Working Paper no. 271

Consumption, house prices and expectations

Orazio Attanasio, Laura Blow, Robert Hamilton
and Andrew Leicester

September 2005

Bank of England

Consumption, house prices and expectations

*Orazio Attanasio**
*Laura Blow***
*Robert Hamilton****
and
*Andrew Leicester***

Working Paper no. 271

* University College London.
** Institute for Fiscal Studies.
*** Structural Economic Analysis Division, Monetary Analysis, Bank of England, EC2R 8AH.
Email: robert.hamilton@bankofengland.co.uk. (Author for correspondence.)

The views expressed in this paper are those of the authors, and not necessarily those of the Bank of England. We would like to thank Peter Andrews, Ian Bond, Neal Hatch, Simon Price, Gabriel Sterne, Simon Whitaker and two anonymous referees for useful comments on earlier drafts. We are also grateful to seminar participants at the Bank of England, the Institute for Fiscal Studies and the UK Treasury. This paper was finalised on 16 June 2005.

The Bank of England's working paper series is externally refereed.

Information on the Bank's working paper series can be found at www.bankofengland.co.uk/publications/workingpapers/index.htm.

Publications Group, Bank of England, Threadneedle Street, London, EC2R 8AH; telephone +44 (0)20 7601 4030, fax +44 (0)20 7601 3298, email mapublications@bankofengland.co.uk.

ISSN 1368-5562

Contents

Abstract

Over much of the past 25 years, the cycles of house price and consumption growth have been closely synchronised. Three main hypotheses for this co-movement have been proposed in the literature. First, that an increase in house prices raises households' wealth, particularly for those in a position to trade down the housing ladder, which increases their desired level of expenditure. Second, that house price growth increases the collateral available to homeowners, reducing credit constraints and thereby facilitating higher consumption. And third, that house prices and consumption have tended to be influenced by common factors. This paper finds that the relationship between house prices and consumption is stronger for younger than older households, and that the consumption of homeowners and renters are equally aligned with the house price cycle. This suggests that neither the wealth nor the collateral channels have been the principal cause of the relationship between house prices and consumption—instead, the most important factor is likely to have been common causality.

Summary

Over much of the past 25 years, the cycles of house price and consumption growth have been closely synchronised. Three main hypotheses for this co-movement have been proposed in the literature. First, that an increase in house prices raises households' wealth, which increases their desired level of expenditure. Second, that house price growth increases the collateral available to homeowners, reducing credit constraints and thereby facilitating higher consumption. And third, that house prices and consumption have tended to be influenced by common factors (eg productivity growth or tax changes), which cause revisions to households' expected lifetime income. This paper uses individual household level data to assess the importance of these different hypotheses. Revisiting this link seems particularly timely, as the housing market has cooled since the end of 2004, generating widespread press speculation about the outlook for prices. In addition, there is the puzzle, discussed in a box in the Bank of England's *Inflation Report* in November 2004, about the recent decline in the correlation between house price and consumption growth, and hence the likely impact of house prices on consumption in the future.

Many previous related studies have focused on the late 1980s consumption and house price booms. Attanasio and Weber recognised that microeconomic data on individual households' expenditure provides a way to distinguish between the competing wealth and common causality hypotheses. If wealth effects were important, older homeowners—who are less likely to demand more housing services in the future—should be the primary beneficiaries of a house price boom and should increase their consumption the most. In contrast, if house prices and consumption are both influenced by common expectations of income growth, younger consumers, with a greater remaining lifespan to realise the gain, should be the ones to raise consumption the most. Their paper argued that common causality was the more likely explanation for the late 1980s correlation. But since then, many other studies, mainly relying on aggregate data, have argued that there is a direct wealth effect.

This paper extends and updates Attanasio and Weber's results, covering data spanning the consumption and house price weakness of the early 1990s, and developments up to and including 2001. We estimate various specifications for individual households' consumption using pseudo-cohorts drawn from 24 years of the Family Expenditure Survey between 1978 and 2001/02. In our baseline specification, the consumption of a household in a given year depends on the cohort to which it belongs, the age of the head of household and various other demographic and household characteristics. We then assess the extent to which adding various house price terms to our baseline model can help explain the consumption patterns over time. By analysing the results for households in different age groups, we determine whether house price movements appear to be a more important determinant of the consumption of younger or older households, or of renters or homeowners, using similar identifying assumptions to those previously used by Attanasio and Weber.

We find several pieces of evidence which suggest that common causality has been the most significant explanation for the co-movements between house price and consumption growth. First, younger cohorts had the largest swings in expenditure during the consumption and housing cycles. Second, the effect of regional house price growth on consumption is found to be stronger for these younger households. Third, the coefficient on the regional level of house prices is as large for younger as for older households, while they had a greater response to the effect of 'unexpected' house price movements. And fourth, the consumption of both homeowners and renters are equally aligned with the house price cycle. Of course, it remains likely that the wealth and collateral channels are important for some households at some points in time. But the evidence in this paper suggests that the main reason for house prices and consumption being correlated in the past is changes in common driving factors—like income expectations.

1 Introduction

Over the past two decades, consumer spending growth in Britain has undergone large shifts: very strong growth in the late 1980s was followed by stagnation in the early 1990s and then sustained strength over the past eight years. This pattern has borne a remarkably close resemblance to house price inflation—although the boom in house prices has been rather more pronounced over recent years (Chart 1). This paper uses microeconomic evidence to improve our understanding of the link between these trends.

Chart 1: Annual real house price and consumption growth

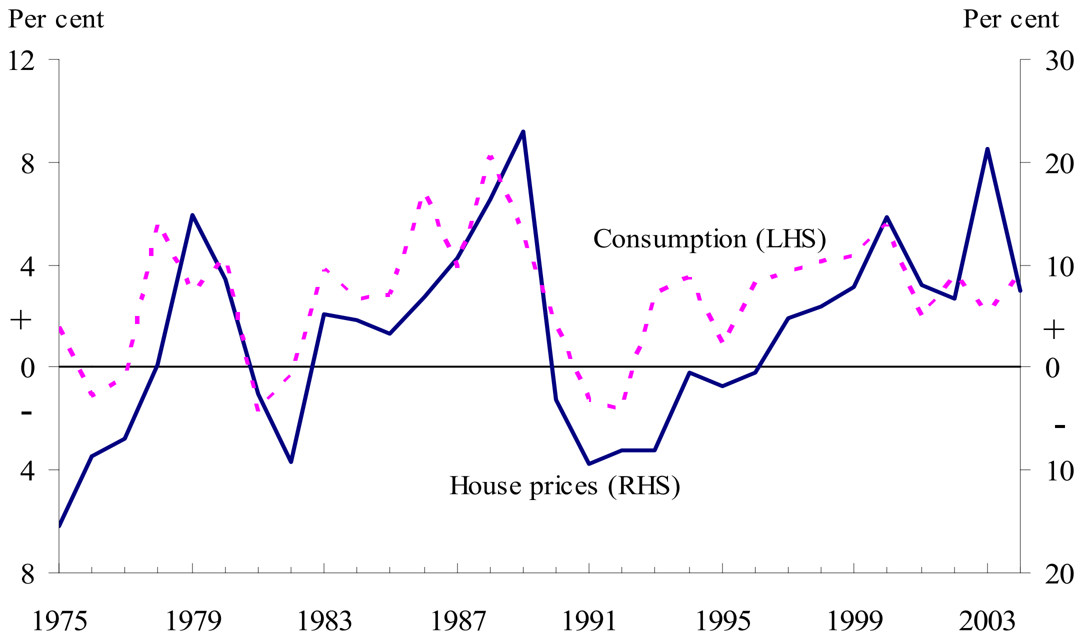

Sources: ODPM, ONS and authors' calculations.

Understanding the relationship between changes in house prices and aggregate consumption is of particular current importance. Strong house price growth in the late 1990s, and first few years of the 2000s, has raised the house price to income ratio above its long-run average, and led to much speculation about the outlook for house prices—with the debate intensified by signs that the housing market was cooling towards the end of 2004 and beginning of 2005.[(1)] Given this uncertainty, it is important to assess the effect that changes in house prices could have on consumption. Related to this, Bank of England (2004) reports that the traditionally strong correlation between house price and consumption growth has weakened since 2001, concluding that 'the association between house

---

[(1)]For example, Barr and Buckley (2004) argue that house prices may be between 10% and 40% overvalued. Bank of England *Inflation Reports* provide a regular commentary on the housing market.

prices and consumer spending has weakened in recent years… [and] that it is likely to be less strong in the future'.

It remains true that housing is an important part of households' balance sheets. An increase in the size, or quality, of the housing stock unambiguously increases wealth, as it increases the flow of housing services that can be derived from the stock. But the real effect of an increase in the nominal price of a unit of the housing stock is less clear. Consider a stylised economy where all houses are owned by a holding company that rents them back to the household sector (which owns the shares in the company). In this case, an increase in house prices would leave households' aggregate wealth unchanged—as the increase in the value of their shares in the property holding company would be offset by the higher discounted future rental costs. Even without this elaborate set-up, house prices should not obviously affect households' *aggregate* wealth, as an increase in house prices would be offset by a rise in the price of housing services ie a higher user cost of housing.

But these stylised arguments may not be the whole story. At the individual household level, house price changes seem likely to have important effects—with these effects varying with current and planned future housing status.[(2)] As discussed in Sinai and Soules (2003), and Campbell and Cocco (2004), homeowners who expect to remain in their current dwelling for a very long time are perfectly hedged against fluctuations in rents and house prices. These fluctuations have no effect on their real wealth and, in the absence of any substitution effects or credit constraints, should not affect consumption choices. However, for homeowners that plan to trade down, or stay put and access their housing wealth through an equity release scheme, the overall wealth effect from an increase in house prices might be positive, facilitating an increase in consumption. For homeowners wishing to trade up in the future, the effect is more ambiguous since the value of their current property has unexpectedly increased, but so too has the price of any future dwelling. For non-homeowners hoping to purchase a house in the future, an increase in prices might be expected to lead to a reduction in current consumption. Similarly, long-term renters are likely to be worse off if rents move in line with house prices. The effects on the younger generation, who are less likely to be owners, might be mitigated by the extent to which they receive bequests from the older generation, who are more likely to own their own home. Differences in preferences and circumstances across these different types of households may mean that house price changes do affect aggregate consumption.

Changes in house prices may also affect consumption by changing the degree to which credit constraints are binding. Aoki *et al* (2001), recognising the strong relationship between net housing

---

[(2)] Statistics from the ODPM reveal that in 2001, 70% of households in the UK were owner-occupiers (including those with mortgages), 10% rented privately and 19% rented from local authorities or registered social landlords. In 1981, the proportions were: 58%, 11% and 31%.

equity and mortgage equity withdrawal, pointed out that a rise in house prices increases the collateral available to homeowners. This may encourage them to borrow more (in the form of mortgage equity withdrawal), enabling them to finance higher consumption.[3] A related issue is that changes in house prices may also affect households' desire for other forms of precautionary saving.

An alternative explanation for the historic relationship between consumption and house prices is that both are driven by common factors. King (1990) and Pagano (1990) argued that the strong correlation between aggregate house price and consumption growth in the late 1980s does not necessarily infer causation. They advanced a competing hypothesis: that an upward revision to expected future incomes (due to either current or anticipated productivity gains) simultaneously increases the demand for housing services (which, with a relatively fixed supply of housing, in turn raises house prices) and consumption.

Attanasio and Weber (AW) (1994) used micro data to investigate the relationship between consumption and house prices in the late 1980s. In particular, they assessed the common causality hypothesis against the possibility that the boom reflected a wealth effect caused by higher real house prices, as suggested by Muellbauer and Murphy (1990).[4] They did this by recognising that the two have different implications for the behaviour of different types of households. Under the wealth explanation, households composed of older people are likely to be the ones that could increase their consumption most after an increase in house prices, since they are more likely to own their home, less likely to want to trade up in the future and have a shorter time-frame to enjoy their wealth gain. Those with a second house may also experience an additional wealth effect. However, strong house price increases could reduce consumption among those, predominantly young, people that are either renting or looking to buy a home—and who now face higher-than-expected future housing costs. In contrast, under the productivity explanation, younger households would be expected to benefit most, as a permanent upward revision to all expected future earnings would be more significant as they have longer remaining working lives.

AW tested these hypotheses by estimating a 'baseline' life-cycle age profile for different cohorts of households from 'pre-boom' data. They compared the actual consumption of these cohorts during the boom years to the baseline profile, finding that the consumption of the younger cohorts was much higher than expected, but it was close to expected for the older. They also found that including regional house price terms in a regression of consumption on age and other covariates for various birth cohorts enabled them to explain the majority of the late 1980s consumption boom for older

---

(3) Credit-constrained consumers will also be sensitive to other factors, including changes in their current income. For example, strong real income growth during the late 1980s boom may have eased credit constraints for some households, enabling them to increase their consumption.
(4) Muellbauer and Murphy (1990) argued that the late 1980s house price boom reflected increased demand following financial liberalisation.

households, but much less of the (rather more pronounced) boom for younger households. They therefore concluded that the productivity hypothesis was the more likely explanation for that consumption boom.

More recently, a number of authors have re-visited the nature of the relationship between house prices and consumption. Case, Quigley and Shiller (2003), using aggregate data, argued that variations in housing wealth have important effects on consumption—a conclusion also reached by Ludwig and Slock (2004). Given the difficulties in distinguishing between the alternative hypotheses using aggregate data, perhaps of even greater significance is that Campbell and Cocco's (2004) micro study also concluded that the wealth channel was the most significant.

This paper adds to this debate by updating and extending the work of AW by considering the evidence for the wealth and collateral explanations for the relationship between house prices and consumption beyond the late 1980s.(5) The main alternative—that the strong historic relationship between the two reflects common causality—also encompasses the possibility that there may be shocks to house prices that do not affect expectations, and thus consumption. Relative to AW, we have 13 additional years of Family Expenditure Survey (FES) data to investigate this relationship, including important episodes such as the fall in house prices in the early 1990s and some of their subsequent boom. Section 2 discusses the empirical background to this work, looking at data on consumption and house prices. Section 3 describes the econometric methodology used to discern the possible relationship between our key variables, while Section 4 presents the results. Section 5 concludes.

2 Data

This section describes the paths of household expenditure and regional house price growth since the mid-1970s, when our sample period begins.

2.1 Consumption

This paper uses household level survey data on consumer spending from the Family Expenditure Survey.(6) This is a micro data set, from which we use details on households' spending, demographics, region of residence and housing status. This enables us to examine the effects of

(5)The paper does not attempt to provide empirical tests of other consumption theories. See Attanasio (1999) for a review of this literature.
(6)The FES is one of the inputs used by the Office for National Statistics when calculating their official expenditure statistics.

house prices on individual households' consumption, helping us to unpick theories which are observationally equivalent at the aggregate level.

We use 24 years of the FES,[7] comprising the surveys conducted between 1978 and 2001/02. The FES has been run on an annual basis since 1968, and each year records detailed information on the expenditure of around 7,000 households. However, comprehensive information on educational attainment is only available from 1978, so the first few waves are omitted from our analysis.

We use data on both durable and non-durable consumption, which are collected by households recording a diary of expenditure over a two-week period, together with information on income and demographics, which are recorded at a subsequent interview.[8] Our measure of consumer spending does not include housing expenditure as there is no measure in the FES which matches the treatment of housing in the National Accounts, where an imputed value of housing services is estimated. However, we include other durable expenditure, as such expenditure has tended to be particularly aligned with the housing cycle in the past (Bank of England (2004)). Nominal expenditure is deflated to December 2001 values using a non-housing Retail Price Index (RPI) measure.

Chart 2 shows average per capita per week (pcpw) expenditure from the FES between 1975 and 2001 at constant (2001) prices. Real expenditure grew only modestly in the early 1980s. The late 1980s saw a consumption boom, with strong growth from around £91 pcpw in 1985 to £115 pcpw in 1990. There then followed a further period of stagnation in the early 1990s. Real per person spending in 1996 was still below its 1990 figure. Consumption growth was robust in the final part of our sample period, with a real increase of 27% between the 1995 trough and 2001.

To ensure the wider relevance of our results, it is important to check that the aggregate information in the FES exhibits the same trends as official data recorded by the Office for National Statistics (ONS). In theory, a grossed-up FES should—given accurate recording of expenditures and an adequately stratified sample—provide a similar total economy-wide spending figure to that in the national accounts. But, in practice, some items are recorded badly in the FES, such as alcohol and tobacco expenditure. Comparing the growth rates of spending from both sets of data therefore

---

[7]From 1993/94, the Family Expenditure Survey data was reported in fiscal rather than calendar year terms; however, our analysis uses calendar years by reclassifying responses according to the month of interview. From 2001/02, the FES has been amalgamated with the National Food Survey to give the Expenditure and Food Survey (EFS).

[8]Since 1989, the interview has also included questions on purchases of large, infrequently purchased items such as durables and holidays. By asking the households what their expenditure had been over, say, the previous 6 or 12 months, it is hoped that better weekly equivalent data (in terms of reduced variance) could be obtained. However, there is evidence that households report a higher average weekly equivalent expenditure—perhaps by rounding-up figures from memory, or recalling spending over a longer period. Given this problem, and that the recall questions were only introduced halfway through our sample, we solely rely on expenditure reported in the diary.

provides a useful check of whether analysis of the FES data might be able to help us understand trends in the aggregate data.

Chart 2: Average per capita weekly expenditure

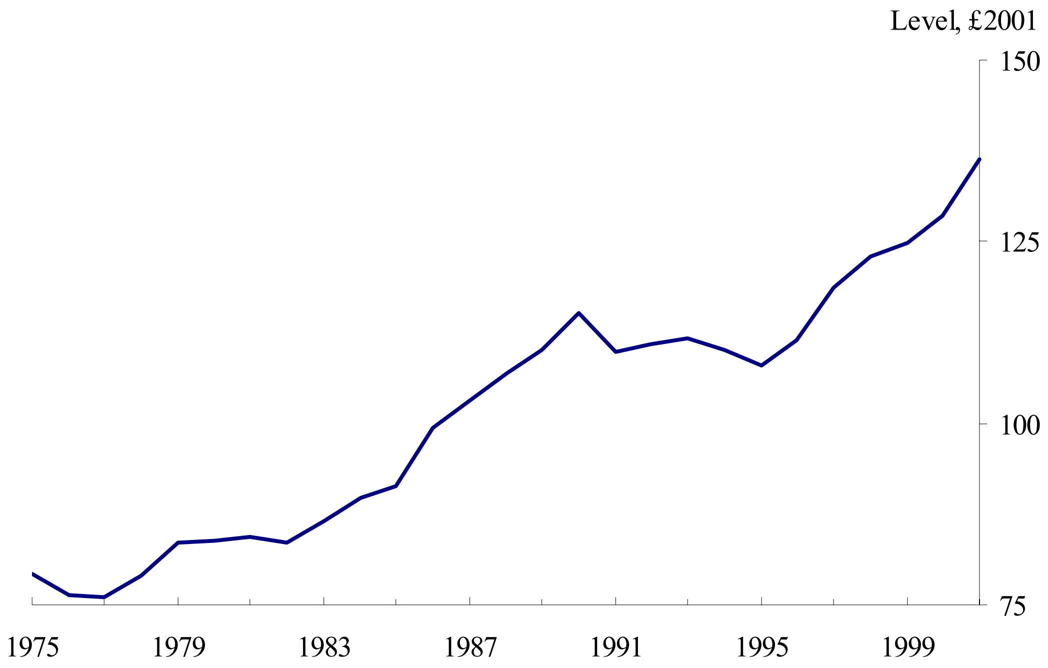

Sources: FES/EFS, ONS and authors' calculations.

Chart 3 shows how real consumption growth, as recorded by the ONS and the grossed-up FES, have evolved since 1975. Excluded from the ONS measure are housing components comparable to those excluded from the FES. Both are reported on a per capita basis to account for population growth. These results show a good correlation between the ONS and FES data.(9) The FES measure is slightly more volatile, though this is not unexpected given sampling variation inherent in survey data. The late 1980s stands out as a period of very high consumption growth in both data sets, although the peak in the FES is in 1986, two years earlier than in the ONS data. The early 1990s decline is also stronger in the FES data, and the recovery slightly slower. Indeed, expenditure in the FES sample declined in 1994/95, while the ONS measure increased. Both sets pick up the strong recovery in the latter 1990s. It is possible that the 2001 discrepancy relates to the changes in the household survey, but we have no explanations for the other divergences.(10)

---

(9)Other robustness studies of the FES have also concluded that it captures national spending data quite well. See, for example, the chapter by Tanner in Banks and Johnson (1998) and Blow, Leicester and Oldfield (2004).
(10)Given the lack of correspondence between the FES and the National Accounts for these periods, we undertook robustness analysis where we dummied out those periods. This did not change the conclusions from this work. Similarly, the results are robust to excluding 2001. Additional results are available on request.

Chart 3: Real per capita expenditure growth

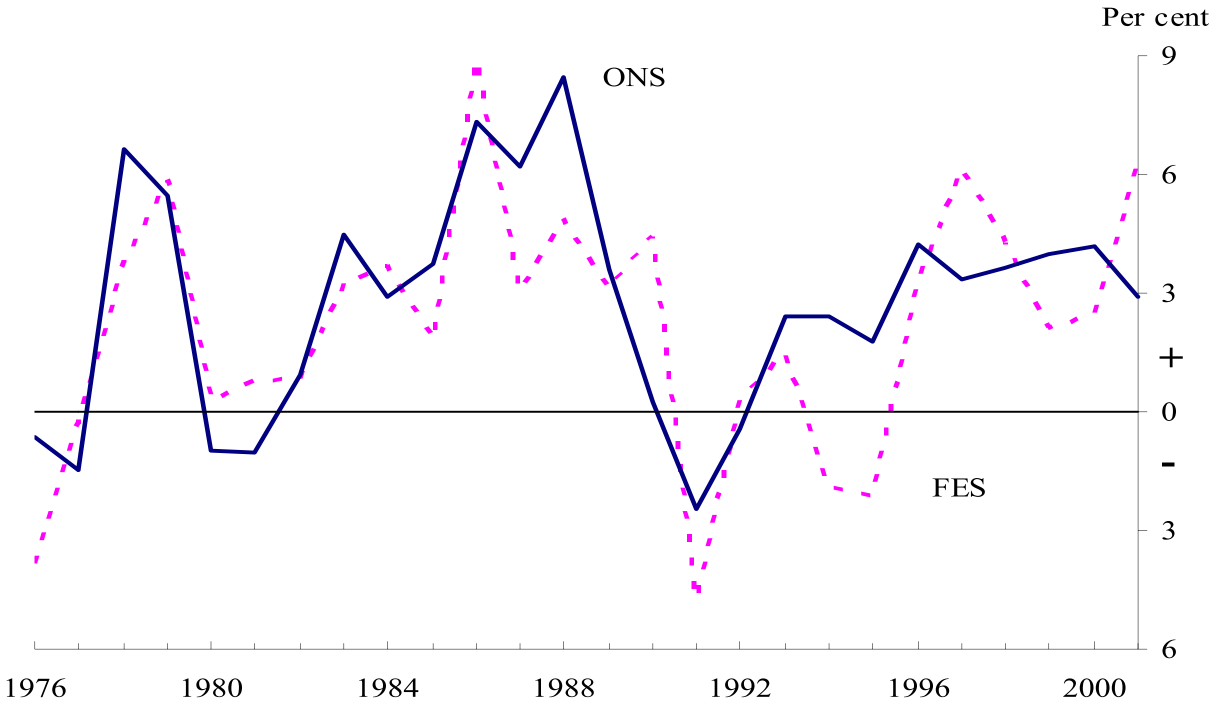

Sources: FES/EFS, ONS and authors' calculations.

As well as looking at these broad trends, it is worth examining how expenditure has fluctuated within the different household types that form the crux of our analysis in this paper. If house prices are an important determinant of consumption, we might expect differences to emerge in the spending patterns of different groups at different points of the house price cycle. In order to smooth out some of the year-on-year fluctuations in spending growth, we average over periods corresponding to high or low consumption growth: 1977-79, 1980-85, 1986-90, 1991-95 and 1996-2001, and plot these averages for three groups of consumers in Chart 4.(11) These groups are defined, both here and in much of the rest of the paper, as the 'young' (aged under 35), the 'middle-aged' (aged 35 – 60) and the 'old' (aged over 60). Part of the aim of this paper is to consider whether the synchronisation of the consumption and housing cycles has been most pronounced among younger people—consistent with changes in expectations and, perhaps, the collateral channel.(12) Or whether house price changes are better able to explain expenditure for older people—who are more likely to own houses and therefore enjoy (suffer) capital gains (losses). From Chart 4, it seems that the consumption

(11)The smoothing is solely to improve the clarity of the graphical presentation: the empirical analysis conducted in the remainder of the paper uses the unmodified raw data.
(12)Consumption of the young may be volatile for other reasons, including that they may be more likely to be credit constrained, so that their consumption may move with fluctuations in income. Habits may also be less important.

behaviour of the young is indeed prone to greater fluctuation, with this pattern following that for house price growth. In what follows, we consider whether this remains the case once we use a more rigorous analytical framework.

Chart 4: Real consumption growth of different age groups

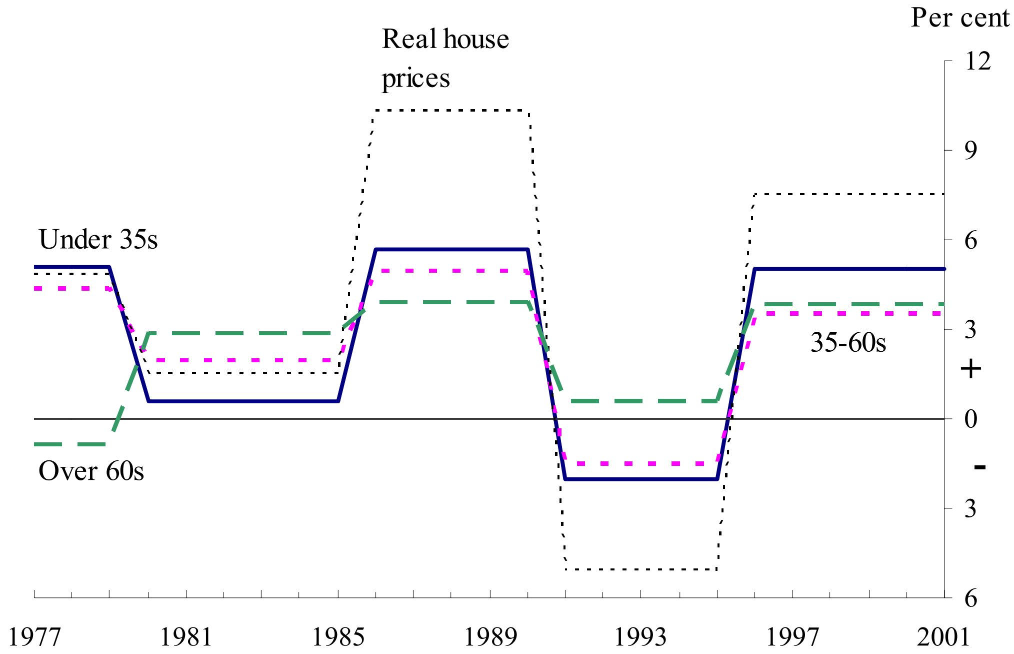

Sources: FES/EFS, ODPM, ONS and authors' calculations.

Homeowners and renters have had a strikingly similar pattern of consumption growth over our sample period (Chart 5). Homeowners increased their spending slightly faster during the late 1980s house price boom, but it declined less sharply during the early 1990s house price falls. More recently, average consumption growth has been almost identical for both groups. These findings constitute *prima-facie* evidence against the wealth hypothesis being the main explanation for the joint movements of consumption and house prices—as under that hypothesis we might have expected renters' consumption to fall in the house price booms and increase in the periods of weakness (as rents might be expected to move in line with house prices, and because renters may hope to move onto the housing ladder). Similarly, the results provide no support for the collateral hypothesis, which would also predict a much stronger alignment for homeowners, and perhaps a negative one for renters.

Chart 5: Real consumption growth of homeowners and renters

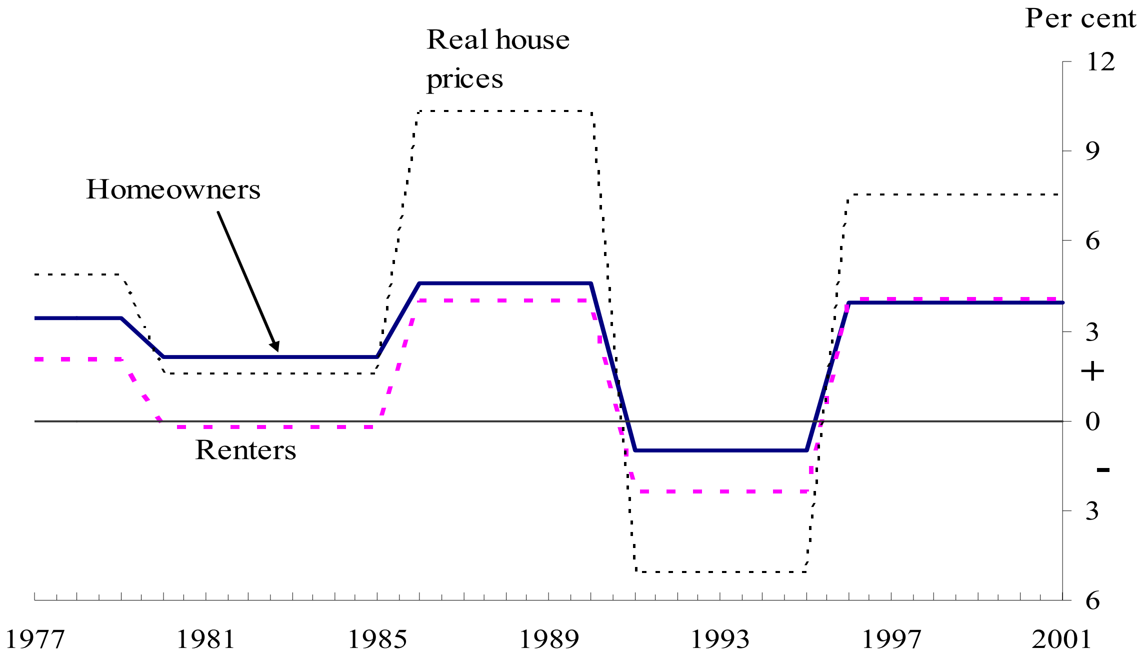

Sources: FES/EFS, ODPM, ONS and authors' calculations.

While the patterns in Chart 5 are interesting, and in what follows we discuss a specification using renters and homeowners, the primary focus of this paper is the analysis by age group. This is because house tenure status, unlike age (for most people), is a decision variable. Therefore, a distinction between renters and homeowners introduces possible selection issues: for example, a household experiencing a positive shock which allows them to increase expenditure might also move from being a renter to a homeowner. Equally, someone experiencing a negative shock might seek to cut back expenditure and may choose to rent rather than own as a result. Moreover, much of the analysis below is based on the use of synthetic panels, which are constructed using time series of cross-sections drawn from a population of fixed composition, and the split between homeowners and renters has not been constant with the former increasing over our sample period. Therefore, the results from pseudo-panel analysis based on these groups could be biased.

2.2 House prices

We use the official house price statistics, as published by the Office of the Deputy Prime Minister (ODPM).(13) A new house price index was launched by the ODPM in 2003, but our data is taken from their earlier series, for which there is a longer time series.(14) Nominal house prices were converted into real values by deflating by the non-housing RPI measure that was used in the construction of the real expenditure measure.

At the national level, annual real house price growth between 1975 and 2001 has been volatile, with the rate fluctuating between +25% and –15% (Chart 1). In real terms, prices increased rapidly in the late 1970s, late 1990s and first few years of the 2000s; while they fell in the early 1980s and, more sharply, in the early 1990s.

These aggregate changes also mask some sharp regional differences (Table A1 in the appendix).(15) In the 1980s, almost all regions saw substantial price inflation, though prices in Scotland lagged behind somewhat and the fastest growing regions were, as in the late 1990s, London, East Anglia and the South. These regions also saw a more pronounced house price contraction in the early 1990s. In the late 1990s and first part of the new century (which coincides with the end of our data sample period in 2001) house price inflation outside London and the South has been much lower: the 15% or so real gains observed nationally in 2000 masked rises of almost 20% in London, East Anglia and the South, but considerably less than 10% in most other regions.

In 2002, house prices began to rise much more rapidly outside London and the South East: growth rates were 10% in London and 16% in the South East, compared to more than 20% in Wales, the South West, East Anglia, the Midlands and Yorkshire. In 2003, annual house price inflation picked up to over 20%, with strong rises in all regions. Annual house price growth slowed in 2004, with prices falling between some months on the Halifax and Nationwide measures.

Given these different patterns of house price movements across Great Britain, we control for price changes at the regional rather than national level.(16) Under the wealth and collateral hypotheses, homeowners who live in regions with high price inflation would be expected to increase consumption more than those with lower house price inflation. Of course, the negative effect of

---

(13)Based on completed transactions prices rather than agreed prices before a sale is completed.
(14)See the ODPM website at
www.odpm.gov.uk/stellent/groups/odpm_control/documents/contentservertemplate/odpm_index.hcst?n=1575&l=3
(15)Results in our main analysis are for Great Britain (ie excluding Northern Ireland) so the regional data for this region is excluded from Table A. However the data in Charts 1, 3, 4 and 5 for aggregate changes do include Northern Ireland.
(16)We also consider house prices in the quarter (as opposed to the year) in which the household was surveyed.

house prices on consumption on those who have aspirations to own homes, or trade up, in the future would also be expected to be exacerbated in regions with high house price growth.

As we discuss below, the influence of house prices on consumption may well depend not only on the change in price but also on its level.(17) From Chart 6 it is clear that, in real terms, the value of houses at the end of our data period was not much above the 1980s peak in the country as a whole. However, real house prices fell by around one third from their 1989 Q3 peak to the 1996 Q2 trough, with this decline subsequently slightly more than reversed. There is also sharp regional variation, with London experiencing the highest overall increase between 1975 and 2001 and Scotland the lowest.

Chart 6: Real house price indices: aggregate and for selected regions

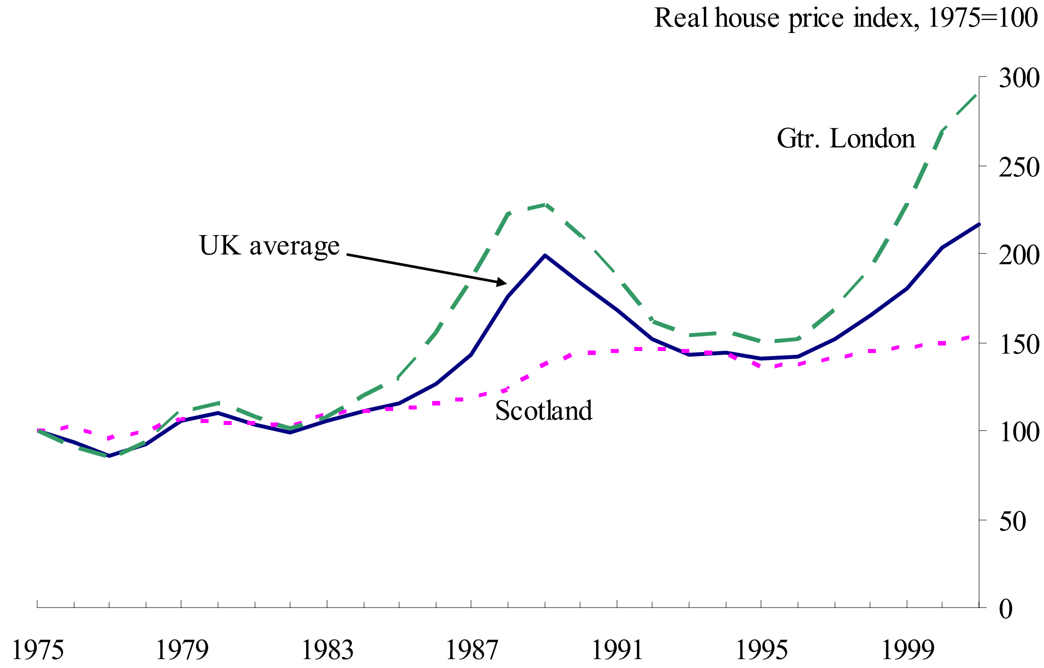

Sources: ODPM, ONS and authors' calculations.

3 Methodology

3.1 General framework

While the methodology we use in this paper is derived from the approach taken by AW, we extend it in two important ways. First, we focus more explicitly on the role played by the level of house

---

(17)The ODPM report that the average UK value of a house (all dwellings) at the end of 2001 was £114,083.

prices. And second, as we have a longer sample period, we are less constrained in the structure we need to put on the estimation of the 'age profiles'. The consumption boom of the late 1980s constituted a large fraction of their sample period but, with the additional years of data we now have, this period can now be treated as just another cyclical episode.

Our framework is derived from a simple version of the life-cycle model. This predicts that consumption in each period is given by a fraction of lifetime wealth that depends only on age. Thus

$$X_t^h = \kappa(age^h) W^h \exp(\varepsilon_t^h) \qquad \textbf{(1)}$$

where X_t^h is the consumption expenditure of household h at time t, W^h is their lifetime wealth and ε_t^h is a residual term. W^h includes human wealth (discounted lifetime earnings), net financial wealth, pension and housing wealth. The function $\kappa(age^h)$ captures several factors, including: the age composition of the household (and therefore the distance from the end of their lives), changes in household needs, changes in discount factors (which could be induced by changes in the probability of death) and so on. We can capture some of these factors using observable variables (such as family size and composition), while we proxy others with a flexible function of age. Variation between different cohorts' non-housing lifetime wealth is captured by cohort dummies and, possibly, other variables that we discuss below. As discussed in the introduction, housing is also an important component of many households' wealth, and we detail our treatment of housing terms in Section 3.2.

As it is typically impossible to obtain a closed-form solution for consumption from a standard life-cycle model, except under very strong and unattractive assumptions (such as quadratic utility), equation (**1**) should only be interpreted as an approximation. In such a framework, the residual term reflects both innovations to permanent income and transitory shocks to current income. These residuals, while correlated across individuals, vary due to differences in circumstances and preferences. Moreover, they also incorporate measurement error in consumption. We assume that the residuals are not correlated with the explanatory variables, that they average out to zero over the estimation period and that they are uncorrelated with deterministic trends. The latter two assumptions enable us to disentangle age, cohort and time effects.

Taking logs of equation (**1**) we get:

$$\ln X_t^h = \ln \kappa(age^h) + \ln W^h + \varepsilon_t^h \qquad \textbf{(2)}$$

By incorporating observable variables (vector z), such as family size and composition, that capture some of the age and other effects on consumption, and using a polynomial in age (function f), this becomes:

$$\ln X_t^h = f(age^h) + \gamma' z_t^h + \ln W^h + \varepsilon_t^h \quad \textbf{(3)}$$

Unfortunately, the FES is not a panel data set, so we cannot follow individual households over time. Instead, we exploit the repeated cross-section nature of the survey to build a pseudo-panel by birth cohort. This technique, first pioneered by Browning *et al* (1985) and Deaton (1985), has become relatively standard in the consumption literature. We create our pseudo-panel data by looking at cohorts of households where the head of household (defined in the FES survey as the principal owner or renter of the property) was born in a particular year. For the period 1975 – 2001 we distinguish 15 cohorts. The oldest, cohort 1, contains all households where the head was born before 1910. Cohort 2 contains households where the head was born between 1910 and 1914, cohort 3 between 1915 and 1919, and so on up to cohort 15, born between 1975 and 1979.

Averaging equation (**3**) over birth cohort *c* gives:

$$\ln X_t^c = f(age^c) + \gamma' z_t^c + \alpha^c + \varepsilon_t^c \quad \textbf{(4)}$$

where $\ln X_t^c$ denotes the average of $\ln X_t^h$ across all households belonging to birth cohort *c*, and so on. Given appropriate normalisation of the coefficients, the average (log) lifetime wealth of households belonging to cohort *c* is α^c.

In the context of a dynamic analysis, based on average cohort techniques, there are additional reasons to control for observable variables, captured in vector *z*. First, controlling for variables that are fixed over the life cycle in the group population but are not fixed in our sample (because of, for example, sampling variation) and that are correlated with consumption, help improve the precision of our estimates. Examples include covariates such as seasonal variables, sex and, perhaps, region and occupation. Second, if there is significant attrition from the group that is correlated with consumption, then our estimates will be biased by sample selection unless we control for variables that capture the compositional changes. An example is the differential mortality by educational level, which also tends to be correlated with lifetime wealth and hence consumption (even if education does not affect preferences in itself). Controlling for education therefore helps prevent biased estimates. And finally, there are some other variables that vary over the life cycle which directly affect consumption decisions, such as household size and composition. If the life-cycle profiles of these variables vary across different cohorts, we would want to control for them to obtain unbiased estimates of the age profile. For example, the simple life-cycle model predicts that households attempt to keep the marginal utility of consumption constant across periods. It is plausible that households of different size face different marginal utilities from the same level of expenditure, thus different lifetime profiles of household size over different cohorts would affect the age profile of the consumption *level*.[18] Since the FES is a random sample of the UK population, using it to calculate the mean log consumption of each cohort over time should give us an unbiased estimate of the expected age profile for a member of a given cohort (with some caveats discussed below).

In estimating equation (**4**) with average cohort data we assume that the age profile of consumption is the same across cohorts except for an intercept (representing differences in lifetime resources) and observable variables related to consumption needs. For the latter set of variables, we assume that the effect they have on consumption decisions is the same across all cohorts. Of course, these variables can generate differences in the age profile of consumption across cohorts if they evolve differently for different cohorts.

Equation (**4**) can be estimated using average cohort data or individual data. In the latter case, the specification we use is:

$$\ln X_t^{ch} = \alpha^c + f(age^c) + \gamma' \mathbf{z}_t^{ch} + \varepsilon_t^c + u_t^{ch} \tag{5}$$

where the superscript ^{ch} denotes household h belonging to cohort c, and u_t is the household's deviation from the cohort average. By observing several cohorts over a long period of time, we assume that the consumption innovations ε_t^c average out to zero over time in our regression. Notice that the flexible function of age and cohort dummies takes care of any deterministic trends in the data.

Estimating equation (**5**) gives us a 'base' consumption profile for each cohort. As in AW, we interpret deviations of observed consumption from such a profile as being determined by innovations to either lifetime resources or to transitory income. By comparing the innovations of different cohorts, we hope to shed some light on changes in aggregate consumption. We then add additional variables to capture some specific changes to resources (such as capital gains to house prices) and check how they alter different cohorts' residuals.

AW estimated the base consumption profiles using data up to 1985. Beyond that date, they added dummy variables for each year and cohort such that the estimate of year-cohort mean consumption after 1985 is completely unconstrained by the functional form imposed in equation (**5**). This ensured that the large residuals of the late 1980s, which were the focus of AW's analysis, did not distort the

---

[18]Of course there are other variables that affect the marginal utility of consumption but that are determined jointly with consumption. Examples include labour supply or housing tenure choices. Controlling for the life-cycle evolution of such variables is obviously more complicated as they and consumption might be affected by the same shocks. One possibility is to use more structural methods, such as those based on Euler equations, to estimate preference parameters and then use these estimates to get life-cycle profiles for marginal utilities. Some would also argue that family size should also be modelled jointly with consumption.

estimates of the base age profiles. In other words, their worry was that those residuals, over a relatively short time horizon, would violate the hypothesis that the innovations to equation (**5**) average out to zero over the sample period.

As we have a much longer sample period, we employ a different strategy. We assume that, even including the late 1980s and the last few years of the sample, consumption innovations do average out to zero in our sample period (which also, for example, includes the early 1990s consumption decline). We therefore estimate the age profile of each cohort by estimating equation (**5**). Deviations from these profiles are identified by the residuals for each cohort and time period. This approach also enables the identification of cohort intercepts for cohorts born between 1965 and 1980—which would not be possible with the AW method since the constituent households were too young to be included in the data before 1985. However, the cost of this is that we are not able to consider the hypothesis of a permanent upward shift in the consumption age profile in the mid-1980s, although there is no sign that there has been a permanent shock to productivity in the official statistics.

3.2 Introducing house price effects

The main aim of this paper is to investigate whether house price movements can help explain the pattern of consumption since the late 1970s. However, as discussed above, house price changes may have different effects for different types of household. A change in house prices has a wealth effect, the size (and even direction) of which reflects the household's current tenure status, future housing needs and the extent to which the household values its descendants' welfare. It may stimulate substitution effects on non-housing consumption or affect credit constraints. There is also the non-causal channel—that house prices may move in line with other determinants of the economy. And a full treatment of housing would also require an allowance for transaction costs, volatility, the cost and ability to finance house purchases, and using housing equity to finance non-housing consumption. Given the potential for complicated interactions, we explore a number of different specifications involving house prices and do not attempt to identify a single specification with a particular channel.

We initially investigate the extent to which annual regional house price growth improves our ability to explain the deviation of a cohort's consumption away from its estimated 'trend' life-cycle profile. This is motivated by the very high correlation between this deviation and annual house price growth—a correlation which is only slightly lower than that between annual house price and consumption growth.[19] At a more theoretical level, house price growth might be important if it helps ease credit constraints, or it may coincide with a more optimistic view of the economy.

We next consider the effect of introducing the level of house prices. This might be expected to better-capture any pure wealth effects, as under the PIH theory it should be the *level* of resources that affects the *level* of consumption. In addition, it might also better capture differences in lifetime resources that are not captured by the cohort dummies in equation (**5**). This could improve the efficiency of our estimates if there is any sampling variation in average cohort wealth in our data which can be captured by including house prices. Moreover, regional house prices levels can also be related to permanent income, as they might be influenced by the level of productivity and economic activity in different regions, and thus help us assess the common causality hypothesis.

To improve the power of our tests, our final two specifications investigate the effect of 'unexpected' changes in regional house prices on consumption. Our first method for predicting 'expected' house price changes uses *ex-post* values of income and real interest rates. For robustness, we consider a second specification—where we assume that households have an *ex-ante* belief about a sustainable rate of house price growth, and consider the extent to which actual house price appreciation since house purchase differs from this belief. Any 'unexpected' changes might be more likely to capture *innovations* to the level of life-cycle resources—caused by the wealth or common causality channels.[20] Such unexpected changes may also affect credit-constrained consumers. Although we place less weight on specifications where we separate homeowners from renters, we also test whether unexpected movements in regional house prices principally affect homeowners, as predicted by the collateral hypothesis, or whether they are also relevant for renters.

Representing these different regional house price channels with the generic function $g(hp^r)$ and interacting these functions with dummies that identify whether the household is in one of three age groups (young, middle-aged or older) gives:

$$\begin{aligned} \ln X_t^{ch} &= \alpha^c + f(age) + \gamma' z_t^{ch} + \theta_{1Y} g(hp_t^r) DY_t^{ch} \\ &\quad + \theta_{1M} g(hp_t^r) DM_t^{ch} + \theta_{1O} g(hp_t^r) DO_t^{ch} + \varepsilon_t^c + u_t^{ch} \end{aligned} \qquad (\mathbf{6})$$

where the dummies *DY*, *DM* and *DO* indicate the age group which the household belongs to. Note that the coefficients on house prices cannot be interpreted as the causal effect of house prices on

[19]Correlation between house price and consumption growth is 0.69 over our sample period, and between house price growth and the deviation of consumption from its exponential trend is 0.55.
[20]In principle there is an additional substitution effect on non-housing consumption, stemming from the fact that housing services have become relatively more expensive.

consumption, as they may also reflect common factors. Instead, we compare the coefficients across age groups, making the identifying assumption that if house prices capture the direct wealth effect, we would expect the coefficient to be larger for older consumers, who have higher rates of homeownership. On the other hand, if they capture differences in economic activity, and in particular expected future income, then we expect the effect to be greater for younger consumers, who stand to gain the most. For our specifications that compare the effect for homeowners and renters, these dummies are replaced by ones for tenure status.

4 Results

In this section we estimate the consumption-age profiles described above to address four main issues:

i. Was the level of consumer expenditure during the booms of the late 1980s and late 1990s/early 2000s higher than can be explained by the consumption-age profiles? Similarly, was the level of consumption in the early 1990s lower than expected given these profiles?
ii. Which age groups have driven the aggregate fluctuations in consumption growth?
iii. Can these fluctuations be explained once house prices are considered?
iv. What causes the relationship between house prices and consumption?

To address these questions we examine the coefficients from, and the fit of, our different specifications. The effects of changing household composition and other observable factors are removed from the life-cycle consumption profiles,[21] which would otherwise tend to give them a pronounced hump-shaped appearance (Attanasio (1999)). Regression outputs for each of the main specifications are given in the appendix.

4.1 Baseline specification – no house price terms

Chart 7 presents the predicted levels of log consumption from the consumption-age framework (equation (**5**)), together with actual outturns. The top left panel gives the profile for all the cohorts, with the remaining ones each showing a third of the results. It indicates that there are periods where the actual level of consumption deviates markedly from its predicted level. Close inspection reveals that this divergence is particularly pronounced among the younger households.

---

[21]The age profile is calculated from the predicted values with family size set to two adults and no children.

Chart 7: Consumption-age profiles – baseline predictions and actual for different cohorts[(a)]

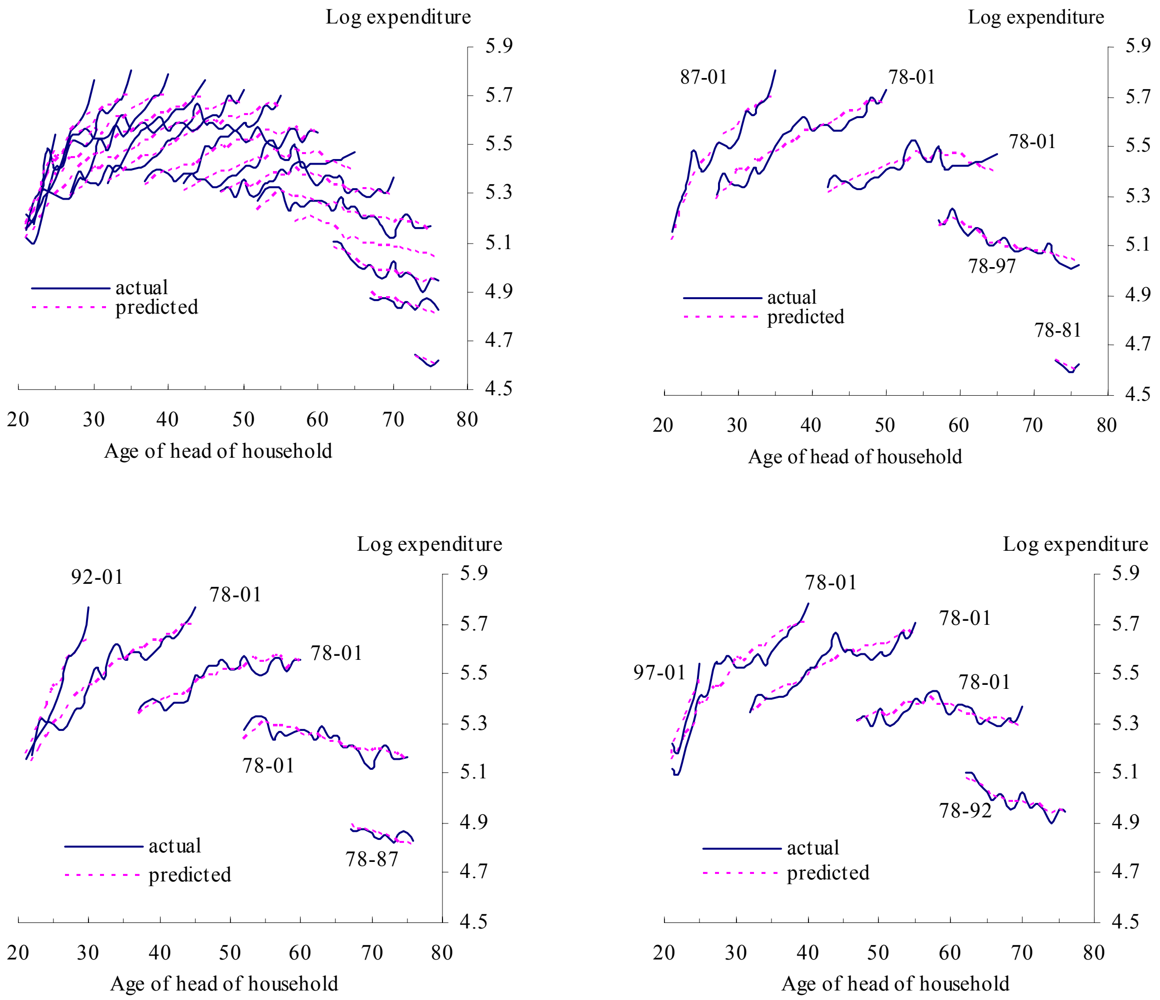

(a) Top left quadrant shows results for all cohorts. Remaining quadrants each show results for a third of cohorts. Dates give time period over which cohorts were observed.

Chart 8 summarises these results, giving the average (across all cohorts) percentage difference between actual and predicted consumption by year. The level of consumption was on average lower than predicted during the early 1980s, with the situation reversed later in the decade. After the mid-1990s dip, consumption was again above its predicted level at the beginning of the 2000s.[22] This analysis considers the gap between the actual and predicted *levels* of consumption, meaning that the peaks and troughs are slightly later than in an analysis based upon consumption *growth* discussed in the data section above. For example, consumption growth was below trend in 1990, but the gap

---

[22] The size of spike in 2001 may partly reflect the survey changing from the FES to the EFS. However, the general patterns of the results is unchanged by excluding 2001.

between the actual and predicted levels of consumption continued to be positive due to the previous years of strong consumption growth.

Chart 8: Gap between actual and predicted levels of consumption

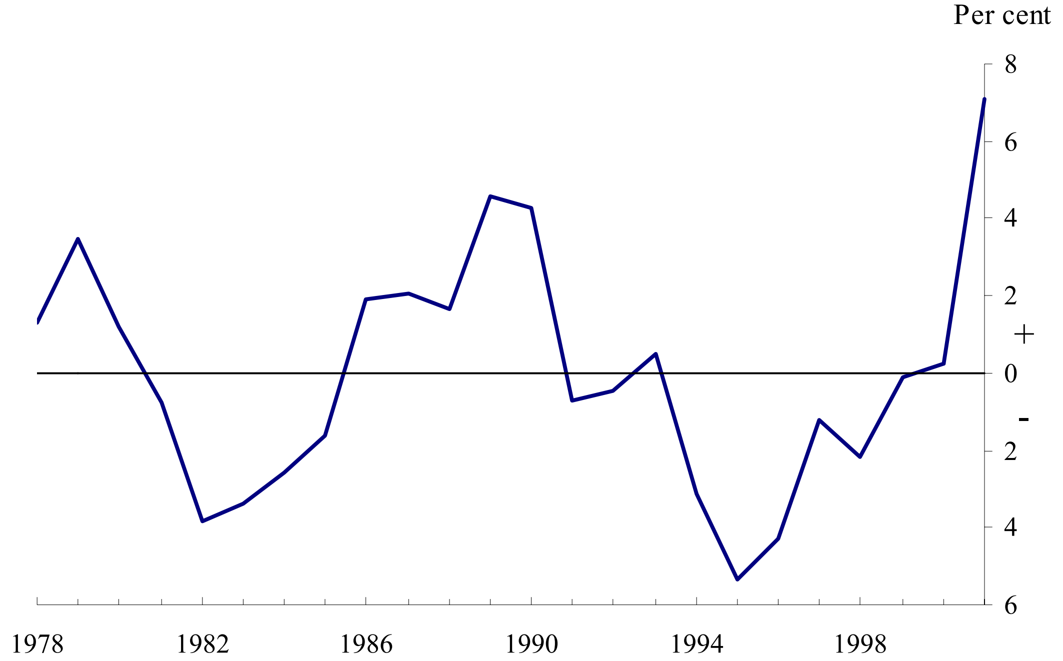

There are marked differences in the spending patterns of the different age groups (Chart 9). The gap between actual and predicted is much less volatile for the older than the younger cohorts, with the pattern for the middle aged between the two extremes. For example, the consumption of the young was almost 10% higher than predicted in the booms of the late 1980s and early 1990s/2000s, but there was little divergence for the older cohorts in the 1980s boom, and their consumption was only 5% greater than expected in 2001.

This evidence is consistent with the results reported in AW, who found that most of the consumption boom of the late 1980s was caused by the youngest cohorts. On that basis, they argued that the boom reflected a change in perceived permanent income, probably induced by a change in expected productivity growth, which would have a greater effect on younger households. This paper shows that younger households also reduced their consumption the most in the early 1990s, and were the primary drivers of the pickup in the late 1990s and early 2000s.

Chart 9: Gap between actual and predicted levels of consumption, by age group

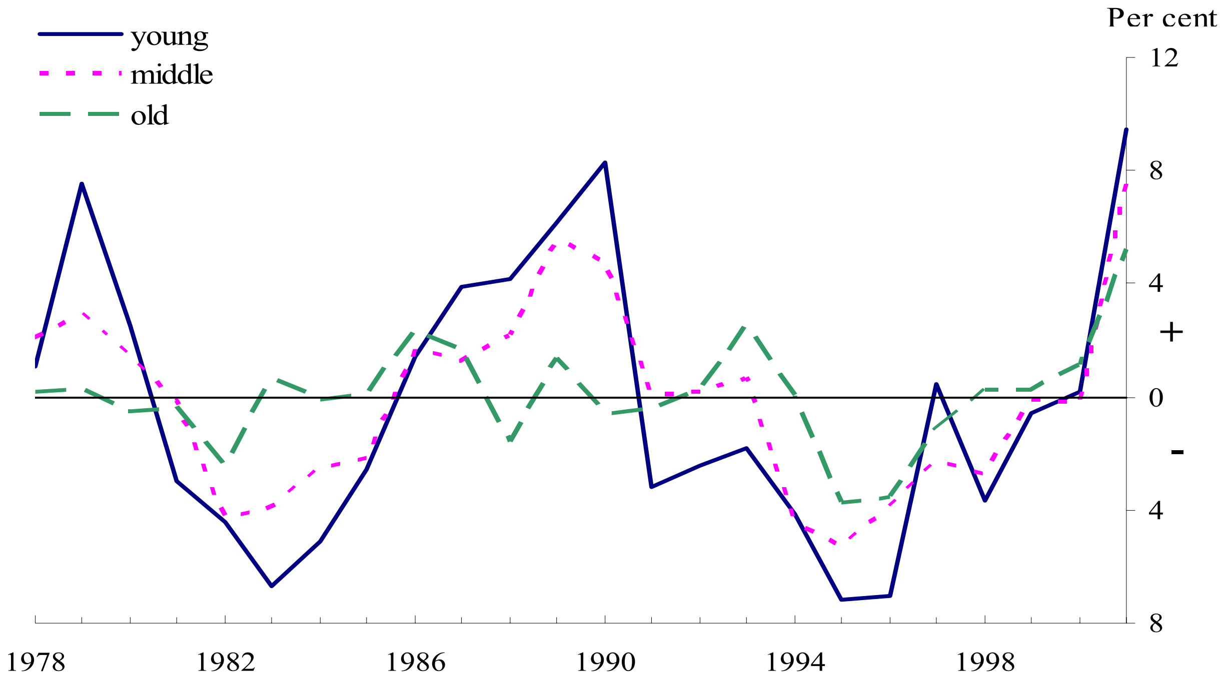

The remainder of Section 4 assesses the extent to which these patterns of consumption are correlated with house price developments. And in particular what the relative strength of the relationship for different types of households tells us about the relative importance of the wealth, collateral and common shock channels.

4.2 Investigating the role of house prices

House price growth

Following AW, we add the regional annual percentage change in real house prices to our baseline specification. We allow the coefficient to vary across the three age groups. The idea is that if changes in house prices are interpreted by households as a gain in wealth, they should be most relevant for homeowners (and particularly those looking to trade down), who will tend to be older. On the other hand, house price changes could be capturing innovations to productivity and income growth in a specific region, which we would expect to be more important for younger households.

The key results are summarised in Table A. The coefficient for the youngest group is quite large. On average, a 1 percentage point increase in annual house price growth is associated with a 0.21% increase in consumption. For the middle-aged group, the increase is smaller, at 0.13%, and for the oldest group smaller still at just 0.04% (indeed this is not significantly different from zero).

Therefore, while house price changes at the regional level seem to be significant 'explanatory' variables for consumption, the pattern of coefficients does not appear to offer support for the wealth hypothesis. We therefore do not interpret the relationship as causal. These results are in sharp contrast with those of Campbell and Cocco (2004), who found that a 1 percentage point increase in house price inflation caused consumption to increase by more than 1%, with even larger effects for the old.

Table A: Effect of a change in house prices on consumption

Age group	Coefficient	Standard Error
Young (21 - 34)	0.209*	0.029
Middle-aged (35 - 59)	0.127*	0.020
Old (60 - 75)	0.042	0.026

* Indicates significant at the 5% level.

In Chart 10 we plot the gap between actual and predicted consumption over time for the new specification, along with the gap for the baseline specification plotted in Chart 8. The difference between the two series is not very large. While it is true that the size of the residuals is most reduced over the late 1980s, when house prices were increasing rapidly, the magnitude of this effect is relatively small. We therefore conclude that, consistent with the results reported by AW, house price changes, while significant, do not explain the periods of consumption strength and weakness observed in the sample.

Chart 10: Actual minus predicted: baseline and including house price growth

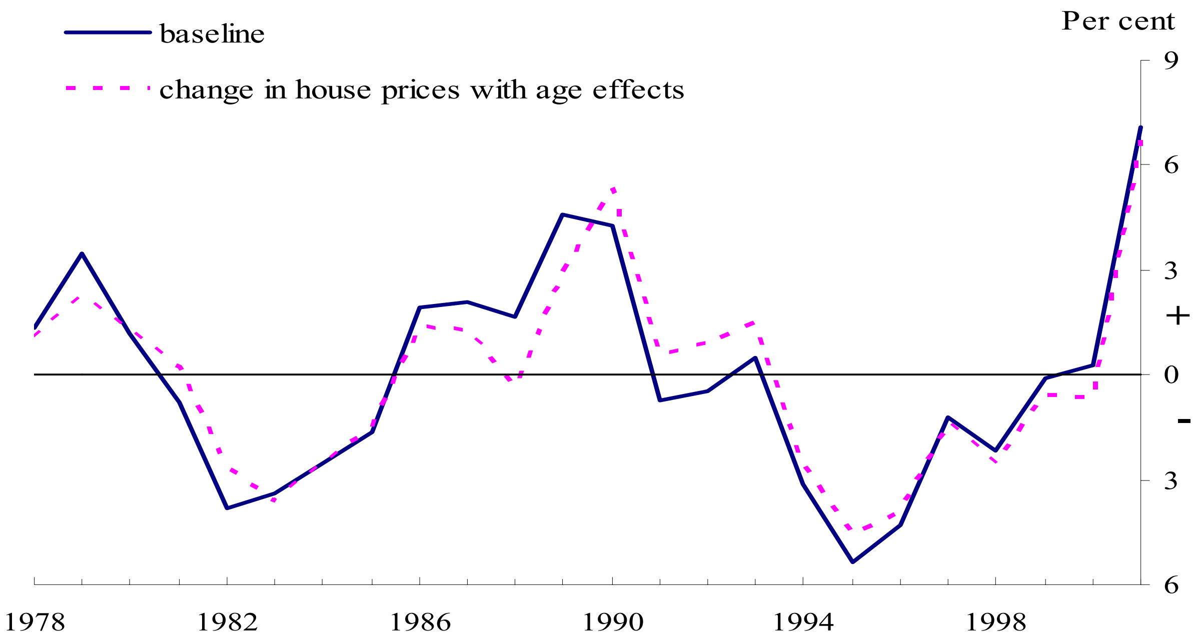

The larger coefficient on house prices for the younger cohort means that the inclusion of the change in house price term helps explain more of their unexplained pattern of consumption than it does for older cohorts (comparison of Chart 11 with Chart 9). However, a substantial part of their cyclicality remains, again suggesting that the peaks and troughs were primarily driven by other factors.

These results are clearly consistent with the common causality hypothesis, as younger households have more to gain from future productivity changes. But they are inconsistent with the wealth hypothesis being the main determinant of the consumption fluctuations. It is less clear whether they provide support either for or against the collateral explanation—the young are less likely to be homeowners, but if they were, they would be more likely to be credit constrained and thus want to borrow against any increase in housing equity.

Chart 11: Actual minus predicted: including the change of house prices by age group

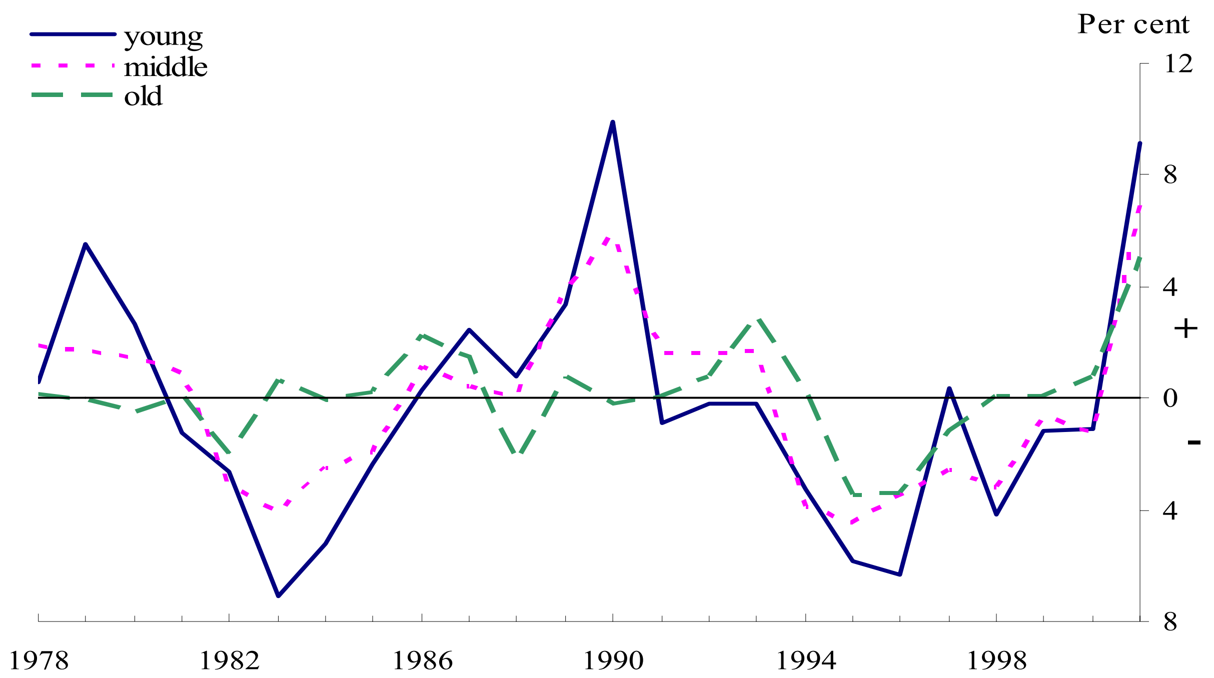

However, the results from an analysis by homeownership status—rather than age—provide less support for the collateral hypothesis. We find similar coefficients on the change in house price term for homeowners and renters (specification VII in Table A2 in the appendix),[23] suggesting that the common causality effect dominates any positive wealth or collateral effects that should benefit

---

[23]In fact, the coefficient for renters is slightly larger, which might reflect that they are more likely to be young.

homeowners.[24] This result is particularly remarkable as any sample selection and endogeneity problems are probably likely to bias the coefficient on renters downwards, and thus accentuate the gap between the two.[25] These results are consistent with, and provide further support for, the findings reported in Chart 5 above that the consumption of renters fluctuated by as much as that of homeowners during periods of consumption and house price booms.

The level of house prices

As discussed above, house prices could act as a proxy for a significant component of household lifetime resources. So instead of changes in house prices, we add their (log) *level* to equation (**5**). We use the level of regional house prices, rather than homeowners' estimate of the current value of their property, since no consistent data were available in the FES.[26] Again, we interact the price level with the three age groups.

Chart 12 shows that including the level of house prices leads to an apparent improvement of the overall fit, especially over the 1980s and 1990s. A substantial unexplained boom remains in 2001 though, as discussed in the data section, this could be due to data problems associated with the move from FES to the EFS survey in 2001.

In Chart A1 in the appendix, we plot the residuals for the three age groups considered previously. The improvement relative to the baseline specification arises because the younger and, particularly, the middle-aged cohorts are better predicted. The overall fit for the older cohorts seems much the same.

While it might be tempting to interpret the improvement in the fit as providing support for the wealth hypothesis, and an important role of house prices in determining consumption, it is likely that this variable is actually capturing other differences at regional level. An indication of this is the fact that the coefficients on the level of house prices (Table B) are remarkably similar for the three age groups we consider: under the wealth hypothesis we would have expected the effect to be much smaller for young consumers, both because the fraction of this group owning houses is smaller and because they are probably net buyers of houses over the rest of their life cycle (so that an increase in house prices would make them worse off). To further investigate the possibility that house prices capture differences in the economic performance of different regions we move to our next specification, in

---

[24]As discussed in Aoki *et al* (2001), it is possible that any collateral channel would be more significant since the mid-1980s, after financial liberalisation (see Attanasio and Weber (1994) for a discussion of liberalisation in the UK). However, we continue to find a similar role for renters when starting the analysis in 1985.
[25]Because the increase in homeownership in the UK has been correlated with the house price cycle, and it is likely that it is the better-off former renters that have moved into homeownership.
[26]An owner's estimate of their house price would anyway suffer from the problems of measurement error and that the price level might be correlated with unobserved heterogeneity and therefore bias the estimation results.

which we decompose changes in house prices into a term that can be explained by earnings and a residual term.

Chart 12: Actual minus predicted: including the level of house prices

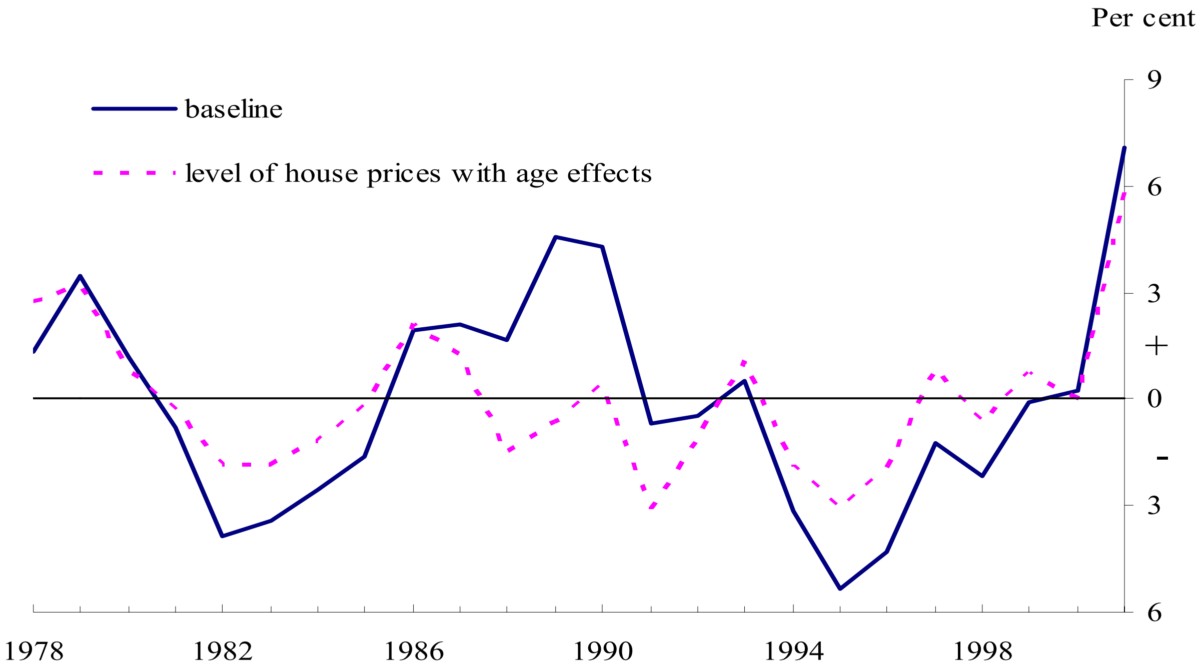

Table B: Coefficient on level of house prices

Age group	Coefficient	Standard Error
Young (21 - 34)	0.161*	0.009
Middle-aged (35 - 59)	0.163*	0.009
Old (60 - 75)	0.165*	0.009

* Indicates significant at the 5% level.

Predicted and unpredicted house prices

In this section, we consider whether it is 'expected' or 'unexpected' movements in house prices that provide the better explanation for consumption growth within our framework. We first describe how we predict an 'expected' level of regional house prices using regional income growth. By subtracting this from actual house prices we get a measure of 'unanticipated' house price changes, which we then add to our standard consumption specification.

The first step was to regress log real house prices (at a monthly regional level) against average log household income (also at a monthly regional level) and the real interest rate (derived from the monthly average of base rates). This captures the component of the house price level that is explained by regional trends in earnings and which might therefore proxy for developments in 'permanent income' when included in our specification. We also calculate the difference between the actual and the predicted levels of house prices, and include this difference in the specification. This 'unanticipated increase', beyond what might have been explained by incomes, may represent a better reflection of the influence of the wealth effect of house prices level on consumption. Table C reports the results of the first-stage regression.

Table C: First-stage regression results – predicting house prices

Dependent variable is log of real regional monthly house prices

	Coefficient	Standard Error[(a)]
Log average real regional monthly income	1.039 *	0.017
Real interest rate	0.007 *	0.001
Constant	5.106 *	0.097
R^2	0.67	
Observations	170,473[(b)]	

* Indicates significant at the 5% level.

(a) Standard errors are clustered by region and month.

(b) While most of the specifications in this paper are run from 1978 to allow use of educational dummies, this regression is run from 1975 to use as much data as possible in estimating the coefficients.

The second step was to include the expected and unanticipated measures of house prices within our consumption-age framework. Table D reports the results, with both components of house prices interacted with the age dummies.[(27)] The results are remarkable. The predicted variable—which we believe provides a proxy for permanent income—appears to have a similar effect on consumption for each of the three age groups, whereas the 'unanticipated' increase has a *greater* effect for the young than the old. This is inconsistent with the wealth hypothesis, where we would have expected a bigger effect for older households. Instead, it provides support for the common causality explanation for the link between consumption and house prices—perhaps because unexpected movements in house prices may contain news about households' expectations that are in addition to the current level of income.[(28)] Charts 13 and A2 show again the residuals from the actual level of consumption generated by this specification. In comparison with the baseline results, the residuals are much smaller, particularly in the late 1980s and early 1990s, although the fit is not much improved from that which we got from simply interacting the observed regional house price with age in the previous

---

(27) The standard errors of this regression are not adjusted for the presence of generated regressors.

(28) Possibilities include that it captures news about the expected growth rate of *future* income, or greater certainty about it.

specification. Of course, this is not really surprising since we have merely attempted to break down the effect of the level of house prices into an unanticipated and an anticipated component.

Table D: Coefficients on predicted level of house price within our consumption framework

Age group	Coefficient	Standard Error
Predicted: young	0.310*	0.011
Predicted: middle-aged	0.312*	0.011
Predicted: old	0.313*	0.011
Actual - predicted: young	0.13*	0.016
Actual - predicted: middle-aged	0.079*	0.012
Actual - predicted: old	0.033	0.015

* Indicates significant at the 5% level.

Chart 13: Residuals based on predicted house price and deviation from actual

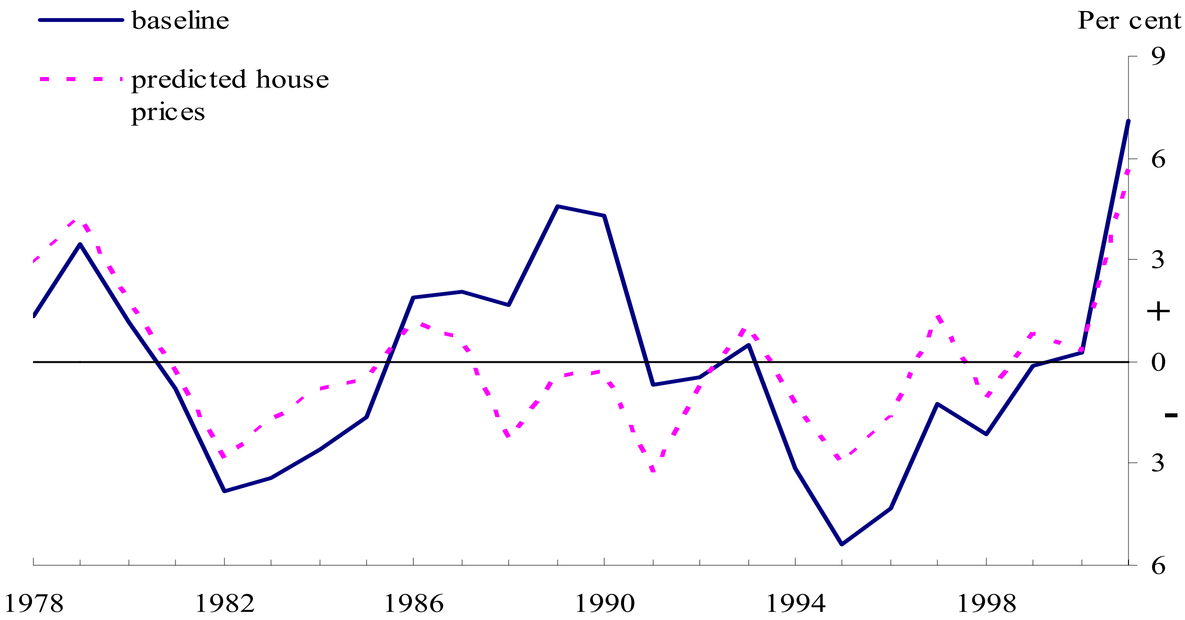

Unexpected movements in house prices since purchase date

An alternative approach to estimating unexpected movements in house prices is to consider evidence for a relationship between the degree of 'unexpected' house price *appreciation* since the date of house purchase and the level of consumption. At the time a household buys a home, they may have expectations about future house price growth. Establishing the extent of house price inflation above and beyond these expectations gives an indication of the innovations to household wealth since house purchase. The sample estimation period for this is from 1985, the date at which a question asking about the date on which households moved into their current residence was added to the FES. We assume the anticipated annual increase in real house prices is a constant 3% over a long horizon.[(29)] The extra term in this specification is therefore the gap between actual house prices and those that would be expected under this belief, which we interact with age group as before. For non-homeowners, this term is zero.

The results are very similar to those we got when predicting house prices using income, with the effects greater for the young than for the old. Charts A4 and A5 in the appendix repeat the residuals analysis. Unsurprisingly, they are similar to those we have already observed: some of the unexplained residual is now removed, but there remains a greater divergence for the young.

We also repeated the analysis, but including the unanticipated change in house prices for renters since they moved into their most recent property (Table E). The motivation behind this is to try to shed more light on the evidence for the collateral/credit-constraints hypothesis. Under this hypothesis, the coefficient on homeowners should be large and positive, but that on renters would be expected to be zero or negative, as an increase in house prices would leave them more credit constrained (if they would have to raise a larger deposit to purchase a house, or if rents move in line with house prices). The interpretation of the results is made more complicated by potential selection bias. This bias seems likely to increases the coefficient on homeowners relative to that on renters, as over time the better-off renters are likely to move into homeownership, meaning that the effect for renters should be considered as a lower bound and that for homeowners as an upper one. But, despite this difficulty, the results are illuminating. The coefficient on young renters is large and positive, while there is a similar age pattern for renters and homeowners: both findings provide support for the common causality channel over the collateral or wealth alternatives.

[(29)]Real house price growth of 3% was chosen as it would appear to reflect a reasonable sustainable long-run rate of growth. For example, it was the same as average real income growth over the sample period. Average annual real house price growth over the sample period 4.0%, and repeating the calculations for this value changes the coefficients to 0.157 for the young, 0.026 for the middle-aged and –0.198 for the old. An alternative assumption that households expected real house price growth of just 2% leads to coefficients of 0.213 for the young, 0.168 for the middle-aged and 0.332 for the old.

Table E: Coefficients on house price change since moved into current house (when included as explanatory variable in equation (6)).

Age group	Coefficient	Standard Error
Homeowner: young	0.195*	0.023
Homeowner: middle-aged	0.105*	0.014
Homeowner: old	0.083*	0.025
Renter: young	0.107*	0.030
Renter: middle-aged	0.056*	0.024
Renter: old	-0.062	0.031

* Indicates significant at the 5% level.

5 Conclusions

Consumer spending growth in the United Kingdom has undergone pronounced cycles over the past two decades. Over much of this time, these changes have mirrored fluctuations in house price growth. Three main competing theories for the close relationship between the two have been proposed in the literature. First, that an unexpected increase in house prices raises households' wealth, and thus their desired consumption (Muellbauer and Murphy (1990)). Second, that both are determined by revisions to expected future income (King (1990) and Pagano (1990)). And third, that house price gains increase housing collateral, which is particularly important to the young as they are more likely to be credit constrained (Aoki *et al* (2001)).

This paper exploits the fact that, although all the theories predict a co-movement between aggregate changes in house prices and consumption, they have different implications for consumption at the individual household level. Under the wealth and collateral hypotheses, an increase in the value of a homeowners' house increases the value of their wealth/collateral. But an increase in house prices might be expected to lead to a reduction in the expected net future wealth of non-owners (who are more likely to be the young) as the rent on their property is likely to ultimately increase in line with house prices. In contrast, under the expectations of future income hypothesis, we might expect a co-movement between house prices and consumption for both homeowners and renters. And this co-movement is likely to be strongest for younger households, as they have more remaining years of working life. An expectation of a higher level or rate of growth of future income could reflect a belief of stronger productivity or, for non-Ricardian consumers, an expectation of lower tax rates. A similar effect would also arise if there was a reduction in uncertainty, reducing households' effective discount rate.

The findings in this paper provide several pieces of evidence which suggest that common causality is the most significant explanation for the co-movements of house price and consumption growth. First, during the consumption and housing cycles, younger cohorts—who are less likely to own a property, and more affected by any revisions to future earnings—had the largest swings in consumption. As they would benefit least from capital gains on housing, this finding is the opposite to what would have been expected under the wealth hypothesis. Second, while house prices seem to be linked to consumption, the relationship between house price growth and consumption is found to be stronger for younger households and not even significant for the older ones. Moreover, house price growth at the regional level did little to improve our ability to explain the consumption cycle, in contrast with the strong aggregate link between house price and consumption growth. Third, the relationship with the regional level of house prices is as strong for young as for older consumers. Moreover, when we split the level of house prices into a part predicted by the level of earnings and a 'surprise' part, we find that the effect of the former is the same for all age groups, while the effect of the latter is stronger, once again, for young households. And fourth, even though homeownership is endogenously determined, the consumption of both homeowners and renters are equally aligned with the house price cycle.

Interestingly, the specifications that included the level rather than the growth of house prices seem to explain a large part of the difference between actual and predicted consumption, even for young consumers. This could be because house price levels are good predictors of future levels of income, or 'permanent income', as perceived by individuals. Revisions to income may be driven by changes in expectations about future productivity.

These results contrast sharply with those found by Campbell and Cocco (2004), who argue that there is a significant wealth effect from house prices to consumption. Although their paper also uses cohort analysis from the BHPS, there are three significant practical differences between the methodologies. First, their sample period started in 1988 rather than 1978. Second, they used a reduced-form regression for analysing cohorts' consumption growth, while we use the permanent income hypothesis as a rationale for considering the level of their consumption. And third, they use self-reported rather than regional house prices, which in the absence of a true panel introduces potential sample selection issues. There are also significant differences between the results. They argue that a 1% increase in house prices leads to 1.2% increase in consumption through a wealth effect, with the estimated elasticity for older homeowners even higher at 1.7%. This is much stronger than the 0.15% relationship found at the aggregate level. In contrast, we find that a 1% increase in house prices is associated with consumption being between 0.21% and 0.04% higher, depending on the age group—although we do not argue that this reflects a causal wealth effect. Campbell and Cocco also find a role for house price increases reducing credit constraints although, somewhat counter-intuitively, this occurs through national rather than regional house price movements. Of course, it is likely that the wealth and collateral channels are important for some households at some points in time. But, on average over the past twenty five years, this paper concludes that these effects have been smaller than the common causality channel.

Appendix

Table A1: Real regional house price increases, 1975 – 2001

	North	Yorks & Humber	N. West	E. Midlands	W. Midlands	E. Anglia	Gtr. London	S. East	S. West	Wales	Scotland
1975	-10.6	-12.0	-14.7	-14.0	-15.1	-17.1	-18.2	-16.6	-17.7	-13.3	-9.4
1976	-4.4	-4.9	-5.9	-7.5	-7.2	-9.1	-8.1	-8.8	-8.7	-5.5	2.3
1977	-2.1	-7.4	-5.6	-8.1	-6.6	-10.6	-7.8	-8.4	-8.4	-9.5	-5.9
1978	3.1	3.7	7.0	4.3	6.5	6.2	9.8	8.1	6.4	6.7	3.9
1979	8.7	11.4	14.8	12.4	13.7	14.5	18.9	18.2	17.1	13.4	6.3
1980	1.5	5.2	5.2	6.1	1.9	8.4	4.7	4.0	7.2	2.5	-1.5
1981	-3.6	1.2	-4.3	-2.7	-5.9	-5.0	-7.0	-5.8	-6.9	-2.2	-1.4
1982	-3.4	-6.9	-4.9	-4.4	-5.8	-4.1	-6.0	-5.9	-4.0	-2.2	-1.2
1983	6.0	7.3	5.0	5.8	3.3	3.9	7.5	7.8	5.2	4.8	6.9
1984	1.3	2.0	1.5	4.9	2.6	5.6	11.4	8.7	3.2	-0.3	1.8
1985	-1.7	1.0	-0.1	2.5	0.0	7.0	8.6	5.2	5.1	2.6	0.8
1986	5.0	5.9	6.6	6.5	6.1	12.9	18.2	13.9	11.6	3.3	2.4
1987	5.5	4.9	3.0	9.3	10.3	15.3	20.8	18.7	13.6	7.3	3.7
1988	7.4	13.1	12.5	23.2	28.3	30.6	18.2	24.7	28.8	15.6	3.7
1989	22.3	32.5	27.5	23.6	24.5	11.0	3.1	9.7	14.9	29.0	12.1
1990	6.1	1.8	7.7	-6.3	-4.9	-16.6	-8.2	-14.4	-15.3	-2.5	4.2
1991	-6.5	-3.3	-5.8	-9.4	-7.1	-9.7	-10.8	-12.0	-10.1	-8.4	0.7
1992	-0.4	-6.1	-3.3	-6.4	-7.2	-9.6	-14.2	-11.2	-9.8	-5.6	-0.3
1993	-0.9	-2.0	-6.7	-5.7	-4.3	-7.7	-4.5	-6.9	-6.9	-1.9	0.7
1994	2.1	-3.8	-0.6	-0.1	-1.8	-0.9	2.8	1.0	0.8	-1.0	-1.2
1995	-7.6	-2.4	-3.5	-2.5	0.0	-0.1	-1.7	-1.8	-1.7	-4.5	-1.6
1996	0.8	0.3	-1.3	2.7	0.1	-2.1	0.1	1.9	1.4	1.7	0.3
1997	4.2	2.9	4.6	1.9	3.6	6.8	12.3	8.6	5.7	3.3	3.3
1998	2.0	1.2	2.0	6.8	5.4	10.5	12.3	12.3	8.7	2.5	3.4
1999	6.2	4.3	5.1	2.1	6.0	3.9	21.3	9.7	9.5	6.1	0.5
2000	2.7	3.4	5.7	9.3	11.3	17.8	16.3	17.6	14.7	3.9	1.5
2001	3.3	5.9	8.8	9.2	7.2	7.0	8.9	9.5	11.4	10.2	2.7

Source: ODPM.

Table A2: Regression results

Dependent variable is log of real non-housing household expenditure

	I	II	III	IV	V	VI	VII
Constant	0.913	1.041 *	-1.396 *	-2.654 *	1.678 *	1.610*	1.532*
	(0.492)	(0.495)	(0.585)	(0.587)	(0.681)	(0.481)	(0.481)
Dummy: cohort 2	0.187 *	0.187 *	0.164 *	0.138 *	-	0.177*	0.180*
	(0.011)	(0.011)	(0.011)	(0.011)		(0.011)	(0.011)
Dummy: cohort 3	0.295 *	0.296 *	0.234 *	0.194 *	0.089 *	0.279*	0.282*
	(0.011)	(0.011)	(0.012)	(0.012)	(0.021)	(0.011)	(0.011)
Dummy: cohort 4	0.392 *	0.393 *	0.311 *	0.257 *	0.189 *	0.360*	0.365*
	(0.011)	(0.011)	(0.012)	(0.012)	(0.021)	(0.011)	(0.011)
Dummy: cohort 5	0.492 *	0.494 *	0.390 *	0.322 *	0.307 *	0.441*	0.446*
	(0.011)	(0.011)	(0.013)	(0.013)	(0.021)	(0.011)	(0.011)
Dummy: cohort 6	0.586 *	0.587 *	0.460 *	0.377 *	0.417 *	0.512*	0.518*
	(0.012)	(0.012)	(0.014)	(0.014)	(0.023)	(0.012)	(0.012)
Dummy: cohort 7	0.672 *	0.676 *	0.524 *	0.425 *	0.515 *	0.583*	0.589*
	(0.012)	(0.012)	(0.015)	(0.016)	(0.024)	(0.012)	(0.012)
Dummy: cohort 8	0.742 *	0.745 *	0.570 *	0.456 *	0.602 *	0.642*	0.648*
	(0.013)	(0.013)	(0.016)	(0.017)	(0.025)	(0.013)	(0.013)
Dummy: cohort 9	0.826 *	0.829 *	0.631 *	0.501 *	0.686 *	0.718*	0.724*
	(0.013)	(0.013)	(0.017)	(0.018)	(0.025)	(0.013)	(0.013)
Dummy: cohort 10	0.867 *	0.872 *	0.650 *	0.506 *	0.733 *	0.763*	0.769*
	(0.013)	(0.014)	(0.018)	(0.019)	(0.026)	(0.013)	(0.013)
Dummy: cohort 11	0.965 *	0.969 *	0.725 *	0.564 *	0.826 *	0.852*	0.859*
	(0.014)	(0.014)	(0.019)	(0.020)	(0.027)	(0.014)	(0.014)
Dummy: cohort 12	1.067 *	1.073 *	0.803 *	0.625 *	0.935 *	0.948*	0.955*
	(0.014)	(0.014)	(0.020)	(0.022)	(0.027)	(0.014)	(0.014)
Dummy: cohort 13	1.128 *	1.134 *	0.845 *	0.652 *	0.997 *	1.015*	1.023*
	(0.015)	(0.015)	(0.022)	(0.023)	(0.028)	(0.015)	(0.015)
Dummy: cohort 14	1.156 *	1.161 *	0.859 *	0.652 *	1.025 *	1.052*	1.059*
	(0.017)	(0.017)	(0.023)	(0.025)	(0.029)	(0.016)	(0.016)
Dummy: cohort 15	1.171 *	1.168 *	0.847 *	0.622 *	1.039 *	1.079*	1.081*
	(0.022)	(0.022)	(0.028)	(0.030)	(0.033)	(0.022)	(0.022)
Age	0.252 *	0.235 *	0.364 *	0.350 *	0.115	0.200*	0.208*
	(0.058)	(0.058)	(0.069)	(0.069)	(0.081)	(0.057)	(0.057)
Age Squared	-0.008 *	-0.008 *	-0.014 *	-0.013 *	-0.002	-0.007*	-0.008*
	(0.003)	(0.003)	(0.003)	(0.003)	(0.004)	(0.003)	(0.003)
Age Cubed	0.000 *	0.000 *	0.000 *	0.000 *	0.000	0.000*	0.000*
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Age ^ 4	-0.000	-0.000	-0.000 *	-0.000 *	-0.000	0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Age ^ 5	0.000	0.000	0.000 *	0.000 *	0.000	0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Kids aged 0	-0.016	-0.016 *	-0.019 *	-0.020 *	-0.023 *	-0.014*	-0.014*
	(0.008)	(0.008)	(0.008)	(0.008)	(0.009)	(0.008)	(0.008)
Kids aged 1	-0.022 *	-0.022 *	-0.020 *	-0.019 *	0.006 *	-0.017*	-0.017*
	(0.006)	(0.006)	(0.006)	(0.006)	(0.009)	(0.006)	(0.006)
Kids aged 2	-0.003	-0.003	-0.002	-0.002	0.007	0.000	0.001
	(0.007)	(0.007)	(0.007)	(0.007)	(0.009)	(0.007)	(0.007)
Kids aged 3 – 4	0.014 *	0.014 *	0.014 *	0.015 *	0.018 *	0.020*	0.020*
	(0.005)	(0.005)	(0.005)	(0.005)	(0.006)	(0.005)	(0.005)
Kids aged 5 – 10	0.041 *	0.041 *	0.041 *	0.042 *	0.042 *	0.051*	0.051*
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
Kids aged 11 – 16	0.085 *	0.085 *	0.085 *	0.085 *	0.085 *	0.091*	0.091*
	(0.003)	(0.003)	(0.003)	(0.003)	(0.004)	(0.003)	(0.003)
Kids aged 17 – 18	0.258 *	0.259 *	0.260 *	0.259 *	0.271 *	0.245*	0.245*
	(0.009)	(0.009)	(0.009)	(0.009)	(0.012)	(0.009)	(0.009)
Number of adults	0.558 *	0.558 *	0.558 *	0.555 *	0.582 *	0.524*	0.524*
	(0.003)	(0.003)	(0.003)	(0.003)	(0.004)	(0.003)	(0.003)

	I	II	III	IV	V	VI	VII
Dummy: more than two adults	-0.270 * (0.006)	-0.270 * (0.006)	-0.270 * (0.006)	-0.269 * (0.006)	-0.329 * (0.009)	-0.220* (0.006)	-0.220* (0.006)
Dummy: A Levels	0.154 * (0.003)	0.154 * (0.003)	0.154 * (0.003)	0.152 * (0.003)	0.158 * (0.004)	0.120* (0.003)	0.119* (0.003)
Dummy: degree	0.275 * (0.005)	0.274 * (0.005)	0.274 * (0.005)	0.270 * (0.005)	0.258 * (0.006)	0.237* (0.005)	0.237* (0.005)
Percentage change in house price: young	-	0.209 * (0.028)	-	-	-	-	-
Percentage change in house price: middle-aged	-	0.127 * (0.020)	-	-	-	-	-
Percentage change in house price: old	-	0.042 (0.026)	-	-	-	-	-
Log real house price level: young	-	-	0.161 * (0.009)	-	-	-	-
Log real house price level: middle-aged	-	-	0.163* (0.009)	-	-	-	-
Log real house price level: old	-	-	0.165* (0.009)	-	-	-	-
Predicted log real house price level: young	-	-	-	0.310 * (0.011)	-	-	-
Predicted log real house price level: middle-aged	-	-	-	0.312 * (0.011)	-	-	-
Predicted log real house price level: old	-	-	-	0.313 * (0.011)	-	-	-
Difference between actual and predicted: young	-	-	-	0.130 * (0.016)	-	-	-
Difference between actual and predicted: middle-aged	-	-	-	0.079 * (0.012)	-	-	-
Difference between actual and predicted: old	-	-	-	0.033 * (0.015)	-	-	-
Appreciation of house price since purchase: young	-	-	-	-	0.192 * (0.023)	-	-
Appreciation of house price since purchase: middle-age	-	-	-	-	0.111 * (0.013)	-	-
Appreciation of house price since purchase: old	-	-	-	-	0.098 * (0.025)	-	-
Homeownership dummy	-	-	-	-	-	0.255* (0.004)	0.256* (0.004)
Mortgage dummy	-	-	-	-	-	-0.005 (0.004)	-0.005 (0.004)
Percentage change in house price: homeowner	-	-	-	-	-	-	0.116* (0.017)
Percentage change in house price: renter	-	-	-	-	-	-	0.133* (0.023)
R^2	0.51	0.51	0.51	0.52	0.50	0.53	0.53
Observations	149,484	149,484	149,484	149,484	87,697	149,484	149,484

Note: Dummies for occupation, region and month included as standard in all specifications but not reported. In specification V, since we can only use data from 1985, the sample size is reduced. Further, by 1985 everyone from cohort 1 has been removed from the sample as they are all over 76 years old; hence cohort 2 becomes the comparison group for this specification and is omitted.
* = significant at 5% level. Standard errors in parentheses.

Specification I: No housing terms.
Specification II: Annual percentage change in house price interacted with age.
Specification III: (Log) level of real house price interacted with age.
Specification IV: Predicted level of real house price and difference between actual and predicted, interacted with age.
Specification V: Appreciation of house price since purchase (data since 1985 only) interacted with age.
Specification VI: No housing terms, but inclusion of a dummy for homeowners and renters.
Specification VII: Annual percentage change in house price interacted with homeownership status.

Chart A1: Actual minus predicted: including the level of house prices, by age

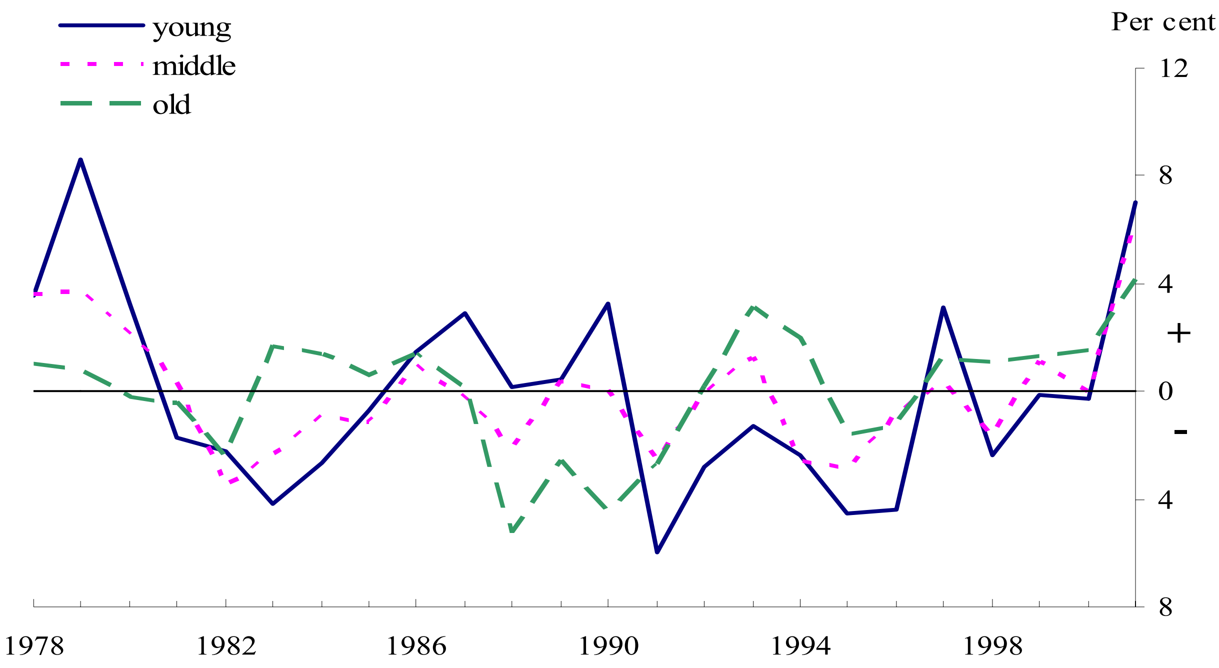

Chart A2: Residuals based on predicted house price and deviation from actual, by age

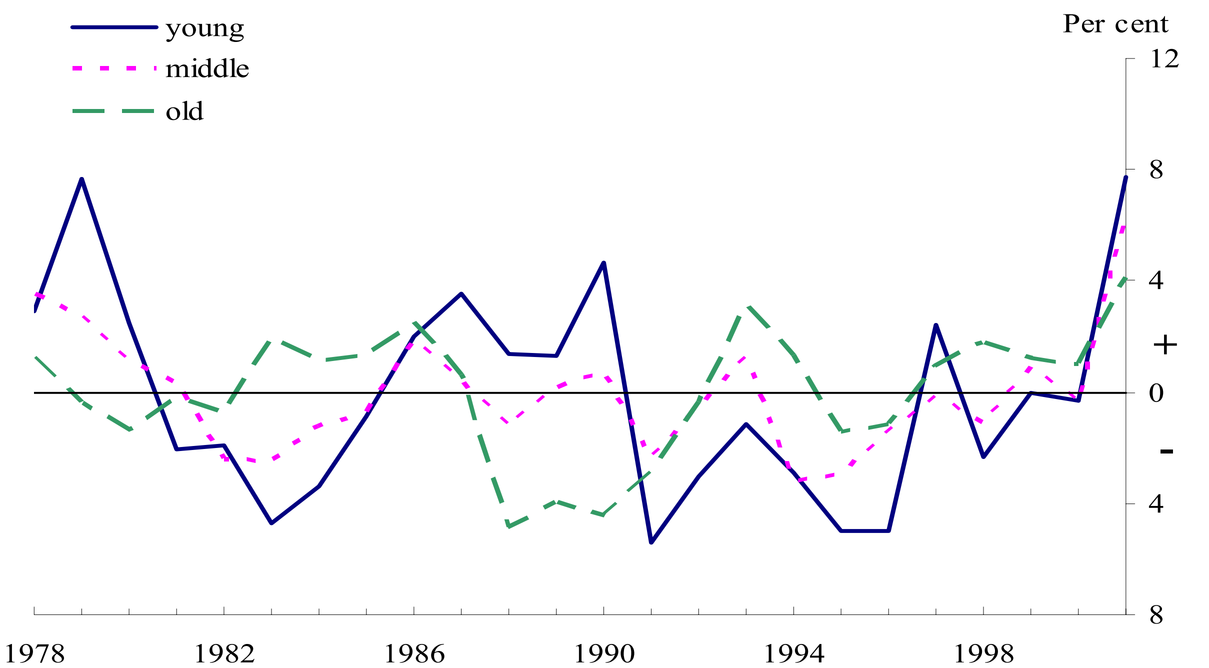

Chart A3: Actual log house prices and those predicted by incomes, by region

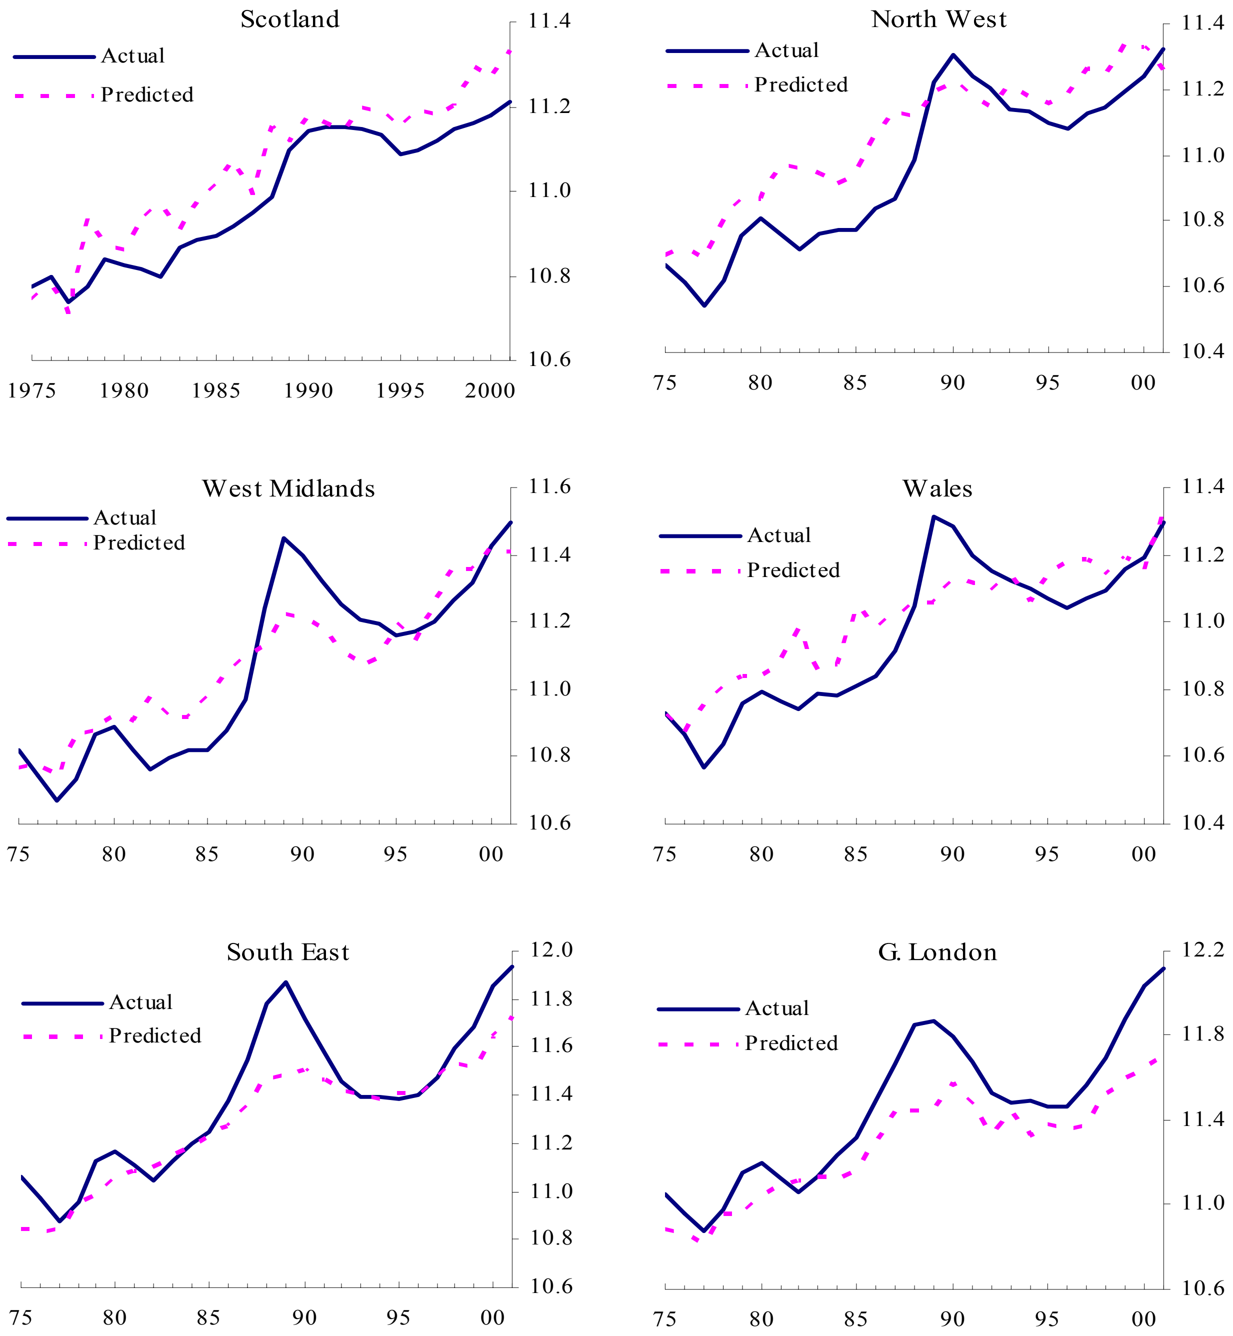

Chart A4: Residuals based on unanticipated increase in house price since purchase

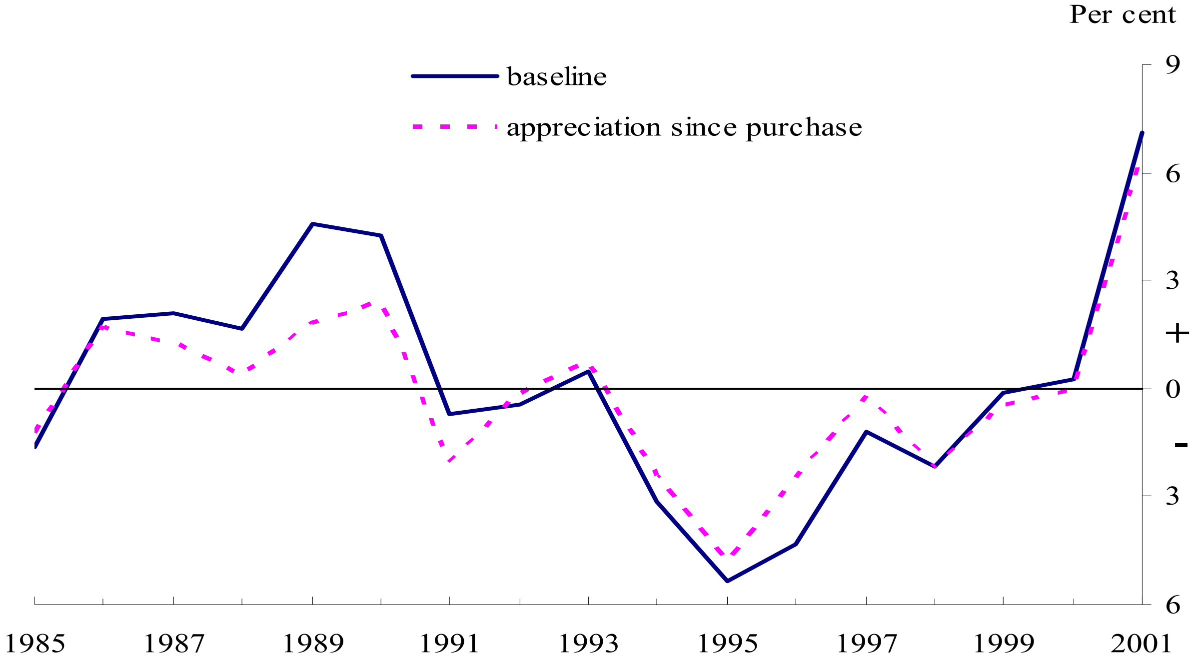

Chart A5: Residuals based on unanticipated increase in house price since purchase, by age

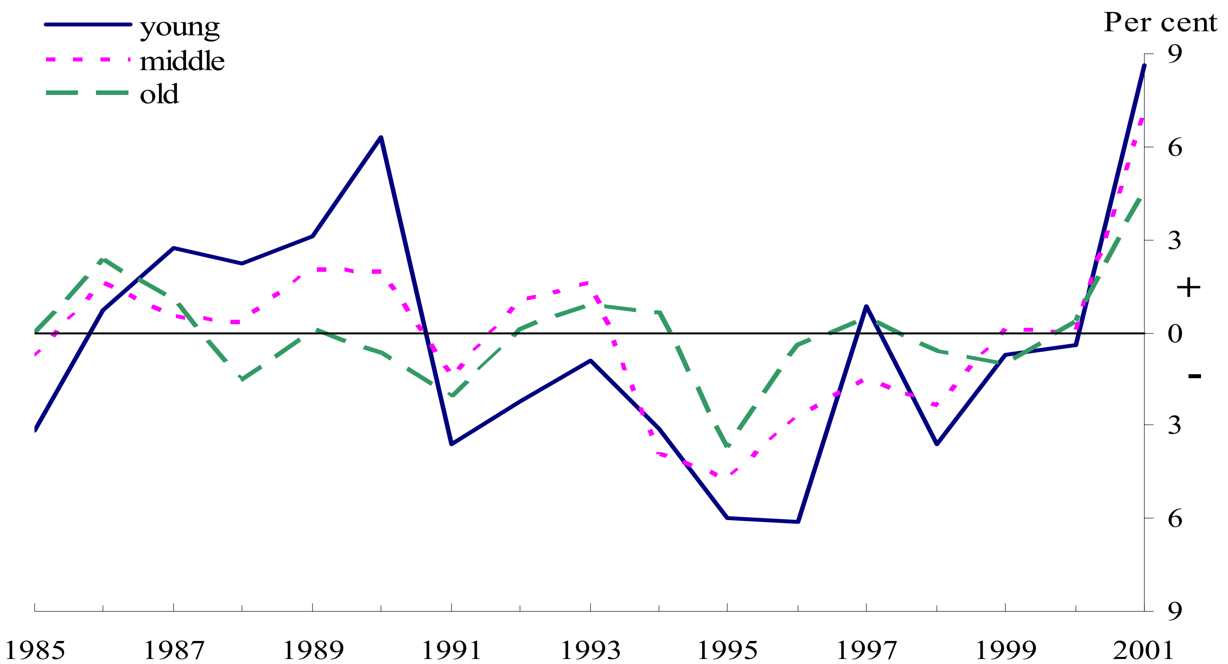

References

Aoki, K, Proudman, J and Vlieghe, G (2001), 'Why house prices matter', *Bank of England Quarterly Bulletin*, Winter, pages 460–68.

Attanasio, O (1999), 'Consumption', in Taylor, J and Woodford, M (eds), *Handbook of macroeconomics*, *Vol.1b.*

Attanasio, O and Weber, G (1994), 'The UK consumption boom of the late 1980s: aggregate implications of microeconomic evidence', *Economic Journal*, Vol. 104, Issue 427 (November), pages 1,269–302.

Bank of England, *Inflation Repor*t, various editions.

Bank of England (2004), *Inflation Repor*t, November, pages 12-13.

Barr, C and Buckley, G (2004), 'UK housing market special: creaking housing', Deutsche Bank Global Markets Research.

Blow, L, Leicester, A and Oldfield, Z (2004), 'Consumption trends in the UK, 1975–99', London: Institute for Fiscal Studies.

Browning, M, Deaton, A and Irish, M (1985), 'A profitable approach to labour supply and commodity demands over the life-cycle', *Econometrica*, Vol. 53, No. 3 (May), pages 503–44.

Campbell, J and Cocco, J (2004), 'How do house prices affect consumption? Evidence from the micro data', unpublished.

Case, K, Quigley, J and Shiller, J (2003), 'Comparing wealth effects: the stock market versus the housing market', unpublished.

Deaton, A (1985), 'Panel data from time series of cross sections', *Journal of Econometrics*, Vol. 30, pages 109–26.

King, M (1990), 'Discussion', *Economic Policy*, Vol. 11, pages 383–87.

Ludwig, and Slock, T (2004), 'The relationship between stock prices, house prices and consumption in OECD countries', *Topics in Macroeconomics*, Vol. 4, Issue 1, Article 4.

Muellbauer, J and Murphy, A (1990), 'Is the UK balance of payments sustainable?', *Economic Policy*, Vol. 11, pages 345–83.

Pagano, M (1990), 'Discussion', *Economic Policy*, Vol. 11, pages 387–90.

Sinai, T and Soules, S (2003), ‘Owner-occupied housing as a hedge against rent risk’, *NBER Working Paper no. 9462*, National Bureau of Economic Research.

Tanner, S (1998), ‘How much do consumers spend? Comparing the FES and National Accounts’, in Banks, J and Johnson, P (eds), *How reliable is the Family Expenditure Survey?*, London: Institute for Fiscal Studies.