth shares

for the intermediate years as follows. Let $\hat{w}_p(t)$ be a statistic for the precise wealth share for time t : for example, the covariance of consumption with the housing wealth share. Similarly let $\hat{w}_a(t)$ be the statistic for the approximate wealth share at time t , which includes mortgage debt, but does not include financial assets. Then for years between 1995 and 2000 (except for 1996) I interpolate the precise wealth share statistic as

$$\hat{w}_p(t) = \hat{w}_a(t) \times \left(\frac{(2000 - t)}{5} \times \frac{\hat{w}_p(1995)}{\hat{w}_a(1995)} + \frac{(t - 1995)}{5} \times \frac{\hat{w}_p(2000)}{\hat{w}_a(2000)} \right)$$

In words I re-weight the approximate wealth share estimate using weights observed in 1995 and 2000. These weights are linearly interpolated. I perform exactly the same procedure for years between 2000 and 2005.

I do not observe $\hat{w}_a(1996)$. Therefore I compute $\hat{w}_p(1996)$ as

$$\hat{w}_p(1996) = \left(\frac{(2000 - t)}{5} \times \hat{w}_p(1995) + \frac{(t - 1995)}{5} \times \hat{w}_p(2000) \right)$$

a standard interpolation.

I compute wealth shares for years before 1995 using data on housing assets and mortgage debt for 1993 and 1994. I first compute $\hat{w}_a(t)$ as the approximate statistic, not including financial assets. I then compute interpolated statistics as follows:

$$\hat{w}_p(t) = \hat{w}_a(t) \times \frac{\hat{w}_p(1995)}{\hat{w}_a(1995)}$$

I perform the same procedure for the years after 2005.

F.2 Results for Wealth Share Statistics Assuming Different Discount Factors on Future Income

Table A3 shows the results of a robustness exercise for the key moments used in section 4. The second and fifth columns reproduce the results shown in table 1. These columns use a net rate of 4% to discount future income. The other columns show the key wealth share moments using a low discount rate of 1% and a high discount rate of 7%.

F.3 Estimating Expected Discounted Lifetime Labour Income Using Different Data Structures

As discussed in section 2, estimating discounted lifetime labour income presents a number of difficulties. For example, accurate estimation requires that the income model presented in equations 5 and 6 is correctly specified. Here I discuss a related version of this challenge caused by the basic identification problem for income: we cannot distinguish permanent from transitory incomes at the household level, but only observe current income. This identification problem arises even in long panels. In this appendix I examine estimation under two different data structures: cross-sectional data, and a short, 2-wave panel. I show that, when estimating the crucial moments,

Table A3: ROBUSTNESS OF WEALTH SHARE ESTIMATES TO COMPUTATION OF DISCOUNTED LIFE-TIME INCOME

disc. rate year		Cov(ψ_t^H, c_t)			Cov(ψ_t^H, y_t)		
		1%	4%	7%	1%	4%	7%
1950s Cohort							
	1995	0.0048 (0.0049)	0.0065 (0.0061)	0.0084 (0.0074)	0.0059 (0.0022)	0.0079 (0.0027)	0.0101 (0.0033)
	2000	0.0120 (0.0056)	0.0156 (0.0067)	0.0184 (0.0078)	0.0052 (0.0028)	0.0079 (0.0034)	0.0097 (0.0039)
	2005	0.0216 (0.0092)	0.0271 (0.0103)	0.0308 (0.0112)	0.0076 (0.0042)	0.0092 (0.0046)	0.0111 (0.0050)
1960s Cohort							
	1995	0.0076 (0.0026)	0.0116 (0.0036)	0.0162 (0.0049)	0.0027 (0.0016)	0.0050 (0.0023)	0.0086 (0.0029)
	2000	0.0078 (0.0030)	0.0106 (0.0041)	0.0150 (0.0051)	0.0041 (0.0015)	0.0061 (0.0020)	0.0089 (0.0026)
	2005	0.0125 (0.0062)	0.0141 (0.0074)	0.0172 (0.0088)	0.0009 (0.0026)	0.0023 (0.0031)	0.0031 (0.0037)

Notes: Standard errors shown in parentheses and obtained using a bootstrap with 1000 draws. For further details see notes to table 1, and accompanying text.

the bias caused by observing current incomes only is second order in the case of the short panel, but somewhat larger when using a cross-section.

For this discussion, I focus on estimating $\text{Cov}(c_t, \psi_t^H)$. Estimating this moment requires estimating discounted lifetime wealth from time $t+1$ onwards. The formulation of discounted lifetime labour income used in appendix A is as follows:

$$L_{it+1} \equiv \sum_{j=0}^{T_w-t-1} \exp[\mathbb{E}_t \ln Y_{it+1+j} - jr]$$

where $r = \ln R \approx R - 1$ is the risk-free discount rate. Given the income model used this gives:

$$L_{it+1} \equiv P_{it} \sum_{j=0}^{T_w-t-1} \exp(X_{it+1+j} \varphi_{t+j} - jr)$$

where L_{it+1} is given in terms of P_{it} because P_{it} is the relevant component of the households time- t information set.

With cross-sectional data we observe current income only. The best we can do is to use time- t income to estimate L_{it+1} . We must therefore form the following mismeasured estimate of discounted expected lifetime wealth:

$$\begin{aligned} L_{it+1}^* &\equiv \exp(\epsilon_{it}) P_{it} \sum_{j=0}^{T_w-t-1} \exp(X_{it+1+j} \varphi_{t+j} - jr) \\ &= \exp(\epsilon_{it}) L_{it+1} \end{aligned}$$

where ϵ_{it} is the time- t transitory shock. Now suppose we have access to a short panel, as in the current analysis. We can now use time- $t+1$ income. Define:

$$\begin{aligned} L_{it+1}^{**} &\equiv \exp(\epsilon_{it+1}) P_{it+1} \sum_{j=0}^{T_w-t-1} \exp(X_{it+1+j} \varphi_{t+j} - jr) \\ &= \exp(\epsilon_{it+1} \nu_{it+1}) L_{it+1} \end{aligned}$$

L_{it+1}^{**} is therefore less close than L_{it+1}^* as an approximation of L_{it+1} . As we shall see, however, it provides useful orthogonality properties. Interest centres on the resulting effect on mismeasurement of the housing wealth share. Accordingly define:

$$\psi_{it}^{H*} \equiv \frac{\exp \mathbb{E}_t \ln \left(p_{t+1}^H \tilde{H}_{it+1} \right)}{\Lambda_{it+1}^*}$$

which uses housing endowments from time t rolled into time $t+1$, and such that

$$\Lambda_{it+1}^* = \exp \mathbb{E}_t \ln \left(p_{t+1}^H \tilde{H}_{it+1} \right) + \exp \mathbb{E}_t \ln (RA_{it}) + L_{it+1}^*$$

is mismeasured total discounted life-time wealth. Analogously define:

$$\psi_{it}^{H**} \equiv \frac{\exp \mathbb{E}_t \ln \left(p_{t+1}^H \tilde{H}_{it+1} \right)}{\Lambda_{it+1}^{**}}$$

which is the housing wealth share using the short panel.

As discussed, the relevant consideration for the results here is the bias in estimating the covariance of the mismeasured version of the housing wealth share with log consumption. I define the following variables, recycling notation from other parts of the appendix. Let:

$$\xi_{it} = \exp(\epsilon_{it}) - \mathbb{E}[\exp(\epsilon_{it})]$$

and

$$\tilde{\xi}_{it+1} = \exp(\epsilon_{it+1} \nu_{it+1}) - \mathbb{E}[\exp(\epsilon_{it+1} \nu_{it+1})]$$

which are both mean-zero variables.

Applying a second-order Taylor series expansion it is simple to show the following relationships:

$$\begin{aligned} \psi_{it}^{H*} &\approx \psi_{it}^H + \xi_{it} \psi_{it}^H \frac{1}{Z_{it}} + \frac{1}{2} \xi_{it}^2 \psi_{it}^H \frac{1}{Z_{it}^2} \\ \psi_{it}^{H**} &\approx \psi_{it}^H + \tilde{\xi}_{it+1} \psi_{it}^H \frac{1}{Z_{it}} + \frac{1}{2} \tilde{\xi}_{it+1}^2 \psi_{it}^H \frac{1}{Z_{it}^2} \end{aligned}$$

where $Z_{it} \equiv 1 + \psi_{it}^H$.

Therefore, using the notation from section [B](#):

$$\begin{aligned} \text{Cov}(c_t, \psi_{it}^{H*}) &\approx \text{Cov}(c_t, \psi_{it}^H) + \text{Cov}\left(c_t, \xi_t \psi_{it}^H \frac{1}{Z_t}\right) + \text{Cov}\left(c_t, \frac{1}{2} \xi_t^2 \psi_{it}^H \frac{1}{Z_t^2}\right) \\ &= \text{Cov}(c_t, \psi_{it}^H) + \mathbb{E}\left(\psi_{it}^H \frac{1}{Z_t}\right) \text{Cov}(c_t, \xi_t) + \text{Cov}\left(c_t, \frac{1}{2} \xi_t^2 \psi_{it}^H \frac{1}{Z_t^2}\right) \end{aligned} \quad (40)$$

where the second line follows from the first by application of Bohrstedt's formula discussed in appendix [B](#), and because $\mathbb{E}(\xi_{it}) = 0$.

Similarly:

$$\text{Cov}(c_t, \psi_{it}^{H**}) \approx \text{Cov}(c_t, \psi_{it}^H) + \mathbb{E}\left(\psi_{it}^H \frac{1}{Z_t}\right) \text{Cov}(c_t, \tilde{\xi}_{t+1}) + \text{Cov}\left(c_t, \frac{1}{2} \tilde{\xi}_{t+1}^2 \psi_{it}^H \frac{1}{Z_t^2}\right) \quad (41)$$

Expressions [40](#) and [41](#) show the advantage of having a panel over a cross-section. In the case of the panel, then $\text{Cov}(c_t, \tilde{\xi}_{t+1}) = 0$, because $\tilde{\xi}_{t+1}$ captures future income shocks. Therefore the error from using current income is second-order in $\tilde{\xi}_{t+1}$ and likely very small. In the case of the cross-section, the bias captured in expression [40](#) is first order and given by $\mathbb{E}\left(\psi_{it}^H \frac{1}{Z_t}\right) \text{Cov}(c_t, \xi_t)$.

I briefly attempt to quantify this latter bias. In the quantitative section 5 I used a value for the variance of transitory shocks, which matches $\text{Var}(\xi_t)$, of 0.05. Therefore the size of the bias is this variance multiplied by $E\left(\psi_t^H \frac{1}{Z_t}\right) \approx E(\psi_t^H)$, and multiplied by the elasticity of consumption with respect to transitory incomes. The framework in the present paper implies this latter elasticity is very small. However, the size of the elasticity is itself a topic of debate. For the sake of this calculation I assume a moderately large value of 0.2. If we then use an average value of $E(\psi_t^H)$ (which in fact differs by age) of 0.3, then the bias is around $0.05 \times 0.3 \times 0.2 = 0.003$. If the central value of (true) $\text{Cov}(c_t, \psi_t^H)$ is around 0.015, then this implies that the bias when using the cross-section is around 20% of the true value.

F.4 The Consumption Response to House Price Shocks in the BHPS

Here I show reduced-form evidence on the consumption response to house price shocks using the panel dimension of the BHPS. The regressions presented here are simple OLS regressions of changes in mortgage holdings and food expenditure on changes to self-reported household-level house value. Standard errors are clustered at the individual level. This evidence is similar to that discussed in [Disney et al. \(2010\)](#) and [Crossley, Levell, and Low \(2017\)](#).

The results are shown in table [A4](#), which is split into two halves. The left hand side shows the re-mortgage elasticity. The regressions shown here include a square in age. I have also removed those aged over 55 because older households are typically unable to access mortgage markets. The regressions also exclude those who have moved home or who say they have extracted equity to improve their homes. I remove these home improvers because house values for these households are endogenous to the response variable. The regressions here are therefore intended to show equity withdrawal intended for non-housing related consumption. The first column shows results from a raw OLS regression with no extra controls. The second column includes year dummies. The third column allows for an interaction of the change in house value with a dummy for whether the household was previously highly levered. Here leverage is defined purely in terms of the mortgage-to-house-value ratio. A household is considered highly levered if this ratio is above 0.7. The third regression implies that highly levered households account for much, but not all, of the re-mortgaging.

The right-hand half of the table shows regressions of changes to household residual food expenditure on changes to house values. All regressions include controls for changes to residual household income. Similarly to the left half of the table, the first column restricts the sample to those who do not move. The food expenditure elasticity, at around 0.05, can be re-interpreted as a consumption elasticity if we divide it by the income elasticity of food expenditure, which is around 0.4. The estimated coefficient therefore implies a raw consumption elasticity of around 0.12. This estimate is fairly low compared to the literature and that used in the rest of this paper. However, in this auxiliary evidence, I do not try to control for measurement error and so the estimate is likely attenuated.

The second column uses only those households that move. As expected, any changes in house values here are not news to the household and so food expenditure does not respond. The third column excludes households who extract equity to finance home improvements. The result in this column is identical to that in the first

column. Finally, the fourth column shows the interaction of changes to house value with the indicator for being highly levered. Here, the consumption response does not depend on leverage, in contrast to the results from the left-hand half of the table.

Table A4: PANEL EVIDENCE FOR HOUSE PRICE EFFECTS FROM THE BHPS

	Re-mortgage elasticity			Food consumption elasticity			
	$\frac{\Delta M_t}{p_{t-1}^H H_{t-1}}$			Δf_t			
$\Delta(\log(p_t^H H_t))$	0.076** (0.034)	0.084** (0.043)	0.037** (0.014)	0.046*** (0.014)	-0.019 (0.027)	0.046*** (0.015)	0.045*** (0.016)
$hilev_{t-1} \times \Delta(\log(p_t^H H_t))$			0.320* (0.194)				0.004 (0.032)
no home improvement	x	x	x			x	
year effects		x	x	x	x	x	x
movers					x		
Observations	17707	17707	17707	23025	1258	22043	23025
Adjusted R^2	0.001	0.001	0.008	0.001	-0.001	0.001	0.001

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Standard errors clustered at the individual level in parenthesis. M_t is (unlogged) mortgage level. f_t is log residual food expenditure. $hilev_t$ is an indicator that the mortgage-to-house-value ratio is greater than 70%. Re-mortgaging regressions exclude those aged over 55 and include a quadratic in age as an extra control. Food expenditure regressions include changes to income as an extra control. See text for more details.

Appendix References

References

- BERGER, D., V. GUERRIERI, G. LORENZONI, AND J. VAVRA (2018): “House Prices and Consumer Spending,” *Review of Economic Studies*, 85, 1502–1542.
- BLUNDELL, R., H. LOW, AND I. PRESTON (2013): “Decomposing Changes in Income Risk Using Consumption Data,” *Quantitative Economics*, 4, 1–37.
- BLUNDELL, R., L. PISTAFERRI, AND I. SAPORTA-EKSTEN (2016): “Consumption inequality and family labor supply,” *American Economic Review*, 106, 387–435.
- BOHRNSTEDT, G. W. AND A. S. GOLDBERGER (1969): “On the exact covariance of products of random variables,” *Journal of the American Statistical Association*, 64, 1439–1442.
- CAMPBELL, J. AND J. COCCO (2007): “How do house prices affect consumption? Evidence from micro data,” *Journal of Monetary Economics*, 54, 591–621.

- CLAUSEN, A. AND C. STRUB (2013): “Envelope Theorems for Non-Smooth and Non-Concave Optimization,” *Mimeo*.
- CROSSLEY, T., P. LEVELL, AND H. LOW (2017): “Consumption Spending, Housing Investments and the Role of Leverage,” *mimeo*.
- DISNEY, R., J. GATHERGOOD, AND A. HENLEY (2010): “House price shocks, negative equity, and household consumption in the united kingdom,” *Journal of the European Economic Association*, 8, 1179–1207.
- FELLA, G. (2014): “A generalized endogenous grid method for non-smooth and non-concave problems,” *Review of Economic Dynamics*, 17, 329–344.
- GOODMAN, L. (1960): “On the exact variance of products,” *Journal of the American Statistical Association*, 55, 708–713.
- ISKHAKOV, F., T. JORGENSEN, J. RUST, AND B. SCHJERNING (2017): “The Endogenous Grid Method For Discrete-Continuous Choice Dynamic Models With (or Without) Taste Shocks,” *Quantitative Economics*, 8, 317–365.
- PIAZZESI, M. AND M. SCHNEIDER (2016): “Housing and Macroeconomics,” *Handbook of Macroeconomics: Volume 2*.
- WOOLDRIDGE, J. (2002): *Econometric analysis of cross section and panel data*, MIT Press.