

Essays in Macroeconomics and Asset Pricing



Jason Lu

Faculty of Economics
University of Cambridge

This dissertation is submitted for the degree of
Doctor of Philosophy

To my family

Declaration

This dissertation is the result of my own work and includes nothing which is the outcome of work done in collaboration, except as declared in the Preface and specified in the text. It is not substantially the same as any that I have submitted, or, is being concurrently submitted for a degree or diploma or other qualification at the University of Cambridge or any other University or similar institution. It does not exceed the prescribed word limit of 60,000 words.

Jason Lu

July 2021

Acknowledgements

First and foremost I would like to thank my supervisor, Charles Brendon, for guiding me as his student and for his constant encouragement and support throughout my time in Cambridge. I am also deeply grateful to my advisor, Pontus Rendahl for his time and numerous discussions thorough out my PhD studies. I would also like to thank Giancarlo Corsetti and Martin Ellison for agreeing to be on my committee and for their time to read and comment on my work.

I would further like to thank the Faculty of Economics, in particular Silvana Dean, and my colleagues in the DAE for sharing many working days, countless coffee breaks and innumerable interesting discussions over the years.

Most of all, I am incredibly grateful to my family for their his continuous support and love throughout the years. I dedicate this thesis to them.

Abstract

The dissertation presents three essays on asset pricing in the context of Macroeconomics. Each chapter develops upon a central theme: that asset price bubbles act to drive fluctuations in the aggregate economy and conversely are themselves a symptom of economic conditions. The dissertation presents arguments and evidence in support of this theme. The implication is that these types of asset price fluctuations are important to study if we wish to develop a full understanding of the macroeconomy as a larger integrated system.

In my first chapter, I study the effects of bubbles on a secular stagnation economy. In such a setting, a negative natural rate of interest drives a deficit of demand and a persistent output gap. I find that bubbles act to increase the natural rate of interest, hence they alleviate the cause of secular stagnation. This suggests a positive role for bubbles, however bubbles are intrinsically unstable. Larger bubbles are shown to be more unstable, and upon collapse the natural rate of interest falls, potentially triggering secular stagnation.

In the second chapter, I offer a demographic explanation to the secular decline in interest rates. Fertility rates fell dramatically in the early 1970s, which created a distortion in the age distribution with the cohort born just before the fertility fall being disproportionately larger than the cohorts born both before and after. As this large cohort accumulates assets for retirement, their savings flood the capital market leading to a collapse in interest rates. The model offers an explanation for the fall in interest rates over the last three decades. Furthermore, it predicts that real interest rates will continue to fall until hitting a trough around the year 2035. Despite the fall in interest rates, the model does not provide a rationale for the rise in land prices. This chapter is a co-authored work with Coen Teulings.

In the final chapter, I develop a model for house price dynamics driven by short and long-term shocks to housing fundamentals. I consider households who have incomplete information and cannot directly observe underlying changes in fundamentals. Instead, they solve a Bayesian signal extraction problem to infer fundamentals from prices. In contrast to the complete information benchmark, my model replicates the observed empirical short-term momentum and long-term reversal in prices, which are often associated with speculation in housing markets. Furthermore, I show that a sequence of positive short term shocks generates an expectations-driven boom and bust in housing that is of far greater magnitude than in the complete information benchmark. Finally, I demonstrate an application of my model to identify episodes of bubbles in US house price data.

Contents

1	Bubbles in a Secular Stagnation Economy	1
1.1	Introduction	1
1.2	Bubbles and the natural rate of interest	3
1.2.1	The model	4
1.2.2	Loan market equilibrium	6
1.3	Bubbles in a secular stagnation economy	9
1.3.1	Price levels and the zero lower bound	9
1.3.2	Firms and the labour market	10
1.3.3	Monetary policy	11
1.3.4	Equilibrium in the secular stagnation economy	13
1.4	Bubble stability and fiscal policy	20
1.5	Conclusion	23
1.6	Appendix	24
1.6.1	Production economy with capital	24
1.6.2	Steady-state equilibrium and aggregate demand and supply	25
1.6.3	Aggregate supply	27
1.6.4	Aggregate demand	28
	References	30
2	The Decline in Fertility Rates and its Effects on Interest Rates and the Value of Land	32
2.1	Introduction	32
2.2	The evolution of age distributions	35
2.2.1	A model of the cohort size transition	37
2.3	OLG model of real interest rates	40
2.4	Application to fertility rate data	46
2.5	Extension to land	48
2.6	Conclusion	54

2.7	Appendix	55
2.8	BGP real interest rate	56
2.8.1	BGP real interest rate in the model with land	59
2.8.2	Adjustment algorithm	61
2.8.3	Alternative parameter specifications	63
	References	67
3	Incomplete Information and Housing Booms and Busts	69
3.1	Introduction	69
3.2	The model	73
3.2.1	Market structure	74
3.2.2	Information structure	76
3.3	Simulation and impulse response functions	80
3.3.1	Calibration of the incomplete information model	81
3.3.2	Simulations	83
3.3.3	Impulse response functions	86
3.3.4	Boom and bust experiment	89
3.4	Empirical application	91
3.4.1	Deconstruction of price data	91
3.4.2	Identifying bubble episodes	95
3.5	Model extensions	96
3.5.1	Intraperiod expectation updating	96
3.5.2	Simple production economy with mortgage contracts	98
3.5.3	Short-term component persistence	100
3.6	Conclusion	101
	References	102

Chapter 1

Bubbles in a Secular Stagnation Economy

1.1 Introduction

The secular stagnation hypothesis argues that low interest rates can be the cause of a persistent fall in output. Specifically, Summers (2013) argues that long-run trends have pushed the natural rate of interest, the real interest rate consistent with full output, to be negative.¹ This negative natural rate of interest poses a problem because monetary policy, when controlling the nominal interest rate, is constrained below by the Zero Lower Bound (ZLB). As a result, the central bank can only deliver at minimum a real interest rate of zero minus the rate of inflation. When this minimum level is above the natural rate of interest, there is necessarily a mismatch in interest rates which results in an output gap. The resulting output gap persists as long as the mismatch remains, which could be until some unknown date in the far future.

Summers (2013) discusses the relevance of bubbles to the secular stagnation hypothesis, however the precise mechanism by which bubbles play a role in secular stagnation is not yet understood, nor formally addressed. In this chapter I propose such a mechanism. Building upon the secular stagnation model by Eggertsson and Mehrotra (2014), I find that bubbles act to increase the natural rate of interest by providing additional assets which absorb excess savings. This directly alleviates the initial cause of secular stagnation, namely a negative natural rate of interest. I find that a bubble of sufficient size can eliminate a secular stagnation equilibrium altogether, although this solution to the problem is not without its trade-offs. I show that the stability of a bubble is negatively associated with its size, and that when a bubble collapses, the natural rate of interest falls potentially resulting in secular stagnation.² Government enforced

¹The real interest rate equals to the natural rate of interest only when there is full output. Nonetheless the fall in real interest rates is indicative of a fall in the natural rate of interest, under the assumption that the economy has been on average operating at full output.

²This mechanism is in line with empirical estimates of the natural rate of interest. Laubach and Williams

fiscal policies, of the type considered in Samuelson (1958), can likewise increase the natural rate of interest by a similar mechanism. These policies can generate the same desirable outcome of a bubbly equilibrium, but without the potential instability that bubbles bring.

I do not look to explain the long-run fall in the natural rate of interest. Instead, I take a negative natural rate of interest as given, and then analyse the effects of bubbles in this environment. The model by Eggertsson and Mehrotra (2014) serves as my point of departure. Within their Keynesian overlapping generations framework, a sufficiently negative natural rate of interest results in an output gap of indefinite duration. In this equilibrium, monetary policy is forever at the ZLB, but nonetheless there is no self-correcting force to full employment. My analysis begins by allowing for the trade of a bubbly asset. I consider rational bubbles of the kind studied by Tirole (1985), where the underlying asset generates no dividends and therefore has a fundamental value of zero. Although this asset has no inherent value, a household is willing to purchase this asset under the rational expectation that they can sell this asset at a positive price in the future. In such a setting, there is a multitude of equilibria, one for each feasible size of the bubble up to an upper bound which I find analytically. I do not attempt to explain the origins of the bubble. While I consider rational bubbles in my model, the same mechanism would apply if the bubble was irrational. Indeed, a bubble regardless of origin, has a positive effect on the natural rate of interest, which has reduces the output gap through the standard secular stagnation channel described in Eggertsson and Mehrotra (2014).

The literature on secular stagnation is closely related to the previous literature on liquidity traps, which differs to secular stagnation mainly in terms of the duration of the output gap (see e.g. Krugman (1998) and Eggertsson and Krugman (2012)). Originally referencing Japan in the 1990s, the liquidity trap literature considers similarly the effect of a negative natural rate of interest, coming to broadly the same conclusions. Their analysis, however, assumes the negative natural rate of interest to be temporary and of known duration, hence when the natural rate of interest returns to normal levels, the output gap will similarly close. Our experience since then has shown that a negative natural rate of interest may not be so short-lived, and certainly it is difficult to forecast when interest rates will return to “normal” levels. For example, Japan’s low interest rates has persisted since the early 1990s. While duration is the primary difference between the two branches of the literature, the implications of the liquidity trap literature can differ greatly once allowing for an indefinitely negative natural rate of interest. In particular, a rational bubble is unsustainable when the natural rate of interest is high. Hence, if an increase in the natural rate of interest is foreseeable or expected, then by backwards induction the bubble must not be feasible under rational expectations today. That is to say, the mechanism for rational

(2016) provide estimates of the natural rate of interest in the U.S., which show a sudden fall in the natural rate of interest following the collapse of the housing bubble in 2007.

bubbles considered in this chapter can only apply under an indefinitely negative natural rate of interest, and will not hold under a temporarily negative natural rate of interest as is considered in the liquidity trap studies.

My analysis relates also to the literature investigating the role of bubbles in driving the business cycle (see e.g. Ventura (2003) and Martin and Ventura (2012)), which looks to resolve the contrary prediction of Samuelson-Tirole models that bubbles are contractionary. Specifically, standard Samuelson-Tirole models find that bubbles increase the real interest rate, which reduces the equilibrium level of capital and therefore lowers output. This is the reverse of our experience, where we witness that episodes of bubbles are associated with investor optimism, high investment (as opposed to low), and also high growth.³ My chapter offers an alternative mechanism of how bubbles can be expansionary. By increasing the natural rate of interest, trade in the bubbly asset can increase output despite the assets having no inherent value. The caveat, however, is that this channel operates only under the condition that the natural rate of interest in the fundamental economy is sufficiently negative. Whereas this is a reasonable premise for recent decades, it will not apply in past business cycles when interest rates were higher.

The structure of the rest of this chapter shall be as follows. In section two I describe a basic endowment economy model which demonstrates the effect of a bubble on the natural rate of interest. In section three I describe a Keynesian production economy, which shows that a bubble increases output in a secular stagnation economy by increasing the natural rate of interest. In section four I consider the stability of bubbles and show that the stability of a bubble is negatively associated with its size. Further I give a brief discussion on the role of fiscal policy. In section five I conclude.

1.2 Bubbles and the natural rate of interest

In this section I show that bubbles act to increase the natural rate of interest by providing an additional store of value for saving generations. This simple mechanism provides the main insight of how bubbles relate to the secular stagnation hypothesis.

I consider an endowment economy that follows closely the endowment model by Eggertsson and Mehrotra (2014). Here, I do not consider production, monetary policy, nor the ZLB. Income is endowed and hence exogenously determined, so the only question that remains to be answered is: what real interest rate clears markets? The market-clearing real interest rate is also the real interest rate consistent with full output in the full model, i.e. the natural rate of interest. Hence, the analysis in this section directly maps to the mechanics for the natural rate of interest in the full model considered later.

³In welfare terms, bubbles are often positive even in these settings as they exist only when the economy is dynamically inefficient.

Eggertsson and Mehrotra (2014) represent the negative natural rate of interest as the result of a reduction in the debt limit, D , which acts to impose loan market frictions that lowers borrowing demand. Instead I propose to think of the parameter D as a capture-all variable representing various forces that have acted to reduce the natural rate of interest over recent decades. These include, for example, ageing demographics, slowing technological progress, or a glut of external savings, to name a few. I prefer this interpretation as it maps closer to the original Summers (2013) discussion, and because these forces are long term and therefore can better support a rational bubble. However, the interpretation is not central to the analysis and is mostly a matter of taste. My main point is that I take a negative natural rate of interest in the fundamental economy as given, and within this setting I analyse the effect of bubbles.

The bubble part of the model draws from the previous literature on rational bubbles. In particular we know that rational bubbles are sustainable in overlapping generations economies when the economy is dynamically inefficient, i.e. when the real interest rate, r , is less than the growth rate of the economy, g . Bubbles act to increase r , which takes on a new meaning when the real interest rate is mapped to the natural rate of interest in the full model.

1.2.1 The model

Households live for three periods. They are initially young, in the next period they are middle aged, and eventually they are old. Generations may borrow and lend through a bond market: the young borrow subject to a debt limit D , the middle aged receive an endowment income some of which they save for retirement, and the old consume their savings.

Households born in period t have the following utility function

$$\log(C_t^y) + \beta \log(C_{t+1}^m) + \beta^2 \log(C_{t+2}^o), \quad (1.1)$$

where $C_t^y, C_{t+1}^m, C_{t+2}^o$ denote respectively the levels of consumption for the cohort born in period t , in period t when young, in period $t + 1$ when middle aged, and in period $t + 2$ when old.

Borrowing and lending occurs through the trade of a one period risk-free bond at the market real interest rate r_t . With three generations, the middle aged accesses the bond market to buy bonds from the young. Next period, the previously middle aged (now old) is repaid for their loan by the previously young (now middle aged).

In addition to the bond market, I allow for the trade of a bubbly asset which generates no dividends. In each period the middle aged purchase the bubbly assets from the old. This rewards the old as they consume the price of the bubbly assets they sell. Likewise the middle aged will resell their bubbly assets to the next generation in retirement. Altogether the bubbly asset provides an additional asset to save with. The budget constraints faced by households

born at period t are given by

$$C_t^y = B_t^y \quad (1.2)$$

$$C_{t+1}^m = Y_{t+1} - (1 + r_t)B_t^y + B_{t+1}^m - p_{t+1}A_{t+1}^m \quad (1.3)$$

$$C_{t+2}^o = -(1 + r_{t+1})B_{t+1}^m + p_{t+2}A_{t+1}^m \quad (1.4)$$

$$(1 + r_t)B_t^i \leq D_t, \quad (1.5)$$

where B_t^i represents the bond issuance (or borrowing) by generation i in period t , A_t^m represents the quantity of bubbly assets purchased at period t by the middle aged, p_t represents the price of the bubble at period t , and D_t represents the exogenous borrowing limit at time t .

The debt limit depends on the young's ability to repay in the next period and it includes debt appreciation. Furthermore I assume that the debt limit is binding for the young, hence

$$B_t^y = \frac{D_t}{(1 + r_t)}. \quad (1.6)$$

Denote by N_t the size of the middle-aged cohort at time t , which evolves according to

$$N_t = (1 + g_t)N_{t-1}, \quad (1.7)$$

where g_t is the growth rate of cohort size from period $t - 1$ to period t . Market clearing for bonds implies

$$(1 + g_t)B_t^y = -B_t^m \quad (1.8)$$

Taking the supply of the bubble to be fixed at $\bar{A}$, market clearing for the bubbly asset implies that the middle aged purchases the entire supply from the old, hence

$$A_t^m = \frac{\bar{A}}{N_t}. \quad (1.9)$$

Bubbles generate no utility, and with perfect foresight they also are not expected to collapse. Hence, owning the bubbly asset is a perfect substitute to owning bonds and no-arbitrage between these two assets gives

$$p_{t+1} = (1 + r_t)p_t. \quad (1.10)$$

The old will always consume all remaining wealth, hence

$$C_t^o = -(1 + r_{t-1})B_{t-1}^m + p_t A_{t-1}^m, \quad (1.11)$$

whereas the middle aged are at an interior solution that follows the Euler equation,

$$\frac{1}{C_t^m} = \frac{\beta(1+r_t)}{C_{t+1}^o}, \quad (1.12)$$

which follows from the utility function given by equation (1.1).

1.2.2 Loan market equilibrium

Definition 1. *An equilibrium in the endowment economy is defined as a path for the collection of quantities $\{C_t^y, C_t^m, C_t^o, B_t^y, B_t^m, A_t^m\}$, and a path for prices $\{r_t, p_t\}$, that solve equations (1.3), (1.4), (1.5), (1.6), (1.7), (1.8), (1.9), (1.10), (1.11), and (1.12) given an exogenous path for $\{Y_t, D_t, g_t\}$.*

To solve for the equilibrium path in the endowment economy model, it is sufficient to consider market clearing in the bond market. This requires that the bond supply of the young matches the bond demand of the middle aged

$$(1+g_t)B_t^y = -B_t^m. \quad (1.13)$$

Denoting the left-hand side of the above equation to be the demand for loans, L_t^d , and the right-hand side of the equation to be the supply for loans, L_t^s , it follows from equation (1.6) that

$$L_t^d = \frac{1+g_t}{1+r_t} D_t, \quad (1.14)$$

and using the Euler equation given by equation (1.12), substituting for C_t^m and C_{t+1}^o using equations (1.3), (1.4), and (1.6) to solve for B_t^m and find that the the supply of loans is given by

$$L_t^s = \frac{\beta}{1+\beta} (Y_t - D_{t-1}) - \mathcal{A}_t, \quad (1.15)$$

where $\mathcal{A}_t \equiv \frac{p_t \bar{A}}{N_t}$, is the value (or size) of the bubble per capita, and $\mathcal{A}_t = 0$ represents the fundamental economy case with the value of the bubbly asset equal to zero.

Note that the introduction of the bubbly asset does not affect the demand for loans. The young borrow the maximum they are able to subject to the borrowing limit which is assumed to be binding. An increase in the size of the bubble per capita reduces the supply of loans one-to-one, as the bubbly asset is a perfect substitute for bonds at any given real interest rate.

Equating the supply and demand for loans in equations (1.14) and (1.15) gives the equilibrium real interest rate

$$1+r_t = \frac{(1+g_t)D_t}{\frac{\beta}{1+\beta}(Y_t - D_{t-1}) - \mathcal{A}_t}, \quad (1.16)$$

which says that the equilibrium real interest rate is an increasing function of the size of the bubble, $\mathcal{A}_t$.

Proposition 1. *Given constant levels for $\{Y_t, D_t, g_t\} = \{Y, D, g\} \forall t$, if*

$$1 + r^f \equiv \frac{(1 + g)D}{\frac{\beta}{1+\beta}(Y - D)} < 1 + g,$$

there is a set of equilibria, with each equilibrium characterised by a particular size of the bubble per capita at time t , $\mathcal{A}_t$, in the following range

$$\mathcal{A}_t \in \left[0, \frac{\beta Y - (1 + 2\beta)D}{1 + \beta}\right].$$

For a given $\mathcal{A}_t$, the equilibrium real interest rate at time t is

$$1 + r_t = \frac{(1 + g)D}{\frac{\beta}{1+\beta}(Y - D) - \mathcal{A}_t},$$

which is an increasing function of the size of the bubble per capita, $\mathcal{A}_t$.

Proof. The proof is similar to that of Tirole (1985). Firstly I claim that for a bubble to be sustainable in equilibrium, r_t must never exceed g .

Consider the path of $\mathcal{A}_t$. Because $\mathcal{A}_t \equiv \frac{p_t \bar{A}}{N_t}$, the size of the bubble per capita grows at rate $r_t - g$. Suppose that this growth rate was positive for any one period, i.e. if $r_t > g$ for some t , then in the next period $\mathcal{A}_{t+1}$ must be strictly greater than $\mathcal{A}_t$. If $\mathcal{A}_{t+1} > \mathcal{A}_t$, then by equation (1.16) it follows that $r_{t+1} > r_t$. Hence by induction, if $r_t > g$ for one period, it follows that $\{r_i\}_{i>t} > r_t > g$.

Knowing that r^f , the equilibrium real interest rate in the fundamental economy (with $\mathcal{A}_t = 0$), is strictly less than g , then $r_t > g$ implies $\mathcal{A}_t > 0$. So if $r_t > g$ for one period, then $\mathcal{A}_t > 0$ and $\{r_i\}_{i>t} > r_t > g$. Then it must be the case that at some point in the future $\mathcal{A}_t$ exceeds $\frac{\beta}{1+\beta}(Y - D)$, i.e. the middle aged are no longer willing to buy up the entire stock of the bubbly asset. At this point the bubble is no longer sustainable, and the final old generation can no longer sell their holdings of the bubble. As households have perfect foresight, that old generation who loses out on their savings would not purchase the bubble in the first place, and by backwards induction it must be the case that no households, even now, would value the bubbly asset. Hence for rational bubbles to emerge, it must be the case that $r_t \leq g$ for all periods. This implies a maximum size for the size of the bubble $\mathcal{A}_t \leq \frac{\beta Y - (2+\beta)D}{1+\beta}$.

Given that $\mathcal{A}_t \leq \frac{\beta Y - (2+\beta)D}{1+\beta}$, then in the next period $r_t \leq g$. Hence the size of the bubble per capita (weakly) gets smaller, and hence r_t is (weakly) decreasing. By induction r_t is decreasing with time. Given this, the size of the bubble is likewise permanently (weakly) getting smaller,




Figure 1.1: Loan market equilibrium following a decrease in D .




Figure 1.2: Loan market equilibrium following an increase in bubble size.

and hence the middle aged will always be able to purchase the entire stock of the bubbly asset. This ensures that the bubble is sustainable.

When $r_t < g$, the bubble shrinks at an increasing rate of $r_t - g$, where asymptotically r_t tends to r^f and $\mathcal{A}_t$ tends to 0. When $r_t = g$, the size of the bubble is constant and therefore r_t is constant at the level of g . $\square$

The natural rate of interest in the fundamental equilibrium of this simple economy is given by $1 + r^f \equiv \frac{(1+g)D}{\frac{\beta}{1+\beta}(Y-D)}$, which is a function of the parameters g , Y , and D . We see in Figure 1.1 that a fall in D leads to a reduction in loan demand and an increase in loan supply. The result is a fall in the real interest rate to a level that is potentially negative.

From Proposition 1 we see that a bubble of some size is sustainable whenever r^f , the real interest rate in the fundamental economy, is less than g , the growth rate of the population. In a situation of secular stagnation, r^f is strongly negative, and hence this condition is more likely fulfilled. Also from Proposition 1, we know that any bubble of size up to a maximal level $\frac{\beta Y - (1+2\beta)D}{1+\beta}$, is sustainable. Figure 1.2 shows the effect of an increase in bubble size on the equilibrium real interest rate.

An increase in the size of the bubble leads to a decrease in loanable fund supply, and hence an increase in the real interest rate. This comparative exercise explains the difference between a bubbly equilibrium versus a fundamental equilibrium. If one introduces a bubble to a fundamental economy, the result would be an increase in the real interest rate.

Likewise, the reverse scenario to Figure 1.2 informs us that a fall in the size of the bubble leads to a fall in real interest rate. In particular, this says that if a bubble collapses, the real interest rate would revert back to the lower fundamental level r^f .

1.3 Bubbles in a secular stagnation economy

Given the effect of bubbles on the natural rate of interest, what are the implications for the wider economy? Eggertsson and Mehrotra (2014) demonstrate that a negative natural rate of interest may lead to secular stagnation in equilibrium. In this section I show that trade in bubbles can lift the natural rate of interest and hence restore full output, however the collapse of a bubble has the opposite effect and can result in a secular stagnation equilibrium.

To formally characterise a secular stagnation economy, a few minimal Keynesian ingredients are necessary. These ingredients generate an environment where a strongly negative natural rate of interest can generate a secular stagnation equilibrium with a binding ZLB and a persistent output gap. I give a description and a brief discussion for each ingredient below.

1.3.1 Price levels and the zero lower bound

A nominal price level is introduced by assuming that the one period bond is denominated in money, the interest rate of which is determined by the central bank. This gives the nominal analog of the Euler equation

$$\frac{1}{C_t^m} = \beta(1 + i_t) \frac{P_t}{P_{t+1}} \frac{1}{C_{t+1}^o}, \quad (1.17)$$

where i_t and P_t are respectively the nominal interest rate and the price level. The ZLB on nominal interest rates, due to no-arbitrage between holding money versus holding bonds, is imposed by

$$i_t \geq 0. \quad (1.18)$$

Euler equations (1.12) and (1.17) together give the Fisher equation

$$1 + r_t = (1 + i_t) \frac{P_t}{P_{t+1}}. \quad (1.19)$$

The ZLB on nominal interest rates implies the following lower bound for real interest rates

$$1 + r_t \geq \frac{1}{\Pi_{t+1}},$$

where Π_{t+1} is the inflation rate in period $t+1$ defined by $\Pi_t \equiv \frac{P_{t+1}}{P_t}$. Hence for a given inflation rate, the ZLB on nominal interest rates implies a corresponding lower bound for real interest rates. In particular if the inflation rate is low, the lower bound imposed on real rates is harsher. When the lower bound on real interest rates is above the natural rate of interest, then there is by definition of the natural rate of interest an output gap.

Already one can understand the problem of the ZLB in the situation of a negative natural rate of interest. However, the inflation rate is endogenously determined in equilibrium by the meeting of aggregate demand with aggregate supply. To understand the effect of ZLB, it yet requires a description of the production side of the economy.

1.3.2 Firms and the labour market

Firms hire labour to produce output. Each firm's production function features decreasing marginal productivity, therefore firms earn a profit. Every middle-aged household is endowed with $\bar{L}$ units of labour, which they supply inelastically to firms, and the ownership of one firm's associated profits. Note that although capital does not feature in this production function, it's inclusion does not change the mechanisms in this model (although it will affect the likelihood of dynamic inefficiency, see the Appendix for a derivation of the model with capital).

The budget constraint for the middle aged becomes

$$C_{t+1}^m = \frac{W_{t+1}}{P_{t+1}}L_{t+1} + \frac{Z_{t+1}}{P_{t+1}} - (1+r_t)B_t^y + B_{t+1}^m - p_{t+1}A_{t+1}^m, \quad (1.20)$$

where Z_t is the nominal profit of the firms, and L_t is the quantity of labour supplied to the firm. Note that although labour supply is inelastic, firms may not be able to hire all of the labour supply due to wage frictions, hence $L_t \leq \bar{L}$. Firms have the following production function

$$Z_t = \max_{L_t} P_t Y_t - W_t L_t$$

subject to

$$Y_t = L_t^\alpha,$$

and the firm's first-order condition with respect to labour hired gives labour demand as the following function of the real wage

$$\frac{W_t}{P_t} = \alpha L_t^{\alpha-1}. \quad (1.21)$$

Supposing that the economy was frictionless, in equilibrium the nominal wage will adjust such that all of the available labour supply is hired, with $\frac{W_t}{P_t} = \alpha \bar{L}^{\alpha-1}$. In this case there is no

unemployment and hence production is always at its maximum (potential) level. This is exactly equivalent to the endowment economy considered in section two, therefore the equilibrium real interest rate without frictions is equal to the natural rate of interest. Following Eggertsson and Mehrotra (2014), a nominal friction in the form of a wage rigidity is introduced. Due to this rigidity it is now possible that the nominal wage may fail to adjust to clear the labour market. Wages follow

$$W_t = \max\{\bar{W}_t, P_t \alpha \bar{L}^{\alpha-1}\}, \text{ where } \bar{W}_t \equiv \gamma W_{t-1} + (1 - \gamma) P_t \alpha \bar{L}^{\alpha-1}. \quad (1.22)$$

The above process says that the nominal wage tries to adjust to clear the labour market, however it is downwardly bounded by a wage norm. The wage norm reflects a compromise between the wage last period, and the market-clearing wage this period. This compromise is summarised by the parameter γ , where full downwards rigidity case is given by $\gamma = 1$, and $\gamma = 0$ gives the flexible wage case.

Note that in order for the above wage norm to be binding in a steady state equilibrium, there needs to be deflation. With any positive inflation rate, the wage last period would be less than the wage this period, hence the wage norm cannot be binding. This is not to say that in general, a secular stagnation equilibrium cannot occur with positive inflation. It is only for simplicity that this particular wage norm does not adjust for nominal wage growth, in which case positive inflation can still be consistent with a binding wage norm.

1.3.3 Monetary policy

The central bank sets the nominal interest rate following a Taylor rule that is constrained by the ZLB

$$1 + i_t = \max \left(1, (1 + i^*) \left(\frac{\Pi_{t+1}}{\Pi^*} \right)^{\phi_\Pi} \right), \quad (1.23)$$

where $\phi_\Pi > 1$, Π^* , and i^* are parameters of the policy rule. The Taylor rule attempts to stabilise inflation around an inflation target Π^* using the nominal interest rate. As $\phi_\Pi > 1$, when inflation increases above Π^* , the nominal interest rate increases by a bigger amount. As the real interest rate is given by the nominal interest rate minus the rate of inflation, the Taylor rule ensures that following an increase in the inflation rate, the net effect on real interest rates is an increase. Conversely suppose that the rate of inflation falls below Π^* , then the nominal interest will fall by a bigger amount to deliver a lower real interest rate. For the second regime, the nominal interest rate when responding to deflation can only fall to zero, after which the nominal interest rate cannot further adjust. Further reductions in Π_{t+1} results in an increase of the real interest rate, as opposed to a decrease which the Taylor rule is designed to achieve.

The level of inflation at which the ZLB becomes binding is denoted to be Π_{kink} . This is




Figure 1.3: Taylor rule subject to the zero lower bound.

given by

$$\Pi_{kink} = \left(\frac{1}{1 + i^*} \right)^{\frac{1}{\phi_{\Pi}}} \Pi^*. \quad (1.24)$$

The ZLB, together with the Taylor rule, has implications on the set of *steady-state* real interest rates that a central bank can possibly deliver. Consider the steady-state inflation rate given by $\Pi = \Pi_{kink}$. For an increase in the inflation rate above Π_{kink} , $\phi_{\Pi} > 1$ means that the Taylor rule delivers a higher steady-state real interest rate than at Π_{kink} . Likewise for a fall in the inflation rate below Π_{kink} , the ZLB means that the resulting steady-state real interest rate is also higher than at Π_{kink} . Hence at an inflation rate of Π_{kink} , the steady-state real interest rate is at its lowest at $\Pi_{kink}^{-1} - 1$. See Figure 1.3 for a representation of the policy function in $(1 + r, \Pi)$ space.

Without the Taylor rule, it is possible to achieve a steady-state real interest rate of below $\Pi_{kink}^{-1} - 1$. For example imagine that the central bank always minimises the real interest rate by setting the nominal interest to zero. Then any steady-state real interest rate, however negative, is potentially possible with a high enough level of inflation. In general, any policy function curve is possible so long as it stays to the right of the line joining A and B in Figure 1.3. It is only because the central bank is unwilling to tolerate high inflation that there is a lower bound to the real interest rate that can be delivered.

Given that there is a minimum steady-state real interest rate that can be delivered by the central bank, by definition we cannot have full output if the steady-state natural rate of interest is below $\Pi_{kink}^{-1} - 1$. Note that this is irrespective of the inflation rate, whereas the lower bound derived earlier is conditional on the rate of inflation.

1.3.4 Equilibrium in the secular stagnation economy

Definition 2. *An equilibrium in the production economy model is defined to be a path for the following quantities $\{C_t^y, C_t^m, C_t^o, B_t^y, B_t^m, A_t^m, L_t, Y_t, Z_t\}$, and a path of the following prices $\{P_t, W_t, p_t, r_t, i_t\}$, that solve equations (1.20), (1.6), (1.7), (1.8), (1.9), (1.10), (1.11), (1.12), (1.19), (1.21), (1.22), and (1.23), given an exogenous path for $\{D_t, g_t\}$.*

In particular, as I want to consider a secular stagnation equilibrium that can hypothetically last indefinitely, it is useful to examine my argument in a steady-state setting. To find a steady-state equilibrium with bubbles, however, it is necessary to adjust my initial assumption of a fixed supply for the bubbly asset.

What would happen if the supply of the bubbly asset was fixed? From section 2, loan supply is given by equation (1.15)

$$L_t^s = \frac{\beta}{1+\beta}(Y_t - D_{t-1}) - \mathcal{A}_t,$$

where $\mathcal{A}_t = \frac{p_t A_t}{N_t}$, and the supply of the bubbly asset is taken to be fixed at $\bar{A}$. The size of the bubble, $\mathcal{A}_t$, grows at the rate of $\frac{1+r_t}{1+g_t}$. In order for the bubble to be sustainable it needs to be the case that $r_t - g_t \leq 0 \forall t$. A steady-state equilibrium requires the equilibrium real interest rate to be constant, i.e. $r_t = r \forall t$, but for a particular level of g it will not generally be true that $r - g = 0$ exactly. In fact, as secular stagnation equilibria generally have $r < 0$, the case of $r - g = 0$ is uninteresting for my analysis (assuming $g > 0$). Hence, the inequality is usually strict. When $r - g < 0$, the size of the bubble per capita is decreasing, and therefore loan supply is increasing. We know that the equilibrium real interest rate in the endowment economy, i.e. the natural rate of interest, is given by equation (1.16)

$$1 + r_t = \frac{(1+g)D_t}{\frac{\beta}{1+\beta}(Y_t - D_{t-1}) - \mathcal{A}_t},$$

which must similarly be decreasing over time. In fact, Tirole (1985) shows that a bubble with $r_t - g_t \leq 0$ must eventually tend to zero. This means that a fixed supply of the bubble is inconsistent with a steady-state equilibrium.

To get around this problem, I now consider the supply of the bubbly asset A_t to grow over time. In particular, the supply of the bubble is to grow at a rate of $\frac{g_t - r_{t-1}}{1+r_{t-1}} \geq 0$. This forces the size of the bubble per capita to be constant, which is a necessary condition for the existence of a steady-state equilibrium. Again the middle aged purchases the entire supply of the bubbly asset. What differs is that the young generation (overall) is given an additional endowment of $\frac{g_t - r_{t-1}}{1+r_{t-1}} A_{t-1}$ units of the bubbly asset. That is to say, they are given the additional units of the bubbly asset, $A_t - A_{t-1}$. The size of the bubble (per capita) is now constant at a level which I call $\mathcal{A}$.

I assume for simplicity that the young are given this endowment of the bubbly asset. The advantage of this is that the young's consumption is already at a corner solution. If we assume that the young's consumption is still at a corner solution after being given this endowment (which will remain the case as long as the debt limit is sufficiently binding), the solution to the household remains simple as the additional bubble endowment will not change the Euler equation from section 2. Note, however, that the qualitative effect of the bubble is robust to where the endowment of the bubble is given. This assumption is made only to simplify the analysis.

Supposing that the young remain at a corner solution, replacing equation (1.6) the consumption of the young increases due to their new endowment

$$C_t^y = \frac{D_t}{1 + r_t} + p_t \frac{g_t - r_{t-1}}{1 + r_{t-1}} \frac{A_{t-1}}{N_{t+1}}. \quad (1.25)$$

With these adjustments, I define a steady-state equilibrium in the production economy model as follows.

Definition 3. A steady-state equilibrium is defined to be a path for the following quantities $\{C_t^y, C_t^m, C_t^o, B_t^y, B_t^m, A_t^m, L_t, Y_t, Z_t\}$, and for the following prices $\{P_t, W_t, p_t, r_t, i_t\}$, where each quantity/price is constant or grows with constant rate, that solve (1.20), (1.25), (1.7), (1.8), (1.9), (1.10), (1.11), (1.12), (1.19), (1.21), (1.22), and (1.23), given constant levels for $\{D_t, g_t\} = \{D, g\}$.

Proposition 2. Given constant levels for $\{D_t, g_t\} = \{D, g\} \forall t$, if

$$1 + r^f \equiv \frac{(1 + g)D}{\frac{\beta}{1+\beta}(Y^n - D)} < 1 + g,$$

where $Y^n \equiv \bar{L}^\alpha$, there exists a set of steady-state equilibria, with each equilibrium characterised by a particular size of the bubble per capita, $\mathcal{A}$, in the following range

$$\mathcal{A} \in \left[0, \frac{\beta Y^n - (1 + 2\beta)D}{1 + \beta}\right].$$

For a given $\mathcal{A}$, the natural rate of interest is defined to be

$$r^n = \frac{(1 + g)D}{\frac{\beta}{1+\beta}(Y^n - D) - \mathcal{A}} - 1,$$

which is an increasing function of the size of the bubble per capita, $\mathcal{A}$. When the natural rate of interest is below $\Pi_{kink}^{-1} - 1$, given by equation (1.24), output is increasing in the natural rate of interest. When the natural rate of interest is above $\Pi_{kink}^{-1} - 1$, output is at its maximum potential level.

Proof. The first part of Proposition 2 follows exactly the proof for Proposition 1. The second part of Proposition 2 follows from the argument given in the discussion of the Taylor rule. $\square$

Proposition 2 states that when the natural rate of interest is below $\Pi_{\text{kin}}^{-1} - 1$, a further fall in the natural rate of interest leads to a fall in output. As bubbles act to increase the natural rate of interest, an increase in the size of the bubble can lead to an increase in output. The reason this occurs is due to the ZLB on nominal interest rates. In order for the central bank to maintain full output, the real interest rate delivered in the economy needs to match the natural rate of interest. However the lowest real interest rate that the central bank can deliver (due to the Taylor rule) is $\Pi_{\text{kin}}^{-1} - 1$. Hence a fall in the natural rate of interest beyond $\Pi_{\text{kin}}^{-1} - 1$ leads to a contraction in output.

To demonstrate Proposition 2 in a aggregate demand and supply framework, I derive the steady-state aggregate supply and aggregate demand curves as a function of the size of the bubble $\mathcal{A}$. Aggregate supply and aggregate demand meet to determine the steady-state equilibrium, and analysing the effect of bubbles on these curves allows one to deduce its effect on output,

1.3.4.1 Steady-state aggregate supply and demand

Steady-state aggregate supply is not affected by the introduction of the bubble. Aggregate supply is divided into two regimes, one when the wage norm in equation (1.22) is not binding and there is full output, and the other when wage is above market clearing and there is an output gap.

For the first regime where there is full output ($Y = \bar{L}^\alpha$), it follows that the wage norm is not binding in the steady state. This gives the following condition

$$\frac{W}{P} \geq \gamma \frac{W}{P} \Pi^{-1} + (1 - \gamma) \alpha \bar{L}^{\alpha-1}.$$

As the wage norm is not binding, $\frac{W}{P} = \alpha \bar{L}^{\alpha-1}$, therefore the condition above is satisfied whenever $\Pi \geq 1$. This implies that the division of the aggregate supply curve into its two regimes is separated by when $\Pi < 1$, and when $\Pi \geq 1$. The first regime is given by

$$Y = \bar{L}^\alpha \equiv Y^n \quad \text{for} \quad \Pi \geq 1. \tag{1.26}$$

Now consider the scenario when the wage norm is binding with $\Pi < 1$, then it follows that $W = \bar{W}$. Therefore the real wage is given by

$$\frac{W}{P} = \frac{(1 - \gamma) \alpha \bar{L}^{\alpha-1}}{1 - \gamma \Pi^{-1}}.$$

For a particular real wage, the quantity of labour hired is given by equation (1.21). Com-

binning this with the production function gives

$$\frac{\gamma}{\Pi} = 1 - (1 - \gamma) \left(\frac{Y}{Y^f} \right)^{\frac{1-\alpha}{\alpha}} \quad \text{for } \Pi < 1, \quad (1.27)$$

which defines the second regime of the aggregate supply curve where there is an output gap. See the Appendix for a full derivation of equation (1.27).

The two regimes given by equations (1.26) and (1.27) together specify the steady-state aggregate supply curve, which is shown in (Y, Π) space in figure 1.4. The first regime represents the vertical part of the aggregate supply curve. Here, $\Pi > 1$, the wage norm is not binding, and therefore there is market clearing in the labour market with $Y = \bar{L}^\alpha$. The second regime has a positive slope and represents when the wage norm is binding. To explain the positive slope, consider a further decrease of $\Pi < 1$. The rate of deflation increases resulting in a higher steady-state real wage. This leads to a lower level of labour employed, and hence a further output reduction. Therefore the fall in Π leads to a fall in Y .

Now I describe the steady-state aggregate demand curve which is affected by the introduction of bubbles. Combining the policy function given by equation (1.23), the Fisher equation given by equation (1.19), and summing up consumption across all three cohorts using equations (1.25), (1.3) and (1.11), I find

$$Y = D + \frac{(1 + \beta)(1 + g)D}{\beta(1 + r)} + \frac{1 + \beta}{\beta} \mathcal{A}. \quad (1.28)$$

Aggregate demand is also separated into two regimes. The first regime corresponds to when the Taylor rule is active with $i > 0$. The second regime corresponds to when the ZLB is binding with $i = 0$. Substituting for r using Π in the first regime gives

$$Y = D + \frac{(1 + \beta)(1 + g)D\Gamma^*}{\beta} \frac{1}{\Pi^{\phi_\Pi - 1}} + \frac{1 + \beta}{\beta} \mathcal{A} \text{ for } i > 0, \quad (1.29)$$

and for the second regime I find

$$Y = D + \frac{(1 + \beta)(1 + g)D}{\beta} \Pi + \frac{1 + \beta}{\beta} \mathcal{A} \text{ for } i = 0, \quad (1.30)$$

where $\Gamma^* \equiv (1 + i^*)^{-1}(\Pi^*)^{\phi_\Pi}$. See the Appendix for a full derivation of the above result.

Equations (1.29) and (1.30) describe the steady-state aggregate demand curve, which is also shown in (Y, Π) space by Figure 1.4. For the first regime when the ZLB is not binding, the Taylor rule ensures that an increase in Π results in an increase in r . This leads to a reduction in aggregate demand, hence the first regime of the curve has a negative slope. In the second regime, the ZLB means that a decrease in Π leads to an increase in r . The increase in r reduces aggregate demand, hence the positive slope in the second regime. The two regimes meet at




Figure 1.4: Fundamental equilibrium following a decrease in D .




Figure 1.5: Effect of a bubble when the natural rate of interest is strongly negative.

the level of inflation when the zero lower bound is weakly binding. This is given by Π_{kink} in equation (1.24).

The exact slopes of the aggregate demand and aggregate supply curves depend on the particular parameter specification. In particular looking at Figure 1.4, the existence of a secular stagnation equilibrium represented by point B depends on the condition that the second regime of the aggregate demand curve is more steep than the second regime of the aggregate supply curve. Eggertsson and Mehrotra (2014) show that this is true for sensible parameter values which they present in their paper.

As seen in Figure 1.4, a fall in D leads to a fall in the natural rate of interest which can result in secular stagnation. In particular, as I am describing the steady-state equilibrium, the output gap at point B will last indefinitely. In the fundamental economy, we see the effect of a fall in D in $\{Y, \Pi\}$ space as a shift to the left of the aggregate demand curve. In this example, the intersection of aggregate supply and aggregate demand moves from full output, at point A , to secular stagnation, at point B .




Figure 1.6: The policy function and the natural rate of interest.

Figure 1.4 shows the effect of a strongly negative natural rate of interest in the fundamental economy. Figure 1.5 shows the effect of introducing a bubble to this economy. The bubble shifts out the aggregate demand curve as it acts to increase the natural rate of interest. By shifting out the aggregate demand curve through a higher natural rate of interest, the equilibrium at the intersection of aggregate demand and aggregate supply now occurs at the full output regime of the aggregate supply curve. Specifically, as there is full output in the bubbly equilibrium it follows that the natural rate of interest is above $\Pi_{kink}^{-1} - 1$. In general, it is possible that there still remains secular stagnation, however the bubble will nonetheless make the output gap smaller.

1.3.4.2 Natural rate of interest versus the real interest rate delivered

To better understand how the natural rate of interest affects the resulting equilibrium, it is insightful look closely at the natural rate of interest versus the real interest rate delivered (by the central bank).

Suppose that the natural rate of interest in the fundamental economy is less than $\Pi_{kink}^{-1} - 1$, what is the resulting steady state equilibrium? Figure 1.6 shows that for a natural rate of interest that is below $\Pi_{kink}^{-1} - 1$, the real interest rate delivered by the central bank cannot possibly be equal to the natural rate of interest. In particular suppose that there are no nominal rigidities. Then wages are always at the market-clearing level and there cannot be an output gap. As the real interest rate required to clear the goods market, r^n , cannot be the same as the real interest rate delivered by the central bank, r , the economy with no frictions would mean no possible equilibrium.

The argument above explains the necessity of $r = r^n$ if we are to have full output, $Y = Y^n$. In a world without rigidities, full output is necessary, so $r \neq r^n$ leads to a contradiction. If one allows for the possibility of less than full output, it is now plausible to have an equilibrium with




Figure 1.7: The policy function and the market-clearing real interest rate.

$r \neq r^n$. The introduction of nominal rigidities leads to the Keynesian IS curve, which says that the size of the output gap is increasing in the mismatch between the real interest rate delivered and the natural rate of interest.

The red line in Figure 1.7 shows how the market-clearing real interest rate depends on Π . A fall in Π beyond $\Pi < 1$ results in a fall of aggregate supply (see the second regime of the aggregate supply curve). Aggregate demand is falling in the real interest rate, therefore the reduction in output leads to a higher market-clearing real interest rate. This gives the negative slope of the market-clearing real interest rate below $\Pi < 1$. The intersection of the policy function and the market-clearing real interest rates curve determines the steady-state equilibrium. This second regime of the market-clearing real interest rate curve (for $\Pi < 1$) is the Keynesian IS curve, generated by the simple nominal friction of the wage rigidity. By suffering an output gap, it is possible for market clearing to occur at a higher real interest rate which is consistent with the policy function. Note again the importance of relative slopes of the two curves. For a contractionary equilibrium to exist it must be the case that the slope of the policy function is more steep than the slope of the market-clearing real interest rate curve. This maps to the analogous condition on relative slopes for aggregate demand versus aggregate supply, which Eggertsson and Mehrotra (2014) show is satisfied for a sensible parametrisation.

Given that the effect of bubbles is to increase the natural rate of interest, this shifts the market-clearing real interest rate curve to the right. As seen in Figure 1.8, the effect of this is a decrease in the equilibrium real interest rate, and therefore a higher level of output. In particular when the natural rate of interest equals to $\Pi_{kink}^{-1} - 1$, the economy no longer requires an output gap to achieve market clearing while remaining consistent with the policy function. This means that whenever the natural rate of interest is above $\Pi_{kink}^{-1} - 1$, there exists a steady-state equilibrium with full output.




Figure 1.8: Effect of bubbles on the market-clearing real interest rate.

1.4 Bubble stability and fiscal policy

While bubbles may have a positive effect through its effect on the natural rate of interest, this generally is considered to be not without trade-offs. In particular, the defining characteristic of bubbles is that its price today depends importantly on investors' expectations for its future resale value. When investors' expectations shift, the bubble that was sustained by such expectations will likewise collapse.

Within my model, there is no aspect of uncertainty. In such a framework it is expected that bubbles continue indefinitely. To analyse stability, I offer to consider the thought experiment of introducing exogenous shocks to economic fundamentals. This allows me to analyse the stability of bubbles without needing to deviate from the existing modelling framework, however a more rigorous approach would allow the agents to internalise the probability of collapse as a feature of the model. To proceed, from Proposition 1 we know that a bubble is unstable whenever the real interest rate greater than the growth rate of the economy g . When $r > g$, the bubble grows relative to the aggregate economy, which means that eventually the bubble will be too large to be sustained. I show in this section that bubbles of larger sizes are relative less stable, as they are more likely to become unstable following a fundamental shock. Following the standard analysis of dynamic efficiency, policies that generate the equivalent intergenerational transfers as bubbles act as a substitute for trade in bubbles, which can produce the same desirable outcome as bubbles without the associated instability.

For simplicity I shall make my above argument in the endowment economy model, although the same argument can be applied to the production economy's steady-state equilibrium. Given that r_t must never exceed g , the upper bound to real interest rates maps to an upper bound to the maximum size of the bubble. Suppose we consider a shock to economic fundamentals, namely Y , D or g . Such a shock threatens the stability of a bubble if it decreases this upper

bound. When the the upper bound falls below the actual size of the bubble, then the bubble must collapse due to forward-looking expectations. The larger the bubble, the closer it is to the initial upper bound, and hence the smaller the shock that is required to burst it. In particular for a bubble that is already at its maximum size, any reduction in the upper bound would result in the bubble collapsing.

According to proposition 1 the upper bound for the size of the bubble is given by

$$\mathcal{A}_t \leq \frac{\beta Y - (2 + \beta)D}{1 + \beta},$$

which is not violated whenever

$$\mathcal{A}_t \leq \frac{\beta}{1 + \beta}Y - \frac{2 + \beta}{1 + \beta}D.$$

Hence the upper bound falls due to either an increase in D or a fall in Y . Both an increase in D and a fall in Y act to reduce the desired savings per capita, $\frac{\beta}{1+\beta}(Y - D)$. As $\mathcal{A}_t$ units of per capita savings is necessarily invested into the bubble, this amount can become too large relative to desired savings per capita after an increase in D or a fall in Y . The larger the bubble $\mathcal{A}_t$, the closer we are to violating this stability condition, and hence the smaller the shock to either Y or D would be sufficient to violate this condition, thereby making the bubble unstable.

Similar to a change in Y or D , a reduction in the growth rate g can achieve the same effect but from the other direction. To demonstrate this I need to however adjust an assumption made earlier. Right now a reduction in g leads to a fall in r_t at a one-to-one rate, hence r_t will never exceed g following a fall in g . However this fact owes entirely to the life-cycle income endowment structure. If instead I allow for a smaller endowment income to the old, which follows the endowment economy specification in Eggertsson and Mehrotra (2014), in this setting r_t may exceed g following a reduction in g .

Allowing for an endowment income to the old of level Y^o , loanable fund supply becomes

$$L_t^s = \frac{\beta}{1 + \beta}(Y^m - D) - \frac{1}{1 + \beta} \frac{Y^o}{1 + r_t} - \mathcal{A}_t,$$

which implies that the market-clearing real interest rate is now

$$1 + r_t = \frac{(1 + g)D + \frac{1}{1+\beta}Y^o}{\frac{\beta}{1+\beta}(Y^m - D) - \mathcal{A}_t}.$$

A bubble's stability requires $1 + r_t \leq 1 + g_t$. This condition may be rewritten as

$$\frac{\frac{1}{1+\beta}Y^o}{\frac{\beta}{1+\beta}(Y^m - D) - \mathcal{A}_t - D} \leq 1 + g.$$

Hence following a negative shock to g , it is possible that a bubble becomes suddenly unstable. Again the size of the bubble is negatively associated with its relative stability. For larger sizes of the bubble, $\mathcal{A}_t$, a smaller shock to g will make the bubble collapse.

Bubbles raise the natural rate of interest because they create intergenerational transfers when the economy is dynamically inefficient. Whether or not the economy is actually dynamically inefficient is a somewhat open question. Whereas Abel et al. (1989) find evidence that the U.S. was never dynamically inefficient, Geerolf (2013) finds evidence to the contrary with access to new data on the return of land. Nonetheless for the purpose of my discussion, I make the assumption that the *fundamental* secular stagnation economy is dynamically inefficient. Whereas this assumption is innocuous for my analysis previously on the effect of bubbles in secular stagnation, the effect of policies in my framework depends on this assumption as they target specifically the transfers that rational bubbles provide.

Following the standard dynamic efficiency logic, trade in bubbles generates transfers away from the middle aged to the old. In every period, the middle aged purchase the stock of the bubbly asset from the old, which the old uses to finance consumption in retirement. The middle aged is willing to do this because they rationally expect that the current young generation will likewise do the same the next period. This transfer facilitates consumption smoothing across the periods of a household's life cycle, as it redistributes income away from when the household is rich (in the middle period when it is working), to when the household is poor (in the final period of retirement). As this transfer occurs at a return of $1 + g$, it is efficient whenever the real interest rate is less than the growth rate of population.

Understanding the effect of rational bubbles in this way, it follows similarly that the government can substitute for bubbles by reproducing the same transfers by force of policy, see e.g. Samuelson (1958) and Diamond (1965). Consider a PAYG transfer scheme enforced by the government where the middle aged as a whole is taxed in period t a quantity of Q_t . This quantity is then redistributed evenly among the old generation. Furthermore suppose also that this quantity Q_t grew by a rate of $1 + r_t$, so that

$$Q_{t+1} = (1 + r_t)Q_t.$$

This guarantees that the middle aged know that they will likewise receive a transfer back in retirement, with the transfer received accumulating an equivalent return. This PAYG transfer

is designed to be exactly equivalent to the same quantity of trade in bubbles, hence it acts as a perfect substitute for the bubbly asset. In this situation, the same allocation is achieved with a bubble that is Q_t units smaller. Therefore, the PAYG transfer reduces the size of the bubble by that same quantity (keeping allocations constant). In the context of the fundamental economy secular stagnation equilibrium, it follows that PAYG transfers reduce the output gap by raising the natural rate of interest with exactly the same mechanism as bubbles. Any other policy that enforces similar transfers will similarly substitute for bubbles, for example the issuance of sovereign debt by the government, where repayments for sovereign debt is funded by the taxation of the working generation, see Diamond (1965). By reducing the size of the bubble, there is a reduction in the impact of its potential collapse. This may make the economy more robust to unpredictable changes in investor sentiments, at least with regards to the possibility of secular stagnation. While a bubble's stability is shown to be negatively associated with its size, a bubble that is smaller due to a substituting policy is not more stable. This is because although the bubble is smaller, the maximum stable size of that same bubble is smaller by the same quantity. Nonetheless by achieving a smaller bubble, the effect of its collapse is dampened.

1.5 Conclusion

Although often associated with negative economic outcomes, in an environment of secular stagnation I show that bubbles can play a positive role by absorbing savings and increasing the natural rate of interest. This alleviates the central cause of secular stagnation, thereby increasing output and employment. While bubbles can do this, it likely comes with associated instability. The collapse of a bubble leads to a reduction in the natural rate of interest, which can act as the trigger to secular stagnation. I show in a stylised thought experiment that the stability of a bubble is likely decreasing in its size, hence the bubble solution to secular stagnation is imperfect. Within the framework of rational bubbles, other policies may have the same effect without the associated instability. The basic premise for analysing a bubble's stability and its relation to government policies is introduced in this chapter, however the canonical examples I give are simplified and lack certain features necessary for a comprehensive analysis of the question. An extension where agents internalise the probability of a bubble's collapse will be insightful, however this is left for future research.

1.6 Appendix

1.6.1 Production economy with capital

The introduction of capital does not qualitatively change the mechanisms described in section 3. Again bubbles act to increase the natural rate of interest, which can increase output when the natural rate of interest was strongly negative to begin with.

With capital, there is now a one-to-one relation between the real interest rate to the level of capital per worker. The equilibrium level of capital per worker is higher for lower real interest rates. Hence, low real interest rates leads to a secondary positive effect on output through a higher level of capital. Capital does not necessarily preclude negative real interest rates. As a savings vehicle, capital generates a return of the $MPK_{t+1} + 1 - \delta$. While the marginal product of capital is certainly greater than zero, once accounting for depreciation it possible for capital as an asset to offer a negative real interest rate.

Firms are perfectly competitive and they produce output following a Cobb-Douglas production function

$$Y_t = K_t^{1-\alpha} L_t^\alpha. \quad (1.31)$$

Households save by purchasing capital. In purchasing one unit of capital, a household must forgo $p^k - r^k$ units of consumption in the current period, which in the next period allows for $p^k(1 - \delta)$ units of additional consumption. The savings return on capital equals to the savings return on bonds by arbitrage

$$\begin{aligned} \frac{p^k(1 - \delta)}{p^k - r^k} &= 1 + r, \\ \implies r^k &= p^k \left(1 - \frac{1 - \delta}{1 + r} \right), \end{aligned}$$

which gives equations (1.33) and (1.39).

The contemporaneous rental rate on capital equals to its marginal product in production

$$r^k = (1 - \alpha) \frac{Y}{K},$$

which gives equations (1.34) and (1.40).

When $\Pi < 1$, $W = \bar{W}$.

$$\begin{aligned} \frac{W}{P} &= \gamma \frac{W}{P} \Pi^{-1} + (1 - \gamma) \alpha P_t K^{1-\alpha} \bar{L}^{\alpha-1}, \\ \implies \frac{W}{P} &= \frac{(1 - \gamma) \alpha K^{1-\alpha} \bar{L}^{\alpha-1}}{1 - \gamma \Pi^{-1}}. \end{aligned}$$

Because $\frac{W}{P}$ equals to the marginal product of labor, equation (1.38) follows

$$\begin{aligned}\alpha K^{1-\alpha} L^{\alpha-1} &= \frac{(1-\gamma)\alpha K^{1-\alpha} \bar{L}^{\alpha-1}}{1-\gamma\Pi^{-1}} \\ \implies L &= \left(\frac{1-\gamma\Pi^{-1}}{1-\gamma} \right)^{\frac{1}{1-\alpha}} \bar{L}.\end{aligned}$$

Bonds yield the same return as capital and the bubbly asset, hence its demand from the middle aged is perfectly crowded out by both capital and the bubbly asset, thus

$$B_t^m = -\frac{\beta}{1+\beta}(Y_t - D_{t-1}) + \frac{p_t A_t}{N_t} + (p^k - r^k)K.$$

For a particular steady-state level of capital per capita, investment is needed to account for depreciation and population growth

$$\begin{aligned}K &= \frac{1-\delta}{1+g}K + \frac{I}{p^k} \\ \implies I &= p^k \left(1 - \frac{1-\delta}{1+g} \right) K,\end{aligned}$$

where I is the investment needed to sustain a particular steady-state level of capital K .

As before, consumption across all generations must sum to output minus investment, it then follows that

$$\begin{aligned}Y - I &= (1+g)C^y + C^m + \frac{C^o}{1+g} \\ \implies Y - p^k \left(1 - \frac{1-\delta}{1+g} \right) K &= \\ (1+g) \left(\frac{D}{1+r} + \frac{g-r}{1+r} \mathcal{A} \frac{1+r}{(1+g)^2} \right) &+ \frac{1}{1+\beta}(Y-D) + \frac{1+r}{1+g} \frac{\beta}{1+\beta}(Y-D) \\ \implies Y = D + \frac{1+\beta}{\beta} \frac{1+g}{1+r} D &+ p^k K \left(1 + \frac{1-\delta}{\beta(1+r)} \right) + \frac{1+\beta}{\beta} \mathcal{A},\end{aligned}$$

which gives equation (1.37), and substituting $1+r = \Pi^{-1}$ when $i = 0$ gives equation (1.42).

1.6.2 Steady-state equilibrium and aggregate demand and supply

A steady-state equilibrium is summarised by the following set of equations below.




Figure 1.9: The effect of a bubble in an economy with capital.

For $\Pi \geq 1$ we have

$$r^k = p^k \left(1 - \frac{1 - \delta}{1 + r} \right), \quad (1.32)$$

$$r^k = (1 - \alpha) \frac{Y}{K}, \quad (1.33)$$

$$1 + r = \frac{\Pi^{\phi_{\Pi}-1}}{\Gamma^*}, \quad (1.34)$$

$$Y = K^{1-\alpha} \bar{L}^{\alpha}, \quad (1.35)$$

$$Y = D + \frac{1 + \beta}{\beta} \frac{1 + g}{1 + r} D + p^k K \left(1 + \frac{1 - \delta}{\beta(1 + r)} \right) + \frac{1 + \beta}{\beta} \mathcal{A}. \quad (1.36)$$

For $\Pi < 1$ we have

$$L = \left(\frac{1 - \frac{\gamma}{\Pi}}{1 - \gamma} \right)^{\frac{1}{1-\alpha}} \bar{L}, \quad (1.37)$$

$$r^k = p^k \left(1 - \frac{1 - \delta}{1 + r} \right), \quad (1.38)$$

$$r^k = (1 - \alpha) \frac{Y}{K}, \quad (1.39)$$

$$Y = K^{1-\alpha} L^{\alpha}, \quad (1.40)$$

$$Y = D + \frac{1 + \beta}{\beta} \Pi(1 + g) D + p^k K \left(1 + \Pi \frac{1 - \delta}{\beta} \right) + \frac{1 + \beta}{\beta} \mathcal{A}. \quad (1.41)$$

For $\Pi \geq 1$, equations (1.32), (1.33), (1.34) and (1.35) can be combined to eliminate K, r^k and r to find the aggregate supply curve, and (1.32), (1.33), (1.34) and (1.36) may be combined to eliminate K, r^k and r to find the aggregate demand curve. For $\Pi < 1$, equations (1.37), (1.38), (1.39) and (1.40) may be combined to eliminate L, K and r^k to find the aggregate supply curve, and (1.37), (1.38), (1.39) and (1.41) may be combined to eliminate L, K and r^k to find the aggregate demand curve. A bubble again acts to shift out the aggregate demand curve. Figure 1.9 shows the effect of a bubble when the natural rate of interest in the fundamental economy




Figure 1.10: Expansionary versus contractionary bubbly episodes.

is strongly negative.

Aggregate demand and the second regime of aggregate supply do not qualitatively change with the introduction of capital. The first regime of aggregate supply representing full employment is no longer vertical but now has a negative slope. The negative slope is due to the secondary effect of real interest rates on output. An increase of Π in the first regime of aggregate supply again corresponds to an increase in the real interest rate. This now means a reduction in the steady-state level of capital per worker, which reduces output, hence the negative slope of the first regime.

Martin and Ventura (2012) identify that models of bubbles from Samuelson (1958) and Tirole (1985) associate bubbly episodes with periods of lower output. A bubble increases real interest rates and reduces the capital stock, hence lowering output. This feature of Samuelson-Tirole models is hard to reconcile with experience. In this model, bubbles can be either expansionary or contractionary depending on the initial natural rate of interest in the fundamental economy.

Suppose that the fundamental economy was already at full output. Then a bubble shifts out aggregate demand from AD_1 to AD_2 , which reduces output through the standard Samuelson-Tirole mechanism. But if the fundamental economy generates a secular stagnation equilibrium, a bubble shifts out aggregate demand from AD_3 to AD_4 , which results in an increase in output. Hence a bubble can be either expansionary or contractionary.

1.6.3 Aggregate supply

When the wage norm is not binding ($\Pi \geq 1$), then $Y = \bar{L}^\alpha = Y^n$.

However when the wage norm is binding, it follows that

$$\begin{aligned}\frac{W}{P} &= \gamma \frac{W}{P} \Pi^{-1} + (1 - \gamma) \alpha \bar{L}^{\alpha-1} \\ \Rightarrow \frac{W}{P} &= \frac{(1 - \gamma) \alpha \bar{L}^{\alpha-1}}{1 - \gamma \Pi^{-1}}.\end{aligned}$$

Using the labour demand function we get

$$\begin{aligned}\frac{W}{P} &= \alpha L^{\alpha-1} \\ \Rightarrow L &= \left(\frac{W}{P} \right)^{\frac{1}{\alpha-1}} \\ \Rightarrow Y &= \left(\frac{W}{P} \right)^{\frac{\alpha}{\alpha-1}}.\end{aligned}$$

Substituting for $\frac{W}{P}$, rearranging, and accounting for $\bar{L}^\alpha = Y^n$, I find:

$$\frac{\gamma}{\Pi} = 1 - (1 - \gamma) \left(\frac{Y}{Y^n} \right)^{\frac{1-\alpha}{\alpha}},$$

which gives equation (1.28) for the second regime.

1.6.4 Aggregate demand

The consumption of each generation at time t is given respectively by equations (1.26), (1.3), and (1.4)

$$\begin{aligned}C_t^y &= B_t^y + p_t \frac{g_t - r_{t-1}}{1 + r_{t-1}} \frac{A_{t-1}}{N_{t+1}}, \\ C_t^m &= Y_t - (1 - r_{t-1}) B_{t-1}^y + B_t^m - p_t A_t^m, \\ C_t^o &= -(1 + r_{t-1}) B_{t-1}^m + p_t A_{t-1}^m,\end{aligned}$$

where $Y_t = \frac{W_{t+1}}{P_{t+1}} L_{t+1} + \frac{Z_{t+1}}{P_{t+1}}$, and $B_t^y = \frac{D_t}{1+r_t}$. Using the Euler equation, we further get

$$B_t^m = -\frac{\beta}{1 + \beta} (Y_t - D_{t-1}) + \frac{p_t A_t}{N_t}.$$

In addition, market clearing for the bubbly asset gives

$$A_t^m = \frac{A_t}{N_t}.$$

At a steady-state equilibrium, aggregate consumption across all three generations totals to

Y , thus

$$\begin{aligned}
Y &= (1+g)C^y + C^m + \frac{C^o}{1+g} \\
\Rightarrow Y &= (1+g) \left(\frac{D}{1+r} + \frac{g-r}{1+r} \mathcal{A} \frac{1+r}{(1+g)^2} \right) + \frac{1}{1+\beta} (Y-D) + \frac{1+r}{1+g} \frac{\beta}{1+\beta} (Y-D) \\
\Rightarrow Y &= D + \frac{(1+\beta)(1+g)D}{\beta(1+r)} + \frac{1+\beta}{\beta} \mathcal{A}.
\end{aligned}$$

Next, substitute for $1+r$ using equations below:

$$\begin{aligned}
1+r &= \frac{\Pi^{\phi_{\Pi}-1}}{\Gamma^*} \text{ for } i > 0, \\
1+r &= \Pi^{-1} \text{ for } i = 0,
\end{aligned}$$

where $\Gamma^* \equiv (1+i^*)^{-1}(\Pi^*)^{\phi_{\Pi}}$. Equations (1.30) and (1.31) then fully specify the aggregate demand curve.

References

- Abel, A. B., Mankiw, N. G., Summers, L. H., and Zeckhauser, R. J. (1989). Assessing dynamic efficiency: Theory and evidence. *The Review of Economic Studies*, 56(1):1–19.
- Baldwin, R. and Teulings, C. (2014). Secular stagnation: facts, causes and cures. *London: Centre for Economic Policy Research-CEPR*.
- Carvalho, V. M., Martin, A., and Ventura, J. (2012). Understanding bubbly episodes. *The American Economic Review*, 102(3):95–100.
- Diamond, P. A. (1965). National debt in a neoclassical growth model. *The American Economic Review*, 55(5):1126–1150.
- Eggertsson, G. B. and Krugman, P. (2012). Debt, deleveraging, and the liquidity trap: A fisher-minsky-koo approach. *The Quarterly Journal of Economics*, 127(3):1469–1513.
- Eggertsson, G. B. and Mehrotra, N. R. (2014). A model of secular stagnation. Technical report, National Bureau of Economic Research.
- Galí, J. (2014). Monetary policy and rational asset price bubb. *The American Economic Review*, 104(3):721–752.
- Geerolf, F. (2013). Reassessing dynamic efficiency. *manuscript, Toulouse School of Economics*.
- Hansen, A. H. (1939). Economic progress and declining population growth. *The American Economic Review*, 29(1):1–15.
- Krugman, P. R., Dominquez, K. M., and Rogoff, K. (1998). It’s baaack: Japan’s slump and the return of the liquidity trap. *Brookings Papers on Economic Activity*, 1998(2):137–205.
- Laubach, T. and Williams, J. C. (2016). Measuring the natural rate of interest redux. *FEDS Working Paper*.
- Martin, A. and Ventura, J. (2012). Economic growth with bubbles. *The American Economic Review*, 102(6):3033–3058.

Samuelson, P. A. (1958). An exact consumption-loan model of interest with or without the social contrivance of money. *The journal of political economy*, pages 467–482.

Summers, L. (2013). Why stagnation might prove to be the new normal. *Financial Times*, 15:12.

Tirole, J. (1985). Asset bubbles and overlapping generations. *Econometrica: Journal of the Econometric Society*, pages 1071–1100.

Chapter 2

The Decline in Fertility Rates and its Effects on Interest Rates and the Value of Land

2.1 Introduction

Real interest rates have fallen substantially in the past few decades across the advanced economies. They have been near zero or negative in the United States and the Euro-zone since late 2008, and for Japan since the mid 1990s. This trend has received intense academic scrutiny in recent years as it has important implications on the wider economy, for example on the operation of monetary policy subject to the zero lower bound and on financial stability. See Teulings and Baldwin (2015) for an overview of the secular stagnation debate. In this paper we show that a decline in fertility rates can drive a large fall in real interest rates, like which we have witnessed. A sudden fall in fertility rates leads to a shifting of the age distribution starting from the early 1970s, where today we witness an age imbalance with a disproportionately large number of middle-aged people, aged about 45-54. The consequence of this imbalance is a disruption of the life-cycle saving patterns across generations; the savings of the large cohort drives down the return of capital and hence interest rates. One might wonder if the surplus of saving can also explain the rise in land prices and consequently the rise in house prices, where land acts as a vehicle that absorbs some part of the additional savings. Our analysis surprisingly finds evidence to the contrary, at least from the perspective of a fully rational model driven by fertility rates and life-cycle saving motives.

Fertility rates underwent a dramatic decline for all advanced economies since the early 1970s, although the timing and the magnitude of this fall varies from country to country. Goldin and Katz (2002) document this fall for the US, where fertility rates were high and well above

replacement in the two baby-boom decades following WWII, and the subsequent fall led to a Total Fertility Rate (TFR) of about 1.8, a level that is below replacement. The timing of this aligns with the contraceptive pill becoming widely available towards the late 1960s and early 1970s, as well as increasing career and higher education opportunities for women. Fertility rates have been relatively stable since then, hence we propose to think about this fall as an exogenous one-time shock. The fall in fertility rates drives a shift in the dynamics for cohort size, with the cohorts born thereafter being smaller in size. Because the population was growing rapidly in most countries in the baby-boom decades before the fertility shock, this large pre-shock cohort is also greater in size than the cohorts born before. Altogether, the cohort currently aged between 45-54 is the largest cohort across all age groups, and this leads to a diamond-shaped age distribution which we find common across the advanced economies. The notable caveat to this pattern is Japan, which has an additional large cohort aged 65-69.

The long-term effects of the fall in fertility rates on savings and interest rates is well understood. The seminal work by Samuelson (1958) analyses the long-term Balanced Growth Path (BGP) equilibrium as a function of the population growth rate. The fall in fertility rates results in a fall in the long-term population growth rate, which lowers the real interest rate along the BGP equilibrium. The subject of this paper, however, considers the shorter-run transitional dynamics following a sudden fall in fertility rates. These dynamics are characterised by the life cycle of the large cohort described previously, which presents the significant deviation from the age distribution of the BGP equilibrium. The transition features a sharp decline in real interest rates, one that overshoots the long-run fall in interest rates along the BGP. Although the overshooting is transitional, it nonetheless lasts several decades and is therefore important to analyse and understand. Our model predicts that real interest rates shall continue to fall until around 2035, and thereafter there will be a partial reversal as we converge to the new BGP.

To analyse the transitional dynamics following a fertility shock, we consider a large OLG model where fertility rates iteratively determine the size of the new-born cohort. Time is discrete at the annual level, and an individual in this economy lives on the order of some 70 periods. We abstract from cross-sectional heterogeneity within birth cohorts and instead focus on heterogeneity across birth cohorts. We do not consider idiosyncratic and uninsurable shocks to income and instead assume that labour income is deterministically derived from the value of the individual's labour endowment. We also do not consider price level fluctuations and instead consider only the full output equilibrium. We believe this is an appropriate simplification as we are interested in the longer-term dynamics of real interest rates and land prices and not medium-term cyclical fluctuations. The objective of the individual is then to consume and save optimally so as to maximise its welfare, subject to its perfect foresight for the path of prices. In the extension where we consider the price determination of land, land serves simultaneously as a store of value

and as a good that generates utility. An individual optimally selects the quantity of land to rent depending on its rental price, and individuals also buy and sell land over its life cycle to facilitate consumption smoothing. As is typical for the life-cycling consumption smoothing for an individual, in our model each cohort saves during the years it is active on the labour market to accumulate assets, which it uses to finance their consumption in retirement. The desired asset holdings by a generation peaks at the age of retirement, which is around 65 in most countries, and this pattern of asset accumulation is smoothly accommodated by the other birth cohorts whom are at different stages of their life cycles along a BGP. The large cohort born before the negative fertility shock disrupts this accommodating process; they are accumulating assets while the usual absorbers of these savings (the retirees and the youngsters) are in relative short supply. We find that the saving of this large cohort drives a sharp collapse in real interest rates and our model predicts near zero or negative real interest rates from 2010 onwards. Further, our model predicts that this period of low interest rates will continue for another 1.5 decades, overshooting the new BGP level until hitting a trough around the year 2035. After which the large cohort begins to retire and deplete their assets, leading to an increase in the real interest rate that tends towards the new BGP level. Despite the accumulation in assets, our model does not predict an increase in land prices but rather a decline.

Our model assumes closed capital markets. Due to the close synchronicity of demographic trends across the advanced economies, international capital markets may offer little help to the problem of demography-driven over-saving. While developing economies experienced a lesser fall in fertility rates, there are greater frictions to those markets which limits the movement of capital. Nonetheless, it is still insightful to consider a global capital market with a global interest rate. Our analysis may be extended to a global scale with open and partially frictional capital markets, however we leave this as an avenue for future research.

We also make the assumption of no technological progress. Hence the long-term growth rate of the economy, g , is simply the growth rate of population. To include a constant rate for technological progress would be equivalent to having a faster rate of effective labour growth. Our results would be qualitatively unchanged with this inclusion, but an increase in g would imply a shift upwards in interest rates by a similar margin. Gordon (2012) in his study of historical U.S. growth makes the argument that innovations of recent decades are comparatively minor when compared to the innovations of the industrial revolution, leading to slowing productivity growth in recent decades. His findings imply a downwards pressure on per-capita output growth through the unobserved technology term, as well as a corresponding additional negative effect on interest rates. Gordon's perspective fits into the broader discussion of long-term stagnant growth. In particular, Summers (2013) proposes the secular stagnation hypothesis, where a persistently low natural rate of interest can lead to a persistent New-Keynesian output gap due

to the zero lower bound on nominal interest rates. Our analysis offers a potential demographic explanation for the decline of the natural rate of interest, which could serve as the premise for a secular stagnation equilibrium.

This chapter builds upon literature that looks to investigate the effect of demographics on real interest rates. Rachel and Smith (2015) find that changes in demography has resulted in changes to the support ratio, the ratio of workers to the overall population. They find that an increase in the support ratio can explain an increase in savings supply accounting for a 0.9% fall of real interest rates. Carvalho, Ferrero and Nechio (2016) make a similar argument and find that they can account for a reduction in the natural rate of interest of 1.5%. Goodhart and Erfurth (2014) consider future support ratios and project that support ratios will fall in the future, thereby reversing the period of high savings and low interest rates. Similarly to this model, they predict that by 2025 real interest rates shall return to their historical norm of 2.5-3%.

While the above studies vary in their approach, we distinguish our analysis by our focus on transitional dynamics. Their studies have in common the fact that they consider a static effect following a shift in demography. Consider for example an analysis based on the level of the support ratio, which argues that within a steady-state equilibrium high support ratios correspond to low interest rates. The contribution of this paper is to highlight the importance of the acute medium-term effects on the age distribution along the transition. It is precisely the concentration of mass towards for the older part of the labour force that generates our results. The implication is a much larger response for real interest rates, which would not be captured in static comparisons.

The structure of the remainder of this chapter is as follows. Section 2 presents the evolution of the age distribution common to all advanced economies in recent decades, and we demonstrate that these trends can be explained by a fall in fertility rates in the early 1970s. In section 3, we present the model that maps a collapse in fertility rates to its effects on saving and real interest rates. In section 4 we present a calibration of our model framework to simulate real interest rates for various advanced economies. In section 5 we present an extension of our model to consider the value of land during the demographic transition. In section 6 we conclude.

2.2 The evolution of age distributions

Figure 2.1 provides the age pyramids for the world's four largest economies, the United States, China, Japan, and Germany, in 2016¹. The same figure for the six largest economies in the Eurozone is included in the Appendix.

¹Figure 2.1 is composed from figures taken from <https://www.cia.gov/library/publications/the-world-factbook/fields/2010.html>




Figure 2.1: Age pyramids of the world's four largest economies.

All countries reveal a clear difference compared to the traditional graphs that applied 50 years ago and that do still apply for many developing countries: the age pyramid been replaced by a diamond shape with a concentration of mass in the middle section, aged 45 to 54. However at a more detailed level, the graphs also reveal clear differences between countries. For the sake of our discussion, we ignore the tops of the pillars above the age of 65, which largely reflects the effect of mortality at old age. Instead, we focus on the size of cohorts at the lower ages, which reflects differences in the fertility rates over time and also subsequent migration flows. First, we may identify the relative size of the large cohort born before the early 1970s, who are aged around 45-54 in 2016. This large cohort is most extreme in Germany, and somewhat less so in the United States. Other countries in the Eurozone look closer to Germany than the United states, and we include their age pyramids in the Appendix. In China the fertility shock comes about 5 years later with the introduction of the One Child Policy, hence the largest cohort is also around 5 years younger. In Japan, we see two large cohorts with an additional large cohort aged 65-69.

The large cohort dominates the age distribution. In the mid-1980s, the large cohort is around age 20, at which point they enter the labour force. In the mid-2000s, the large cohort is around age 40 and they are in the peak of their working careers. In 2016, the large cohort is around age 50, approaching retirement. The goal of our analysis is to understand the impact of this large

cohort and the overall age distribution transition on the evolution of real interest rates and the value of land.




Figure 2.2: TFR for the world's four largest economies.

Fertility rate data across the four largest economies explains the origin of the large cohorts. The Total Fertility Rate (TFR) is defined as the average number of children a woman is expected to have in her lifetime. In each country in the decades preceding the large cohort, population was growing rapidly. Following a sudden collapse in the TFR, the final cohort before the fertility shock become large relative to both the cohorts born before and after. As the large cohorts approach middle age, they generate the diamond shape in the age pyramid that we are familiar with today. For all countries, there is an echo effect of the elderly large cohort some 20 to 25 years later, when the large cohort of fertile women gives rise to a boom in newborn babies. In Japan today, the echo cohort of the original large cohort, aged 40 to 44, is close in size to the cohort of their parents due to mortality at the higher ages.

2.2.1 A model of the cohort size transition

We develop a simple model of cohort size and its evolution. Suppose that an individual lives for a total of J years (periods). Define the size of the cohort born at period t to be N_t . The total

population of all those alive at period t is P_t , which is given by

$$P_t = \sum_{i=0}^{J-1} N_{t-i}. \quad (2.1)$$

Those between the ages of $\underline{F}$ and $\overline{F}$ are fertile. They determine the size of the newborn cohort according to

$$N_t = b_t \sum_{i=\underline{F}}^{\overline{F}-1} N_{t-1-i}, \quad (2.2)$$

where b_t is the birth rate at time t , where the birth rate is derived from the TFR by $b_t = TFR/(\overline{F} - \underline{F})$. From these equations, we may derive the constant rate of population (and cohort size) growth for a given fixed birth rate b . By equation 2.1, we see that if population grows at a constant rate of g_P and cohort size grows at a constant rate of g_N , then it must be the case that $g_P = g_N$. From here we shall refer to this common growth rate as simply g . Rewriting equation (2.2) to

$$b = \left(\sum_{i=\underline{F}}^{\overline{F}-1} (1+g)^{-i} \right)^{-1}, \quad (2.3)$$

which gives us a monotonically increasing relationship between the birth rate b and the constant rate of population growth g .

We consider the case of a sudden fall in fertility rates. Suppose that the population was initially growing at the high constant growth rate associated with $b = b^H$, given by equation (2.3). The effect of the fertility shock leads to a reduction in birth rates described by

$$b_t = \begin{cases} b^H, & t < t^* \\ b^L & t \geq t^* \end{cases}, \quad \text{where } b^H > b^L. \quad (2.4)$$

Using equation (2.2), and noting that the initial demographic state was that of constant growth, we may derive the demographic transition thereafter. The transition occurs over many periods, however we find that the growth rate of the population will eventually stabilise at the new constant growth rate given by equation (2.3) with $b = b^L$.

2.2.1.1 Application to Germany

Germany, out of the four largest economies, comes closest to the case of a sudden negative fertility shock. Taking the fertile sub-population to be those between the ages of 20 to 30, we begin under the assumption that the population was initially growing at a constant rate consistent with a TFR of 2.5. We model the fertility shock as a sudden drop in TFR from 2.5




Figure 2.3: Germany's fertility shock and the simulated age distribution in 2016.

to 1.4 in the year of 1970 (see figure 2.3). Figure 2.3 shows the fit of our model. We see that our assumption of a sudden fertility shock captures the decline in fertility rates. In addition, we see that our model of the age distribution fits the age distribution of Germany well. We note that the success of this fit is despite using only two numbers, 2.5 before and 1.4 thereafter, to fit the fertility rate decline.

An important point to note is that the current age distribution for Germany, in terms of its ratio between those approaching retirement to those who are young, is more extreme than the long-run state of constant rate population decay. There is considerably less mass in the young cohorts' part of the Germany's age distribution. This highlights the relative size of the large cohort born before the fertility shock, and the relative scarcity of the absorbers of their savings, the younger cohorts.

2.3 OLG model of real interest rates

We take our model of the cohort size transition as our point of departure. In this section we embed our model of the age distribution in a large but otherwise standard OLG model. Using our OLG model we derive the implied path for real interest rates.

Households live for a total of J periods (years). Following Barro and Becker (1989) and Curtis et al. (2015), we consider each household to be headed by a singular adult, who provides for their offspring due to their parental altruism. An offspring reaches adulthood at age χ . Before an offspring reaches adulthood they make no economic decisions, but they have an effect on the aggregate economy due to their indirect effect on their parent's expenditure. Upon reaching adulthood, an offspring leaves the care of their parent(s) and enters the labour force. A specification of this kind generates a "hump-shaped" consumption pattern for (single-parent) households over their life cycles, which is observed empirically (see Attanasio et al. (1999)).

The utility function for a household born in period τ is of the following form:

$$U_{\tau} = \begin{cases} \sum_{i=\chi}^{J-1} \beta^i \left(\mu(n_{\tau,i})^{\eta} \frac{(c_{\tau,i}^c)^{1-\theta}}{1-\theta} + \frac{(c_{\tau,i}^a)^{1-\theta}}{1-\theta} \right) & \text{for } \theta \neq 1 \\ \sum_{i=\chi}^{J-1} \beta^i \left(\mu(n_{\tau,i})^{\eta} \log(c_{\tau,i}^c) + \log(c_{\tau,i}^a) \right) & \text{for } \theta = 1, \end{cases}$$

where $c_{\tau,i}^a$ denotes the expenditure on consumption dedicated to herself/himself, $c_{\tau,i}^c$ denotes the per-child expenditure on consumption for her/his offspring, θ measures the elasticity of intertemporal substitution, β measures the rate of time discounting, η measures the elasticity of a single child's consumption with respect to the number of children, μ measures the overall weighting on offspring consumption, and finally $n_{\tau,i}$ measures the number of dependent children a household born in period τ has at age i (which we allow to be continuous).

The number of dependent children for a household born in period τ at age i is given by

$$n_{\tau,i} = \sum_{j=\underline{F}}^{\min\{\bar{F},j\}} b_{\tau,j} - \sum_{j=\underline{F}}^{\min\{\bar{F},i-\underline{F}\}} b_{\tau,j}. \quad (2.5)$$

The optimal distribution between offspring consumption versus the adult's own personal consumption is given by

$$\frac{c_{\tau,i}^c}{c_{\tau,i}^a} = (\mu n_{\tau,i}^{\eta-1})^{1/\theta}.$$

Using the above optimality relation, we may rewrite the original utility function in terms of the household's overall expenditure on consumption:

$$U_\tau = \begin{cases} \sum_{i=\chi}^{J-1} \beta^i \frac{(c_{\tau,i})^{1-\theta}}{1-\theta} (1 + \mu^{1/\theta} n_{\tau,i}^{\eta/\theta-1/\theta+1}) & \text{for } \theta \neq 1 \\ \sum_{i=\chi}^{J-1} \beta^i \log(c_{\tau,i}) (1 + \mu^{1/\theta} n_{\tau,i}^{\eta/\theta-1/\theta+1}) & \text{for } \theta = 1, \end{cases} \quad (2.6)$$

where $c_{\tau,i}$ is the household's overall expenditure on consumption, defined as $c_{\tau,i} \equiv c_{\tau,i}^c n_{\tau,i} + c_{\tau,i}^a$.

The rewritten utility function allows us to derive an Euler equation. Households may borrow or save at the market real interest rate, the utility function given by equation (2.6) generates the following Euler equation

$$c_{\tau,i+1} = \beta^{1/\theta} (1 + r_{\tau+i+1})^{1/\theta} \left(\frac{1 + \mu^{1/\theta} n_{\tau,i+1}^{\eta/\theta-1/\theta+1}}{1 + \mu^{1/\theta} n_{\tau,i}^{\eta/\theta-1/\theta+1}} \right)^{1/\theta} c_{\tau,i}, \quad (2.7)$$

which differs from the standard Euler equation due to the third term that is a function of the number of dependent children. Suppose that if the number of dependent children is zero, or if the number of the dependent children does not change from this period to the next, then the new Euler equation collapses to the familiar standard Euler equation.

Each household receives a unitary labour endowment each year it is active on the labour market, which begins at age χ and ends at age ψ . Households derive no utility from leisure and hence inelastically supply their labour to firms at the going market wage. A household born in period τ has the following budget constraint:

$$\sum_{i=\chi}^{J-1} c_{\tau,i} \prod_{j=0}^{i-1} (1 + r_{\tau+j+1})^{-1} \leq \sum_{i=\chi}^{\psi-1} w_{\tau+i} \prod_{j=0}^{i-1} (1 + r_{\tau+j+1})^{-1}, \quad (2.8)$$

where $w_{\tau+i}$ is the wage at period $\tau + i$.

Furthermore the aggregate labour supply is given by

$$L_t = \sum_{i=\chi}^{\psi-1} N_{t-i}, \quad (2.9)$$

where L_t is the aggregate labour supply. Firms produce subject to a Constant Elasticity of Substitution (CES) production function using two inputs, capital and labour

$$Y_t = \begin{cases} \left(\alpha K_t^{\frac{\sigma-1}{\sigma}} + (1-\alpha) L_t^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}} & \text{for } \sigma \neq 1 \\ K_t^\alpha L_t^{1-\alpha} & \text{for } \sigma = 1, \end{cases} \quad (2.10)$$

where σ measures the elasticity of capital labour substitution, and the limiting case of σ tending to 1 gives the Cobb-Douglas production function. Furthermore, the limiting case of α equal to 0 gives the endowment income special case.

We complete the model with some market-clearing conditions. Capital depreciates at rate $\delta \in (0, 1]$. The equation for capital accumulation is given by

$$K_{t+1} = (1 - \delta)K_t + I_t. \quad (2.11)$$

Further the aggregate resource constraint is given by

$$Y_t = I_t + \sum_{i=\chi}^{J-1} c_{t-i,i} N_{t-i}. \quad (2.12)$$

Perfect competition in the capital rental market further gives

$$r_t + \delta = \begin{cases} \alpha k_t^{-\frac{1}{\sigma}} \left(\alpha k_t^{\frac{\sigma-1}{\sigma}} + 1 - \alpha \right)^{\frac{1}{\sigma-1}} & \text{for } \sigma \neq 1 \\ \alpha k_t^{\alpha-1} & \text{for } \sigma = 1, \end{cases} \quad (2.13)$$

and perfect competition in the labour market gives

$$w_t = \begin{cases} (1 - \alpha) \left(\alpha k_t^{\frac{\sigma-1}{\sigma}} + 1 - \alpha \right)^{\frac{1}{\sigma-1}} & \text{for } \sigma \neq 1 \\ (1 - \alpha) k_t^\alpha & \text{for } \sigma = 1, \end{cases} \quad (2.14)$$

where k_t denotes the level of capital per worker.

Definition 4. *An equilibrium in this economy is defined as a path for the following quantities:*

$$\{c_{t,\chi}, \dots, c_{t,J-1}, r_t, w_t, k_t, L_t, K_t, I_t, Y_t\},$$

given an exogenous path for:

$$\{b_t, n_{t,\chi}, \dots, n_{t,J-1}, N_t\},$$

that satisfies equations (2.2) and (2.5), and given parameters $\{\alpha, \beta, \gamma, \eta, \mu, \theta\}$, so that the following market-clearing equations: (2.7), (2.8), (2.9), (2.10), (2.11), (2.12), (2.13) and (2.14) are satisfied in all periods. Furthermore, a Balanced Growth Path equilibrium is a special case of an equilibrium, as defined above, where N_t grows at a constant rate g , and where the real interest rate and hence all per capita variables are constant.

To characterise the BGP equilibrium, it suffices to find the BGP real interest rate. To do this, we first solve the household and firm problems as a function of the unknown real interest rate. Next we find the implied size of the aggregate capital stock by aggregating the individual households' savings. This corresponds to the demand for assets as a function of the real interest rate. The implied size of the aggregate capital stock must also be consistent with the level of capital per worker derived from the real interest rate. This corresponds to the supply of assets as a function of the real interest rate. Combining these two relations determines the BGP real interest rate uniquely. We derive analytically the equation determining the BGP real interest rate in the Appendix.

We apply our OLG environment to the fertility application of Germany from section 2. We begin with the assumption that the population was initially growing at the high constant rate consistent with a TFR of 2.5. We suppose further that the economy was initially following a BGP equilibrium consistent with a TFR of 2.5 until the moment of the shock, and that households did not anticipate the shock. Upon the shock, households are tasked with deriving their optimal consumption/saving decisions for all future periods in their horizon in response to the unexpected shock. We consider a perfect foresight equilibrium thereafter, where households correctly evaluate the future trajectories for all economic variables and optimally consume with respect to this. To calculate the economic transition following the shock, we employ a computational algorithm that iteratively calculates the entire real interest rate path following the shock. Upon convergence of the iterative procedure, we stop our search and check directly the market-clearing conditions. We describe our iterative algorithm in the Appendix.

We show the results of the baseline specification for the Germany simulation. Following section 2 we take J , χ , ψ , $\underline{F}$ and $\bar{F}$ to be 75, 20, 65, 20 and 30 respectively. Furthermore, we pin down household and firm parameters, with $\alpha = 0.5$, $\beta = 0.99$, $\delta = 0.1$, $\eta = 0.76$, $\mu = 0.65$,

$\theta = 2$ (i.e. an elasticity of intertemporal substitution of 0.5) and $\sigma = 0.4$.²




Figure 2.4: Germany application's real interest rate over a long horizon.




Figure 2.5: Application to Germany.

²See Chirinko (2004), Havránek (2015), Curtis et al. (2015) and Nadiri (1996) for the rationale of the baseline specification parameters.

Figure 2.5 shows the evolution of cohort size after a negative fertility shock in 1970, over a very long horizon. Before the shock, population was growing at a constant rate of 1%. The growth rate shall eventually converge to a constant rate of decline at 1.5% per annum. In the process of converging to the future long-run growth rate, the age distribution undergoes a transition period of many decades. Following the shock there is an immediate decline in cohort size. This is followed by a short period of gradual increase, as the size of the fertile sub-population is still increasing. This gradual increase is exhausted after 20 years, and following on from this the first smaller cohorts begin entering the fertile sub-population. This leads to an echo effect of the original fertility shock and hence cohort size falls again. There are additional echos of the original shock, which die out over time as the age distribution tends to the long-run distribution. The periodicity of the echos is 25 years, which corresponds to the average age of a mother.

Figure 2.4 shows the trajectory of the equilibrium real interest rate for Germany under the baseline parameter specification. Initially the real interest rate moves up. The composition of the population is hardly affected at that time since the number of post-shock cohorts is still low, however the older cohorts feel richer because they realise that in subsequent years the size of future cohorts will be smaller, relaxing the tension on the goods market and making their claim of future output more valuable. Hence, they increase consumption, which drives up the real interest rate. From about 1985 onwards, the large cohort enters the labour market and consequently the real interest rate starts declining steeply, from nearly 1% to -1.5% in 2035, a decline of some 2.5%, crossing 0% around the year 2000. What is striking is the duration of the period for which the real interest rate is declining. The downward pressure continues into the 2030s, just as the large cohort begins to retire. The depth of the trough significantly overshoots below the new BGP level.

After 2035, there is some recovery in real interest rates as the large cohort depletes their savings in retirement. This leads to a reduction in the capital stock and hence an increase in the return of capital investment/saving. The echos of the cohort size distribution, discussed above, generates residual cyclicity in the real interest rate path. Because the first post-shock cohort was much smaller in size than the previous cohorts, when they enter the fertile sub-population their smaller size leads to a reduction in the number of births that year. Subsequent generations of the first post-shock cohorts are likewise also smaller, although the magnitude of this disruption fades out over time. With the birth of each new generation of the original post-pill cohort, cohort size undergoes a dampened residual shock similar to that of the original fertility shock. This leads to a fall and a subsequent uptick in real interest rates, which leads to the cyclicity of real interest rates. As noted before the periodicity of the cohort size distribution is 25 years, the average age of a fertile mother. Likewise the periodicity of the real interest rate cycles is

also 25 years. As the magnitude of the shock fades out over time, as does magnitude of the real interest rate cycles. By the year 2250 the transition to the new BGP is essentially complete.

Real interest rates generates market clearing by shifting consumption across time. Through the Euler equation, low interest rates shifts consumption forward in time whereas high interest rates rewards the postponement of consumption. We may understand the fall in interest rates since the mid-1980s as the response to the overwhelming demand for the postponement of consumption from the large cohort. Hence in order to clear markets the response is low interest rates. A simple comparison sums up this channel. Whereas the growth rate of the labour force measures approximately the growth rate of consumption supply, the growth rate of the adult population measures approximately the growth rate of consumption demand. The difference between the two leads to a mismatch of consumption supply and demand, and market clearing requires interest rates to bridge the gap. In equilibrium we see a clear contemporaneous correlation between this difference in growth rates and the real interest rates.

Regarding welfare, we see that the cohort hurt most due to the fertility shock is the cohort born in 1950. This is because they proceed the largest cohort by exactly 20 years. Hence they save during many a period of steep interest rate decline. Furthermore, although they retire during a resource-rich period while the large cohort is still in the labour force, they are unable to trade for the resources from the large cohort at a good price as they do not offer the in-demand resource for the large cohort, i.e. retirement consumption for the large cohort. Altogether, we see a life-time 2 percent fall in welfare for this cohort.

The Appendix shows the real interest rate trajectory for some alternative parameter specifications, however we summarise the intuition here. An increase in life expectancy or a decrease in the retirement age acts to increase the number of years the large cohort needs to finance consumption for in retirement, which leads to a deeper trough in real interest rates. A decrease in the elasticity of intertemporal substitution deepens the trough in real interest rates, as the elasticity measures the degree to which the large cohort substitutes to saving more when the returns to saving falls. Likewise a decrease in the elasticity of capital labour substitution also deepens the trough in real interest rates, as this means that capital is less able to absorb savings at any given real interest rate.

2.4 Application to fertility rate data

We apply our OLG model developed in the previous section and calibrate it to fertility rate data for the four largest economies in the world. The calibration exercise for the six largest Eurozone countries is included in the Appendix. Again we take our model of the cohort size transition as our starting point. However instead of the stylised sudden fertility shock we utilised before, we

take recorded fertility rate data to calibrate our demographic model.

Our new birth rate specification is as follows

$$b_t = \begin{cases} b^H, & t < t^* \\ TFR_t/(\bar{F} - \underline{F}) & t \geq t^* \end{cases}, \text{ where } TFR_t/(\bar{F} - \underline{F}) \rightarrow b^L < b^H. \quad (2.15)$$

Here TFR_t is derived from the fertility data from a particular country. b^H is calibrated so that it matches the initial population growth rate of a country, and b^L is calibrated so that it matches the projected terminal population growth rate of a country. For the US we consider a TFR fall from 3 to 1.8, for China from 3.5 to 1.7, for Japan from 2.1 to 1.4, and for Germany from 2.5 to 1.4.

We solve the model using the same computational method as before. Again we assume that households do not expect the change in fertility rates prior to period t^* . We include our simulation of real interest rates based on the calibrated birth rates below.




Figure 2.6: Simulated real interest rates for the four largest economies.

We see that the same broad pattern for real interest rates from the simulation based on a sudden fertility shock is true also for our calibrated simulation for the four largest economies. The same exercise for the six largest Eurozone economies is included in the Appendix.

There are however differences in timing and magnitude for the real interest rate response across the four largest economies. Out of the four, the United States suffered the smallest fertility fall. Likewise the trough for real interest rates in the United States simulation is the shallowest. China begins with the highest interest rate as China experienced the fastest population growth prior to the fertility fall. China's trough is however of similar depth due to the low fertility rates today. Furthermore China's real interest rate trough is delayed by some 5 to 10 years relative to the other three countries as the fertility fall came 5 to 10 years later. Japan's earlier large cohort means that the decline of real interest rates began some 20 years before the other three countries. Its decline is extended and the final trough occurs at a similar time as Germany and the United States, around 2035. The figure for Germany is closely reminiscent of the stylised simulation from the previous section. Other Eurozone countries share a similar demographic and simulated real interest rate transition as Germany.

Note that our simulation assumes that each economy begins at the BGP equilibrium. That is, each economy is assumed to begin on its steady-state level of capital per worker, consistent with the initial population growth rate. This assumption is unrealistic especially for China, which underwent a process of capital accumulation in recent decades. The implication is that China had far less capital saturation than is assumed in our simulation, and in which case real interest rates will be much higher due to the higher returns to capital. As China approaches saturation in capital investment, this would imply China's real interest rates to approach a similar level to that of the other economies in the top four.

2.5 Extension to land

We extend our OLG model from sections 3 and 4 to include the asset land. We develop this extension to evaluate the following hypothesis: can the demographic shift towards saving explain the rising value of land in recent decades? To the extent that land serves as a store of value, this might be a plausible explanation for the appreciation of land and consequently also housing.

We model land as an asset in addition to capital, which is available in fixed supply in the economy. Households derive utility from the utilisation of land, which they rent in each period in a competitive and frictionless rental market. The intraperiod utility function is Cobb-Douglas with respect to non-durable consumption and land utilisation. The previous lifetime utility function given by equation (2.6) is replaced by

$$U_{\tau} = \begin{cases} \sum_{i=\chi}^{J-1} \beta^i \frac{(c_{\tau,i}^{1-\lambda} l_{\tau,i}^{\lambda})^{1-\theta}}{1-\theta} (1 + \mu^{1/\theta} n_{\tau,i}^{\eta/\theta - 1/\theta + 1}) & \text{for } \theta \neq 1 \\ \sum_{i=\chi}^{J-1} \beta^i \log(c_{\tau,i}^{1-\lambda} l_{\tau,i}^{\lambda}) (1 + \mu^{1/\theta} n_{\tau,i}^{\eta/\theta - 1/\theta + 1}) & \text{for } \theta = 1. \end{cases}$$

Owing to the Cobb-Douglas form of the intraperiod utility function, optimal distribution between consumption expenditure and land rental expenditure is given by

$$\frac{c_{\tau,i}}{1-\lambda} = \frac{l_{\tau,i}R_{\tau+i}}{\lambda}. \quad (2.16)$$

Hence it follows that

$$\lambda e_{\tau,i} = l_{\tau,i}R_{\tau+i}^L, \quad (2.17)$$

where $e_{\tau,i}$ is the overall household expenditure by the cohort born in period τ at age i , defined to be $c_{\tau,i} + R_{\tau+i}l_{\tau,i}$.

Using the above optimality relation, we may rewrite the final extended utility function above to

$$U_{\tau} = \begin{cases} \sum_{i=\chi}^{J-1} \beta^i \frac{(e_{\tau,i})^{1-\theta}}{1-\theta} (1 + \mu^{1/\theta} n_{\tau,i}^{\eta/\theta-1/\theta+1}) & \text{for } \theta \neq 1 \\ \sum_{i=\chi}^{J-1} \beta^i \log(e_{\tau,i}) (1 + \mu^{1/\theta} n_{\tau,i}^{\eta/\theta-1/\theta+1}) & \text{for } \theta = 1. \end{cases} \quad (2.18)$$

Note that the utility function in terms of expenditure is the same as the previous utility function in terms of consumption. The new Euler equation relating intertemporal substitution for expenditure is given by

$$e_{\tau,i+1} = \beta^{1/\theta} (1 + r_{\tau+i+1})^{1/\theta} \left(\frac{1 + \mu^{1/\theta} n_{\tau,i+1}^{\eta/\theta-1/\theta+1}}{1 + \mu^{1/\theta} n_{\tau,i}^{\eta/\theta-1/\theta+1}} \right)^{1/\theta} e_{\tau,i}. \quad (2.19)$$

The budget constraint for a household born in period τ at age i is given by

$$e_{\tau,i} + L_{\tau,i+1}p_{\tau+i} + K_{\tau,i+1} + B_{\tau,i+1} \leq L_{\tau,i}(p_{\tau+i} + R_{\tau+i}^L) + K_{\tau,i}(R_{\tau+i}^K + 1 - \delta) + B_{\tau,i}(1 + r_{\tau+i}) + w_{\tau+i}l_i \quad (2.20)$$

where $L_{\tau,i}$, $K_{\tau,i}$ and $B_{\tau,i}$ are the holdings of land, capital and the risk-free bond respectively for cohort τ entering age i , $R_{\tau+i}^L$ and $R_{\tau+i}^K$ are the rental rates of land and capital respectively, r is the interest rate of the risk-free bond, p is the price of land, δ is the depreciation rate of capital, w is wage rate, and l is the labour endowment of a household at i .

Substitution between land ownership and the risk-free bond gives

$$1 + r_{t+1} = \frac{R_{t+1}^L + p_{t+1}}{p_t},$$

whereas substitution between capital investment and the risk-free bond is already given.

Hence, the price of one unit of land is given by the sum of its future discounted dividends

$$p_t = \sum_{i=1}^{\infty} R_{t+i} \prod_{j=1}^i (1 + r_{t+j})^{-1}. \quad (2.21)$$

Substitution across assets means that we may aggregate across asset classes

$$e_{\tau,i} + A_{\tau,i+1} \leq A_{\tau,i}(1 + r_{\tau+i}) + w_{\tau+i}l_i,$$

where total assets for cohort τ at age i is given by

$$A_{\tau,i} = L_{\tau,i}p_{\tau} + K_{\tau,i} + B_{\tau,i}.$$

The budget constraints across time can be summarised as the following intertemporal budget constraint

$$\sum_{i=\chi}^{J-1} e_{\tau,i} \prod_{j=0}^{i-1} (1 + r_{\tau+j+1})^{-1} \leq \sum_{i=\chi}^{\psi-1} w_{\tau+i} \prod_{j=0}^{i-1} (1 + r_{\tau+j+1})^{-1}. \quad (2.22)$$

Market clearing for the land rental use and land ownership is given

$$\sum_{i=\chi}^{J-1} N_{t-i} l_{t-i,i} = \sum_{i=\chi}^{J-1} N_{t-i} L_{t-i,i} = \bar{L}, \quad (2.23)$$

where $\bar{L}$ is the aggregate fixed supply for land.

Aggregating eq (2.17) across cohorts, the rental price of land is pinned down by the supply of land and aggregate expenditure

$$\lambda E_t = \bar{L} R_t^L, \quad (2.24)$$

where $E_t \equiv \sum_{i=\chi}^{J-1} (1 - \gamma) e_{t-i,i} N_{t-i}$ is the economy's aggregate expenditure.

Aggregating eq (2.20) across cohorts in period t gives

$$E_t + \bar{L} p_t + K_{t+1} = \bar{L} (p_t + R_t^L) + K_t (R_t^K + 1 - \delta) + w_t L_t,$$

which once rearranged gives

$$E_t - \bar{L} R_t^L + I_t = K_t R_t^K + w_t L_t,$$

where L_t is the aggregate labour force.

The RHS of the above equation is aggregate output given that the production function has constant returns to scale. Hence, substituting using eq (2.24) gives the following aggregate resource constraint

$$(1 - \lambda) E_t + I_t = Y_t. \quad (2.25)$$

Definition 5. *An equilibrium in this economy is defined as a path for the following quantities:*

$$\{e_{t,\chi}, \dots, e_{t,J-1}, r_t, p_t, w_t, k_t, L_t, K_t, I_t, Y_t\},$$

given an exogenous path for:

$$\{b_t, n_{t,\chi}, \dots, n_{t,J-1}, N_t\},$$

that satisfies equations (2.2) and (2.5), and given parameters $\{\alpha, \beta, \gamma, \eta, \mu, \theta, \lambda\}$, so that the following market-clearing equations: (2.19), (2.22), (2.9), (2.10), (2.11), (2.25), (2.21), (2.13) and (2.14) are satisfied in all periods. Furthermore, a *Balanced Growth Path equilibrium* is a special case of an equilibrium, as defined above, where p_t and N_t grows at a constant rate g , and where the real interest rate and hence all per capita variables are constant.

Our model is solved using a similar computational algorithm to our previous model. The main difference in this extended model is the need to calculate and track the price of land and its evolution through time. In each iteration of our adjustment algorithm, we also need to update the transition path for land prices. We present the results of our simulation, again for a sudden TFR fall from 2.5 to 1.4, as in section 2.




Figure 2.7: Simulated real interest rates for the model with land.

We see that the same pattern for interest rates holds in the model with land. However, notably interest rates are higher in the model with land, consistently above zero in this simulation. This is due to the availability of another store of value to facilitate saving.




Figure 2.8: Simulated value for land.

Contrary to our initial hypothesis, the value or price of land is not predicted to increase following the fertility shock. There is a small sudden fall in the value of land at the moment of the shock, due to the sudden arrival of new information. Following from this, the value of land rises for two decades, however the value of land reaches a peak well before the year 2000. Thereafter the value of land falls consistently, at a rate that approaches the new BGP level.

To understand this transition, we decompose the drivers that affect land prices. In our perfect foresight model where land is valued as the sum of discounted dividends, there are two forces influencing the value of land. Firstly the value of land is affected by the path for real interest rates, which determines the discount factor for future dividends. A fall in the real interest rate, *ceteris paribus*, increases the value of land. Secondly, the value of land is affected by aggregate expenditure. A constant proportion of consumption expenditure is diverted towards land rent, and this determines the stream of dividends that land entitles one to in the future. A rise in consumption expenditure, *ceteris paribus*, increases the value of land. Below we show the evolution of the value of land, scaled by GDP.




Figure 2.9: Simulated value of land relative to output.

We see that following the fertility fall, the value of land relative to GDP immediately falls. It adjusts to a level around the BGP, at 1.55, by around 2015. There is a short uptick in the value of land relative to GDP around the year 1985, just as the large cohort enters the labor force and begins renting land. This channel is similar to the mechanism considered in Mankiw and Weil (1989). Nonetheless, the overall effect on the value of land as a proportion of output is negative, which we understand in the following way.

The BGP equation for the value of land as a proportion to GDP can be written in the following way

$$\frac{\bar{L}p_t}{Y_t} = \frac{\lambda}{1-\lambda} \frac{Y_t - I_t}{Y_t} \frac{1+g}{r-g}.$$

The fall in fertility rates leads to a fall in g and a simultaneous fall in r . Furthermore, the fall in r is smaller than the fall in g , due to the ability for capital to absorb capital at an increasing rate as r falls. The largest source of variation in the above equation is the third fractional term of the right-hand side. A fall in g and an increase in $r - g$ leads to an unambiguous fall in the value of land relative to GDP. Importantly, the rise in $r - g$ has a large effect as $r - g$ is close to zero. The dependence on $(r - g)^{-1}$ measures the discounting of future dividends, once incorporating the growth rate of consumption expenditure. The fall in the interest rate is dominated by the fall in growth rates, hence the overall effect on land value as a proportion of GDP is negative.

2.6 Conclusion

In this chapter we demonstrate that changes in the age distribution of the population, common across most advanced economies in recent decades, can explain the pattern of high savings and low interest rates witnessed today. These changes were driven by a sharp decline in fertility rates in the early 1970s, and the consequences of this event are intuitive and foreseeable. We evaluate the hypothesis of whether this surplus of savings can explain inflated asset prices, in particular for land which we model as extracting dividends in each period proportional to overall consumption expenditure. Our conclusion is however in conflict with this hypothesis; despite the surplus of savings the value of land should fall from the late 1990s onwards.

We note the broader discussion of how low interest rates pose a problem for monetary policy due to the ZLB on policy rates. Beyond an increase in the retirement age, which would reduce the targeted savings of current working cohorts and thereby raise interest rates, there is little guidance for policy responses. An increase in government debt may offset the excess savings. In particular how the debt shall be funded in repayment shall determine how households respond to these changes. Further analysis is required to better understand the impacts of such an option, which we feel to be an insightful avenue for future research.

2.7 Appendix




Figure 2.10: Age pyramids of the Eurozone's six largest economies.




Figure 2.11: Simulated real interest rates for the Eurozone's six largest economies.

2.8 BGP real interest rate

To find the BGP equilibrium real interest rate in the OLG economy without land, we need to find the supply and demand of assets (i.e. capital) as a function of the hypothetical real interest rate. Equating supply and demand and solving for the hypothetical real interest rate gives the BGP equilibrium real interest rate.

Consider a BGP equilibrium, $r_t = r, w_t = w, k_t = k$ and

$$N_t = (1 + g)^t,$$

where N_0 is normalized to unity without loss of generality. Define:

$$P_i(r) \equiv (1 + r)^{-i} \quad (2.26)$$

$$\implies P_i(g) = N_{-i}$$

$$R_i(r) \equiv \sum_{j=\chi}^{i-1} P_j(r) = \begin{cases} r^{-1} (P_{\chi-1}(r) - P_{i-1}(r)) & \text{for } r \neq 0 \\ i - \chi & \text{for } r = 0, \end{cases},$$

$$W \equiv w \sum_{i=\chi}^{\psi-1} P_i(r) = w R_{\psi}(r), \quad (2.27)$$

where W is lifetime wealth of the representative household at birth, and $P_i(r)$ is the intertemporal price of period $t + i$ consumption in terms of period t income, which does not depend on t since $r_t = r$ does not depend on t . Using this notation, the Euler equation given by equation 2.7 can be written as

$$c_{j+1} = \beta^{1/\theta} (1 + r)^{1/\theta} \left(\frac{1 + \mu^{1/\theta} n_{j+1}^{\eta/\theta-1/\theta+1}(b)}{1 + \mu^{1/\theta} n_j^{\eta/\theta-1/\theta+1}(b)} \right)^{1/\theta} c_j.$$

where c_j is the consumption of the representative household at age j , and $n_i(b)$ is the number of dependent children a representative parent has at age i , given by

$$n_i(b) \equiv \sum_{j=\chi}^{\min\{\bar{F}, i\}} b - \sum_{j=\chi}^{\min\{\bar{F}, i-\chi\}} b.$$

Define:

$$g_j^c(b) \equiv \left(\frac{1 + \mu^{1/\theta} n_{j+1}^{\eta/\theta-1/\theta+1}(b)}{1 + \mu^{1/\theta} n_j^{\eta/\theta-1/\theta+1}(b)} \right)^{1/\theta}.$$

The Euler equation before implies

$$c_j = \beta^{j/\theta} P_j(r)^{-1/\theta} \left(\prod_{l=\chi}^{j-1} g_l^c(b) \right) c,$$

where c is a constant to be determined.

We next define:

$$\begin{aligned}
B_i(r, b) &\equiv \sum_{j=\chi}^i \beta^{j/\theta} P_j(r)^{(\theta-1)/\theta} \left(g_l^c(b) \right), \\
\sum_{j=\chi}^i P_j(r) c_j &= \sum_{j=\chi}^i P_j(r) \beta^{j/\theta} P_j(r)^{-1/\theta} \left(g_l^c(b) \right) c \\
&= B_i(r, b) c \\
\implies W &= \sum_{i=\chi}^{J-1} P_i(r) c_i = B_{J-1}(r, b) c \\
\implies c &= \frac{W}{B_{J-1}(r, b)}.
\end{aligned}$$

Let $A_i(r)$ be the stock of assets accumulated by the representative agent at age i . It is equal to the present value of the labour income up to age i , minus the present value of consumption up to that age i :

$$A_i(r) = \sum_{j=\chi}^i P_{j-i}(r) (w I_{j \in [\chi, \psi-1]} - c_j).$$

where $I_x = 1$ if $x = \text{true}$ and $I_x = 0$ if $x = \text{false}$. This equation is intuitive. The factor $P_{j-i}(r)$ accounts for the discounting of savings and debts. The first term, for the case $i \in [\chi, \psi-1]$, is the present value of labour income up to age i . The second term is the present value of consumption up to age i . The exhaustion of the budget constraint gives

$$w \sum_{i=\chi}^{\psi-1} P_i(r) = \sum_{i=\chi}^{J-1} P_i(r) c_i,$$

therefore we have $A_J(r) = 0$, i.e. households spend their full life time wealth on consumption.

We can now calculate the aggregate net demand for assets for all cohorts at period 0, $S(r)$:

$$\begin{aligned}
S(r) &= \sum_{i=\chi}^{J-1} N_{-i} A_i(r) = w \sum_{i=\chi}^{\psi-1} N_{-i} \sum_{j=\chi}^i P_{j-i}(r) + w \sum_{i=\psi}^{J-1} N_{-i} \sum_{j=\chi}^{\psi-1} P_{j-i}(r) - \sum_{i=\chi}^{J-1} N_{-i} \sum_{j=\chi}^i P_{j-i}(r) c_j \\
&= w \sum_{i=\chi}^{\psi-1} \frac{P_i(g)}{P_i(r)} \sum_{j=\chi}^i P_j(r) + w \sum_{i=\psi}^{J-1} \frac{P_i(g)}{P_i(r)} \sum_{j=\chi}^{\psi-1} P_j(r) - W \sum_{i=\chi}^{J-1} \frac{B_i(r, b)}{B_{J-1}(r, b)} \frac{P_i(g)}{P_i(r)} \\
&= w \sum_{i=\chi}^{\psi-1} \frac{P_i(g)}{P_i(r)} R_{i+1}(r) + w \sum_{i=\psi}^{J-1} \frac{P_i(g)}{P_i(r)} R_{\psi}(r) - w R_{\psi}(r) \sum_{i=\chi}^{J-1} \frac{B_i(r, b)}{B_{J-1}(r, b)} \frac{P_i(g)}{P_i(r)}
\end{aligned}$$

Since the capital stock satisfies $K_t = kL_t$, the total net demand for assets minus the next period's capital stock in period 0, $S^*(r)$, is given by

$$S^*(r) = S(r) - K_1 = S - (1 + g) kL_0.$$

To find the BGP equilibrium real interest rate, it suffices to solve $S^*(r) = 0$. This is given by

$$\sum_{i=\chi}^{\psi-1} \frac{P_i(g)}{P_i(r)} R_{i+1}(r) + \sum_{i=\psi}^{J-1} \frac{P_i(g)}{P_i(r)} R_{\psi}(r) - R_{\psi}(r) \sum_{i=\chi}^{J-1} \frac{B_i(r, b)}{B_{J-1}(r, b)} \frac{P_i(g)}{P_i(r)} - (1+g) \frac{k}{w} R_{\psi}(g) = 0,$$

where $\frac{k}{w}$ is given by

$$\frac{k}{w} = \frac{1}{r + \delta} \frac{\alpha}{\left(\frac{r+\delta}{\alpha}\right)^{\sigma-1} - \alpha}.$$

2.8.1 BGP real interest rate in the model with land

Consider a general BGP equilibrium. Here $r_t = r$, $w_t = w$, and $k_t = k$ and

$$N_t = (1+g)^t,$$

where N_0 is normalised to unity without loss of generality. Analogous to section 4, define:

$$\begin{aligned} P_i(r) &\equiv (1+r)^{-i} \\ \implies P_i(g) &= N_{-i}, \\ R_i(r) &\equiv \sum_{j=\chi}^{i-1} P_j(r) = \begin{cases} r^{-1} (P_{\chi-1}(r) - P_{i-1}(r)) & \text{for } r \neq 0 \\ i - \chi & \text{for } r = 0 \end{cases}, \\ W &\equiv w \sum_{i=\chi}^{\psi-1} P_i(r) = w R_{\psi}(r), \end{aligned} \tag{2.28}$$

where W is lifetime wealth at birth and $P_i(r)$ is the intertemporal price of period $t+i$ consumption in terms of period t income, which does not depend on t since $r_t = r$ does not depend on t . Using this notation, the Euler equation can be written as

$$c_{j+1} = \beta^{1/\theta} (1+r)^{1/\theta} \left(\frac{1 + \mu^{1/\theta} n_{j+1}^{\eta/\theta-1/\theta+1}(b)}{1 + \mu^{1/\theta} n_j^{\eta/\theta-1/\theta+1}(b)} \right)^{1/\theta} c_j.$$

Define:

$$\begin{aligned} g_j^c(b) &\equiv \left(\frac{1 + \mu^{1/\theta} n_{j+1}^{\eta/\theta-1/\theta+1}(b)}{1 + \mu^{1/\theta} n_j^{\eta/\theta-1/\theta+1}(b)} \right)^{1/\theta} \\ \Rightarrow c_j &= \beta^{j/\theta} P_j(r)^{-1/\theta} \left(g_{l=\chi}^{j-1}(b) \right) c, \end{aligned}$$

where c is a constant to be determined. We firstly define

$$\begin{aligned}
B_i(r, b) &\equiv \sum_{j=\chi}^i \beta^{j/\theta} P_j(r)^{(\theta-1)/\theta} \left(g_{l=\chi}^{j-1}(b) \right), \\
\sum_{j=\chi}^i P_j(r) c_j &= \sum_{j=\chi}^i P_j(r) \beta^{j/\theta} P_j(r)^{-1/\theta} \left(g_{l=\chi}^{j-1}(b) \right) c \\
&= B_i(r, b) c \\
\implies W &= \sum_{i=\chi}^{J-1} P_i(r) c_i = B_{J-1}(r, b) c \\
\implies c &= \frac{W}{B_{J-1}(r, b)}.
\end{aligned}$$

Let $A_i(r)$ be the stock of assets accumulated by a person of age i at time $t = 0$. It is equal to the present value of the labour product up to age i minus the present value of consumption up to that age

$$A_i(r) = \sum_{j=\chi}^i P_{j-i}(r) (w I_{j \in [\chi, \psi-1]} - c_j).$$

where $I_x = 1$ if $x = \text{true}$ and $I_x = 0$ if $x = \text{false}$. The factor $P_{j-i}(r)$ accounts for the discounting of savings and debts to time $t = 0$. The first term, for the case $i \in [\chi, \psi - 1]$, is the present value of labour income. The second term is the present value of consumption. Since the exhaustion of the budget constraint gives

$$w \sum_{i=\chi}^{\psi-1} P_i(r) = \sum_{i=\chi}^{J-1} P_i(r) c_i,$$

we have $A_J(r) = 0$: households spend their full life time wealth on consumption. We can now calculate the aggregate net demand for assets $S(r)$ for all cohorts at time $t = 0$:

$$\begin{aligned}
S(r) &= \sum_{i=\chi}^{J-1} N_{-i} A_i(r) = w \sum_{i=\chi}^{\psi-1} N_{-i} \sum_{j=\chi}^i P_{j-i}(r) + w \sum_{i=\psi}^{J-1} N_{-i} \sum_{j=\chi}^{\psi-1} P_{j-i}(r) - \sum_{i=\chi}^{J-1} N_{-i} \sum_{j=\chi}^i P_{j-i}(r) c_j \\
&= w \sum_{i=\chi}^{\psi-1} \frac{P_i(g)}{P_i(r)} \sum_{j=\chi}^i P_j(r) + w \sum_{i=\psi}^{J-1} \frac{P_i(g)}{P_i(r)} \sum_{j=\chi}^{\psi-1} P_j(r) - W \sum_{i=\chi}^{J-1} \frac{B_i(r, b)}{B_{J-1}(r, b)} \frac{P_i(g)}{P_i(r)} \\
&= w \sum_{i=\chi}^{\psi-1} \frac{P_i(g)}{P_i(r)} R_{i+1}(r) + w \sum_{i=\psi}^{J-1} \frac{P_i(g)}{P_i(r)} R_{\psi}(r) - w R_{\psi}(r) \sum_{i=\chi}^{J-1} \frac{B_i(r, b)}{B_{J-1}(r, b)} \frac{P_i(g)}{P_i(r)}
\end{aligned}$$

The total value of land is

$$\begin{aligned}
\bar{L}p_0^L &= \sum_{i=1}^{\infty} \frac{\gamma}{1-\gamma} \frac{Y_i - I_i}{(1+r)^i} \\
&= \frac{\gamma}{1-\gamma} (Y_0 - I_0) \frac{1+g}{r-g} \\
&= \frac{\gamma}{1-\gamma} \frac{1+g}{r-g} R_{\psi}(g)(y - (g + \delta)k)
\end{aligned}$$

Since the capital stock satisfies $K_t = kL_t$, the total net demand for assets minus next period's capital stock, S^* (still at $t = 0$), reads

$$S^*(r) = S(r) - K_1 - \bar{L}p_0^L = S - (1+g)K_0 - \bar{L}p_0^L.$$

Asset balance is given by $S^* = 0$, that is, if

$$\begin{aligned}
w \sum_{i=\chi}^{\psi-1} \frac{P_i(g)}{P_i(r)} R_{i+1}(r) + w \sum_{i=\psi}^{J-1} \frac{P_i(g)}{P_i(r)} R_{\psi}(r) - w R_{\psi}(r) \sum_{i=\chi}^{J-1} \frac{B_i(r, b)}{B_{J-1}(r, b)} \frac{P_i(g)}{P_i(r)} &= \\
(1+g)k R_{\psi}(g) + \bar{L}p_0^L & \\
\Leftrightarrow \sum_{i=\chi}^{\psi-1} \frac{P_i(g)}{P_i(r)} R_{i+1}(r) + \sum_{i=\psi}^{J-1} \frac{P_i(g)}{P_i(r)} R_{\psi}(r) - R_{\psi}(r) \sum_{i=\chi}^{J-1} \frac{B_i(r, b)}{B_{J-1}(r, b)} \frac{P_i(g)}{P_i(r)} &= \\
(1+g) \frac{k}{w} R_{\psi}(g) + \frac{\gamma}{1-\gamma} \frac{1+g}{r-g} R_{\psi}(g) \left(\frac{y}{w} - (g + \delta) \frac{k}{w} \right) &
\end{aligned}$$

Since

$$\begin{aligned}
r + \delta &= \alpha k^{-\frac{1}{\sigma}} \left(\alpha k^{\frac{\sigma-1}{\sigma}} + 1 - \alpha \right)^{\frac{1}{\sigma-1}}, \\
w &= (1 - \alpha) \left(\alpha k^{\frac{\sigma-1}{\sigma}} + 1 - \alpha \right)^{\frac{1}{\sigma-1}},
\end{aligned}$$

k and k/w satisfy

$$\begin{aligned}
k^{\frac{\sigma-1}{\sigma}} &= \frac{1 - \alpha}{\left(\frac{r+\delta}{\alpha} \right)^{\sigma-1} - \alpha}, \\
\frac{k}{w} &= \frac{\alpha}{1 - \alpha} \frac{k^{\frac{\sigma-1}{\sigma}}}{r + \delta} = \frac{1}{r + \delta} \frac{\alpha}{\left(\frac{r+\delta}{\alpha} \right)^{\sigma-1} - \alpha}.
\end{aligned}$$

2.8.2 Adjustment algorithm

As demography tends from an initial state of high population growth to an eventual state of low population growth, the equilibrium real interest rate must likewise transition from a high level to a low level (the BGP real interest rate is calculated using the equation described above). Note that real interest rates characterise the entire transition, as all allocations may be derived as a function of the real interest rate path (although the allocation in any period shall depend on

future interest rates due to forward-looking agents). For the transition between the two BGP's, the real interest rate transition is calculated iteratively using an adjustment mechanism that we describe below.

We assume that the transition period lasts a maximum of $10J$ periods (where J is the life span of a household), i.e. after $10J$ periods the economy has converged to the terminal BGP. For the first iteration, we make the assumption that the real interest rate transition is a linear interpolation between the initial level to the final level. Each successive iteration is then calculated in the following way.

In the $k + 1$ th iteration of our adjustment algorithm, we begin with the past iteration's real interest rate path, which we denote to be $\{r_t^k\}$ (the initial linear interpolation is denoted by $\{r_t^0\}$), and derive from this the implied path for wages, output, and so on. Using $\{r_t^k\}$ and the implied path of wages, we solve the household problem in *every* period for *every* cohort. This allows us to calculate aggregate consumption demand in each period t :

$$C_t^{k+1} = \sum_{i=0}^{J-1} N_{t-i} c_{t-i,i}^{k+1},$$

where $c_{t-i,i}^{k+1}$ is the demanded consumption in period t by the cohort born in period $t - i$, in iteration $k + 1$.

Using the path we calculate for aggregate consumption demand, we can calculate aggregate savings supply by

$$S_t^{k+1} = w_t^{k+1} L_t - C_t^{k+1}.$$

Note the distinction between savings and investment. Savings equals to investment minus capital returns, i.e. savings is the additional quantity devoted to asset accumulation after capital returns is already reinvested. Using savings supply, we can calculate the desired aggregate asset position through time by

$$A_{t+1}^{k+1} = A_t^{k+1}(1 + r_{t+1}^k) + S_t^{k+1},$$

where r_{t+1}^k denotes the real interest rate in period t , iteration k , and the initial asset position is given by $A_{t^*-1}^{k+1} = K_{t^*}$. This allows us to derive the implied capital stock, consistent with the desired asset position, by

$$K_{t+1}^{k+1} = A_t^{k+1}.$$

All assets in our economy are real, i.e. in equilibrium it follows that $A_{t+1} = K_t$. In principle there could be financial assets, e.g. money, which would mean that the total amount of assets accumulated by households exceeds the level of the capital stock. Note that the implied capital stock is derived from the *desired* asset accumulation of households. The *actual* quantity of asset

accumulated may differ as $1 + r_t^k \neq R_t^K$, the return on capital investment. For the purpose of our algorithm however, we are interested in the *desired* asset position of households, hence we assume that assets grow at the real interest rate of the past iteration, $1 + r_t^k$.

Using the implied capital stock, we may calculate the implied level of capital per worker, and hence induce the implied real interest rate path. We call the path of the implied real interest rates $\{r_t^I\}$. The real interest rate path for the next iteration is finally given by

$$r_t^{k+1} = (1 - \phi)r_t^k + \phi r_t^I,$$


where ϕ is a parameter that controls the degree to which we adjust per iteration. Note that for a stable adjustment, we use a conservative level for ϕ , around the level of 0.05. A larger ϕ leads to over-adjustment and explosive behavior of the algorithm.



Finally, we compare $\{r_t^k\}$ and $\{r_t^I\}$. If the Euclidean distance between these two vectors is sufficiently small, we have reached convergence and we end our search for the equilibrium transition path at $\{r_t^{k+1}\}$. Note that market clearing is directly guaranteed by convergence, as when $r_t^k = r_t^I \forall t$, it must be the case that in each period aggregate consumption demand equals to aggregate labour income, plus aggregate capital income, minus aggregate capital investment.



2.8.3 Alternative parameter specifications



We offer alternative parameter specifications for our simulations in section 3.

OLG model









References

- Kenneth J Arrow, Hollis B Chenery, Bagicha S Minhas, and Robert M Solow. Capital-labor substitution and economic efficiency. *The Review of Economics and Statistics*, pages 225–250, 1961.
- Orazio P Attanasio, James Banks, Costas Meghir, and Guglielmo Weber. Humps and bumps in lifetime consumption. *Journal of Business & Economic Statistics*, 17(1):22–35, 1999.
- Richard Baldwin and Coen Teulings. Secular stagnation: facts, causes and cures. *London: Centre for Economic Policy Research-CEPR*, 2014.
- Robert J Barro and Gary S Becker. Fertility choice in a model of economic growth. *Econometrica: journal of the Econometric Society*, pages 481–501, 1989.
- Carlos Carvalho, Andrea Ferrero, and Fernanda Nechio. Demographics and real interest rates: Inspecting the mechanism. *Available at SSRN 2713443*, 2016.
- Robert S Chirinko. σ : The long and short of it. *Journal of Macroeconomics*, 30(2):671–686, 2008.
- Robert S Chirinko, Steven M Fazzari, and Andrew P Meyer. That elusive elasticity: a long-panel approach to estimating the capital-labor substitution elasticity. 2004.
- Chadwick C Curtis, Steven Lugauer, and Nelson C Mark. Demographic patterns and household saving in china. *American Economic Journal: Macroeconomics*, 7(2):58–94, 2015.
- Peter A Diamond. National debt in a neoclassical growth model. *The American Economic Review*, 55(5):1126–1150, 1965.
- Gauti B Eggertsson and Neil R Mehrotra. A model of secular stagnation. Technical report, National Bureau of Economic Research, 2014.
- Claudia Goldin and Lawrence F Katz. Career and marriage in the age of the pill. *The American Economic Review*, 90(2):461–465, 2000.

- Charles AE. Goodhart and Philipp Erfurth. Demography and economics: Look past the past. URL <http://voxeu.org/article/demography-and-economics-look-past-past>, 2014.
- Robert J Gordon. Is us economic growth over? faltering innovation confronts the six headwinds. Technical report, National Bureau of Economic Research, 2012.
- Axel Gottfries and Coen Teulings. Can demography explain secular stagnation? URL <http://voxeu.org/article/role-demography-explaining-secular-stagnation>, 2015.
- Tomáš Havránek. Measuring intertemporal substitution: The importance of method choices and selective reporting. *Journal of the European Economic Association*, 13(6):1180–1204, 2015.
- N Gregory Mankiw and David N Weil. The baby boom, the baby bust, and the housing market. *Regional science and urban economics*, 19(2):235–258, 1989.
- M Ishaq Nadiri and Ingmar R Prucha. Estimation of the depreciation rate of physical and r&d capital in the us total manufacturing sector. *Economic Inquiry*, 34(1):43–56, 1996.
- Lukasz Rachel and Thomas Smith. Secular drivers of the global real interest rate. 2015.
- Paul A Samuelson. An exact consumption-loan model of interest with or without the social contrivance of money. *The journal of political economy*, pages 467–482, 1958.
- Lawrence Summers. Why stagnation might prove to be the new normal. *Financial Times*, 15, 2013.

Chapter 3

Incomplete Information and Housing Booms and Busts

3.1 Introduction

Extensive research has studied the effect of the housing market on the aggregate economy. The consensus of this research finds that house price variation has widespread effects on household spending, household borrowing and default, and local employment (see e.g. Campbell and Cocco (2007,2015), Cloyne et.al. (2019), Guren et.al. (2018), and Mian and Sufi (2011), Mian et.al. (2013), Mian and Sufi (2014)). In contrast, there is disagreement over the source of the house price variation. One trend of literature views price changes in the housing market as primarily efficient (see e.g. Himmelber et.al. (2005), Wu et.al. (2012), and Landvoigt et.al. (2015)), whereas others view the housing market as primarily inefficient (see e.g. Shiller (2007), Case et.al. (2012), Glaeser and Nathanson (2017)). This poses a challenge as depending on one's views on this matter, the design of models and potentially the results of one's analysis can differ. Overall, although there is compelling empirical evidence in favour of the latter view of an inefficient housing market, see Piazzesi and Schneider (2016) for a summary, models that contend with inefficient housing markets often need to impose restrictions on the decision-making of households. For example, some papers consider extrapolative expectations on the part of homebuyers, as opposed to the standard rational expectations framework. Assumptions such as these nevertheless offer the scope for study of housing in a wider setting, however they invite debate about the validity of the assumptions and consequently the validity of the results. Much of the theoretical work on housing simply goes around this issue by considering exogenous movements in prices or fundamentals, however there is a question of whether it is possible to generate the observed house price behaviour without needing to impose on the decision-making of households.

In this chapter I develop a search and matching model of housing markets where the key innovation is that households do not observe the underlying fundamentals and instead must infer fundamentals from prices. This leads to a Bayesian signal extraction problem that households solve optimally. Importantly, I do not impose assumptions on the decision-making of households. My model reproduces the main empirical features associated with housing market inefficiency, namely short-term price momentum, long-term price reversal, and boom and bust cycles in house prices. Furthermore, I show that these features are directly generated by the price inference mechanism: in the complete information benchmark, where households directly observe fundamentals, prices resemble a random walk and hence do not exhibit features associated with market inefficiency.

The fundamentals in my model are summarised by short- and long-term shocks that realise in every period. These stochastic shocks drive housing demand and serve as the basis for a dynamic model of house prices. Importantly, long-term shocks have a larger effect on the housing market because they are persistent. The long-term component to housing demand, defined as the aggregated contributions of the history of long-term shocks, is of particular interest to households as it pertains to the value of continued search on the housing market both from the perspective of a home buyer and a seller. The Bayesian signal extraction problem can then be restated as the problem of how to infer this long-term component based on information from the path of prices.

At the start of each period, every buyer is matched with a house that is for sale and this generates a quality draw that measures the utility flow of the buyer living in that matched house. While the quality draw is revealed, this draw reflects the sum of the short-term shock, the long-term component to housing demand, as well a stochastic idiosyncratic draw, and the contributions of each component is not known. Nash bargaining occurs over the transaction price if there is a joint surplus. At the end of the period the average transaction price across all transactions is revealed to the households. I show that the average transaction price is a monotonic function of the sum of the short and long-term components, which I define as the price signal. This result is the solution to the Bayesian signal extraction problem, and it states that the contribution of the stochastic idiosyncratic draw is averaged out once households learn of the average transaction price. Households then optimally update their expectations of the long-term component based on this price signal. However, as the price signal is a noisy measure of the long-term component due to the contribution of the short-term shock, households are conservative and adjust their expectations towards the price signal only partially. Persistent changes to the price signal do nonetheless become fully internalised as expectations adjust and converge towards the new level.

This type of updating of expectations based on prices is novel to this paper. The updating

process leads to “adaptive expectations” for the long-term component that is generated endogenously as a result of optimal household decision-making, which contrasts to much of the existing literature where adaptive expectations are simply assumed. Furthermore, I show that expectations for the long-term component can be written explicitly as a backward-looking average of the price signal. This similarly provides a rationale for “backward-looking expectations” that is often considered a central feature of housing markets. The implication of this type of expectation formation is the possibility of bubble episodes. A long period of sustained short-term shocks can deceive households, leading them to believe a persistent long-term shift has occurred when this is not the case. Specifically I show that the same sequence of short-term shocks generates a large boom and bust in prices in the incomplete information model, while there is little response in prices for the complete information model.

This chapter relates to the broader discussion of efficiency in financial markets. The Efficient Market Hypothesis (EMH) states that at any given time in a market prices reflect all available information, see Malkiel and Fama (1970). Empirically, there has been extensive investigation regarding whether the EMH holds in the stock market, see e.g. Lim and Brooks (2011) for a summary. The conclusions vary and depend on the time-period, dataset, asset pricing model used for calculating the market return, and the exact EMH testing methodology. One of the main critics of the EMH, Robert Shiller, received his Nobel prize for his work showing that markets move too much to be explained by rational investors responding to changing fundamentals. In particular, Shiller highlights that the housing market is especially inefficient, where he claims the market is speculative and prone to bubbles. The following empirical works document various “puzzles” which are in conflict with the EMH in the housing market. Case and Shiller (1990), Englund and Ioannides (1997) and Glaeser and Nathanson (2015), Titman et.al. (2014), Guren et.al. (2018) and Armona et.al. (2018) document short-term momentum and long-term reversal in prices, which contrasts with the EMH where prices form a random walk. Piazzesi and Schneider (2016) document the excess volatility puzzle of house prices. Flavin and Yamashita (2002), Burnside et.al. (2016) document boom-bust cycles in housing. Martínez-García and Grossman (2018) formally detect bubbles in housing prices with a formal test of Phillips et.al. (2015). For a broader discussion of housing market efficiency see also Pavlidis et.al. (2017); Engsted et.al. (2016); Shi (2017); Hu and Oxley (2018).

Many factors affect the housing market and house prices. These factors are varied in the context of how they play a role and they differ also in their persistence. In my model I simplify by categorising any shock to these factors as either being short versus long term, however I do not go beyond this to consider these factors individually in their original contexts. Demography is one example of a long-term factor that affects the housing market. Faster population growth leads to increased demand and prices. Furthermore, the decision to purchase housing is a life-cycle

one. In order to purchase a house, one is required to make a large financial commitment in the form of a down payment that often requires many years to accumulate. Chambers et.al. (2009) shows that the homeownership rate increases from roughly 40% for young households (aged 20-34 years) to twice that share for older households (aged 65-74 years). The homeownership rate then declines slightly for very old households. Furthermore, housing is special in that it provides simultaneously a consumption good, namely housing services, while also serving as an asset for investment. The type and size of housing demanded evolves over time and with the age of the household, see for example Morrow-Jones and Wenning (2005). As an investment vehicle housing also serves the purpose of precautionary savings through various means of equity extraction, see for example Ong et.al. (2013) and Bhutta and Keys (2016). Consequently, long term socioeconomic trends in demography, family structure, labour force structure, to name a few, all contribute to determining the current state as well as the future trajectory of the housing market. The above are all examples of factors that drive housing appetite and I look to summarise them as a long-term shock or the long-term component to demand. In addition, changes in household income and savings, the employment rate, and household prospects for the future, each serve to drive the housing market in the short run, see Hwang and Quigley (2006). I look to summarise these factors as the short-term shock/component to housing demand. Jointly, these two components generate the basis for my dynamic model of house prices.

This chapter also relates to the broader behavioural finance literature, specifically in relation to the study of housing markets. Barberis et.al. (2018) examine the ways extrapolative expectations can generate short term bubbles of different magnitudes and durations. Glaeser and Nathanson (2017) consider a similar framework where households take boundedly rational approximations to approximate fundamentals, which they show can generate bubbly housing markets. Kaplan et.al. (2018) find that shifts in beliefs about future housing demand were the dominant force behind the observed swings in house prices and the rent-price ratio around the Great Recession. Case and Shiller (2012) and Kuchler and Zafar (2015) document correlations between past home price changes and subjective home price expectations. Burnside et.al. (2016) develop a model where agents optimistic beliefs about future house prices can spread like an infectious disease. Boz and Mendoza (2014) show that a model with Bayesian learning about a regime shifting loan-to-value limit can produce a pronounced run-up in credit and land prices followed by a sharp and sudden drop. Gete (2018) shows that introducing the backward-looking survey expectations from Case et.al. (2012) into a DSGE model with housing can help account for movements in U.S. house prices over the period 1994 to 2012. Gelain et.al. (2018) show that specification with random walk expectations outperforms one with rational expectations in plausibly matching the patterns in the data.

The rest of the chapter shall be organised as follows. In section 3.2 I describe the basic setup

of my model. In section 3.3 I present the results of my model and how the price behaviour of the incomplete information model contrasts to the complete information benchmark. In section 3.4 I consider the application of model to identify bubbles in US house price data. Finally, in section 3.5 I consider several extensions which generalise various assumptions made in the basic setup in section 3.5. Section 3.6 concludes.

3.2 The model

In this section I describe the setup of the model and three main theoretical results that summarise the solution to the Bayesian signal extraction problem.

The housing market is made up of a measure 1 of households and a measure 1 of the indivisible housing asset. Each house is owned by some household, that is to say in any period there is an ownership function from the measure 1 of houses to the measure 1 of households. It is possible for one household to own multiple houses so this function is not generically injective. In order for a particular household to enjoy the services from a house it must own the house and occupy that house. Otherwise, the households will be a renter and earn a baseline level of rental utility. Each household can only occupy one house at a time, which implies that there is a bijective function from the subset of *owner-occupant* households to the subset of owner-occupied houses. Any household that does not occupy a house they own is a prospective home buyer in the market and any house that is not owner-occupied is available for sale. Given that the function between owner-occupant households and owner-occupied houses is bijective, this means that the measure of home buyers and houses for sale is equal. An owner-occupancy relationship begins when a home buyer is matched with a house that is for sale and decides to buy it from a seller. An owner-occupancy relationship ends stochastically with exogenous probability ϕ in every period after the initial match.




Figure 3.1: Several cases of the household to housing mapping.

3.2.1 Market structure

In every time period there is a sequence of events detailed as follows. At the start of the period, a proportion ϕ of all owner-occupant households are separated from their houses. These separated households become house buyers and house sellers simultaneously. Next, each home buyer is matched with a house that is for sale. The matching gives each home buyer indexed by i a stochastic draw $\epsilon_{i,t}$, which shall be described next, that measures the quality of this potential match. The quality draw gives the utility flow in each period for this home buyer to occupy this particular house. A transaction occurs whenever the joint surplus is above the zero threshold, and Nash bargaining occurs to determine the transaction price. After this round of housing transaction, any household that remains unmatched is allocated at random to an unmatched house for rental use. This rental use generates a baseline level of utility u_{rent} but incurs a rental charge of R and a search cost of c_{search} . Likewise, each for-sale house entitles the seller to receive R in rental income however it must pay c_{post} to continue its sell posting. At the end of the period, the prices of all transactions are averaged and this is revealed to all households, which may or may not provide a price signal depending on the information set of the households.

The quality draw is made up of an aggregate short-term component, A_t , an aggregate long-term component, χ_t , and an idiosyncratic stochastic component, $\eta_{i,t}$

$$\epsilon_{i,t} = A_t + \chi_t + \eta_{i,t}. \quad (3.1)$$

In addition, components in the above equation have the following distributions:

$$A_t \sim N(0, \sigma_A^2), \quad (3.2)$$

$$\eta_{i,t} \sim U(-M_\eta, M_\eta), \quad (3.3)$$

$$\chi_t = \rho\chi_{t-1} + e_t, \quad (3.4)$$

$$\text{with } e_t \sim N(0, \sigma_e^2), \quad (3.5)$$

where $\rho \in [0, 1)$ is the AR(1) coefficient. The key distinction between the two aggregate components is the difference in persistence. Here I take the position that the short-term component is noise in every period, however it is possible to reproduce my results where the short-term component is persistent but less nonetheless persistent than the long-term component. See section 3.5 on extensions for a derivation of this alternative model.

I assume that the exchange of money in a housing transaction translates directly to a transfer of utility. This can be for example generated as the result of optimal behaviour from households having quasilinear utility with respect to consumption or quasilinear disutility from providing labour, see for example Lagos and Wright (2005). In section 3.5 I consider a version of my model where there is endowment income for households and there is quasi-linear utility over consumption; this version of the model generates these value functions endogenously. For now, I simply state the Bellman equations that describe the household problem, in a manner similar to Ngai and Tenreyro (2014). The Bellman equation for a home buyer after observing the realisation of the quality draw is given by

$$\begin{aligned} V_B(\Omega, \epsilon) = \max \Big\{ & u_{rent} - R - c_{search} + \beta[V_B(\Omega', \epsilon')], \\ & \epsilon - p(\Omega, \epsilon) + \beta[\phi(V_B(\Omega', \epsilon') + V_S(\Omega', \epsilon')) + (1 - \phi)V_M(\Omega', \epsilon)] \Big\}, \end{aligned} \quad (3.6)$$

where Ω is the set of observable aggregate states (the contents of which depends on the information structure under consideration), V_B , V_S and V_M are the value functions for buying, selling and matched states respectively, and p is the Nash bargaining price function.

The Bellman equation for a seller when considering the value of a particular house for sale is given by

$$V_S(\Omega, \epsilon) = \max \Big\{ R - c_{post} + \beta[V_S(\Omega', \epsilon')], \quad p(\Omega, \epsilon) \Big\}. \quad (3.7)$$

Note that this is the Bellman equation for an additional house that is for sale. One household may own multiple houses however given that my model is linear it is sufficient to consider the selling problem from the perspective of a single house for sale.

The Bellman equation for a owner-occupant household is given by

$$V_M(\Omega, \epsilon) = \epsilon + \beta[\phi(V_B(\Omega', \epsilon') + V_S(\Omega', \epsilon')) + (1 - \phi)V_M(\Omega', \epsilon)]. \quad (3.8)$$

3.2.2 Information structure

In order to contrast the results of the incomplete information model, I consider a complete information case where both the short-term and long-term components are observable at the start of the period. That is to say $\Omega = \{A, \chi\}$. In this situation the average price revealed at the end of the period provides no additional information as a price signal, and indeed this price is perfectly predictable at the start of the period, given Ω , by taking averages over all realisations of the idiosyncratic draw.

In the incomplete information model, I suppose that households in period t have priors (or expectations) for the long-term component *last period* equal to χ_{t-1}^e . This is the only information brought into the start of the period t and hence $\Omega_t = \chi_{t-1}^e$. Next, housing is transacted and the rental market is cleared which reveals the average transaction price p_t at the end of the period. Note that because the innovation to the long-term component and the short-term shock are both mean zero, the priors for the long-term component overlap with the expectations for the long-term component. Figure 3.2 below graphically shows how the average transaction price is formed depending on the signal. Here ϵ^* is the threshold for the idiosyncratic draw such that there is zero surplus, and the average transaction price is the mean surplus given to seller over the transaction region.




Figure 3.2: Average transaction price depending on the signal.

In the incomplete information framework, revealed prices have informational content and in

particular one can state the following result.

Proposition 3. *Given expectations for long-term component last period, the average transaction price is a strictly increasing function of the sum of the short-term and long-term components, $A_t + \chi_t$, which I shall refer to as the price signal, s_t . Hence the knowledge of p_t at the end of the period t directly maps to the period t price signal.*

Proof. The surplus of a potential match as a function of the match quality ϵ is given by

$$\begin{aligned} S(\epsilon) = & \epsilon + \beta[\phi(V_B(\Omega', \epsilon') + V_S(\Omega', \epsilon')) + (1 - \phi)V_M(\Omega', \epsilon)] \\ & - (R - c_{post} + \beta[V_S(\Omega', \epsilon')] + u_{rent} - R - c_{search} + \beta[V_B(\Omega', \epsilon')]). \end{aligned} \quad (3.9)$$

Let ϵ^* be the threshold at which the above surplus is zero, which we assume always exist due to assumptions on the parameters of the idiosyncratic shock, M_η . It follows that the Nash bargaining price for this transaction is given by

$$p(\Omega, \epsilon) = \frac{1}{2}S(\epsilon) \quad \text{if } \epsilon > \epsilon^*. \quad (3.10)$$

Recall also that

$$\begin{aligned} \epsilon &= A_t + \chi_t + \eta_{i,t} \\ &= s_t + \eta_{i,t}. \end{aligned}$$

Hence the average transaction price at the end of the period is given by

$$p_t = \frac{1}{s_t + M_\eta - \epsilon^*} \int_{\epsilon^*}^{s_t + M_\eta} p(\Omega, \epsilon) \, d\epsilon. \quad (3.11)$$

Consider now a higher signal s'_t , such that $s'_t = s_t + \delta$ for some $\delta > 0$. For this higher signal, the average transaction price is given by

$$p'_t = \frac{1}{s_t + \delta + M_\eta - \epsilon^*} \int_{\epsilon^*}^{s_t + \delta + M_\eta} p(\Omega, \epsilon) \, d\epsilon. \quad (3.12)$$

This average transaction p'_t is higher than p_t as $V_M(\Omega', \epsilon)$ and consequently $p(\Omega, \epsilon)$ is strictly increasing in ϵ . This establishes the claim in Proposition 3. □

The price signal is a noisy measure for the long-term component. It matters for home buyers and sellers alike how much of the price signal is comprised of the long-term component as opposed to the short-term component due to the fact that long-term component is persistent. Specifically, the long-term component affects the value of continued search on the behalf of home buyers and continued posting on the behalf of sellers. Note also that information realised in period t is partially revealed only at the end of period t , after transactions have already been completed at the start of the period. Hence households at best can only respond with a lag to new information arrival.

After the extraction of the price signal, how should households update their priors or expectations? I derive the analytical expression for optimal χ_t^e updating which is formally stated in Proposition 4 below, however in fact this problem can be restated as special case of the Kalman filtering problem (see Ljungqvist and Sargent (2018) for an overview of the application of Kalman filters in economics).

Proposition 4. *Conditional on period t priors of the long-term component of $\chi_{t-1} = \chi_{t-1}^e$, the maximum likelihood updating of χ_t^e based on a revealed price signal in period t of s_t is*

$$\chi_t^e = \rho\chi_{t-1}^e + \left(\frac{\sigma_e^2}{\sigma_A^2 + \sigma_e^2}\right)(s_t - \rho\chi_{t-1}^e). \quad (3.13)$$

Proof. Given expectations of the long-term component last period, the signal this period from the perspective of the households is given by

$$s_t = \rho\chi_{t-1}^e + e_t + A_t. \quad (3.14)$$

The joint density of A_t and e_t is given by

$$f(A_t, e_t) = \frac{1}{2\pi\sigma_A\sigma_e} \exp(-A_t^2/2\sigma_A - e_t^2/2\sigma_e). \quad (3.15)$$

Given knowledge of the signal, we may therefore rewrite the likelihood function in terms of e_t in the following way

$$\mathcal{L}(e_t) = \frac{1}{2\pi\sigma_A\sigma_e} \exp(-(s_t - e_t))^2/2\sigma_A - e_t^2/2\sigma_e). \quad (3.16)$$

The first-order condition is given by

$$\frac{1}{2\pi\sigma_A\sigma_e} \exp(-(s_t - e_t^{ML}))^2/2\sigma_A - (e_t^{ML})^2/2\sigma_e)((s_t - e_t^{ML})/\sigma_A - e_t^{ML}/\sigma_e) = 0 \quad (3.17)$$

Solving for e_t^{ML} gives

$$e_t^{ML} = \left(\frac{\sigma_e^2}{\sigma_A^2 + \sigma_e^2} \right) (s_t - \rho\chi_{t-1}^e), \quad (3.18)$$

which establishes the claim in Proposition 4. □

To visualise the above result, consider Figure 3.3 below:




Figure 3.3: Updating of χ^e .

Any deviation of s_t from $\rho\chi_{t-1}^e$ is a “surprise” to the household as the short-term shock and the innovation to the long-term component are both mean zero. Given the close relationship between the price signal and the long-term component, Proposition 4 states intuitively that households will update their priors positively with s_t . Households update χ_t^e towards the signal, however they are conservative in their updating. The updating proportion, $\sigma_e^2/(\sigma_A^2 + \sigma_e^2) < 1$, depends positively on σ_e and negatively on σ_A . The final result derives immediately from Proposition 4:

Proposition 5. *Supposing households update priors according to Proposition 4, priors at the end of period t is a backward-looking average of the history of price signals given by*

$$\chi_t^e = \sum_{i=0}^{t-1} \left(\frac{\sigma_e^2}{\sigma_A^2 + \sigma_e^2} \right) \left(\frac{\sigma_A^2}{\sigma_A^2 + \sigma_e^2} \right)^i \rho^i s_{t-i} + \left(\frac{\sigma_A^2}{\sigma_A^2 + \sigma_e^2} \right)^t \rho^t \chi_0^e. \quad (3.19)$$

Proof. The proof immediately follows from Proposition 4 by backward substitution of χ_{t-1}^e . □

Proposition 5 summarises how the history of price signals affects today’s priors/expectations. Consequently the coefficient of persistence for the price signal is given by $\rho\sigma_A^2/(\sigma_A^2 + \sigma_e^2) < 1$. The implication is that priors are persistent due to their backward-looking nature, and furthermore they are slow in adjusting to new information due to conservatism in response to a noisy signal.

3.3 Simulation and impulse response functions

I solve the model via standard value function iteration techniques, however one needs to obtain the Nash bargaining price function, which maps the state vector to the average transaction price, as well as the three value functions.

For the incomplete information model, the aggregate state vector is simply the expectation of the long-term component from the last period. In each value function iteration, for all points on a grid for the aggregate state home buyers and sellers match over all possible realisations of the idiosyncratic draw. For each level of the idiosyncratic draw, households evaluate the transaction surplus conditional on the value functions from the previous value function iteration. This yields the entire distribution of transaction prices, for all points on a grid for the aggregate state. I use equations (3.6) - (3.8) to update the three value functions; the expectation of next period values is obtained by taking expectations of the last iteration's value functions evaluated over all possible innovations to the long-term component *today*, and the contemporaneous contributions to the value functions are already obtained from solving the household problem above. Finally, using the distribution of transaction prices I am able to calculate the mean transaction price as a function of the state vector. I iterate upon the value functions and the price function until convergence.

Solving the complete information model is entirely analogous to solving the incomplete information model. In this case, the aggregate state vector is the realised long-term component for the this period. In each value function iteration, for all points on a grid for the aggregate state home buyers and sellers match over all possible realisations of the idiosyncratic draw. For each level of the idiosyncratic draw, households similarly evaluate the transaction surplus conditional on the value functions from the previous value function iteration. This again yields the entire distribution of transaction prices, for all points on a grid for the aggregate state. I use equations (3.6) - (3.8) to update the three value functions, however this time the expectation of next period values is obtained by taking expectations of the last iteration's value functions evaluated over all possible innovations to the long-term component *tomorrow*. The contemporaneous contributions to the value functions are again already obtained from solving the household problem above. Finally, I calculate the mean transaction price and I iterate this entire procedure until convergence.

Next I describe the calibration of my model before showing the results of my simulations and the impulse response functions.

3.3.1 Calibration of the incomplete information model

The calibration of my model utilises two datasets. The first dataset is the annual US national house price index data from the Federal Housing Finance Agency between 1975 and 2018. This house price index is based on repeated sales of houses which ensures comparability of the index over time as the composition of houses that are transacted is liable to change.

Furthermore, I also use the annual US national homeownership rate data provided by the United States Census Bureau between 1965 and 2018. Although homeownership rate is not a direct equilibrium outcome in the model, it is possible to track the homeownership rate by tracking the proportion of home buyers that become owner occupants. This allows me to further discipline the dynamics of my model by matching moments of my model to the data.

Table 3.1: List of Parameters

Type	Parameter	Value	Meaning
Externally Calibrated	β	0.98	discount rate of households
	ϕ	1/13	rate of ownership separation
Internally Calibrated	σ_A	1.1	standard deviation of short-term shock
	σ_e	1.2	standard deviation of long-term component's innovation
	M_η	3	dispersion of idiosyncratic draws
	ρ	0.9	persistence of long-term component

Table 3.1 describes the set of parameters that are necessary to calibrate. Specifically, the duration of homeownership based on the assumption of my model follows a geometric distribution with parameter ϕ . The rate of ownership separation is selected to match the average duration of homeownership, which according to National Association of Home Builders is 13 years in the USA. Hence ϕ is chosen to be 1/13.

Table 3.2: Targeted Moments for Incomplete Information Model

Type	Moment	Empirical Value	Model Value
Targeted	volatility of price levels	0.04	0.03
	autocorrelation coefficient of price levels	0.9	0.91
	mean ownership rate	0.66	0.68
	standard deviation of homeownership rate	0.02	0.03

The four internally calibrated parameters are used to jointly match the four targeted moments in my incomplete information model to the equivalent four moments from the data described in Table 3.2. Intuitively, the standard deviation of the short-term shock versus the standard deviation of the long-term shock affects the updating parameter described in Proposition 4. When the standard deviation of the short-term shock is large relative to the long-term shock, households are conservative in updating expectations and this affects both the average volatility of price levels and the mean homeownership rate. The dispersion of the idiosyncratic draws controls the importance of the two aggregate shocks; when the dispersion of the idiosyncratic draw is large, the relative importance of both aggregate shocks is smaller. This parameter closely relates to the standard deviation of the homeownership rate as when the idiosyncratic draw is dominant, homeownership rates are more stable over time. Finally, the persistence of the long-term component directly relates to the autocorrelation coefficient (i.e. persistence) of price levels as the parameter ρ is the only source of persistence in the model.

The relationship between the parameters and the moments is highly nonlinear, hence it is difficult to match all four moments exactly. I solve and simulate the model to obtain the moments at each point in a large grid over the four parameters. I then select the point on the grid that most closely matches the moments in my data and this provides the values of my internally calibrated parameters described in Table 3.1.

Table 3.3: Non-targeted Moments

Type	Moment	Empirical Value	Model Value
Incomplete information	One-period autocorrelation in price changes	0.65	0.41
	Ten-period autocorrelation in price changes	-0.14	-0.09
Complete information	One-period autocorrelation in price changes	0.65	0
	Ten-period autocorrelation in price changes	-0.14	-0.01

The calibration of my model results in the following non-targeted moments in my model which I also compare to the data. Specifically, these moments relate to the short-term momentum in house prices, namely one-period autocorrelation in price changes, and the long-term reversal of house prices, namely ten-period autocorrelation in house price changes. I show these non-targeted moments for both the incomplete information model, as well as the complete information model, with the same parameters. We see that the incomplete information model generates both these features of house price data to a quantitatively similar magnitude, whereas the complete information model generates essentially zero house price momentum short term and near zero house price reversal long term.

Note that the parameters for the model are based on the calibration of incomplete information model to data, hence Table 3.3 is a somewhat unfair comparison between the performance of the incomplete information model versus the complete information model. However, the calibration of the complete information model to the same moments in data generates broadly similar parameter values and likewise the resulting untargeted moments are robust insofar as generating zero house price momentum and near zero house price reversal. Indeed, the rationale for these findings is better understood in the later subsection on impulse response functions shown later.

3.3.2 Simulations

To examine the behaviour of the incomplete information model versus the complete information model, I simulate generate $T = 100$ periods (years) of the short-term and long-term shocks. These add together to give the price signal.

For the incomplete model, Proposition 4 gives the updated long-term component expectations given the last period's long-term component expectations and the price signal today. Hence I am able to calculate the entire path of long-term component expectations, subject to an initial level for the long-term component expectations which I set to the unconditional mean of zero for simplicity.

For the complete information model all information is observed and there is no need for any additional tracking of the aggregate state. In each period, I explicitly calculate the distribution of transaction prices over all realisations of the idiosyncratic draw for both versions of the model. I take the means of these price distributions to obtain the average transaction prices and furthermore I calculate the proportion of matches that generate transactions. The proportion of matches that generate transactions is used to update the homeownership rate, given an initial level which I set to the long-term average of 0.68 for simplicity. I show the results of the simulation in Figure 3.4 for both the incomplete information model versus the complete information benchmark for comparison.




Figure 3.4: Simulation results of incomplete and complete information models.

Figure 3.4 is organised in the following way. Both figures on the left-hand side are relating to the dynamics of the incomplete information model and how it relates to the signal, which is the sum of the short and long term components. Both figures on the right-hand side are relating to the dynamics of the incomplete model and how it compares to the dynamics of the complete information model. Both figures on the upper half relate to the behaviour of prices whereas both figures on the lower half relates to the behaviour of homeownership rates. The blue lines identify the incomplete information model whereas the red lines identify the complete information model. The black dashed line represents the price signal.

Prices in the incomplete information model follow closely the price signal. Homeownership rates also follow the price signal although this co-movement is less strong. Prices following the signal shows intuitively that the average joint surplus given a transaction occurring is increasing in the signal. This is also a consequence of Proposition 4, which says that price is a positive monotonic function of the signal. Homeownership rates following the signal indicates that the probability of a match to generate a owner-occupancy relationship is increasing in the signal. This is intuitive as the average quality draw is increasing in the signal and high quality draws generate transactions.

Prices in the complete information model follow closely prices in the incomplete information model. This shows that prices in the complete information model also follow closely the signal, as a higher signal similarly drives a higher average joint surplus of transactions. Given that the prices co-move closely, the difference between prices in the two models primarily stems from differences in fluctuations around a trend that is driven by a common signal. Nonetheless this difference in fluctuations around this trend generates the drastically different behaviour of the models with respect to short-term momentum and long-term reversal. Overall, prices in the incomplete information model are smoother as households are optimally conservative in updating their expectations. Similarly to prices, homeownership rates in the complete information model co-moves closely with that of the incomplete information model. Both are driven by the signal with a higher signal increasing the likelihood of a match resulting in a transaction.




Figure 3.5: Signal versus (lead) expected signal.

Figure 3.5 shows the relationship between the signal and the expected long-term component/signal¹ by households in the incomplete information model. The expected long-term component is able to closely follow the path of the realised signal on a grand scale, hence households in the incomplete information model are able to extract information from prices in an overall successful manner. However, due to the conservatism of households in the Bayesian signal extraction problem, the expected signal is slow to adjust and less volatile. We see that sharp movements in the signal are updated in expectations only in the next period and furthermore that expectations are more smooth than the path for the signal.

3.3.3 Impulse response functions

I consider price and homeownership rate impulse response functions for i) a one-period one standard deviation shock to A and ii) a one-period one standard deviation shock to e . In both cases, the short-term effect on the signal is an increase of a similar magnitude. However, due to the persistence of the increased signal in the case of a shock to e , the resulting responses for prices and homeownership rates differ. Figure 3.6 demonstrates these impulse response functions.

The two figures to the left are the impulse response functions for an A shock, which is evident

¹The expected long-term component equals the expected signal as the short-term shock and the innovation to the long-term component are both have zero mean.

by the short-term effect to the price signal. The two figures to the right are the impulse response functions for an e shock, which is evident by the persistent effect to the price signal. The two figures in the upper half contrast the behaviour of prices of the two models for both shocks. The two figures in the lower half show the behaviour of homeownership rates. The blue lines indicate the incomplete information model and the red lines indicate the complete information model. The black dashed line represents the price signal.




Figure 3.6: IRF's of A and e shocks.

For the A shock, we see that there is a small effect on prices that immediately reverses in the complete information model, while for the incomplete information model, there is a similar contemporaneous effect to prices but there is a further increase of prices in the next period. This delayed effect to prices for the incomplete information model is driven by the change in expectations for the long-term component. The increase in the signal leads to an increase in expectations, which in this case is mistaken. This raises the value of continued search and this leads to a further increase in prices despite the short-term shock having already reversed. Thereafter, expectations gradually decline back towards the true level and likewise prices gradually decline to the long-term level. In the homeownership rate figure we see that there is a short-term increase in homeownership rate that quickly declines as expectations for the long-term component gradually correct itself after the shock.

For the e shock, we see that there is a larger long-term effect on prices in the complete information model that dies out slowly due to the persistence of the long-term component. For the incomplete information model, there is a smaller effect on prices at the moment that the shock hits, however after a few more periods prices increase to almost exactly match the level of the complete information model. This is intuitive as households are not immediately sure that the change to the signal is a long-term change hence it takes several periods for households to be convinced of the fact. For homeownership rates, we see interestingly that there is again a short-term increase to homeownership rates that quickly declines in a manner very similar to the A shock. This shows that the effect of a long-term shock is to primarily increase reservation values rather than increasing the rate of matches leading to a transaction. This can be rationalised as the long-term component is very persistent and hence there is no immediate urgency of finding a house to occupy. Conversely if the long-term component was less persistent, then households would have an incentive to find a house to occupy today to “lock in” the high level of the signal today.

3.3.4 Boom and bust experiment

The distinction between the A shock versus the e shock from the perspective of a household in the incomplete information model is entirely reflected in the transitory change of the signal versus the persistent change of the signal respectively. This motivates the following boom and bust experiment which considers sustained shock to A . Figure 3.7 shows the results for prices in both models (incomplete information versus complete information) for a sustained one standard deviation shock to A for five periods between $t = 5$ to $t = 10$.




Figure 3.7: Price response of a sustained increase in A .

The five period increase in the signal leads to drastically different responses in prices in the two models. In the complete information model, prices increase from the moment of the shock begins until the end of the shock by exactly the same magnitude as in the A impulse response function. In the incomplete information model, prices increase by a small magnitude at the moment the shock hits by the same magnitude as in the A impulse response function. However the sustained increase in the signal leads households to continuously update their expectations of the long-term component upwards as they become increasingly sure that this increase in the signal is long-term. At the end of the period, they are suddenly surprised by the reversal of the signal, which leads to the reversal in expectations and hence a reversal in prices. Indeed the response in prices until period 10 before the reversal in prices looks similar to the first five periods of the e impulse response function after the shock hits.

Importantly the dynamics shown in the above boom and bust experiment are driven entirely by expectations in the incomplete information model. This happens endogenously due to the price inference of households, and in this case it misleads households to firstly believe that there was a long-term change leading to a price boom, and later to quickly reverse their beliefs leading to a price bust.

3.4 Empirical application

In this section I structurally deconstruct house prices ex-post and identify episodes of bubbles, where households were overly optimistic about future long-term trends in the housing market. To do this I use my model to analyse historical annual house prices in the US. The mechanism of a bubble is the same as in the boom and bust experiment considered previously. A sequence of positive short-term shocks leads households to believe that there was a long-term change in demand, which turns out to be mistaken when the short-term shocks reverse leading to a reversal in expectations. The broad strategy is to utilise Proposition 3 to extract the price signal from price data, and use this price signal to identify episodes of bubbles.

3.4.1 Deconstruction of price data

The deconstruction procedure can be summarised in the following steps:

- STEP 1. Detrend log prices to obtain stationary log real house prices;
- STEP 2. Map price data to simulated prices by matching quantiles;
- STEP 3. Extract price signal from mapped simulated prices using Proposition 3;
- STEP 4. Obtain long-term component expectations;
- STEP 5. Optimally ex-post extract the long-term component from the price signal;
- STEP 6. Compare expectations with extracted long-term component to identify bubbles/over-optimism.

The results of these summarised steps are demonstrated in Figure [3.8](#) below.




Figure 3.8: Empirical application summarised steps.

The first step of detrending log house price data is necessary to make the price data stationary. Note that house prices in my model are stationary as the long-term component is stationary with $\rho < 1$.

Price data and prices in my model are on different scales, hence it is necessary to map the prior to the latter before I can apply Proposition 3. The mapping is done via a quantile to quantile mapping. I simulate my model for $T = 10000$ periods which gives me an accurate estimate for the statistical distribution of simulated prices. Similarly, my sample of price data gives me 44 data points to form an empirical distribution for actual prices. Each point of the price data is then compared to the whole sample to find its rank, which gives its empirical quantile. This is then mapped to the same quantile in simulated prices. The upper left part of Figure 3.8 shows the quantile plot of the price data versus simulated prices. The upper right half of figure 8 shows this mapping applied to the price data to give the series of *mapped simulated prices*.

The lower left part of Figure 3.8 shows the result of signal extraction from mapped simulated prices by Proposition 3. The algorithm for signal extraction is as follows:

- STEP 1. Begin with long-term component expectations last period χ_{t-1}^e ;
- STEP 2. Conditional on χ_{t-1}^e , invert the mapped simulated price today p_t to find the extracted signal today s_t ;
- STEP 3. Update χ_t^e using χ_{t-1}^e and s_t according to Proposition 4.

While Proposition 3 establishes the existence of a monotonic function between the signal and the price, this function is not known in analytic form. In previous simulations the signal was generated and hence known, and therefore it was sufficient to suppose that this function was known to households and that the price extraction is performed correctly automatically. In this empirical exercise, it is necessary to uncover the monotonic function and use it to perform the signal extraction. To do this, in each period I obtain by simulation the average transaction price at each level of the signal, conditional on χ_{t-1}^e . The relationship between the average price and the signal gives the monotonic function. I use this function to invert the mapped simulated price to the signal this period. Note that the monotonic function changes every period as it depends on χ_{t-1}^e which changes. Furthermore, I choose for simplicity that $\chi_0^e = 0$, although my results are robust to this choice. The signal extraction algorithm also simultaneously gives the path for long-term component expectations, χ_t^e .

Obtaining the long-term component is also different to simulations where I simply observe the path of the long-term component. In this application I need to optimally extract the long-term component in a manner that is analogous to the household signal extraction problem.

Proposition 4 describes how to optimally extract the long-term component with data *up to time* t . However with ex-post data, i.e. hindsight, one is able to better extract the long-term component than households with information up to period t . The signal gives the long-term component plus noise introduced by the short-term component, hence if one observes an increase in the signal in period t that reverses in period $t + 1$, one downwardly adjusts their best guess for the long-term component in period t . This hindsight gives us an advantage to uncovering the long-term component relative to households who have data only up to time t .

Formally, I solve for the optimal extraction of the long-term component via maximum likelihood, although this problem is similarly a special case of the Kalman filtering problem. My data has T periods of the price signal, and I need to choose a path of length T of the long-term component. Hence, my maximisation problem is exactly identified.

The likelihood function for the whole path is given by

$$\mathcal{L} = \prod_{t=1}^T f_A(s_t - \chi_t) f_e(\chi_t - \rho\chi_{t-1}), \quad (3.20)$$

where $f_A(\cdot)$ and $f_e(\cdot)$ are the density functions of A and e respectively, and χ_t is the long-term component in period t .

The log-likelihood function is (up to scaling) given by

$$\log \mathcal{L} = \sum_{t=1}^T \frac{(s_t - \chi_t)^2}{\sigma_A^2} + \frac{(\chi_t - \rho\chi_{t-1})^2}{\sigma_e^2}.$$

First-order conditions for χ_t for t between 1 and $T - 1$ are given by

$$-\sigma_e^2 s_t + \chi_t((1 + \rho^2)\sigma_A^2 + \sigma_e^2) - \rho\sigma_A^2 \chi_{t-1} - \rho\sigma_A^2 \chi_{t+1} = 0.$$

This gives essentially a second-order Euler equation in terms of χ_t .

In addition, the first-order condition for the terminal χ_T is given by

$$-\sigma_e^2 s_T + \chi_T(\sigma_A^2 + \sigma_e^2) - \rho\sigma_A^2 \chi_{T-1} = 0.$$

Solving for χ_t as a function of χ_{t-1} and χ_{t+1} gives

$$\chi_t = \frac{\sigma_e^2 s_T + \rho\sigma_A^2 \chi_{t-1} + \rho\sigma_A^2 \chi_{t+1}}{(1 + \rho^2)\sigma_A^2 + \sigma_e^2}. \quad (3.21)$$

Solving for the optimal χ_T as a function of χ_{T-1} is given by

$$\chi_T = \frac{\sigma_e^2 s_T + \rho \sigma_A^2 \chi_{T-1}}{\sigma_A^2 + \sigma_e^2}. \quad (3.22)$$

Intuitively, the optimal χ_t is a weighted average of s_t , χ_{t-1} and χ_{t+1} . This reflects the optimal maximum likelihood trade-off that desires χ_t to be close to s_t to minimise A_t , to be close to χ_{t-1} to minimise e_t , and to be close to χ_{t+1} to minimise e_{t+1} . Similarly, χ_T is a weighted average of s_T and χ_{T-1} .

It is non-trivial to analytically solve for the optimal path using the above equations as the optimal χ_t in each period depends on χ_{t-1} and χ_{t+1} which are not known. Instead I use the following algorithm to iteratively solve for the optimal path:

- STEP 1. Begin with a path for $\{\chi_t\}_{t=1,\dots,T}$;
- STEP 2. Update χ_t using equation (21) and substituting for χ_{t-1} and χ_{t+1} using the previous path;
- STEP 3. Update χ_T using equation (22) and substituting for χ_{T-1} using the previous path;
- STEP 4. Check for convergence $\{\chi_t\}_{t=1,\dots,T}$ and continue if convergence is not achieved.

The lower right part of Figure 3.8 shows the results of solving for the extracted long-term component using the algorithm described above. Specifically, the extracted long-term component is more smooth than contemporaneous expectations of the long-term component. This is due to the benefit of hindsight; the extracted long-term component does not respond as sharply to short-term changes to the signal that are later reversed.

3.4.2 Identifying bubble episodes

Bubbles, when defined as a period of irrational price movements, do not exist in my framework. Nonetheless, I can uncover periods where households were overly (but not irrationally) optimistic about the long-term component. This is precisely the mechanism discussed in the boom and bust experiment considered in the earlier section.

Based on the boom and bust experiment example, I define a bubble to be *a period of continued long-term component expectations increase that exceeds the extracted long-term component and at its peak exceeds by at least one standard deviation*.




Figure 3.9: Bubble episodes between 1975-2018.

In Figure 3.9 I show the extracted long-term component and long-term component expectations. The episodes of bubbles are highlighted and they correspond to 1976-1979, 1986-1987, and 2002-2006. The timing of these episodes map broadly to the housing bubbles documented in the U.S. Specifically, the housing bubble in the early 2000's stand out as being longer in duration, a total of five years, as well as being larger in magnitude. If we consider short-term shocks to the housing market as being driven primarily by the business cycle, this period maps also to the period of the Great Moderation.

3.5 Model extensions

In this section I consider three extensions of my existing model. They enrich the model in various ways but the basic mechanisms and insights from the basic model are preserved.

3.5.1 Intraperiod expectation updating

In the model described earlier, households update their expectations only at the end of the period when they observe the price signal. In this extension I allow for each buyer and seller to update their expectations based on the realisation of ϵ , as ϵ is the sum of the price signal plus some idiosyncratic noise. Allowing for this means that households during the transaction period

may have heterogeneous expectations due to receiving different realisations of the idiosyncratic shock. However, at the end of period the price signal is revealed and as this is more informative than ϵ , the distribution of expectations collapses and every household will again be in agreement.

The joint density of the three shocks is as follows

$$f(A, e, \eta) = \frac{1}{2\pi\sigma_A\sigma_e} \exp(-A^2/2\sigma_A - e^2/2\sigma_e) \frac{1}{2M_\eta}. \quad (3.23)$$

Due to the uniform distribution being flat, the maximum likelihood values for A , e and η , given a value for ϵ , is always at a corner solution. This is unlike the previous example of updating subject to the price signal at the end of the period where the optimum is always at an interior solution. Specifically, in maximising the above density subject to ϵ taking a particular value, it is always desirable to maximise the flexibility of η subject to the constraint of being within its domain. In particular, whenever

$$|\epsilon - \rho\chi_{t-1}^e| \leq M_\eta \quad (3.24)$$

the maximum likelihood values for A and e are both zero, and the maximum likelihood value for η is $\epsilon - \rho\chi_{t-1}^e$. This is evident by considering the marginal densities of A, e and η , which are jointly maximised at $(0, 0, \epsilon - \rho\chi_{t-1}^e)$.

That indicates that the optimal behaviour is for households to be conservative and take the view that a particular realisation of ϵ is due entirely to the idiosyncratic component rather than a shift in the fundamentals.




Figure 3.10: Intraproduct updating scheme.

For ϵ outside of the above range, expectations are updated in the following way in order to maximise the likelihood function

$$\chi_t^e = \begin{cases} \rho\chi_{t-1}^e + (\frac{\sigma_e^2}{\sigma_A^2 + \sigma_e^2})(\epsilon - \rho\chi_{t-1}^e - M_\eta) & \text{for } \epsilon > \rho\chi_{t-1}^e + M_\eta \\ \rho\chi_{t-1}^e - (\frac{\sigma_e^2}{\sigma_A^2 + \sigma_e^2})(\rho\chi_{t-1}^e - M_\eta - \epsilon) & \text{for } \epsilon < \rho\chi_{t-1}^e - M_\eta \end{cases}. \quad (3.25)$$

This is evident by again considering the marginal densities. Without loss of generality, consider

$\epsilon_t > \rho\chi_{t-1}^e + M_\eta$, then the solution for η_t is at the corner of M_η as this minimises the sum of A_t and e_t , which maximises the overall likelihood. Conditional on $A_t + e_t = \epsilon_t - \rho\chi_{t-1}^e - M_\eta$, the first-order condition implies the following maximum likelihood e_t given by

$$e_t^{ML} = \left(\frac{\sigma_e^2}{\sigma_A^2 + \sigma_e^2} \right) (\epsilon - \rho\chi_{t-1}^e - M_\eta), \quad (3.26)$$

which establishes the above result for the case of $\epsilon_t > \rho\chi_{t-1}^e + M_\eta$.

The implications of this is that a large proportion of households do not update their expectations during the transaction period, and for those that do they update it only by a small margin, see Figure 3.10 for a visual representation. Nonetheless, the overall effect of this is to generate some small adjustment of expectations on average towards the signal in the intraperiod, which slightly reduce the lag at which households update their expectations. After the transaction period, households observe the price signal and the distribution of expectations collapses. Otherwise the model follows the same steps.

3.5.2 Simple production economy with mortgage contracts

In the stylised model I present immediately the value functions after describing the shocks to the economy. In this section I describe a simple production economy which generates the same Bellman equations as a result of optimal households behaviour. Households supply labour and earn a wage. Households have a per period utility function of the following form

$$u(\epsilon, c, h) = \begin{cases} c - f(h) + \epsilon & \text{if matched to a house} \\ c - f(h) + u_{rent} & \text{if unmatched and renting} \end{cases}, \quad (3.27)$$

where $f(\cdot)$ is a convex twice continuously differentiable function, c is consumption, h is the quantity of labour supplied, and ϵ is the quality of the match between the household and the house it is living in. Furthermore I assume that there exists a h^* such that $f'(h^*) = 1$. The fact that utility is quasilinear with respect to consumption is important for the simple form of the household Bellman equations.




Figure 3.11: Timing

Similarly to Lagos and Wright (2005), I formally characterise the sequence of events within each period to a day sub-period and a night sub-period, see Figure 3.11. During the day, households meet and transact in housing in exactly the same manner as previously described. Again Nash bargaining occurs over the price of the house, however the unit of denomination is now in units of consumption, and this is to be repaid through a bilateral mortgage contract that shall be described next. During the night, the average transaction price is revealed and households update their priors in the same manner. Furthermore, households supply labour and consume the consumption good by supplying labour. I assume that labour is exchanged one-to-one to the consumption good, i.e. the marginal product of labour is one. Intraproduct optimality between labour and consumption is summarised by

$$h_{i,t} = h^*, \quad (3.28)$$

which also gives the labour income in each period.

Supposing that a match between a buyer and seller generates a transaction of price p , the households engage in a bilateral mortgage contracts which repay the sum of p in present value over a fixed term of N periods. Specifically, the equal quantity to be repaid in each period is given by

$$m = \frac{p(1 - \beta)}{1 - \beta^N}, \quad (3.29)$$

which is strictly greater than p/N due to accounting for discounting by β . Consequently $\beta^{-1} - 1$ is the implicit mortgage interest rate.

In summary, households trade in future labour income with the seller receiving a stream of future labour income from the buyer. The transaction leads to an expansion in utility for the seller by

$$\sum_{i=0}^{N-1} m\beta^i = p, \quad (3.30)$$

and similarly a reduction in the buyer's utility by the same amount, due to the quasilinearity of utility with respect to consumption. The Bellman equation for a buyer after observing the

realisation of the quality draw is therefore given by

$$V_B(\Omega, \epsilon) = h^* + \max \left\{ u_{rent} - R - c_{search} + \beta[V_B(\Omega', \epsilon')], \epsilon - p(\Omega, \epsilon) + \beta[\phi(V_B(\Omega', \epsilon') + V_S(\Omega', \epsilon')) + (1 - \phi)V_M(\Omega', \epsilon)] \right\}, \quad (3.31)$$

where again Ω is the set of observable aggregate states. The Bellman equation for a seller is given by

$$V_S(\Omega, \epsilon) = h^* + \max \left\{ R - c_{post} + \beta[V_S(\Omega', \epsilon')], p(\Omega, \epsilon) \right\}. \quad (3.32)$$

Finally the Bellman equation for a matched household is given by

$$V_M(\Omega, \epsilon) = h^* \epsilon + \beta[\phi(V_B(\Omega', \epsilon') + V_S(\Omega', \epsilon')) + (1 - \phi)V_M(\Omega', \epsilon)]. \quad (3.33)$$

Normalising the above three Bellman equations by subtracting h^* obtains the original Bellman equations of the original model.

3.5.3 Short-term component persistence

The existing model considers the short-term component as noise and the long-term component as a persistent $AR(1)$ process. This assumption simplifies the analysis, however it is possible to consider the short-term component to be also persistent. Importantly, I maintain that the persistence of the short-term component is strictly less than that of the long-term component. Recall that the long-term component follows the following process

$$\chi_t = \rho\chi_{t-1} + e_t,$$

where $\rho \in [0, 1)$ is the $AR(1)$ coefficient and e_t is normal with mean 0 and variance σ_e^2 . Next, I consider the short-term component to follow

$$A_t = \theta A_{t-1} + u_t, \quad (3.34)$$

where $\theta \in [0, \rho)$ is the $AR(1)$ coefficient and u_t is normal with mean 0 and variance σ_u^2 . This change necessitates a change in the state space for the households to $\Omega_t = \{\chi_{t-1}^e, A_{t-1}^e\}$ as the short-term component is also of relevance for future periods. As a consequence, households now formulate expectations by

$$\chi_t^e = \rho\chi_{t-1}^e + \left(\frac{\sigma_e^2}{\sigma_u^2 + \sigma_e^2} \right) (s_t - \rho\chi_{t-1}^e - \theta A_{t-1}^e). \quad (3.35)$$

$$A_t^e = \theta A_{t-1}^e + \left(\frac{\sigma_u^2}{\sigma_u^2 + \sigma_e^2} \right) (s_t - \rho\chi_{t-1}^e - \theta A_{t-1}^e). \quad (3.36)$$

Similarly to Proposition 3, household expectations can be written as the history of price signals

$$\begin{pmatrix} \chi_t^e \\ A_t^e \end{pmatrix} = \mathbf{M}^t \begin{pmatrix} \chi_0^e \\ A_0^e \end{pmatrix} + \sum_{j=0}^{t-1} \mathbf{M}^j \begin{pmatrix} q \\ 1-q \end{pmatrix} s_{t-j}, \quad (3.37)$$

where $q = \frac{\sigma_e^2}{\sigma_u^2 + \sigma_e^2}$ and $\mathbf{M} = \begin{pmatrix} \rho(1-q) & -\theta q \\ -\rho(1-q) & \theta q \end{pmatrix}$.

3.6 Conclusion

This paper demonstrates that short-term momentum and long-term reversal in house prices, well documented in the empirical literature, can be readily explained in a rational expectations framework with incomplete information. The key signal extraction problem for the households to solve is how to optimally interpret the ambiguous signal provided by prices, which results in conservative adjustment in expectations in the same direction as prices and endogenously provides a rationale for backward-looking expectations over past prices. The resulting behaviour in prices is broadly similar to the complete information model benchmark, however it differs in the short/medium term fluctuations around the trend. This generates the momentum and reversal mentioned earlier and also generates endogenous boom and bust episodes. The application of my model to US house prices offers a new perspective on how to structurally identify bubbles ex-post.

References

Daron Acemoglu, David Autor, David Dorn, Gordon H Hanson, and Brendan Price. Import competition and the great us employment sag of the 2000s. *Journal of Labor Economics*, 34(S1):S141–S198, 2016.

Manuel Adelino, Antoinette Schoar, and Felipe Severino. House prices, collateral, and self-employment. *Journal of Financial Economics*, 117(2):288–306, 2015.

Sumit Agarwal, Gene Amromin, Souphala Chomsisengphet, Tim Landvoigt, Tomasz Piskorski, Amit Seru, and Vincent Yao. Mortgage refinancing, consumer spending, and competition: Evidence from the home affordable refinancing program. Technical report, National Bureau of Economic Research, 2017.

Aditya Aladangady. Housing wealth and consumption: Evidence from geographically-linked microdata. *American Economic Review*, 107(11):3415–46, 2017.

Luis Armona, Andreas Fuster, and Basit Zafar. Home price expectations and behaviour: Evidence from a randomized information experiment. *The Review of Economic Studies*, 86(4):1371–1410, 2018.

Pierfederico Asdrubali, Bent E Sørensen, and Oved Yosha. Channels of interstate risk sharing: United states 1963–1990. *The Quarterly Journal of Economics*, 111(4):1081–1110, 1996.

Orazio Attanasio, Andrew Leicester, and Matthew Wakefield. Do house prices drive consumption growth? the coincident cycles of house prices and consumption in the uk. *Journal of the European Economic Association*, 9(3):399–435, 2011.

Orazio P Attanasio, Laura Blow, Robert Hamilton, and Andrew Leicester. Booms and busts: Consumption, house prices and expectations. *Economica*, 76(301):20–50, 2009.

Michael Bailey, Ruiqing Cao, Theresa Kuchler, and Johannes Stroebel. The economic effects of social networks: Evidence from the housing market. *Journal of Political Economy*, 126(6):2224–2276, 2018.

- Nicholas Barberis, Robin Greenwood, Lawrence Jin, and Andrei Shleifer. Extrapolation and bubbles. *Journal of Financial Economics*, 129(2):203–227, 2018.
- David Berger, Veronica Guerrieri, Guido Lorenzoni, and Joseph Vavra. House prices and consumer spending. *The Review of Economic Studies*, 85(3):1502–1542, 2017.
- Neil Bhutta and Benjamin J Keys. Interest rates and equity extraction during the housing boom. *American Economic Review*, 106(7):1742–74, 2016.
- Emine Boz and Enrique G Mendoza. Financial innovation, the discovery of risk, and the us credit crisis. *Journal of Monetary Economics*, 62:1–22, 2014.
- Craig Burnside, Martin Eichenbaum, and Sergio Rebelo. Understanding booms and busts in housing markets. *Journal of Political Economy*, 124(4):1088–1147, 2016.
- Charles W Calomiris, Stanley D Longhofer, and William Miles. The housing wealth effect: The crucial roles of demographics, wealth distribution and wealth shares. Technical report, National Bureau of Economic Research, 2012.
- John Y Campbell and Joao F Cocco. Household risk management and optimal mortgage choice. *The Quarterly Journal of Economics*, 118(4):1449–1494, 2003.
- John Y Campbell and Joao F Cocco. How do house prices affect consumption? evidence from micro data. *Journal of monetary Economics*, 54(3):591–621, 2007.
- John Y Campbell and Joao F Cocco. A model of mortgage default. *The Journal of Finance*, 70(4):1495–1554, 2015.
- Christopher D Carroll, Misuzu Otsuka, and Jiri Slacalek. How large are housing and financial wealth effects? a new approach. *Journal of Money, Credit and Banking*, 43(1):55–79, 2011.
- Karl E Case, John M Quigley, and Robert J Shiller. Comparing wealth effects: the stock market versus the housing market. *Advances in macroeconomics*, 5(1), 2005.
- Karl E Case and Robert J Shiller. Forecasting prices and excess returns in the housing market. *Real Estate Economics*, 18(3):253–273, 1990.
- Karl E Case, Robert J Shiller, and Anne Thompson. What have they been thinking? home buyer behavior in hot and cold markets. Technical report, National Bureau of Economic Research, 2012.
- KE Case, JM Quigley, and RJ Shiller. Wealth effects revisited: 1975-2012 (no. w18667), 2013.

- Matthew Chambers, Carlos Garriga, and Don E Schlagenhauf. Accounting for changes in the homeownership rate. *International Economic Review*, 50(3):677–726, 2009.
- Thomas Chaney, David Sraer, and David Thesmar. The collateral channel: How real estate shocks affect corporate investment. *American Economic Review*, 102(6):2381–2409, 2012.
- Hui Chen, Michael Michaux, and Nikolai Roussanov. Houses as atms? mortgage refinancing and macroeconomic uncertainty. Technical report, National Bureau of Economic Research, 2013.
- James Cloyne, Kilian Huber, Ethan Ilzetzki, and Henrik Kleven. The effect of house prices on household borrowing: a new approach. *American Economic Review*, 109(6):2104–36, 2019.
- Daniel Cooper. House price fluctuations: the role of housing wealth as borrowing collateral. *Review of Economics and Statistics*, 95(4):1183–1197, 2013.
- Thomas Davidoff et al. Supply constraints are not valid instrumental variables for home prices because they are correlated with many demand factors. *Critical Finance Review*, 5(2):177–206, 2016.
- Morris A Davis and Stijn Van Nieuwerburgh. Housing, finance, and the macroeconomy. In *Handbook of regional and urban economics*, volume 5, pages 753–811. Elsevier, 2015.
- Anthony A DeFusco. Homeowner borrowing and housing collateral: New evidence from expiring price controls. *The Journal of Finance*, 73(2):523–573, 2018.
- Yongheng Deng, John M Quigley, and Robert Van Order. Mortgage terminations, heterogeneity and the exercise of mortgage options. *Econometrica*, 68(2):275–307, 2000.
- Peter Englund and Yannis M Ioannides. House price dynamics: an international empirical perspective. *Journal of Housing Economics*, 6(2):119–136, 1997.
- Tom Engsted, Simon J Hviid, and Thomas Q Pedersen. Explosive bubbles in house prices? evidence from the oecd countries. *Journal of International Financial Markets, Institutions and Money*, 40:14–25, 2016.
- Giovanni Favara and Jean Imbs. Credit supply and the price of housing. *American Economic Review*, 105(3):958–92, 2015.
- Jack Favilukis, Sydney C Ludvigson, and Stijn Van Nieuwerburgh. The macroeconomic effects of housing wealth, housing finance, and limited risk sharing in general equilibrium. *Journal of Political Economy*, 125(1):140–223, 2017.

- Marjorie Flavin and Shinobu Nakagawa. A model of housing in the presence of adjustment costs: A structural interpretation of habit persistence. *American Economic Review*, 98(1):474–95, 2008.
- Marjorie Flavin and Takashi Yamashita. Owner-occupied housing and the composition of the household portfolio. *American Economic Review*, 92(1):345–362, 2002.
- Christopher L Foote, Kristopher S Gerardi, and Paul S Willen. Why did so many people make so many ex post bad decisions? the causes of the foreclosure crisis. Technical report, National Bureau of Economic Research, 2012.
- Peter Ganong and Pascal Noel. The effect of debt on default and consumption: Evidence from housing policy in the great recession. *Unpublished Working Paper*, 2017.
- Thomas A Garrett, Rubén Hernández-Murillo, Michael T Owyang, et al. Does consumer sentiment predict regional consumption. *Federal reserve Bank of st. Louis review*, 87(2):123–135, 2005.
- Paolo Gelain, Kevin J Lansing, and Gisle J Natvik. Explaining the boom–bust cycle in the us housing market: A reverse-engineering approach. *Journal of Money, Credit and Banking*, 50(8):1751–1783, 2018.
- Pedro Gete. Expectations and the housing boom and bust. an open economy view. *An Open Economy View (December 2018)*, 2018.
- Simon Gilchrist and Egon Zakrajšek. Credit spreads and business cycle fluctuations. *American Economic Review*, 102(4):1692–1720, 2012.
- Edward L Glaeser and Joseph Gyourko. Urban decline and durable housing. *Journal of political economy*, 113(2):345–375, 2005.
- Edward L Glaeser and Charles G Nathanson. Housing bubbles. In *Handbook of regional and urban economics*, volume 5, pages 701–751. Elsevier, 2015.
- Edward L Glaeser and Charles G Nathanson. An extrapolative model of house price dynamics. *Journal of Financial Economics*, 126(1):147–170, 2017.
- Denis Goria and Virgiliu Midrigan. Liquidity constraints in the us housing market. Technical report, National Bureau of Economic Research, 2017.
- Luca Guerrieri and Matteo Iacoviello. Collateral constraints and macroeconomic asymmetries. *Journal of Monetary Economics*, 90:28–49, 2017.

Veronica Guerrieri and Guido Lorenzoni. Credit crises, precautionary savings, and the liquidity trap. *The Quarterly Journal of Economics*, 132(3):1427–1467, 2017.

Nezih Guner, Remzi Kaygusuz, and Gustavo Ventura. Income taxation of us households: Facts and parametric estimates. *Review of Economic Dynamics*, 17(4):559–581, 2014.

Adam M Guren, Arvind Krishnamurthy, and Timothy J McQuade. Mortgage design in an equilibrium model of the housing market. Technical report, National Bureau of Economic Research, 2018.

Adam M Guren, Alisdair McKay, Emi Nakamura, and Jón Steinsson. Housing wealth effects: The long view. Technical report, National Bureau of Economic Research, 2018.

Adam M Guren, Alisdair McKay, Emi Nakamura, and Jón Steinsson. What do we learn from cross-sectional empirical estimates in macroeconomics?, 2018.

Fatih Guvenen, Fatih Karahan, Serdar Ozkan, and Jae Song. What do data on millions of us workers reveal about life-cycle earnings risk? Technical report, National Bureau of Economic Research, 2016.

Fatih Guvenen and Anthony A Smith. Inferring labor income risk and partial insurance from economic choices. *Econometrica*, 82(6):2085–2129, 2014.

Jonathan Heathcote, Fabrizio Perri, and Giovanni L Violante. Unequal we stand: An empirical analysis of economic inequality in the united states, 1967–2006. *Review of Economic dynamics*, 13(1):15–51, 2010.

Charles Himmelberg, Christopher Mayer, and Todd Sinai. Assessing high house prices: Bubbles, fundamentals and misperceptions. *Journal of Economic Perspectives*, 19(4):67–92, 2005.

Yang Hu and Les Oxley. Bubbles in us regional house prices: evidence from house price–income ratios at the state level. *Applied Economics*, 50(29):3196–3229, 2018.

Erik Hurst and Frank Stafford. Home is where the equity is: Mortgage refinancing and household consumption. *Journal of Money, credit and Banking*, pages 985–1014, 2004.

Min Hwang and John M Quigley. Economic fundamentals in local housing markets: evidence from us metropolitan regions. *Journal of regional science*, 46(3):425–453, 2006.

G Kaplan. Violante gl weidner j.(2014),the wealthy hand-to-mouth,. *Brookings Papers On Economic Activity*.

Greg Kaplan, Kurt Mitman, and Gianluca Violante. Consumption and house prices in the great recession: Model meets evidence. *Manuscript, New York University*, 2015.

- Greg Kaplan, Kurt Mitman, and Giovanni L Violante. Non-durable consumption and housing net worth in the great recession: Evidence from easily accessible data. Technical report, National Bureau of Economic Research, 2016.
- Greg Kaplan, Kurt Mitman, and Giovanni L Violante. The housing boom and bust: Model meets evidence. Technical report, National Bureau of Economic Research, 2017.
- Theresa Kuchler and Basit Zafar. Personal experiences and expectations about aggregate outcomes. *The Journal of Finance*, 2015.
- M Kuhn, M Schularick, and UI Steins. The great american debt boom, 1949-2013. Technical report, Discussion paper, University of Bonn, 2017.
- Ricardo Lagos and Randall Wright. A unified framework for monetary theory and policy analysis. *Journal of political Economy*, 113(3):463–484, 2005.
- Tim Landvoigt, Monika Piazzesi, and Martin Schneider. The housing market (s) of san diego. *American Economic Review*, 105(4):1371–1407, 2015.
- Wenli Li and Rui Yao. The life-cycle effects of house price changes. *Journal of Money, Credit and banking*, 39(6):1375–1409, 2007.
- Carl Liebersohn. Housing demand, regional house prices and consumption. *Regional House Prices and Consumption (April 24, 2017)*, 2017.
- Kian-Ping Lim and Robert Brooks. The evolution of stock market efficiency over time: A survey of the empirical literature. *Journal of Economic Surveys*, 25(1):69–108, 2011.
- Lars Ljungqvist and Thomas J Sargent. *Recursive macroeconomic theory*. MIT press, 2018.
- Burton G Malkiel and Eugene F Fama. Efficient capital markets: A review of theory and empirical work. *The journal of Finance*, 25(2):383–417, 1970.
- Sebastian Mallaby. *The Man Who Knew: The Life & Times of Alan Greenspan*. Bloomsbury Publishing, 2016.
- Enrique Martínez-García and Valerie Grossman. Explosive dynamics in house prices? an exploration of financial market spillovers in housing markets around the world. 2018.
- Atif Mian, Kamalesh Rao, and Amir Sufi. Household balance sheets, consumption, and the economic slump. *The Quarterly Journal of Economics*, 128(4):1687–1726, 2013.
- Atif Mian and Amir Sufi. House prices, home equity-based borrowing, and the us household leverage crisis. *American Economic Review*, 101(5):2132–56, 2011.

Atif Mian and Amir Sufi. What explains the 2007–2009 drop in employment? *Econometrica*, 82(6):2197–2223, 2014.

Atif Mian and Amir Sufi. *House of debt: How they (and you) caused the Great Recession, and how we can prevent it from happening again*. University of Chicago Press, 2015.

Hazel A Morrow-Jones and Mary V Wenning. The housing ladder, the housing life-cycle and the housing life-course: Upward and downward movement among repeat home-buyers in a us metropolitan housing market. *Urban Studies*, 42(10):1739–1754, 2005.

Emi Nakamura and Jon Steinsson. Fiscal stimulus in a monetary union: Evidence from us regions. *American Economic Review*, 104(3):753–92, 2014.

Charles G Nathanson and Eric Zwick. Arrested development: Theory and evidence of supply-side speculation in the housing market. *The Journal of Finance*, 73(6):2587–2633, 2018.

L Rachel Ngai and Silvana Tenreyro. Hot and cold seasons in the housing market. *American Economic Review*, 104(12):3991–4026, 2014.

Rachel Ong, Sharon Parkinson, Beverley A Searle, Susan J Smith, and Gavin A Wood. Channels from housing wealth to consumption. *Housing Studies*, 28(7):1012–1036, 2013.

Christopher Palmer. Why did so many subprime borrowers default during the crisis: Loose credit or plummeting prices? *Available at SSRN 2665762*, 2015.

Efthymios Pavlidis, Ivan Paya, David Alan Peel, and Alisa Yevgenyevna Yusupova. Exuberance in the uk regional housing markets. 2017.

Peter CB Phillips, Shuping Shi, and Jun Yu. Testing for multiple bubbles: Historical episodes of exuberance and collapse in the s&p 500. *International economic review*, 56(4):1043–1078, 2015.

Monika Piazzesi and Martin Schneider. Housing and macroeconomics. In *Handbook of macroeconomics*, volume 2, pages 1547–1640. Elsevier, 2016.

Monika Piazzesi, Martin Schneider, and Selale Tuzel. Housing, consumption and asset pricing. *Journal of Financial Economics*, 83(3):531–569, 2007.

James D Rodgers and Judy A Temple. Sales taxes, income taxes, and other non-property tax revenues. *Management Policies in Local Government Finance, Municipal Management Series, Fourth Edition*. Washington, DC: International City/County Management Association for the ICMA University, pages 229–59, 1996.

- Albert Saiz. The geographic determinants of housing supply. *The Quarterly Journal of Economics*, 125(3):1253–1296, 2010.
- Shuping Shi. Speculative bubbles or market fundamentals? an investigation of us regional housing markets. *Economic Modelling*, 66:101–111, 2017.
- Robert J Shiller. Understanding recent trends in house prices and home ownership. Technical report, National Bureau of Economic Research, 2007.
- Todd Sinai and Nicholas S Souleles. Owner-occupied housing as a hedge against rent risk. *The Quarterly Journal of Economics*, 120(2):763–789, 2005.
- Sheridan Titman, Ko Wang, and Jing Yang. The dynamics of housing prices. *Journal of Real Estate Research*, 36(3):283–317, 2014.
- Arlene Wong. Population aging and the transmission of monetary policy to consumption and population aging. 2018.
- Jing Wu, Joseph Gyourko, and Yongheng Deng. Evaluating conditions in major chinese housing markets. *Regional Science and Urban Economics*, 42(3):531–543, 2012.
- Xia Zhou. *Essays on US state-level financial wealth data and consumption data*. PhD thesis, jhu, 2010.