The Coincident Cycles of House Prices and Consumption in the U.K.: Do House Prices Drive Consumption Growth?*

Orazio Attanasio

(Institute for Fiscal Studies, University College, London, NBER and CEPR)

Andrew Leicester

(Institute for Fiscal Studies and University College, London)

and

Matthew Wakefield[†]

(Institute for Fiscal Studies, and University College, London)

January 2008

Abstract

Three mechanisms have been suggested to explain the strong correlation between house prices and consumption in the UK: *wealth effects* and *collateral effects* involve house prices driving consumption changes, but *common causes* may drive both house price and consumption growth simultaneously. Our contribution lies in our systematic use of a realistic structural lifecycle model of consumption and housing decisions, to help us understand how data might distinguish between these mechanisms. We carefully match our model to recent episodes in the UK housing market and compare analyses of these episodes to analyses of counterfactual scenarios in which the mechanisms identified above are shut down. We show that our model can provide a good match to U.K. data on home-ownership and aggregate consumption growth, and provide a firmer theoretical footing for past studies that have argued that since the observed correlation between house price growth and consumption is particularly strong for the young, it is unlikely that wealth effects have been the key factor driving this correlation.

Keywords: Consumption, house prices, lifecycle models

JEL codes: D91 E21

The Coincident Cycles of House Prices and Consumption in the U.K.: Do House Prices Drive Consumption Growth?

1   Introduction

Over much of the past thirty-five years, the cycles of house prices and consumption have been relatively closely synchronised in the UK. Figure 1 shows the remarkable comovements of house price growth and consumption growth (notice however that the scale of house price movements is two and half time that of consumption growth) . The correlation between house prices and consumption makes changes in house prices an important indicator for those wishing to judge inflationary pressures within the economy. Indeed, this indicator is closely watched by the Bank of England's Monetary Policy Committee precisely because, in the words of Nickell (2004), ". . . The evidence suggests that house price inflation is significantly related to household consumption growth and hence to aggregate demand growth and future consumer price inflation in the economy". The perceived importance of this indicator is also illustrated by recent press discussion of what could happen if house prices were to decline or stagnate from current levels after an extended period of real growth (see Economist, 2005, Financial Times, 2008).
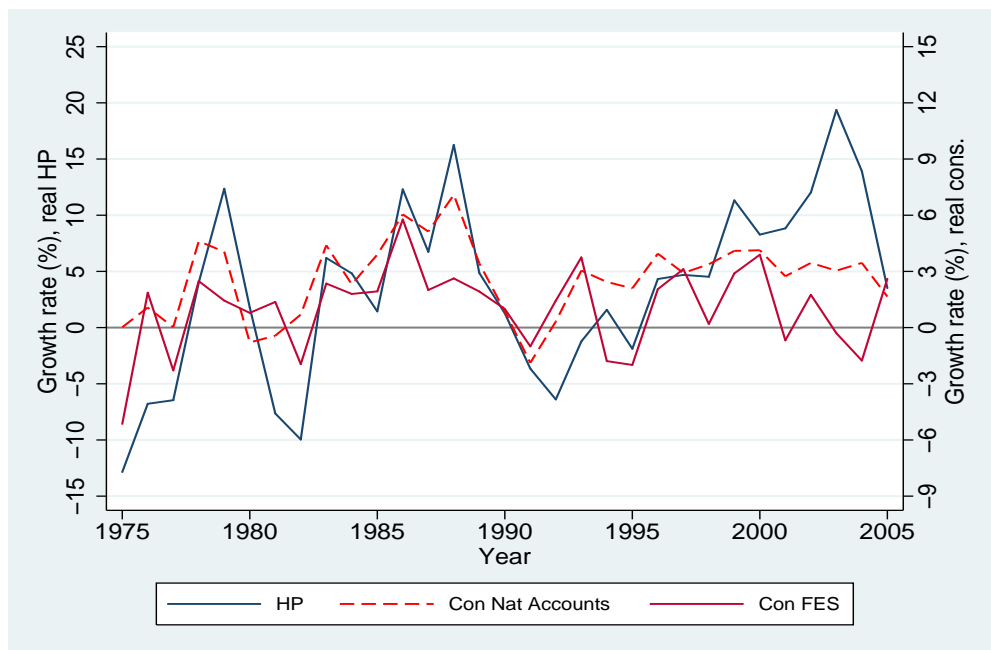

Figure 1: Growth in real consumption and house prices

To understand the implications of house price movements for consumption levels in the economy, and therefore the appropriate policy response to conditions in the housing market, it is necessary to understand what drives the link between house prices and consumption. Three main mechanisms to explain this co-movement have been proposed in the literature:

1. A "wealth effect": increases in house prices raise households' wealth, raising their desired level of expenditure.

2. A "credit constraints" channel: house price growth increases the collateral available to home-owners, thus loosening borrowing constraints and facilitating higher consumption (see Muellbauer and Murphy (1990) for a discussion of mechanisms 1. and 2.).

3. A "common causality" model: factors such as changes in expected income growth, tax changes or changes in credit market conditions lead to increases in both households' expenditure and house prices (see King (1990) and Pagano (1990)).

It has been argued that these different mechanisms might have effects that are observationally equivalent if one looks only at macro-data, and so researchers must use data on the behaviour of individuals or households in order to disentangle the different effects. Attanasio and Weber (AW, 1994) used such micro-data in order to examine the consumption boom of the late nineteen-eighties, which coincided with a rise on real house prices of slightly more than 40% in the space of four years. In order to identify what features of the micro-data they should look for, these authors drew on insights from a life-cycle model. AW (1994) argued that if wealth effects were important, then these were likely to have the biggest effects on the consumption of older individuals who are most likely to have equity in any housing assets, and who have a relatively short time horizon over which to distribute the consumption of the extra wealth. They also simulated a simple lifecycle setup (with no housing asset) and showed that an upward revision of expectations of future income (productivity) could lead to the consumption of the young (with longer to enjoy the higher income stream) responding more strongly than that of the old. The finding that the consumption boom was driven in a large part by strong consumption by the young, was therefore a key component of an argument that this boom, and the simultaneous boom in house prices, may have owed much to common causes. AW's exercise was recently extended in a variety of ways by Attanasio, Blow, Hamilton and Leicester (2007, hereafter ABHL), who, by and large confirmed the results in AW

(1994).

The complicated nature of the relationship between housing assets and non-durable consumption means that it may be perilous to rely on a stylised model for insights. AW themselves acknowledged this when mentioning that the consideration of credit constraints, and modeling of housing decisions and house price booms, would be interesting extensions that were beyond the scope of their model. The purpose of this paper is to derive the implications of changes to house prices and earning innovations for the consumption of different groups in the population. We can therefore check in a rigorous fashion whether such a structural model confirms intuitions about how and why the consumption of different groups might correlate with house prices, which have formed the basis of past empirical tests. We will also explore whether our model can be used to add to our understanding of what has driven consumption growth during fluctuations over the last 35 years in the U.K. economy, and by performing counterfactual simulations attempt to quantify the effects of different factors that move consumption.

Our central contribution comes from our systematic use of a realistic structural lifecycle model of consumption, savings and housing choices, to inform the interpretation of empirical analyses intended to distinguish between the three mechanisms proposed to explain the correlation of house prices and consumption. The first part of the exercise that we undertake involves constructing and numerically solving the structural model which is the main tool of our analysis, and calibrating this model to match U.K. data. The model, which we describe in detail in section 3, includes some innovative features. We model the financial markets available to the agents in our model to be a realistic representation of the UK mortgage market. We calibrate the stochastic processes faced by our agents to include both idiosyncratic and aggregate components. The former are calibrated using micro data. The latter include *aggregate* shocks to house prices and incomes (and as such are experienced by everyone in the economy at the same point in time). We estimate the parameters of the time series processes for house prices and cohort level earnings from actual time series data on house prices and individual earnings aggregated at the cohort level covering the last 35 years in the UK. We are not aware of other studies that use this combination of aggregate and micro data in the calibration of a individual level model and yet it is, in our opinion, important given that we want to understand aggregate fluctuations by aggregating individual consumption.

We choose the parameters of our individual model (such as the preference for housing services)

4

so to reproduce some cohort level facts, such as the level of home ownership and its evolution over the life cycle. Having successfully matched these moments, we can simulate individual behaviour given the set of aggregate innovations to house prices we estimate on our time series data and check the extent to which our model is able to reproduce the features of aggregate consumption growth in the UK. Notice that the moments of aggregate data are *not* used to calibrate our individual level model.

The next step of our analysis is then to simulate behaviour under a set of counterfactual scenarios in which the mechanisms that might drive the link between house prices and consumption are shut down in turn. The construction of these counterfactuals is, by definition, an exercise that cannot be undertaken using data and (as explained in section 4) forms our main means for exploring how we might disentangle the influence of the different mechanisms that might drive the link between house price shocks and consumption growth.

Our results (see section 4) provide a firmer theoretical grounding for reduced form empirical analyses of the type conducted by AW and ABHL. We show that in a model with credit constraints and simultaneous housing and consumption choices, a house price shock that drives consumption changes through a wealth effect will lead to the biggest consumption responses from older groups. In addition to this basic intuition, we can quantify the size of these effects with a model we show fits the data in a number of dimensions. As documented in AW and ABHL (whose results are described in more detail in section 2), this age pattern of responses is the reverse of that observed in data. Our model also shows that the pattern in the data could instead be explained by a shock to aggregate incomes of the kind that has been suggested as a possible common cause of house price and consumption growth.

The structure of the paper is as follows. Section 2 provides a fuller description of evidence which points towards the conclusion that factors other than wealth effects have driven the recent correlation between house price shocks and consumption growth in the U.K. Section 3 describes the lifecycle model that we use, and how we calibrated this model to match recent U.K. history. Section 4 contains our main analyses of different scenarios within the model, and discusses how the analysis of these scenarios relates to empirical results. This section also briefly considers how the model might be exploited to examine forward looking scenarios. Section 5 concludes.

2 Existing evidence

Amongst recent papers that have used micro-data to assess the relationship between house prices and consumption in the UK, there appears to be little consensus. Campbell and Cocco (2007), argue that their empirical results suggest that wealth effects are the most likely explanation of the correlation between house prices and consumption. They find that increases in consumption observed during recent house price booms, were mainly driven by increases in the consumption of home owners (rather than renters) and older consumers rather than younger ones. Campbell and Cocco (2007) also use a structural model to assess whether endogeneity of the home-ownership decision might bias their empirical results. By contrast, ABHL find the common causality channel to be most important. Both papers use micro-data from the UK Family Expenditure Survey / Expenditure and Food Survey and yet reach very different conclusions.[1]

To be more precise about the findings, we summarise the exercise of ABHL. Their baseline specification is the following:

$$\ln X_t^{ch} = \alpha^c + f(age) + \gamma' \mathbf{z}_t^{ch} + \varepsilon_t^c + u_t^{ch} \tag{1}$$

where c stands for cohort and h stands for household; X is non housing expenditure, α^c a cohort-specific intercept, $f(age)$ is a quintic in age, $\mathbf{z}_t^{ch}$ are a set of demographics, and ε_t^c and u_t^{ch} are cohort-specific and household-specific error terms.

The baseline specification is then augmented with (regional) house price terms, as follows:

$$\ln X_t^{ch} = \alpha^c + f(age) + \gamma' \mathbf{z}_t^{ch} + \theta'_{ag} g(hp) + \varepsilon_t^c + u_t^{ch} \tag{2}$$

where ag are three age groups: young (aged less than 35), middle-aged (aged between 35 and 60) and old (aged over 60).

Given the potential for complicated interactions in the relationships being analysed using equation 2, ABHL try a number of different specifications for the function $g(hp)$ in an attempt to capture the relationship between house prices and consumption in a flexible fashion. Table 1 reports the coefficients of interest for two of their specifications. The 'shocks' specification reported in

---

[1]It is unclear why this is the case, though there are several methodological differences between the papers. A comparative study by Cristini and Sevilla (2007) attempted to replicate both studies as closely as possible. The ABHL results were found to be robust.

table 1 includes as regressors predicted (log) house prices[2] and the difference between predicted and realised (log) house prices. The latter term is intended to capture the effect of unexpected house price growth for these three age groups. The house price terms in the 'growth' specification reported in table 1 measure the effect on consumption of proportionate growth in house prices. The strongest associations between unexpected and proportionate house price growth and consumption are seen to show up in for young individuals. Since this group contains individuals who are unlikely to hold large amounts of housing wealth, and this is also the group for which (due to the planning horizon) wealth would have the smallest immediate effect on consumption, these results have been interpreted as suggesting that wealth effects have not been the main driver of the correlation between house prices and consumption.

Table 1: ABHL house price terms by age groups: predicted price specification

	Estimated Coefficients	
House Price Terms	*Shocks spec*	*Growth spec*
Predicted log house price		
Young	0.291*	
Mid-age	0.292*	N/A
Old	0.294*	
Shocks to log house price		
Young	0.188*	
Mid-age	0.088*	N/A
Old	-0.012	
Proportional growth in house price		
Young		0.209*
Mid-age	N/A	0.127*
Old		0.042

3   The model

As discussed in section 1, our central contribution lies in our systematic use of a realistic structural lifecycle model of consumption, savings and housing choices, to disentangle and distinguish the

---

[2]Where house prices are predicted from real interest rates and regional incomes.

three mechanisms that have been proposed to explain the correlation of house price shocks and consumption growth. As such our work is building on several recent contributions that have examined the relationship between housing (or durables) and consumption in structural frameworks.[3] Such analyses have found that the nature of housing as consumption good, asset, source of collateral, and potential intergenerational heirloom, makes the relationship complex. Fernandez-Villarverde and Krueger (2001), for example, argue persuasively that a housing asset may have an important role in explaining the observed "hump shape" of profiles of consumption and durable consumption over the lifecycle. Flavin and Nakagawa (2004) discuss the fact that when the amount of housing can only be adjusted at cost, this may explain why both housing and non-durable consumption tend to be smooth over time apart from at infrequent periods of large adjustment. Similarly, the analysis of Li and Yao (2007) shows how a housing asset can result in consumption behaviour (including for housing consumption) that is very insensitive to income in some ranges, but very sensitive in others.

Our contribution to this literature comes in a large part from the way we apply our model to match features of aggregate data, as we simulate the behaviour of a series of cohorts designed to resemble cohorts of the UK economy. The model is constructed with some of the stochastic elements faced by the agents- the process for house prices, and a component of the income generating process - that are thought of as aggregate in that they affect all members of the population at the same time. As well as using data to estimate the parameters of the processes generating shocks to these aggregate features, we also have estimates of the path of aggregate shocks that have been experienced in the U.K. economy since the beginning of the 1970s, and we use this path as an input into our simulations. More precisely, we take the data on actual growth rates in house prices and aggregate cohort level earnings in the economy each year, and input these into our model at the correct age for the particular cohort that we are simulating. We repeat this process for each five year cohort born between the 1910s and the 1970s, and create simulated data for a large number of different households (i.e. realisations of the idiosyncratic shocks) in each cohort given the house prices and the aggregate income that they actually faced. These simulated data are the basis for our comparisons to data on the U.K. economy, including to household survey data

---

[3]As mentioned in the previous section, Campbell and Cocco (2007) makes some use of a structural model, while Cocco (2005) is a contribution to another growing literature on portfolio choice in the presence of housing

which incorporate the same set of cohorts. It is important to stress that our simulations take as an input the actual path of prices and incomes in the U.K. economy. Since the relationships between choices, prices and income in our model are not linear, simulating for a specific path of prices and aggregate income as we have done gives different results from simulating behaviour on average (across different levels of the house price and aggregate income) for particular shocks. Our simulations based on actual shocks to house prices and incomes also form the "baseline" against which we can compare counterfactual scenarios in which the mechanisms that might drive the link between house prices and consumption are shut down in turn; these scenarios and how they compare to our calibrated baseline are the subject of section 4, below.

In order to carry out the exercise described in the previous paragraph, our model had to be carefully constructed with certain novel features, and is therefore a contribution in its own right to the understanding of lifecycle choices in the face of uncertainty. The model is of a household's lifecycle choices concerning consumption and saving, and whether or not to own housing. Since we want to capture how aggregate shocks feed in to aggregate consumption, we are careful to model aggregate income and house price uncertainty in a way that reflects the data. We are particularly careful to model correlation in the shocks in the two processes since house prices and incomes both tend to go up (down) when the economy is performing strongly (poorly), and this must be reflected in agents (rational) expectations. To our knowledge, building these aggregate shocks, and their empirical correlation, into a model of the kind we construct has not been done previously. In addition, building on the models presented in Bottazzi, Low and Wakefield (2007) and Attanasio et al (2007), we are careful to model mortgage related borrowing constraints as realistically as possible since these are likely to have first order effects on behaviour. We now describe the model, and its calibration to UK data, in detail.

3.1 The household maximisation problem

The households in our model are (ex-ante) heterogeneous in one dimension: their level of education. In practice this will imply the calibration of different earning processes for households with different level of education. In addition, different idiosyncratic shocks. And of course at a given point in time, households of different ages are present.

A household lives $T = 59$ periods (ages 22-80). In every period $t \leq T$, the household maximises

utility by choosing consumption, with $c_t \in R_+$, and whether to own a flat, a house, or no housing (which may be thought of as costless rental), with $h_t \in \{0, 1, 2\}$. The household value function in period t is given by

$$V_t\left(A_t, h_{t-1}, P_t, Y_t, Z_t\right) = \max_{\{c_t, h_t\}} u\left(c_t, h_t\right) + \beta E V_{t+1}\left(A_{t+1}, h_t, P_{t+1}, Y_{t+1}, Z_{t+1}\right) \tag{3}$$

subject to

$$A_{t+1} = R_{t+1} \begin{cases} A_t + W_t - c_t - \kappa P_t(1+F)I(h_t = 1) - P_t(1+F)I(h_t = 2) & \text{if } h_{t-1} = 0 \\ A_t + W_t - c_t + \kappa P_t(1-F)I(h_t \neq 1) - P_t(1+F)I(h_t = 2) & \text{if } h_{t-1} = 1 \\ A_t + W_t - c_t - \kappa P_t(1+F)I(h_t = 1) + P_t(1-F)I(h_t \neq 2) & \text{if } h_{t-1} = 2 \end{cases} \tag{4}$$

where A_t is the start of period asset stock and $R_{t+1} = 1 + r_{t+1}$ and r_{t+1} is the (real) interest rate on the liquid asset; P_t is the price of housing which is realised at the start of period t; F is the cost of selling or buying a house, which is proportional to the price; W_t is household income in period t, while Y_t and Z_t are the idiosyncratic and aggregate components of the income.[4] The number of state variables in this problem (with four continuous states plus the current home ownership and time) means it is computationally demanding to solve; for more details on the solution method see appendix A.

We only allow for collateralised debt, i.e. households are only able to have negative financial assets when they are home owners, so that when they do not own a house ($h_t = 0$) they are subject to the constraint

$$A_t \geq 0 \tag{5}$$

Home owners can borrow, and when they do so they are subject both to a terminal asset condition that translates into an implicit borrowing constraint, and to two explicit borrowing constraints. In particular, we impose the terminal condition $A_{T+1} = 0$. The specification of marginal utility becoming infinite at 0 consumption means this terminal condition prevents households borrowing more than they can repay with certainty. In addition to this implicit borrowing constraint, we allow for two explicit constraints. The first is a function of the value of the house and the second is a function of the household annual income. They determine how much a household is able to borrow at the time of purchase or when remortgaging, and translate into the following

---

[4]The household is assumed to be fully aware of the separate stochastic processes that generate these components.

constraints in the period after the new mortgage is agreed:

$$A_{t+1} \geq -\lambda_h \kappa P_t (1+r), \quad \kappa = \begin{cases} 0 < \kappa < 1 & \text{if } h_t = 1 \\ 1 & \text{if } h_t = 2 \end{cases} \tag{6}$$

The value $(1 - \lambda_h)$ can be thought of as a downpayment requirement.

$$A_{t+1} \geq -\lambda_w W_t (1+r) \tag{7}$$

The explicit constraints on the downpayment and the debt to income ratio only apply when households buy the property or remortgage. Formulating the constraints in this way makes the model more complicated to solve numerically than, for example, having only an income-related constraint that must be satisfied every period (as, for example, in Campbell and Cocco, 2007, Li and Yao, 2007, Cocco, 2005, Campbell and Hercowitz, 2004). However it seems important to us to capture the institutional features of the UK mortgage market, since these are likely to affect how house price shocks feed into consumption. The correlate of only applying the constraints in periods of buying or remortgaging is that when a household continues owning without remortgaging, they can keep their existing debt if they have negative financial assets:

$$A_{t+1} \geq A_t \tag{8}$$

Although there is no mortgage repayment schedule, the household does have to pay off mortgage interest each year in which it does not remortgage.

3.1.1 Utility and bequest functions

Households get utility from consumption, from home-ownership, and from leaving bequests.

The within period utility function is CRRA. This is augmented by a term reflecting the value of home-ownership:

$$u(c_t, h_t) = \exp(\theta \phi h_t) \frac{c^{1-\gamma}}{1-\gamma} \quad \begin{cases} \theta, \ \phi \in \mathbb{R} \backslash \{0\} & \text{if } h_t = 0 \\ \theta \in \mathbb{R}, \ 0 < \phi < 1 & \text{if } h_t = 1 \\ \theta \in \mathbb{R}, \ \phi = 1 & \text{if } h_t = 2 \end{cases} \tag{9}$$

The parameter θ is a housing preference parameter which determines the utility that households obtain from owning a house rather renting it. ϕ determines the relative utility from owning a flat versus a house. These parameters are calibrated in our model.

11

The utility from leaving a bequest is described by a second iso-elastic function:

$$b(W) = \tau * \frac{(W/\tau)^{1-\gamma}}{1-\gamma} \tag{10}$$

Where W is the value of wealth (both financial and housing wealth, net of the fixed cost of selling a property) left over at the end of life T, after all shocks to resources (income and the house price) have been realised, and all consumption decisions have been made.[5] The parameter τ is calibrated. A bequest motive, although not central to our analysis, is crucial to match certain features of the life cycle profile of home ownership. Further discussion of the preferences just discussed can be found in Attanasio et al. (2007).

3.2 The environment: exogenous stochastic processes

Households face three dimensions of uncertainty: shocks to house prices, which are aggregate (i.e. common across all properties in the economy); aggregate shocks to income; and idiosyncratic shocks to income. In the present version of the model, the interest rate on liquid assets and debt is fixed.[6]

If we take the income generating process first, this may be thought of as being composed of three parts:

$$\ln W_t = d_t + y_t + z_t \tag{11}$$

where lower case has been used for logs, and d_t is a deterministic part to the income generating process, y_t is a persistent idiosyncratic stochastic element; and, z_t is the aggregate stochastic component.

The deterministic component d_t is hump-shaped over the working lifetime and is captured using (the log of) a polynomial. The coefficients of these polynomial are calibrated for the two education groups we consider. In practice, to fit the observed data we need a cubic specification (see Appendix B).

The idiosyncratic stochastic component y_t is modelled as an AR(1) process:

$$y_t = \rho_y y_{t-1} + \xi_t, \quad \xi_t \sim N\left(0, \sigma_\xi^2\right) \tag{12}$$

---

[5]For more motivation of the modelling of bequests, see, for example, De Nardi (2004) [16].

[6]It is a relatively simple extension to add i.i.d shocks to the interest rate. Early experiments indicated that such a change has no qualitative impact on our results.

The stochastic aggregate component to income is modelled jointly with the stochastic house price price using a first-order vector auto-regression with correlated innovations. When we used data to estimate this process, we could not reject the hypothesis that lagged income does not affect current house prices and similarly that lagged house prices do not explain current income. However, we find a significant and sizeable correlation between house price and earnings innovations. As for the effects of own lags, we could not reject that these coefficients were equal to unity: we therefore impose this value.

If we let HP stand for the house price and again use lower case letters to represent logs, this can be written as:[7]

$$
\begin{bmatrix} z_t \\ hp_t \end{bmatrix} = \begin{bmatrix} \alpha_0^z \\ \alpha_0^h \end{bmatrix} + \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} z_{t-1} \\ hp_{t-1} \end{bmatrix} + \begin{bmatrix} u_t^z \\ u_t^h \end{bmatrix} \tag{13}
$$

The value of modelling these aggregate processes jointly as a VAR, rather than as separate autoregressions, is that the joint distribution of the error terms will capture correlation between aggregate shocks to house prices and incomes. Capturing this correlation is valuable since it will affect the degree to which individuals will choose to modify their asset accumulation as a means of self insurance against shocks. The joint distribution of the shocks is assumed to be normal:

$$
\underline{u}_t \sim \mathbb{N}(0, \underline{\Omega}) \tag{14}
$$

where

$$
\underline{\Omega} = \begin{bmatrix} \sigma2_z & \pi\sigma_z\sigma_h \\ \pi\sigma_h\sigma_z & \sigma2_h \end{bmatrix}
$$

with π measuring the correlation between the shocks.

3.3 Calibration and estimation

The parameters of our model can be divided into three categories: those we take from other studies, those we estimate outside the model and those we calibrate to fit some moments of the micro data we consider. We estimate the parameters of the exogenous stochastic processes faced by the agents

---

[7]Given the unit persistence in this equation, the constant terms α_0^z and α_0^h capture drift over time even though there is no time trend in equation 13, see Davidson and MacKinnon, pp.606f. Since the process is in logs, the constant terms measure the trend growth rate.

in our model using time series data. Our calibration exercise, instead, involves matching moments for life cycle levels of home ownership status in the UK for those aged 26-60 in 1990-2006. Since we simulate cohorts born between the 1910s and 1960s, we match the ownership levels for the cohorts in the relevant age range in the 1990s and early 2000s.

3.3.1 Parameters fixed or estimated outside of the model

For inputs into the calibrated model, we need to use data on earnings, the house price process and the interest rate on liquid assets. Values for parameters fixed or estimated outside the model are summarised in table 2.

Table 2: Estimated / Fixed Parameters

Parameter	Value		Source
	Estimated/Fixed Parameters		
	Utility Parameters		
γ	1.43		(Attanasio and Weber, 1995)
β	1.0358^{-1}		
	Aggregate House Price and Income Process		
α_0^z	1.66%		FES
σ_z	0.033		FES
α_0^h	3.58%		DCLG
σ_h	0.091		DCLG
τ	0.645		FES / DCLG
κ	0.6		BHPS
	Idiosyncratic Income Process		
Deterministic component: cubic in age			BHPS
	High Edu	Low edu	
ρ_y	0.76	0.77	BHPS
σ_ξ	0.39	0.41	BHPS
$\frac{MedianP_{22}}{MedianY_{22}}$	3.3	4.4	BHPS
	Credit market Institutions		
λ_y	3.0		
λ_h	0.9		
$\bar{r}$	0.03		B.o.E.

Utility function

The preference parameter γ in the utility function is set to match the consumption elasticity of intertemporal substitution of 0.7 found in data (see Attanasio and Weber, 1995). This corresponds to a curvature $\gamma = 1.43$ for our within period utility function. The parameter ϕ indicating the relative utility value of a flat to a house is set at the same level as the price ratio of flats to houses, κ, so that $\phi = 0.6$. We also choose a baseline discount rate of 3.58%, which matches the expected rate of return on housing.

Aggregate Processes: The house price and aggregate income shocks

As described by equations 13 and 14, the aggregate house price and income processes are specified in logs as a vector autoregression of order one (a VAR(1)), with drift. To estimate this process we use: data on the national mix-adjusted house price series for the UK[8], which are the same data underlying figure 1 and are freely available from the Department for Communities and Local Government (DCLG); and, data on the UK Average Earnings. Since our model is set up in terms of real prices, we deflate both series to prices for the latest year using the all item Retail Prices Index, and we use data for the years 1969-2006. As anticipated in section 3.2, estimation of this process does not reject unit persistence, and so we impose that aggregate shocks are permanent as in equation 13. The other parameters returned from this estimation are trend growth rates for the house prices (3.46%) and income (1.65%), and standard deviations of the shocks to the processes of 0.094 and 0.027 respectively for the house price and income, as well as the correlation of the shocks for the two processes of approximately 0.4.

It is important to realise that in our simulation exercises in addition to inputing the parameters we have been discussing to solve the model, we use as inputs the estimated income shocks that are derived from the estimation of the aggregate time series processes in equation (13). These processes are treated as aggregate so that the same shocks hit all individuals in a given period, and thus affect different cohorts at different ages. We feel that it is important that our simulations reflect actual realizations of aggregate shocks to drivers of consumption, since the complicated nature of the relationship between house prices and consumption choices is likely to mean that how this shows up in data will depend on the actual path that prices follow and not just the

---

[8]We use the series reporting average house prices for all dwellings.

processes generating prices. The growth rates in house prices and aggregate income that we input into our model as aggregate shocks are shown in figure 15 in the appendix.

Idiosyncratic income process

On top of the aggregate income process, there are two further elements to the household income generating process (see equation 11). These are a deterministic process, which captures the hump-shaped profile of household incomes over the working life and a 60% replacement rate in retirement, plus a persistent stochastic component which is described by a first order autogeression (equation 12). The deterministic component plus the process generating shocks to the stochastic component are both education group specific, and the realised shocks to the stochastic component are idiosyncratic at the household level.

Parameters for both processes are estimated together as described in Appendix A. Since this involves estimation of a dynamic process at the household level, we require panel data and so use the British Household Panel Survey (BHPS), which is available for the years 1991-2005. The estimation yields that the deterministic component can be approximated as a cubic that shows a hump-shape over the working life which is slightly more pronounced (in particular with a steeper slope at the beginning of the working life) for the college educated group than for those with only compulsory level education. The process generating stochastic innovations is quite similar across groups with a persistence parameter of slightly more than 0.75 and a variance of the shock of around 0.16 (exact values are in Table 2 or Appendix B).

Credit market institutions

The parameters that determine the fraction of the house price (λ_h) and the multiple of earnings (λ_y) that households can borrow are chosen to match the UK institutional features. At the time of taking out a new mortgage (i.e. of buying or remortgaging) households can borrow whichever amount is lower between three times household earnings ($\lambda_y = 3$) and 90% of the house price ($\lambda_h = 0.9$). The interest rate on the liquid asset / debt is taken to be fixed and is set to match the average of 3% for the real interest rate on 90 day treasury bills for the period since 1990 (i.e. the period for our calibration).

3.3.2 Calibrated parameters

We select the preference parameter for housing, θ,[9] plus the parameter τ that determines the bequest motive, and the fixed cost of housing, F, by to matching average life-cycle home-ownership rates between ages 26 and 35, and 36 and 60, for our two education groups. We assume that the parameters are common across the two education groups. As reported in table 3 we obtain values such that owning a house raises utility by approximately 1.5% and owning a flat raises utility by approximately 1%, the fixed cost of buying and selling is 3% of the property price, and the parameter τ takes a value of 4. Table 4 compares home-ownership rates predicted by the model to those observed in the data[10]; we do a reasonably good job of matching the moments of interest, particularly for the high-education group.

Table 3: Calibrated Parameters

Parameter	Value
θ	-0.015
ϕ	$\frac{2}{3}$
τ	4
$F_b = F_s$	0.03

Table 4: Calibration Statistics

Statistic	High Education		Low Education	
	Data	Model	Data	Model
Ownership rate				
Age 26 - 35	0.558	0.570	0.474	0.412
Age 36 - 60	0.794	0.826	0.632	0.689

Notes: The data figures for home-ownership rates are based on the years 1990-2006 of the FES.

---

[9]Since the estimate of the intertemporal elasticity is taken from a paper which does not condition on home-ownership, there is a possible bias.

[10]Data come from the years 1990-2006, as years prior to 1990 are affected by the large-scale selling off of local authority housing.

3.4 Model fit

Since we do not match any moments of consumption, the properties of the evolution of consumption provide a useful check on whether our calibrated model of overlapping cohorts is doing a good job of matching data on the U.K. economy. It is also sensible to check the properties of simulated consumption since the question we wish to address concerns the evolution of consumption. In fact the model does a remarkably good job of matching the consumption growth rate. This is shown in figure 2 which compares the actual growth rate of aggregate consumption in the U.K. economy to that in our simulated data for the years 1975-2006.
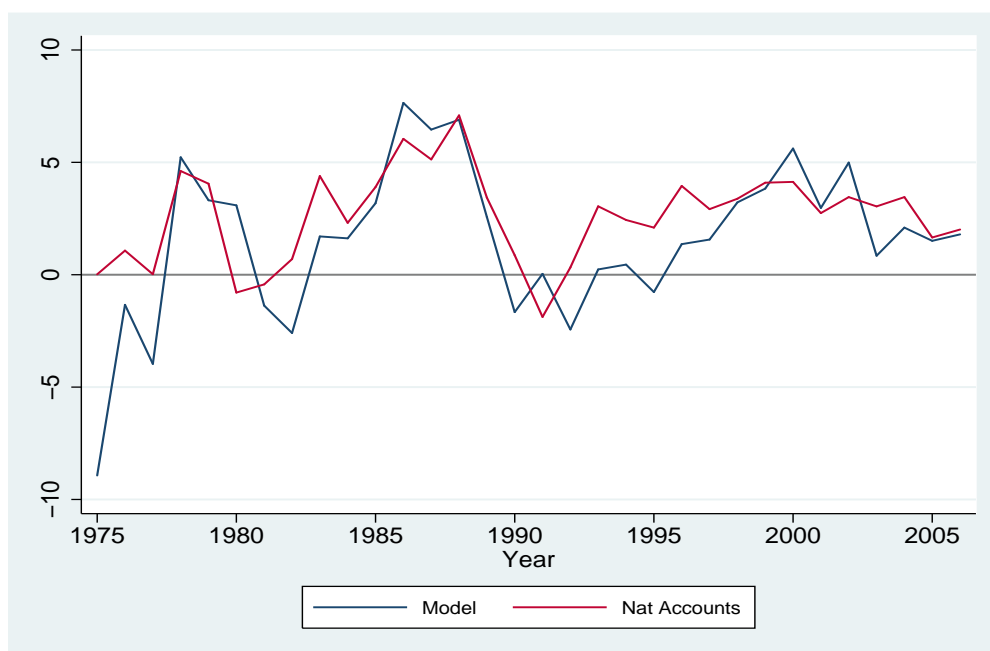
Figure 2: Growth in real consumption: Data versus model

This good match for the consumption growth rate is an indicator that our model might be useful for drawing quantitative, as well as qualitative conclusions. However, unfortunately the model does less well capturing patterns at a more disaggregated level, particularly by cohort. For example, in the late 1980s consumption boom we know (from Attanasio and Weber, 1984) that the strongest consumption growth was from young groups. Figure 3 shows modelled consumption growth during this boom for different cohorts. During the middle years of the decade the model does indeed predict the fastest consumption growth for the young, but in 1988, the year of strongest consumption and house price growth, all cohorts are seen to have strong consumption growth and

the eldest cohort shown (aged in their late 50s - the grey line) is predicted to have consumption growing just as rapidly as those cohorts in their 30s. This counterfactual pattern leads us to be cautious in drawing quantitative conclusions from our model about the extent to which house prices have caused shifts in aggregate consumption. Nonetheless the model is still useful for exploring the qualitative properties of consumption growth driven by increased housing wealth or by increased permanent income.
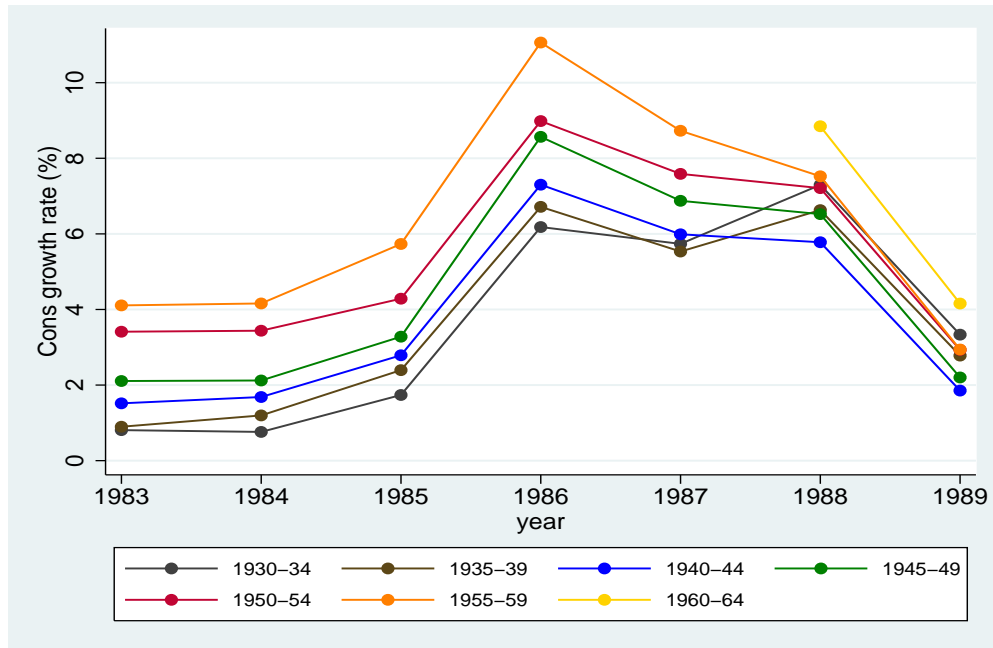

Figure 3: Model consumption growth by cohort: 1980s boom

4   Booms and busts in house prices and consumption

The main motivation of our analysis using the calibrated life-cycle model is to explore how we can distinguish between the mechanisms that have been proposed to explain the correlation of house price shocks and consumption growth. The first part of this subsection reports our investigation of this issue through using our model to explore counterfactual scenarios in which the possible mechanisms are shut off in turn. The final subsection briefly examines how the model might be useful for considering different forward looking scenarios.

4.1 Using the model to disentangle mechanisms

To disentangle the possible mechanisms underlying the correlation between house-price and consumption growth, we compare a baseline run of our model to runs in which aggregate shocks to house prices and to incomes are switched off in turn. These aggregate shocks are the only drivers of changes in aggregate consumption growth in our model, so we start by comparing aggregate consumption growth across simulations. This approach will give us a good idea of which factor is more important in driving swings in aggregate consumption growth in the model economy. We then examine the simulated data in a more disaggregate way to explore what micro-data on consumption can tell us about the relative importance of house prices in driving aggregate consumption growth.
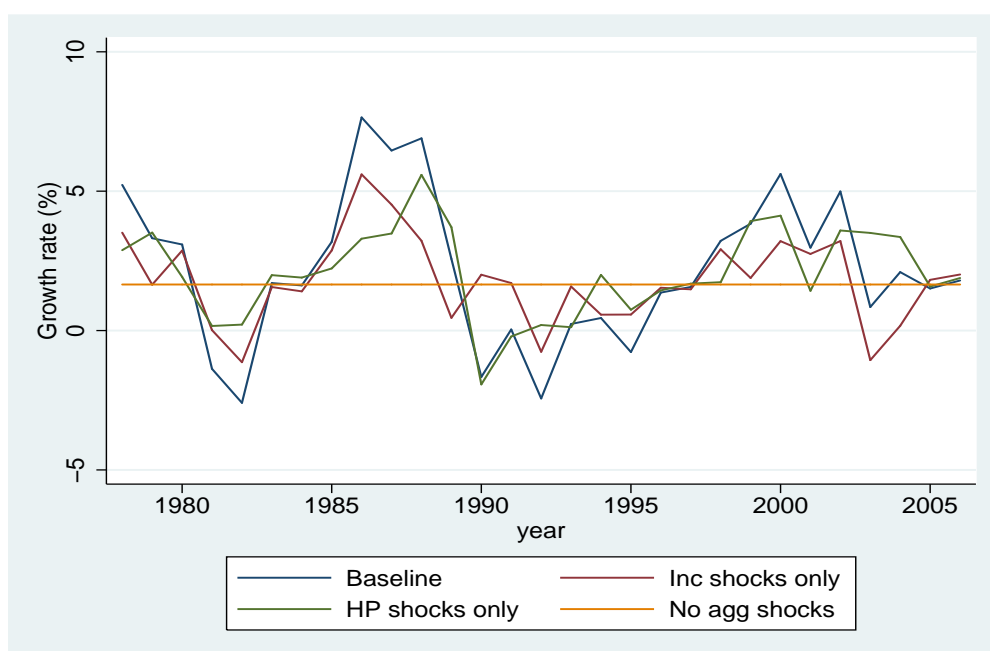
Figure 4: Model consumption growth: different scenarios

We saw in the previous section that our model appeared to generate a series for aggregate consumption growth that closely tracked that seen in aggregate data. Figure 4 below shows the growth rate of aggregate consumption using different runs of the model in which various shocks are switched off. The"baseline" results are the growth rates we saw previously in figure 2. When both house price and income shocks are shut down, consumption grows at the constant rate of approximately 1.7%. The other lines show consumption growth rates in the presence of only income shocks and only house price shocks. In the late 1970s and early 1980s, it is clear that income shocks

are driving much of the movement in consumption growth: the growth rates when only house price shocks hit are typically closer to the constant level whilst growth rates when only income shocks hit are closer to the baseline growth rates. However, by the latter 1980s, the effect of changes in house prices on consumption growth becomes more substantial. By the house price boom of the late 1980s, more of the total consumption growth is driven by house prices than by incomes; similarly the house price bust of the early 1990s drives almost all of the negative consumption growth rates observed. Since then, house price shocks appear to have been the predominant factor driving aggregate consumption growth in the model. The relatively modest influence of income shocks in the later years modelled is due to the fact that aggregate income shocks, set to match the data shown in figure 15 in the appendix, have become rather weak since the late 1990s, compared to the volatile years of the 1970s and early 1980s.

One way to quantify "how much" of consumption growth is being driven by house price or income shocks, is to look at the proportion of the deviation of consumption growth in our baseline run from that in our run with no shocks, that is accounted for by the deviation in the modelled economy with only a single aggregate shock. Doing this over the whole period shown in figure 4, we see find that house price shocks explain away about 55% of the deviation, while income shocks explain away about 35%.[11] In line with the descriptive analysis of the previous paragraph, this difference is stronger after 1990 (for which period the proportions are 73% and 17% respectively), but is reversed (40% and 50%) for the 1978 to 1990.

The above analysis looked at aggregate trends in consumption growth under different versions of the model where house price and income shocks were present or absent. A more disaggregate analysis looking at these results for different year of birth cohorts will allow us to examine how the response to shocks varies for households at different points in the life cycle. Instead of looking at consumption growth rates, figures 5 to 7 show for different cohorts the 'excess' consumption growth (that is, consumption growth above that generated when both shocks are switched off) for different cohorts, first covering the period 1978 to 1989, then 1989 to 2000, and finally 1999 to 2006. In these results, both house price and aggregate income shocks are present.

The now familiar patterns of consumption are evident: relatively large swings from boom to

---

[11]Though the income and house price shocks are the only aggregate shocks in the model, these two factors might also have a joint effect on consumption, and so there is no reason why these to numbers should sum to 100%.

bust in the late 1970s and early 1980s, the consumption boom of the later 1980s followed by the pronounced bust of the early 1990s, a return to strong growth in the latter 1990s and early 2000s before moderation in consumption growth by the middle of the 2000s. These trends are evident in the paths of 'excess' consumption growth for all cohorts, although there is some variation in the magnitude of swings across cohorts. In particular, the swings are seen to more more pronounced for the youngest cohorts in the late 1970s, a period when consumption growth is largely driven by the income shock in the model. On the other hand in the consumption boom around 1988, the bust of the early 1990s, and again in the period of strong growth around 2004, it is the older cohorts who had the largest magnitude of excess consumption growth. These three periods are all episodes when strong house price growth is important in driving modelled consumption growth.
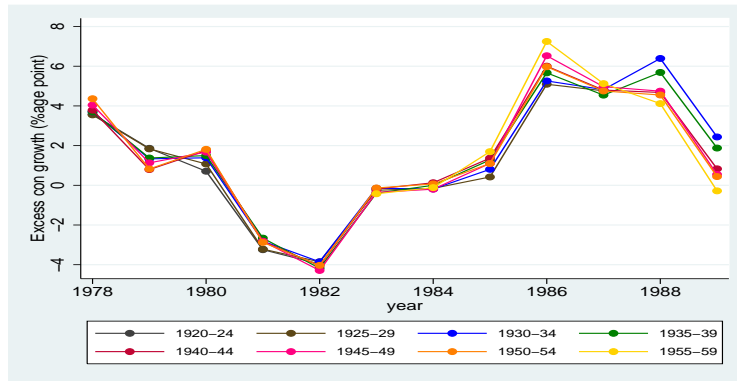
Figure 5: Excess annual consumption growth by cohort, shocks to income and house prices, 1980s
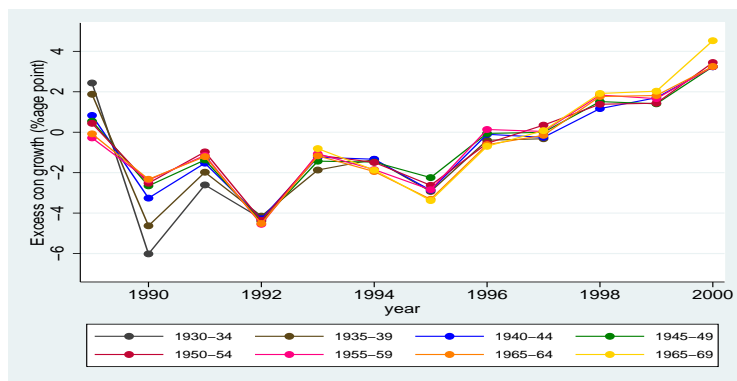

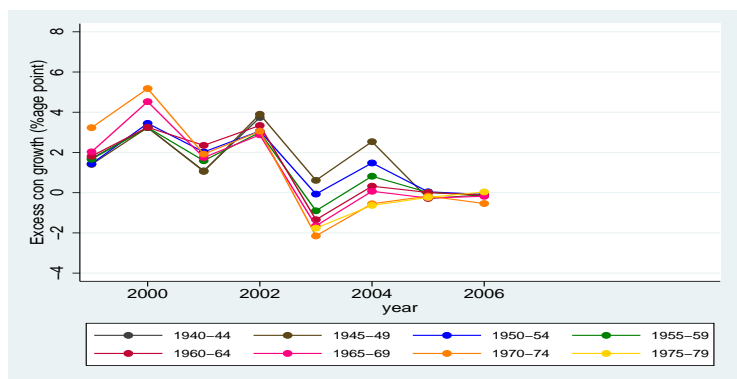Figure 6: Excess annual consumption growth by cohort, shocks to income and house prices, 1990s

Figure 7: Excess annual consumption growth by cohort, shocks to income and house prices, 2000s

Figures 8 to 10 show excess annual consumption growth at the cohort level when the house price shocks have been removed, meaning the only aggregate shocks are to income. Again, the

patterns across cohorts are very similar, but there is a clear tendency for the younger cohorts to have the strongest response to aggregate income shocks: in the 1980s, for example, cohorts born in the 1950s react more strongly than cohorts born in the 1920s or 1930s; similarly, in the 1990s the greatest responses to income shocks are from those cohorts born in the 1960s with relatively weaker consumption responses from those born in the 1930s. These figures also make clear that relatively little consumption response was generated through aggregate income shocks in the 2000s: excess consumption growth is roughly zero for all cohorts in the period 2000 to 2005.

Figures 11 to 13 repeat the analysis when aggregate income shocks have been removed and only house price shocks remain. Again we see that movements in consumption in response to house price shocks are strongly correlated across cohorts, though there is more dispersion than was the case when only aggregate income shocks were in play. By constrast to the previous figures, it is clear that older cohorts are those that respond most to house price shocks: the late 1980s house price boom, for example, generates an excess consumption growth rate at its peak of around 4-6% for cohorts born between 1930 and 1934, but only around half this magnitude for those born later. Similarly the house price slump of the early 1990s has a much stronger negative effect on consumption for older cohorts: by 1992, those born in the 1930s see excess consumption growth in the order of -4% to -6% whilst those born in the 1950s see excess consumption growth no worse than -3%. More recent rapid house price increases also appear to have fed through to the consumption of older cohorts more strongly: in 2004, for example, those born in the late 1940s had excess consumption growth of some 6% whilst those born in the 1970s had excess consumption growth of around 3 - 4%.
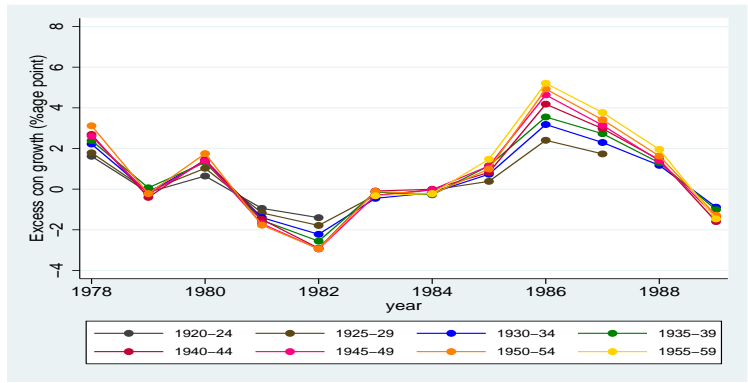
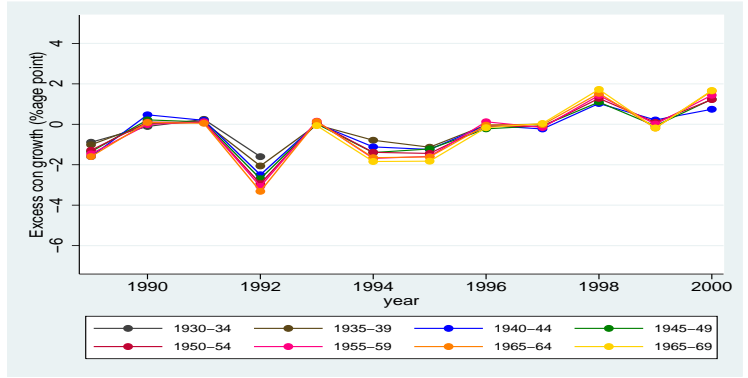Figure 8: Excess annual consumption growth by cohort, aggregate shocks to income only, 1980s

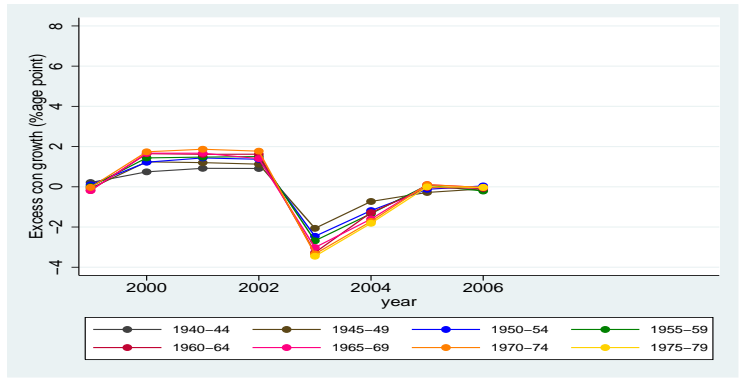Figure 9: Excess annual consumption growth by cohort, aggregate shocks to income only, 1990s

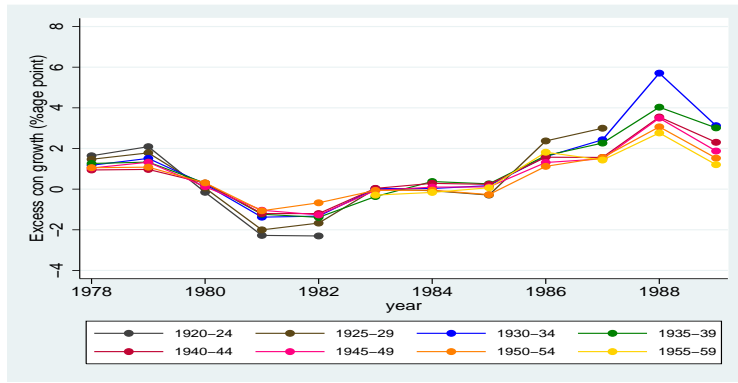Figure 10: Excess annual consumption growth by cohort, aggregate shocks to income only, 2000s

Figure 11: Excess annual consumption growth by cohort, agg. shocks to house price only, 1980s
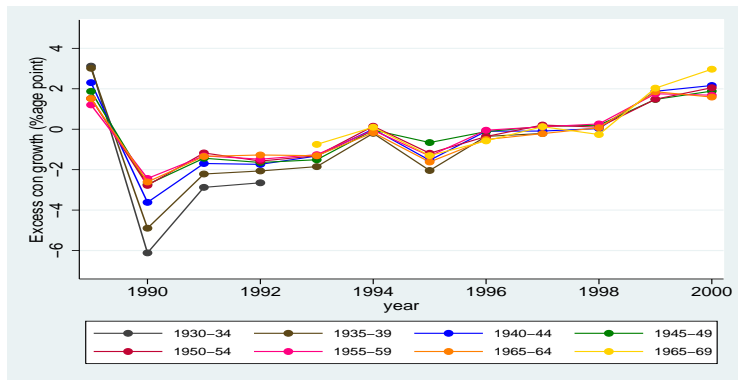

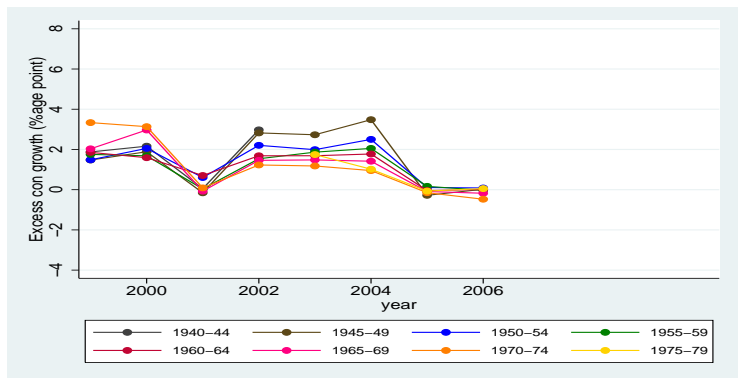Figure 12: Excess annual consumption growth by cohort, agg. shocks to house price only, 1990s

Figure 13: Excess annual consumption growth by cohort, agg. shocks to house price only, 2000s

These cohort analyses have shown that different shocks translate into consumption responses that differ across cohorts. Younger cohorts respond more to aggregate income shocks whilst older

cohorts respond more to house price shocks. In a stylised lifecycle model, it is possible to see an intuition for why this occurs. A permanent income shock effectively acts as a shock to lifetime wealth, and younger cohorts have a longer horizon over which to enjoy a positive shock (or suffer a negative shock) and so will adjust their consumption by more. Older cohorts respond more to house price shocks as they are more likely to have positive equity in their homes and so benefit from a wealth channel that allows them to adjust consumption. They would also respond more to this one off wealth shock simply because they expect fewer periods of life over which they can spread extra consumption.

However, without the kind of analysis that we have conducted, one could not be sure that these results would carry over to a more complicated model with credit constraints and simultaneous housing and non-durable consumption decisions. In particular, the possibility that house price shocks might affect consumption through altering collateral and so borrowing possibilities, would seem more pertinent for the young than the old. This implies that relative to the case of just a wealth effect, collateral effects might complicate the pattern across cohorts of changes in consumption following a house price shock. Our analysis shows that even with this factor in play, it is the old whose consumption is most responsive to the house price shocks.

What we can draw from our analyses is the importance of using micro-data to try to infer whether income shocks or house price shocks are responsible for consumption growth that we observe. If consumption shocks appear particularly strong for the young, that would suggest changes to aggregate incomes are most important in explaining consumption movements whereas if the shocks are strongest for the old, the housing wealth channel is likely to be most important. The earlier reduced for studies considered in section 2 by Attanasio and Weber (1994) and Attanasio et al (2008) indicated that excess consumption growth has been strongest for the young, which was taken to imply that a wealth effect was not the main driver of the observed correlation between house-price growth and consumption growth. Our results give this interpretation on a stronger theoretical footing.

4.2 Using the model to look forward

Having set up our model which captures the growth rate of aggregate consumption in the face of shocks to income and house prices, we can use it to consider what might happen to consumption

over the next few years, if house prices and incomes follow a certain path. Figure 14 shows just such an exercise where the counterfactual scenarios involve house price and aggregate income either growing at their trend rates between 2007 and 2009, or they involve negative shocks of magnitude one standard deviation, to either (or both) of these aggregate variables. These shocks correspond to income falling approximately one-percent for each of the three years, or house prices falling almost six percent each year. The model predicts that the effect of either of these negative shocks in isolation is that aggregate consumption stagnates, with zero real growth each year. However, our assessment of our model has suggested that we might be more accurate at assessing the effects of considering both shocks together, and in that case the model predicts aggregate consumption declining by around 1.7% in each of the three years.
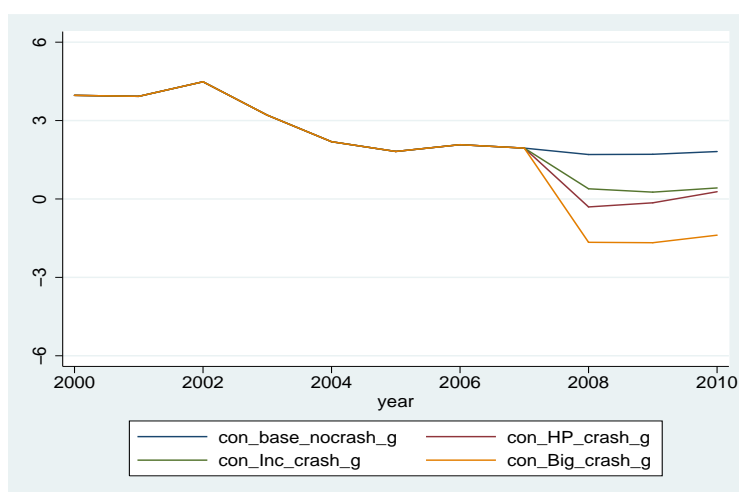

Figure 14: Aggregate consumption growth if prices and/or incomes fall

5 Conclusions and further work

Our analysis has aimed to provide a deeper understanding of how we can distinguish between the mechanisms that have been proposed to explain the correlation of house price shocks and consumption growth. In particular, through our systematic use of a realistic lifecycle model, we have investigated whether structural modelling confirms the intuitions about how and why the consumption of different groups might correlate with house prices that have formed the basis of past, reduced form, empirical tests. Our analysis has involved some methodological innovations in the way we have constructed our model and particularly in the way we have applied it to simulate the behaviour of a series of cohorts, designed to resemble cohorts of the UK economy.

Our analyses have in fact confirmed the intuitions in question. In particular, we have seen that if a house price shock is driving changes in aggregate consumption growth rates, then these changes will be most evident in the consumption paths of older groups. In our simulations the deviations in consumption growth rates *caused by* a house price shock were twice as strong for cohorts in their 50s as for cohorts in their 30s. To get the opposite pattern of stronger deviations for younger groups required some other type of aggregate shock. In the model, and plausibly for many episodes in the data that have coincided with periods of house price booms or busts, this could be a shock to incomes and expected permanent incomes throughout the economy. Thus we have given a more solid footing to the view that micro-data can be useful for disentangling the possible mechanisms underlying the correlation between house price and consumption growth, and to the interpretation of Attanasio and Weber (1994) and ABHL that a stronger correlation for young groups is powerful evidence against the hypothesis that wealth effects from house price changes have been the main mechanism driving the correlation.

On the basis that our model provided a good match to U.K. data on home-ownership rates and consumption growth rates, we have also explored what the model would predict for a scenario in which real house prices and incomes both fell after 2006. For a scenario in which both incomes fell by 1% and house prices by almost 6% for each of three years, the model projected real consumption declines of between 1.5% and 2% per year.

The exercise we have constructed has returned interesting results, but can also be seen as a first step of a research agenda that aims at explaining the behaviour of house prices as an equilibrium process. We have, in effect, modelled the demand side of the housing market. In order to do this

we have made a set of assumptions about the process generating house prices which have facilitated an empirical specification that could be approximated in our simulations. In doing this we have learnt that the precise shape and paramterisation of this process is an extremely important driver of the evolution of the joint decisions of housing demand and consumption expenditure, both at the household and the economy wide level. It is a major research challenge to assess whether the volatile house price process we observe, and other features of aggregate consumption and housing market behaviour, can be obtained from a unified framework in which house prices are determined endogenously.

Appendices

A   Solution method

The setup of our model, with a discrete choice concerning home-ownership, coupled with fixed costs and the particular form of the borrowing constraints, mean that the functions of the household's optimisation problem will not be 'well behaved' and we cannot rely on the existence of smooth first-order conditions that could otherwise have been exploited to improve efficiency in solving the problem.[12] Instead we rely on robust techniques developed in previous ESRC-funded research (ESRC grant RES-000-23-0283, led by Attanasio) which involve solving using iteration on the value function (rather than the first order condition), and finding different "conditional value functions" (one for each of the current choices of house ownership, flat ownership, and non ownership) which can be compared in order to determine the discrete choice.

As is standard for these dynamic problems, the solution for consumption and home-ownership is found recursively from the last period of life, T, backwards. In the final period of life the value function consists of current utility from home ownership and consumption, plus the utility from leaving a bequest, and behaviour in this period is constrained by the necessity that assets at the end of life (after leaving the bequest) be non-negative. Given the optimal choices at $t + 1$, $t < T$, the backwards recursion then proceeds to choose home ownership, consumption and saving that maximise period t's value function, subject to the borrowing constraints.

In order to compute the solution, we solve at a finite number of points in the asset dimension. We store optimal decisions and value functions at grid points but in our simulations households' choices are not restricted to coincide with these points. We perform linear interpolation in all the cases in which choices lie between points.

We also use discrete approximations to the specified continuous processes for idiosyncratic income, and for the house price and aggregate income. This involves modelling these processes using finite state Markov chains that mimic the underlying continuous-valued univariate or bivariate processes. This is done as described in Tauchen (1986). We preferred Tauchen's method of equally

---

[12]The combination of the two borrowing constraints mean that we can show not only that the value function will not be universally concave, but also that the derivative (with respect to assets) will not be defined at all points in the support of this function.

spaced nodes over the quadrature based method proposed by Tauchen and Hussey (1991), because this has been shown to be more robust to very high persistence in the modelled processes (see Floden, 2007).

For the runs presented in this draft of the paper, we have used 105 nodes in each 'conditional' asset grid (we have separate grids underlying each conditional value function, since assets are limited by different borrowing constraints depending on the home-ownership choice). Points are more dense in the lower range of the asset grid, to make sure that non-convexities in the value function are not overlooked in the maximisation process. Idiosyncratic income is represented by a grid of eight nodes, while there are 13 nodes in each of the grids for the house price and aggregate income (or, effectively 169 nodes for the joint process). Monte-Carlo experiments showed that these grid sizes were sufficient to capture the modelled processes to a high degree of accuracy. With this set up, the model solution and simulation takes around 30 hours.

As explained in the paper, the profiles of behaviour that we simulate are obtained by simulating a series of cohorts designed to resemble cohorts of the U.K. population. Every member of a particular cohort experiences the same shocks to aggregate income and the house price at each age. Members of subsequent or earlier cohorts experience the same series of aggregate shocks, but at the relevant ages so that every cohort experiences the same shocks at the same point in time. The aggregate shocks that are input in our baseline run are designed to match actual shocks to house prices and aggregate incomes in the U.K., and these correspond to the growth rates shown in Figure 15.

The profiles of behaviour reported in the paper are obtained by simulating 13 cohorts (born between the 1910s and the 1970s) each containing 2000 individuals split evenly between two education groups. Within education group, households differ according to their initial financial wealth and the idiosyncratic income shock that they face at each period.

In some of the figures in the paper, we plot deviations of average consumption growth from the expected growth rate if the aggregate shocks happen to be zero in each period. With zero shocks, the benchmark expected growth rate of consumption at each age is as shown in Figure 16.
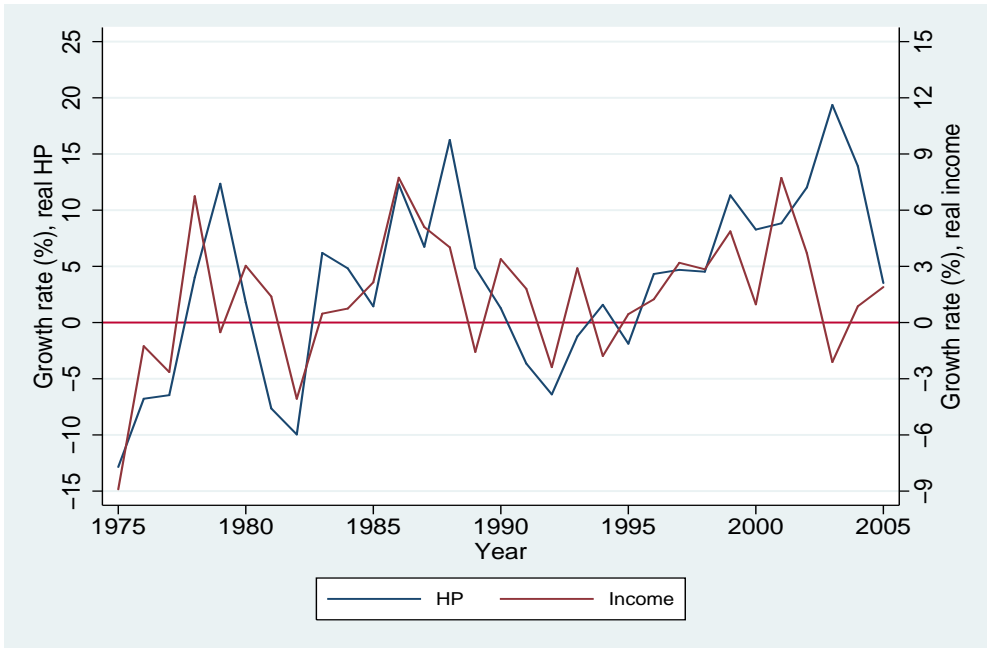
32

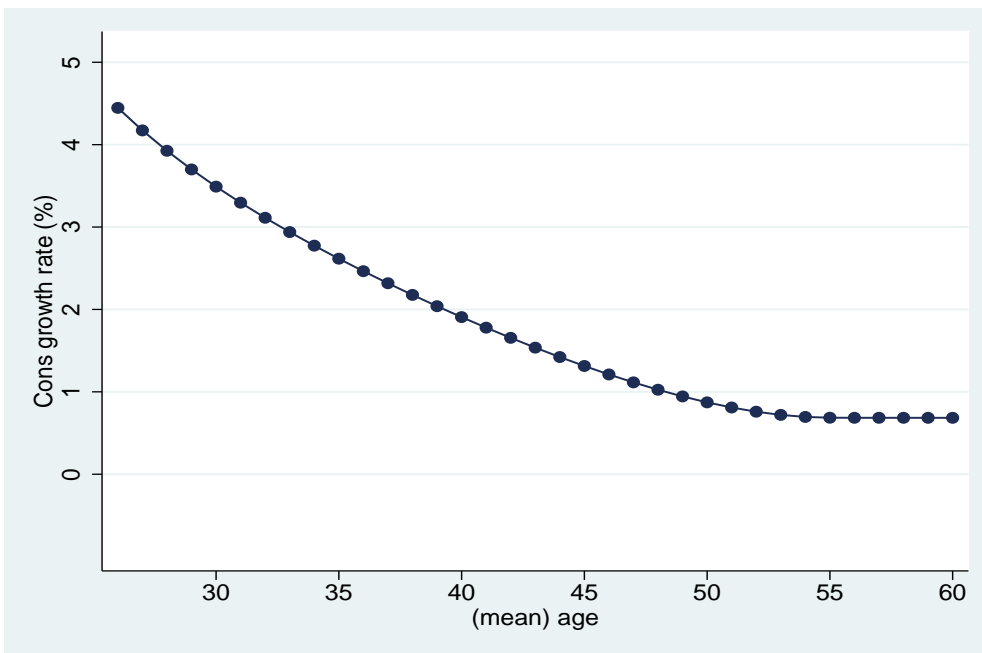Figure 15: Growth in real house prices and "aggregate income"

Figure 16: Growth in consumption with zero realised shocks

B   Estimating the idiosyncratic part of the income process

The estimation of the idiosyncratic element of the income process required panel data on family incomes. We used data from the British Household Panel Survey 1991-2005, and a measure of family (or, more properly, tax unit) non-investment income.

The estimation proceeded in two steps and was carried out separately for each education group. The first step was to regress income on a polynomial in age, and it turned out that in order to capture noticeable differences across education groups a cubic specification was required, but higher order terms did not enter significantly. Having fitted this regression, the persistent element of the process was then estimated as an AR(1) on the residuals from this regression. Results from this exercise are summarised in table 5.

Table 5: Estimated parameters of the idiosyncratic income process

	Low Education	High Education
age	-0.01068	0.01724
age squared	0.00184	0.00037
age cubed	-0.00005	-0.00002
ρ_y	0.76771	0.76365
σ_ξ^2	0.16580	0.15471

BHPS data were also used to fixed the initial level of the house price relative to expected income (this was matched to the ratio of average income to the average house value for those aged 22-26). The 2000 BHPS dataset was our source for the distribution of financial assets for the two education groups at the beginning of adult life. This distribution was matched to that of individuals aged 22-26 in the relevant education group and we assumed that households have zero housing endowments at age 22.

References

[1] ATTANASIO, O., J. BANKS, C. MEGHIR AND G. WEBER (1999), "Humps and bumps in lifetime consumption", *Journal of Business and Economic Statistics*, 17(1): 22-35.

[2] ATTANASIO, O., L. BLOW, R. HAMILTON AND A. LEICESTER (2005), "Consumption, House Prices and Expectations," IFS working paper WP05/24.

[3] ATTANASIO, O., R. BOTTAZZI, H. LOW, L. NESHEIM and M.WAKEFIELD (2007), *Explaining life-cycle patterns of home ownership*, mimeo: IFS.

[4] ATTANASIO, O. AND M. BROWNING (1995), "Consumption over the life cycle and over the business cycle", *American Economic Review*, 85(5): 1118-1137.

[5] ATTANASIO, O. AND G. WEBER (1994), "The UK Consumption boom of the late 1980's: aggregate implications of microeconomic evidence," *The Economic Journal*, 104(427): 1269-1302.

[6] ATTANASIO, O. AND G. WEBER (1995), "Is consumption growth consistent with intertemporal optimization? Evidence for the Consumer Expenditure Survey", *Journal of Political Economy*, 103, 1121-1157.

[7] BANKS, J., R. BLUNDELL, J. SMITH, AND Z. SMITH (2004), "House price volatility and housing ownership over the lifecycle," University College London, Department of Economics Discussion Paper 04-09.

[8] BANKS, J., Z. SMITH AND M. WAKEFIELD (2002), "The distribution of financial wealth in the UK: evidence from 2000 BHPS data", IFS Working Paper no. 02/21.

[9] BLUNDELL, R., L. PISTAFERRI AND I. PRESTON (2004), "Consumption inequality and partial insurance", Working Paper n.04/28, Institute for Fiscal Studies, London.

[10] BOTTAZZI, R., H. LOW and M.WAKEFIELD (2007), *Why do home-owners work longer hours?*, IFS working paper 07/10 (http://www.ifs.org.uk/wps/wp1007.pdf).

[11] CAMPBELL J.R., and Z. HERCOWITZ, (2004) "The Dynamics of Work and Debt", NBER working paper 10201.

[12] CAMPBELL, J.Y. AND J. COCCO (2007), "How Do House Prices Affect Consumption? Evidence From Micro Data", *Journal of Monetary Economics.*

[13] J. COCCO (2005), "Portfolio choice in the presence of housing", *The Review of Financial Studies.*

[14] CRISTINI, A. AND A. SEVILLA (2007), "Comparison Exercise", mimeo

[15] DAVIDSON, R. AND J. G. MACKINNON (2004), "*Econometric Theory and Methods* ", Oxford: Oxford University Press.

[16] De NARDI, M. (2004), "Wealth Inequality and Intergenerational Links," *Review of Economic Studies*, 71(3): 743-68.

[17] DEATON, A. (1991), "Saving and Liquidity Constraints", *Econometrica*, 59(5): 1221-1248.

[18] FLAVIN, M. AND S. NAKAGAWA (2004), "A model of housing in the presence of adjustment costs: a structural interpretation of habit persistence", NBER Working Paper no. 10458.

[19] FLAVIN, M. AND T. YAMASHITA (2002), "Owner-occupied housing and the composition of hte household portfolio", *American Economic Review*, 92(1): 345-362.

[20] FLODEN, M. (2007), *A Note on the Accuracy of Markov-Chain Approximations to Highly Persistent AR(1)-Processes*, SSE/EFI Working Paper Series in Economics and Finance No. 656 (http://swopec.hhs.se/hastef/papers/hastef0656.pdf)

[21] GOURINCHAS, P.O. AND J. PARKER (2002), "Consumption over the life cycle", *Econometrica*, 70(1): 47-89.

[22] GROSSMAN, S. AND G. LAROQUE (1990), "Asset pricing and optimal portfolio choice in the presence of illiquid housing consumption goods", *Econometrica*, 58(1): 25-51.

[23] LI W. and R. YAO (2007), "The life-cycle effects of house price shocks" *Journal of Money, Credit and Banking.*

[24] MUELLBAUER, J. AND A. MURPHY (1997), "Booms and Busts in the UK Housing Market," *The Economic Journal*, 107(445): 1701-1727.

[25] ORTALO-MAGNÉ, F. AND S. RADY (2002), "Tenure Choice and the riskiness of non-housing consumption," *Journal of Housing Economics*, 11: 266-279.

[26] ORTALO-MAGNÉ, F. AND S. RADY (2005), "Housing market dynamics : on the contribution of income shocks and credit constraints", *Review of Economic Studies*, 73(2): 459-485

[27] TAUCHEN, George (1986), "Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions", *Economics Letters* 20, 177-181

[28] TAUCHEN, George and Robert HUSSEY (1991), "Quadrature-Based Methods for Obtaining Approximate Solutions to Nonlinear Asset Pricing Models", *Econometrica* 59, 371-396