Essays on Macroeconomics and Household Heterogeneity

Isaac Gross

Keble College

A thesis submitted for the degree of

Doctor of Philosophy in Economics

Trinity Term 2018

Acknowledgments

I would like to thank my partner Monique for all her help and support over the past five years. I am also incredibly grateful to my friends and family back in Australia.

This work would not have been possible without the support of my supervisor Associate Professor Andrea Ferrero, and the other members of the department I have talked to about this project. While their help and teaching has been invaluable, any errors that remain in this thesis are solely my own.

Finally, I am grateful to the Commonwealth Scholarship Commission for funding my studies over the past three years.

Zac Gross

Supervisor: Professor Andrea Ferrero.

Word count: Approximately 65,000 words, taking p.10 as a typical page.

Statement of authorship: All chapters of this thesis are single-authored and are fully my own work.

Statement of funding: I acknowledge the generous financial support of the Commonwealth Scholarship Commission for funding across the duration of my DPhil, from 2015-2018.

Contents

List of Figures

List of Tables

Introduction

The goal of this thesis is to explore how household heterogeneity propagates and amplifies macroeconomic shocks within the economy using both economic theory and empirical data.

The assumption of a single "representative" household has been a mainstay of macroeconomic research over the past half-century. However recent work suggests that not only is there a considerable degree of heterogeneity among households, but that these differences have a significant impact on a range of macroeconomic issues such as the effectiveness of fiscal stimulus (Kaplan et al., 2014; Broda and Parker, 2014), monetary policy (Auclert, 2017; Kaplan et al., 2016), the housing market (Attanasio et al., 2012; Blundell et al., 2008; Guerrieri and Iacoviello, 2017; Ngai et al., 2016; Mian et al., 2013), consumption (Ahn et al., 2017a; Blundell and Preston, 1998; Campbell and Cocco, 2007; Engelhardt, 1996) and employment (Ravn and Sterk, 2016; McKay and Reis, 2016; Abo-Zaid, 2013a) among many others. This literature has highlighted how households respond differently to aggregate shocks or changes in policy and how simply aggregating or averaging across them can obscure important truths about the economy.

However, relaxing this assumption poses several challenges. The first is choosing the degree and manner in which households differ. While in reality households can differ along many dimensions, in practice it is only feasible to include a small number of these in any given model. Thus one must choose the most salient dimensions along which households differ and the structural reasons behind such differences. For example, when examining the dynamics behind the housing market is it important to model differences in income, wealth, age, tastes or composition? No single model will be able to incorporate all these differences and so it is incumbent on researchers to proritise and justify their choices. In this thesis I will show why household heterogeneity in the housing and labour markets is both empirically relevant and an important consideration when considering the problem of optimal policy.

The second challenge is a computational one. While models can be structured such that differentiated households make identical decisions, in general these differences will cause choices, and thus outcomes, across households to diverge. This produces a non-degenerate distribution of households across their specific state variables. This raises the problem of how this potentially infinite-dimension distribution is incorporated within the model. Previous literature has developed a range of options for handling this problem including approximating the distribution with a small handful of moments (Krusell and Smith, 1998) and approximating it with projection and perturbation methods (Reiter, 2009). In this thesis I will outline two different methods for dealing with this computational problem. The first, set out in Chapter 1, shows how market clearing prices can be feasibly calculated by aggregating over the distribution of households. The second approach involves simulating the model with aggregate uncertainty using numerical derivatives based on impulse response functions.

The first chapter of this thesis will examine how heterogeneity in wealth and income affects households' decision to purchase housing and the implications for their consumption of non-durable goods. It constructs an Aiyagari-Bewley-Huggett model in which households are subject to an idiosyncratic income shock and thus hold different amounts of liquid wealth and illiquid housing. I then evaluate how the anticipated changes in household debt associated with the leveraged purchase of housing affect the consumption of non-durable goods. I show that the differences in income and wealth lead to significant variance in marginal propensities to consume among households. I show that households that are saving for a house deposit can have negative marginal propensities to consume as they lower their consumption in anticipation of being credit constrained as the probability that they will buy a house increases. This result has important implications for the design of fiscal policy, as it shows that payments to first time home buyers, which was a common policy response to the Global Financial Crisis, can lead to falls in aggregate consumption rather than stimulating growth.

The second and third chapters examine how the combination of heterogeneity in workers' wages and downward nominal wage rigidity affects the transmission and design of different aspects of monetary policy. In Chapter 2 I show that in this environment there is a trade-off between a higher rate of inflation which gives workers more flexibility when setting real wages, at the cost of greater price dispersion in the goods market. After outlining a numerical algorithm to solve the model I use micro-data on the distribution of workers' change in wages to calibrate the nominal wage rigidity. I show that downward nominal wage rigidities bend the Phillips curve constraining the inflation rate from falling in times of low demand. This indicates that an inflation rate that is only moderately below its target can mask large falls in the output gap. Finally, I find that the monetary policy rule can be implemented by placing a higher weight on wage inflation, relative to a symmetric nominal wage rigidity. In Chapter 3 I discuss how downwardly rigid wages can amplify or mitigate the welfare loss caused by the zero lower bound on nominal interest rates and how this varies with the parameterisation of the model. I find that the optimal rate of inflation is increased by the presence of both nominal interest rate and wage rigidities, when modeled either separately or in tandem, and is 3 per cent in the baseline calibration of the model.

1 Anticipated Changes in Household Debt and Consumption[1]

Abstract

This chapter evaluates how anticipated changes in debt associated with the leveraged purchase of housing affect consumption. I build an Aiyagari-Bewley-Huggett model in which households save in liquid assets and illiquid housing, where the latter can be used as collateral for borrowing. I show that the model is able to replicate the empirical distributions of income and wealth within the US economy. In the model there is substantial heterogeneity in households' marginal propensities to consume. Households with a high probability of buying housing stock lower their consumption of non-durable goods in anticipation of being credit constrained after their purchase. This results in households with low, even negative, marginal propensities to consume. I verify the model's predictions using micro-data from the Panel Study of Income Dynamics to show that (i) consumption falls in anticipation of, and after, changes in the stock of housing and (ii) households who are planning on purchasing housing have negative marginal propensities to consume. Finally, I use this model to examine the general equilibrium effects of tax credits for first home buyers and show that they lead to decreases in aggregate consumption.

JEL Classification: D14, E21, E62.

Keywords: Household Consumption, Heterogeneous Households, Housing Demand, Leverage.

1 I would like to thank seminar participants at the University of Oxford, and the Workshop on Dynamic Macroeconomics for useful comments. I would like to acknowledge the use of the University of Oxford Advanced Research Computing (ARC) facility in carrying out this work (http://dx.doi.org/10.5281/zenodo.22558).

1.1 Introduction

The recent financial crisis has highlighted the impact that household debt has on consumption. A burgeoning literature has documented how higher levels of household leverage were associated with deeper falls in consumption. Debt overhang has held back consumer spending and the subsequent recovery.[2]

In contrast to the unexpected deterioration of balance sheets during the crisis, most variation in household debt revolves around the decision to purchase a house, which is generally anticipated years in advance. However if transactions in the housing market are anticipated households may be able to smooth their consumption of non-durable goods even if the changes to the balance sheet are quantitatively larger than unexpected cyclical shocks. In this chapter, I show that despite being anticipated these increases in debt lead to a decrease in consumption of non-durable goods. I further show that counter-cyclical policies aimed at increasing home ownership can reduce overall consumption.

In this chapter I make three contributions. First, I extend the multiple asset Aiyagari-Bewley-Huggett model to include a market for housing. My model builds on the two-asset framework from Kaplan, Moll, and Violante (2016), by allowing households to use illiquid assets as collateral for borrowing. I show that households, anticipating that they will be credit constrained following the leveraged purchase of housing, reduce their consumption of non-durable goods in advance of, and after, their purchase.[3] This anticipatory saving causes households which have a high likelihood of purchasing housing to have negative marginal propensities to consume.[4] Second, I validate these results using micro-data from the Panel Study of Income Dynamics (PSID) to show that households who expect to adjust their stock of housing in the

2 Mian, Rao, and Sufi (2013) and Mian and Sufi (2014) highlight how heterogeneity in households' balance sheets affected consumption dynamics during the recent recession. Carroll and Dunn (1997), Campbell and Cocco (2007), and Attanasio et al. (2012) also examine the relationship between the two using rich general equilibrium models. Finally, the empirical link between marginal propensities to consume and household leverage is explored by Broda and Parker (2014), Fagereng et al. (2016) and Kaplan et al. (2014).

3 By contrast, in the standard two-asset model without leverage, the decision to buy illiquid assets triggers an increase in non-durable consumption (Kaplan and Violante, 2014).

4 This anticipatory saving effect is distinct from precautionary saving, which occurs in response to uncertainty with respect to future income. The anticipatory saving effect exists in the absence of either aggregate or idiosyncratic uncertainty.

next two years have lower, even negative, marginal propensities to consume. Finally, I use my model to estimate the impact of tax credits for first home buyers' consumption of non-durable goods. I show that the tax credits cause aggregate consumption to fall as households have a larger incentive to save the necessary house deposit. After the tax credit expires, the share of highly leveraged households is elevated which further suppresses consumption.

Structural models in which heterogeneous agents save in a single risk-free asset are unable to match the empirical results on households' marginal propensity to consume.[5] Kaplan and Violante (2014) resolve this tension by modeling an economy in which households choose between a liquid and an illiquid asset, which is subject to a transaction cost. In their model, the existence of wealthy households that have few liquid assets generates higher marginal propensities to consume which are consistent with the empirical results. However their approach only models households' net asset position, and does not allow for significant levels of debt. This assumption is incongruous with Misra and Surico (2014), who find that the composition of balance sheets has a large impact on households' marginal propensities to consume.

My first contribution is to relax this assumption by developing a heterogeneous agent model à la Aiyagari-Bewley-Huggett, in which households can save in either liquid assets or illiquid housing, and where the latter can be used as collateral for liquid debt. Households save in order to self-insure against fluctuations in labour income, expand their access to credit and enjoy services from housing. Building on Caballero and Farhi (2014), Kaplan and Violante (2014) and Achdou et al. (2017), I introduce a housing market subject to several frictions. Households face borrowing constraints, financial frictions and transaction costs when trading houses. These frictions in the housing market are what generate a negative marginal propensity to consume for households that plan on purchasing a house. In a model without transaction costs, households optimise their consumption of housing and non-durable goods by keeping the marginal rate of substitution between the two constant. This implies a positive co-variance between consumption and housing, as income fluctuates over the life of the household. However, transaction costs induce households to make lumpy,

5 The collective evidence suggests a mean marginal propensity to consume of around 0.25.

debt-financed purchases of housing. When a household chooses to borrow in order to finance the purchase of housing stock, the opportunity cost of non-durable consumption in the present increases for two reasons. First, the household faces a higher interest rate, since the rate they pay to borrow is higher than the rate they earned on their liquid savings. This increases the cost of consumption today relative to saving for consumption tomorrow. Second, borrowing to pay for housing increases the probability that households will be credit constrained in the future, which could force them to forego high marginal utility consumption. Thus consumption today involves the possibility of foregoing a large amount of consumption utility tomorrow in the event that the household is credit constrained.

The combination of a household having to pay a higher interest rate on their debt, and the desire to reduce the probability that they will hit their borrowing constraint, leads them to lower their level of consumption upon purchasing housing. Forward-looking households smooth this transition by lowering their level of consumption as the probability that they will choose to buy a house in the near future rises. In my model, a transitory rise in income increases the probability that a household will find it optimal to purchase housing. This increases their incentive to save in anticipation of being credit constrained post-purchase. For the majority of households, the wealth and liquidity effects outweigh the anticipatory saving effect such that the increase in income leads to a rise in consumption. However, for households that are close to the point at which they would choose to purchase housing, a small increase in income can lead to a large rise in the probability that they will do so. For these households, the anticipatory saving outweighs the wealth and liquidity effects and thus produces a negative co-variance between income and consumption.

My second contribution is to estimate the effects of expected changes in household debt on non-durable consumption. I pursue two empirical strategies. First, I estimate households' marginal propensities to consume from transitory income shocks using the semi-parametric method developed by Blundell, Pistaferri, and Preston (2008) and Kaplan and Violante (2010). I show that average marginal propensities to consume are lower for households that are more likely to buy additional stock, as measured by households' self-reported expectations that they will move in the next two years.[6] I find that households which report that they will definitely move have a marginal propensity to consume that is negative. Second, I estimate a series of models of consumption growth following the approach of Dynan (2012) to show that anticipated changes in household debt are associated with declines in household consumption in magnitudes that are consistent with the results from the model. I estimate these regression models using an instrumental variable approach which I show provides a consistent estimator of the impact of expected changes in debt and home ownership.

Finally, the existence of households with a negative marginal propensity to consume has important implications for public policy. During the financial crisis several governments implemented tax credits aimed at encouraging first home buyers to enter the market. My third contribution is to use my model to analyse the impact of these temporary tax credits on house prices, home ownership and aggregate consumption. I show that while the tax credits do boost prices and turnover in the housing market, the effect is concentrated at the time the tax credit expires rather than when they are implemented. The tax credit "pulls forward" sales that would otherwise have occurred in the future. By encouraging households to purchase more housing, it also increases the share of highly leveraged households within the economy which leads to lower aggregate consumption. This result contrasts with both the stated intention and the previous literature on such tax credits, which assumed that they would stimulate the economy and increase aggregate consumption.

The chapter is organised as follows. Section 1.2 discusses how this chapter contributes to the literature. Section 1.3 presents my general equilibrium model of the housing market under incomplete markets. Section 1.4 discusses the calibration of the model, analyses the implications for household consumption dynamics and discusses alternative calibrations. Section 1.5 validates the main findings of the model using micro-data from the PSID. Section 1.6 applies the model to analyse the impact of tax credits for first home buyers. Section 1.7 concludes.

6 These expectations are surveyed by the PSID and are positively correlated with future changes in both home ownership and increases in household debt.

1.2 Related Literature

My analysis relates to several strands of the literature.

First, I add to the large theoretical literature that models how heterogeneous agents behave in the presence of incomplete markets. Introduced by Bewley (1983) and with early explorations by Huggett (1993) and Aiyagari (1994), these models explore how idiosyncratic shocks affect the distribution of wealth and income within the economy and how households achieve partial insurance by saving in a single risk-free asset. This approach solves for the market-clearing prices by aggregating asset demand over the distribution of households. Further development of this framework include extensions to multiple assets with different degrees of liquidity (Kaplan and Violante, 2014), and heterogeneity among household preferences (Iacoviello and Pavan, 2013). Previous literature on the two-asset model assumes that households are able to borrow up to an exogenous limit. I add to this literature by extending the two-asset model to allow for borrowing limits to be endogenously determined by a household's stock of housing and the market-clearing price.

More recent work by Achdou et al. (2017) shows that by solving the model in continuous time, aggregate shocks can be solved keeping a high-dimensional approximation of the wealth distribution as a state variable. Guerrieri and Iacoviello (2017) highlight the importance of this approach, showing that aggregate shocks have quantitatively different impacts on agents in different regions of the distribution.

Second, my approach to modeling the mortgage market builds on studies of housing markets. Given their computational complexity, life-cycle models typically assume that house prices are purely exogenous (Fernandez-Villaverde and Krueger (2011); Iacoviello and Pavan (2013); Yang (2009)). A handful of papers develop models with aggregate shocks that impact the equilibrium price of housing, but they require either environments with household heterogeneity limited to two types of agent (Justiniano et al. (2015); Iacoviello and Pavan (2013)) or approximations which reduce their dimensionality to a single moment as pioneered by Krusell and Smith (1998). However, for more complex models the number of moments required for an accurate approximation can be orders of magnitude higher (Ahn et al., 2017b). I add to this literature by solving a model which contains both an endogenously determined price of housing and a rich level of household heterogeneity capable of matching the empirical distribution of earnings.

This chapter also contributes to the empirical literature studying the link between a household's balance sheet and non-durable consumption. A number of papers have used increasingly rich models with housing and debt to address aggregate questions and make cross-sectional predictions, e.g. Carroll and Dunn (1997), Campbell and Cocco (2007), and Attanasio et al. (2012). The impact of exogenous house price movements on consumption and debt over the life-cycle has also been extensively examined using the heterogeneous agent framework, mostly notably by Garriga and Hedlund (2016), Berger et al. (2015) and Iacoviello and Pavan (2013). By contrast, my chapter focuses on how transaction costs affect non-durable consumption before and after a household adjusts its stock of housing. Benmelech et al. (2017) and Best and Kleven (2013) both estimate the impact of a house purchase on consumption. Their approach compares how consumption, as measured by the Consumer Expenditure Survey and the UK Living Costs and Food Survey respectively, varies with housing tenure.[7] They find that consumption of durable goods rises for households who have recently purchased a house, while consumption of non-durable goods falls. By contrast I use data from the PSID which follows households as they move and surveys households' expectations. I extend their analysis by using lagged household expectations as an instrumental variable to control for potential endogeneity within the model. This approach controls for contemporaneous shocks which may affect both consumption and the decision to purchase a house. Another approach in the empirical literature is Martin (2003), who shows that households increase their spending on food before moving to a smaller house and decrease it before moving to a larger house. I replicate this analysis with data on broader consumption and extend it to measuring a household's marginal propensity to consume prior to purchasing housing.

This chapter also contributes to the fast expanding literature that estimates heterogeneity in marginal propensities to consume across households. There is a wealth

7 Since neither survey follows households that move, Benmelech et al. (2017) and Best and Kleven (2013) focus on how consumption is affected after a house purchase.

of literature that analyses how debt affects households' marginal propensities to consume, using self-reported values (Auclert, 2017), exogenous variation in fiscal transfers (Greer, Parker, and Souleles (2006) and Parker et al. (2013)), lottery winnings (Fagereng et al., 2016) and a semi-structural approach (Blundell, Pistaferri, and Preston 2008). These methods consistently find that households with high levels of household debt and low access to liquid assets have high marginal propensities to consume. I contribute to this literature by using data on households' self-reported expectation of moving to show that households with a high probability of buying additional housing have lower, even negative, marginal propensities to consume.

Finally, in an application of the model I estimate the impact of tax credits for first home buyers. Dynan et al. (2013) find that while there is some evidence that the tax credit helped support house prices, there was no discernible positive impact of the tax credits on broader economic activity. My work suggests these tax credits have a contradictory impact on aggregate consumption, resolving this puzzle.

1.3 The Model

In this section I introduce a two-asset, incomplete markets model. My main innovation is to augment Kaplan, Moll, and Violante (2016)'s rich representation of household consumption and saving behavior with the ability of households to borrow against the nominal value of their stock of illiquid housing. Households face uninsurable idiosyncratic shocks to labour income and can self-insure through saving in either housing and liquid assets. Outside of the household and housing sector, the rest of the model is kept relatively parsimonious. To economise on computational time I solve the model in continuous time using the upwind approximation outlined in Achdou et al. (2017).

1.3.1 Households

The economy is populated by a continuum of households who are differentiated by their holdings of liquid assets, b, their ownership of illiquid housing, a, and their idiosyncratic labour productivity, z. I assume that time is continuous and that there is no aggregate uncertainty. At each point in time t the economy is governed by the joint distribution of the households $\mu_t(b, a, z)$. Households die off with a fixed probability λ and are replaced by new households with zero wealth $(a = b = 0)$ and the mean level of income.[8]

Households must optimise the expected present value of their utility flow, $u(c_t, h_t, l_t)$, by choosing the optimal path of their labour supply, l_t, non-durable consumption c_t and their ownership of housing h_t. Each household takes prices, wages, interest rates and governmental transfers, $\Phi = \Phi\left\{w_t, r_t^b, a_t^p, \tau_t, T_t\right\}$, as given. Preferences are time-separable, and conditional on surviving, the future is discounted at a rate $\rho \geq 0$:

$$U_t = E_t \int_t^{\infty} e^{-(\rho+\lambda)t} u(c_t, h_t, l_t)\, dt. \tag{1.3.1}$$

The flow utility is increasing and strictly concave in c and h and decreasing and strictly convex in l. I assume the functional form

$$u_t = \frac{\left\{\{c_t - z_{it}G(l_t)\}^{1-\zeta}\{h_t\}^{\zeta}\right\}^{1-\gamma} - 1}{1-\gamma} \qquad G(l_t) = \phi z_t \frac{l_t^{1+\frac{1}{\eta}}}{1+\frac{1}{\eta}}, \tag{1.3.2}$$

where the curvature parameter, γ, determines households' risk aversion and intertemporal elasticity of substitution, ζ is the weight placed on housing relative to non-durable consumption, and η is the Frisch elasticity of labour supply. I assume a functional form for flow utility that follows Greenwood et al. (1988), which cancels the wealth effect on the supply of labour, l_t. Thus all households, regardless of their assets, will supply the same amount of labour which is solely a function of the wage per effective unit of labour, w_t, and does not vary with idiosyncratic productivity shocks, z_t.[9]

Households maximise the present value of expected utility subject to four constraints. The first is the budget constraint, which governs the evolution of liquid assets, b_t, in the form of a risk-free bond

$$\dot{b_t} = z_t w_t l_t + r_t b_t - c_t - c_t^{rent} p_t^{rent} - k\left(a_t, a_t'\right) - p_t^a \alpha_t - T_t. \tag{1.3.3}$$

8 The stochastic death of households is not required to solve the model, but is critical to matching the high share of renters with no assets observed within the economy.

9 This assumption creates a direct mapping between productivity and earnings, which facilitates the calibration of the exogenous productivity process. I relax this assumption in Appendix 1.C.

Housing assets are illiquid as households must pay a fee $k(a_t, a_t')$ in order to buy or sell an amount of housing. I set this friction as a fixed percentage of the current level of housing , a_t, and new level of housing, a_t'. I denote any new purchases or sales of housing $\alpha_t = a_t' - a_t$. Note that while the model is written in continuous time, any purchase of housing stock and the accompanying change in liquid assets occur instantaneously. Each household's stock of housing depreciates over time at a rate δ, and thus the stock of housing evolves according to

$$\dot{a}_t = (1-\delta)\, a_t + \alpha_t. \tag{1.3.4}$$

The interest rate paid on the liquid asset varies depending on whether the household is a saver who deposits their funds with the financial firms, or a net borrower. The spread between saving and borrowing occurs due to the cost of financial intermediation and is detailed below.

$$r_t = \begin{cases} r_t^{+} & \text{if } b_t > 0 \\ r_t^{+} + r_t^{spread} & \text{if } b_t \leq 0 \end{cases} \tag{1.3.5}$$

Each household's consumption of housing services is defined by the stock of housing that they own. Households that do not own housing must lease some via the rental market. However, rental units only grant a fraction, θ^r, of the benefit to households relative to owned stock.

$$h_t = \begin{cases} \theta^r c_t^r & \text{if } c_t^r > 0 \\ a_t & \text{if } c_t^r \leq 0 \end{cases} \tag{1.3.6}$$

Finally, each household can borrow against a proportion of the total value of the stock of housing they own.

$$b_t \geq -M a_t p_t^a \tag{1.3.7}$$

$$a_t > 0 \tag{1.3.8}$$

Households maximise 1.3.1 subject to (1.3.3)-(1.3.8), taking as given the path for the real wage $\{w_t\}_{t\geq 0}$, the interest rate on liquid assets $\left\{r_t^b\right\}_{t\geq 0}$, the price of housing relative to non-durable goods $\{p_t^a\}_{t\geq 0}$, and taxes and fiscal transfers $\{\tau_t, T_t\}_{t\geq 0}$. As described below, the paths for the price variables are determined by the market-clearing conditions for labour, liquid assets and housing. In Appendix 1.A I describe how the households' optimisation problem can be solved recursively with a Hamilton-Jacobi-Bellman equation. The resulting steady state solution yields decision rules governing household behavior as a function of liquid assets, illiquid housing, productivity and the relative prices for non-durable consumption $c\left(a, b, z; \Phi\right)$, labour supply $l\left(a, b, z; \Phi\right)$, and the housing adjustment rule $\alpha\left(a, b, z; \Phi\right)$. Outside of the steady state these rules will be time-varying and respond to the time path of aggregate prices $\Phi_t = \left\{w_t, r_t^b, p_t^a, \tau_t, T_t\right\}_{t\geq 0}$. These optimal decision rules imply drift paths for households' holdings of liquid assets and housing, which when combined with the exogenous stochastic process for z, can be used to calculate the stationary joint distribution of assets and income $\mu_t\left(a, b, z, \Phi\right)$. In Appendix 1.A I outline the Kolmogorov Forward equation that defines the evolution of this distribution over time.

1.3.2 Firms

Final Goods Firms

A continuum of final goods firms produce goods for consumption, y_t, by using effective units of labour, n_t, in a linear production function. These firms are perfectly competitive, take the real effective wage, w_t, as given and have zero profit. Their prices are perfectly flexible and are defined as the numaire. They maximise profits

$$\Pi_t^{goods} = y_t - w_t n_t, \tag{1.3.9}$$

subject to the production function

$$y_t = \zeta_t n_t, \tag{1.3.10}$$

where ζ_t is an aggregate labour productivity shock.

Real Estate Firms

Real estate firms borrow from financial firms to fund the purchase of housing, which is rented out to households that do not own their own home. Real estate firms operate in a perfectly competitive market, and thus the rental price is pinned down by the cost of borrowing, a fixed cost of providing rental services, the price of purchasing house stock and any capital appreciation driven by the drift in the price of housing. I assume that real estate firms are able to borrow against the full value of their housing stock. The zero profit condition pins down the price of rental stock as the function of the cost of housing, capital gains from housing and the interest rate

$$p_t^{rent} = \phi^{rent} + \left(r_t^{b-} + \delta\right) p_t^a - \dot{p}_t^a, \tag{1.3.11}$$

where ϕ^{rent} is the fixed cost of providing rental services.

Financial Firms

Financial firms intermediate saving and borrowing within the economy subject to a fixed cost, r_t^{spread}, in the form of a spread between the interest rate on deposits and loans. Households make liquid deposits in financial firms, which are then loaned out to households who require a mortgage or to real estate firms. Finally, I assume that these financial firms operate in a competitive environment and thus earn zero profits

$$\Pi_t^{financial} = \left(r_t^b + r_t^{spread}\right)\left(\int_{b<0} 1\, d\mu_t + \int_{a=0} c_t^{rent}\, d\mu_t\right) - r_t^b \int_{b>0} d\mu_t. \tag{1.3.12}$$

Construction Firms

Finally, a continuum of construction firms produce new units of housing for sale to both households and real estate firms. These firms are perfectly competitive and produce houses using labour, n_t^a, and housing permits, $\overline{L}$, purchased from the government. This ensures that the price elasticity of supply of housing is finite, and is analogous to assuming investment adjustment costs in the housing sector. They maximise profits

$$\Pi_t^{construction} = p_t^a IA_t - w_t n_t^a - p_t^L \bar{L} \tag{1.3.13}$$

subject to the production function

$$IA_t = \zeta_t^a \left(n_t^a\right)^{\omega} \bar{L}^{1-\omega}. \tag{1.3.14}$$

Thus the price elasticity of supply is given by $\frac{\omega}{1-\omega}$.

I follow Favilukis et al. (2017) by assuming that the government sells the housing permits at the market-clearing price. This ensures that all rents from the fixed supply of housing permits accrue to the government and the construction sector makes no profit in equilibrium.

Maximizing 1.3.13 subject to 1.3.14 generates the housing investment function

$$IA_t = \left(\omega p_t^a\right)^{\frac{\omega}{1-\omega}} \overline{L} \tag{1.3.15}$$

1.3.3 Government

The fiscal authority funds an exogenous level of government spending and fiscal transfers to households via a combination of lump-sum payments and taxes on labour income

$$T_t = \tau_t w_t l_t + \tilde{T}_t. \tag{1.3.16}$$

Any fiscal deficit is financed by issuing debt to households in the form of liquid assets. I assume that the government are able to borrow directly from savers and are not subject to the friction generated by the financial intermediaries. Thus the government's budget position is given (in deficit terms) by

$$\dot{B}_t^{gov} = r_t^b B_t^{gov} - \int T_t \, d\mu_t - p_t^L \bar{L} + G_t. \tag{1.3.17}$$

1.3.4 Market Clearing Conditions

An equilibrium in this economy is defined as the paths for individual household and firm choice variables $\left\{a_t, b_t, c_t, c_t^{rent}, \alpha_t, l_t, n_t, n_t^a\right\}_{t\geq 0}$, and aggregate prices

$\{w_t, r_t^b, p_t^a, \tau_t, T_t\}_{t\geq 0}$ such that households and firms optimise their objective functions subject to their respective constraints, taking prices as given, for all time t and where the distribution of households, $\mu_t(a_t, b_t, l_t)$, is such that (i) the goods market, (ii) the market for liquid assets, (iii) the market for housing and (iv) the labour market all clear.

The final goods market clears if the total production by final goods firms is equal to consumption, government services, and the total resources lost to the frictions related to the housing sector, the financial firms and the real estate sector

$$Y_t = \int_{a_t \neq a_t'} k\left(a_t, a_t'\right)\, d\mu_t + \int_{b<0} r_t^{spread}\, d\mu_t + \int_{a=0} \phi^{rent}\, d\mu_t + C_t + G_t. \tag{1.3.18}$$

The liquid asset market clears when net household savings, B_t^{net}, are equal to the funds borrowed by the real estate sector and the fiscal authority

$$\int b\, d\mu_t = B_t^{net} = p_t^A C_t^{rent} + B_t^{gov}. \tag{1.3.19}$$

The market for housing must clear when total inflows from additional construction are equal to the aggregate depreciation plus the change in total demand by both home owners and renters

$$IA_t = \delta A_t + \int \dot{a}_t + \dot{c}_t^{rent}\, d\mu_t. \tag{1.3.20}$$

Finally, the labour market clears when the total labour demanded by the final goods and construction sectors equals the effective supply from households

$$N_t + N_t^c = \int zl\left(a, b, z\right)\, d\mu_t. \tag{1.3.21}$$

1.4 Calibration

1.4.1 Household Preferences

I calibrate the household utility function to match key moments in the data. I set the disutility of labour ϕ such that hours worked is equal to $1/3$ of the time endowment in equilibrium. I set the weight on housing in the utility function to match the average expenditure share on housing at 16 per cent. I set the curvature parameter, γ, which governs the risk aversion and inter-temporal elasticity of substitution, to 1. Given these values, the average value for the inter-temporal elasticity of substitution is approximately 0.7.[10] Finally, I set the rate of household deaths λ to $\frac{1}{180}$ such that the expected lifespan for a newly created household is 45 years.

1.4.2 Income Process

A critical determinant of the distribution of liquid and illiquid assets is the frequency and size of earnings shocks that households are subject to. An environment in which incomes are subject to small, but frequent, shocks will encourage households to hold a high level of liquid stocks to better smooth consumption. By contrast when shocks are infrequent, but large, households will be more willing to hold illiquid assets, paying the adjustment cost only occasionally when a large shock occurs. It is thus important to match the higher moments of the earnings process faced by households if we are to replicate the empirical distributions of wealth and income. One advantage of the continuous time environment is that we can calibrate shocks by their frequency of arrival, size and persistence. Conversely in a discrete time model shocks are assumed to arrive at the rate of one per unit of time.

Table 1.1: Earnings Process Moments

Moment	Data	Process - Continuous Time	Process - Discreteised Grid
Variance - log earning	0.70	0.70	0.72
Variance - 1yr change	0.23	0.23	0.21
Variance - 5yr change	0.46	0.46	0.48
Kurtosis - 1yr change	17.8	16.5	17.3
Kurtosis - 5yr change	11.6	12.1	11.6
Share of 1yr change $<$ 10%	0.54	0.56	0.59
Share of 1yr change $<$ 20%	0.71	0.67	0.61
Share of 1yr change $<$ 50%	0.86	0.85	0.89

As shown by Guvenen et al. (2015) the distribution of the change in log-earnings has a high degree of kurtosis (Table 1.1). In order to replicate this non-Gaussian

10 The inter-temporal elasticity of substitution is defined as $\frac{-u_c}{c.u_{cc}}$. Given our assumption on the utility function, this corresponds to $\frac{c-G(l)}{(\gamma+\zeta(1-\gamma))c}$ which varies across the distribution of households.

distribution, I follow Kaplan et al. (2016) in assuming that household earnings follow a "jump-diffusion" process. This process allows us to generate a distribution for changes in log-earnings which match the high kurtosis seen in the data.

Since all workers face the same wage per effective unit of labour and choose to supply the same amount of labour, a set of moments of households labour income can be used to estimate the labour productivity process z_{it}. I split the aggregate log-earnings process, z_{it}, into two components

$$z_{it} = z_{1,it} + z_{2,it}, \tag{1.4.1}$$

where each component, $z_{j,it}$, evolves according to

$$dz_{j,it} = -\beta_j z_{j,it} dt + dJ_{it}, \tag{1.4.2}$$

where dJ_{it} is the jump process which arrives at a rate α_j, such that over a small time period dt the probability that a jump occurs is $\alpha_j dt$ and the probability that a jump does not occur is $(1 - \alpha_j)\, dt$. Conditional on a jump occurring, the new level for the component, $z_{j,it}$, is drawn from a normal distribution with zero mean and variance σ_j^2. Thus

$$dJ_{j,it} = -z_{j,it} + \epsilon_{j,it} \quad \text{with} \quad \epsilon_{j,it} \sim N\left(0, \sigma_j^2\right) \tag{1.4.3}$$

To calibrate the parameters of the earnings process I use a simulated method of moments to match data from Guvenen et al. (2015).

The estimates of the parameters suggest the presence of two kinds of shocks. The first is an infrequent, but persistent career shock which occurs on average every 38 years and has a half-life of 18 years. The second is a more temporary shock which arrives roughly every 3 years, but largely dissipates within a quarter.

1.4.3 Housing Market

I calibrate the parameters related to the housing market to match key long-run relationships in the US housing market. The fixed cost faced by households when adjusting their housing stock is set to 7 per cent of the value of the house price. This

Table 1.2: Earnings Process Parameters

Parameter	Persistent Component $j = 1$	Transitory Component $j = 2$
Arrival rate	0.08	0.007
Mean reversion	0.76	0.009
St deviation of jump process	1.76	1.54

Note. Rates are expressed as quarterly values.

produces a housing market in which 9.2 per cent of the housing stock is turned over annually. This compares with roughly 10 per cent in the data as estimated by Ngai et al. (2016).

The elasticity of housing supply is determined by the exponent in the production function of the construction sector, ω, which I set to generate an elasticity of 1.5 per cent, which is the median value of elasticities estimated by Saiz (2010) across US cities. I calibrate the size of the construction sector to be 5 per cent on total output within the economy, matching its long-run share of gross value added.

The remaining parameters $\left(p^a, \theta^r, \rho, s^h, \phi^{rent}\right)$ are set to match five key moments in the data.

Table 1.3: Calibrated Moments

	Model	Data
Share of owners with a mortgage	0.55	0.57
Share of renters	0.38	0.38
Mean net worth - income ratio	5.2	5.5
Average ratio earnings owners - renters	2.1	2.2
Median loan-to-value ratio	0.66	0.63

I match these moments by simulating the model over a grid of parameters, and choosing the set which minimises the difference between the five parameters.

Given that the majority of lending within the model occurs between savers and mortgagees, I calibrate the financial friction to match the long-run average spread between 30-year fixed rate mortgages and the 3-month Treasury Bill. Since 1985 this spread has averaged 3.5 per cent. Finally, the real interest rate and the real wage are internally calibrated to clear the labour markets.

1.4.4 Equilibrium

The households' optimal consumption and saving decisions and the evolution of the joint distribution of their income, housing and liquid assets can be summarised by a Hamilton-Jacobi-Bellman equation and a Kolmogorov Forward equation. The method of approximating these equations is outlined in Appendix 1.A.

How well does this model match the distribution of wealth within the economy? In Table 1.4 I show some key moments related to the ownership of housing and total household net worth in the model.

Table 1.4: Selected Untargeted Moments of the Wealth Distribution

	Data	Model
Housing owned by top 1%	19.43%	8.5%
Housing owned by top 10%	64.7%	55.3%
Housing owned by top 20%	79.4%	73.0%
Gini coefficient - housing	0.73	0.65
Net worth owned by top 1%	35%	19.5%
Net worth owned by top 10%	88%	76.3%
Net worth owned by top 20%	96%	91.2%
Gini coefficient - total net worth	0.85	0.79
Aggregate debt to house value	0.42	0.40

Source: 2004 Survey of Consumer Finance

On the whole, the model matches the data fairly well with the notable exception of the top end of the distribution. While the model generates significant degrees of inequality across both liquid assets and housing, it fails to match the large shares of wealth owned by the top 1 per cent of households.

1.4.5 Marginal Propensities to Consume and Household Leverage

How well does this heterogeneous approach to households match the empirical findings on households' marginal propensities to consume? Of the range of empirical studies that examine this topic, the most compelling use exogenous variation in fiscal payments (Kaplan and Violante (2014); Broda and Parker (2014)), and lottery winnings (Fagereng et al. (2016)). These studies suggest that households spend 15-25 per cent of one-off payments in the quarter that they are received.

A marginal propensity to consume can be thought of as the impact of a one-off increase in liquid wealth. I thus define the marginal propensity to consume of a payment of x additional dollars over a length τ periods as

$$MPC_{\tau}^{x}\left(a,b,z\right)=\frac{C_{\tau}\left(a,b+x,z\right)-C_{\tau}\left(a,b,z\right)}{x}, \tag{1.4.4}$$

where $C_{\tau}\left(a,b,z\right)$ is the sum of expected consumption of an individual household over the next τ periods.

$$C_{\tau}\left(a,b,z\right)=E\left[\int_{0}^{\tau}c\left(a_t,b_t,z_t\right)dt|a_0=a,b_0=b,z_0=z\right] \tag{1.4.5}$$

This formula can be solved using the Feynman-Kac formula as outlined in Appendix 1.B.[11]

The unconditional quarterly marginal propensity to consume in the model is 19 per cent, which is consistent with the empirical results. However this unconditional mean masks a high level of heterogeneity across the distribution of wealth among households. Figure 1.1 shows the distribution of the marginal propensities to consume as a function of households' liquid and illiquid assets.

Figure 1.1: Household Heterogeneity

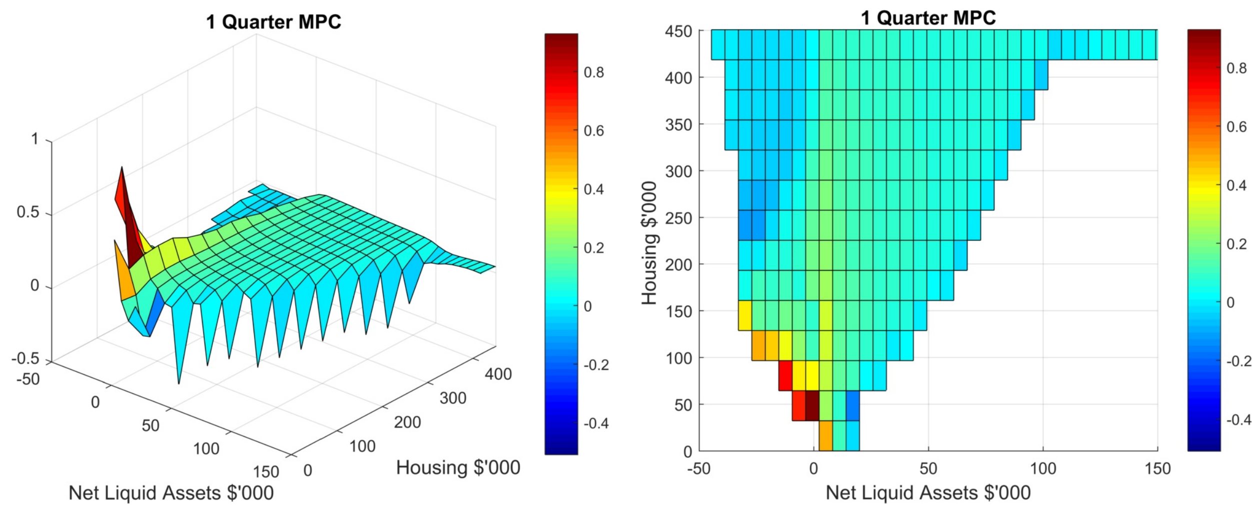

Note. This figure shows the distribution of marginal propensities to consume across households from two different viewpoints, over the subsequent quarter for households with the median income. For a two dimensional slice of this distribution see Figure 1.3.

11 An alternative strategy would be to estimate the impact of a small increase in the transitory component of the households' productivity process. This alternative strategy produces qualitatively similar results, but with larger simulation error due to the coarser grid used to approximate the productivity process.

There are three notable features of the distribution in Figure 1.1. The first is the spike as households approach the borrowing constraint which is governed by the nominal value of their housing stock. The second is the high marginal propensity to consume for households that own some amount of housing stock but have zero liquid wealth. The third notable feature is the fall in marginal propensities to consume, often to negative levels, of households who are close to the point at which they would optimally choose to adjust their stock of housing.

The first two features are well documented in previous research. Broda and Parker (2014) find that households with limited access to liquid funds had a significant increase in consumption in response to the 2008 fiscal stimulus payments in the US. In addition, Kaplan, Violante, and Weidner (2014) use a semi-structural approach to measure marginal propensities to consume and find that households with limited access to liquid assets have significantly higher marginal propensities to consume, even if they have high levels of illiquid assets. These results are consistent with the existence of wealthy hand-to-mouth households. Fagereng et al. (2016) examine the response of household consumption to exogenous lottery prizes using Norwegian administrative data. They find that marginal propensities to consume vary with the amount of households' liquid assets and that households with close to zero liquid assets have high propensities to consume even if they are wealthy in terms of their illiquid asset position.

The existence of households with a negative marginal propensity to consume is a more novel finding. The key to understanding this result is examining how a household's decision to adjust its stock of housing affects its incentive to consume non-durable goods.

Since changes in a household's ownership of housing is subject to transaction costs, it is optimal for households to make lumpy purchases, waiting longer before deciding to interact with the market and adjusting by a larger amount when they do so. These large, lumpy purchases require financing in the form of mortgages, which will place post-purchase households closer to the endogenous credit constraint and subject to the higher interest rate faced by borrowers relative to their previous position as savers.

These two effects are partially offset by the complementarity between housing and the consumption of non-durable goods in the households' utility function. However the effect of this complementarity is quantitatively small.

The combination of a higher probability of being credit constrained and the increase in interest rates faced by the household will increase the incentive for households to repair their balance sheet by reducing spending and paying back debt. Thus, on average, households decrease their consumption by 6 per cent conditional on buying additional housing stock (Figure 1.2).

Figure 1.2: Change in Non-durable Consumption Conditional on Buying Additional Housing Stock

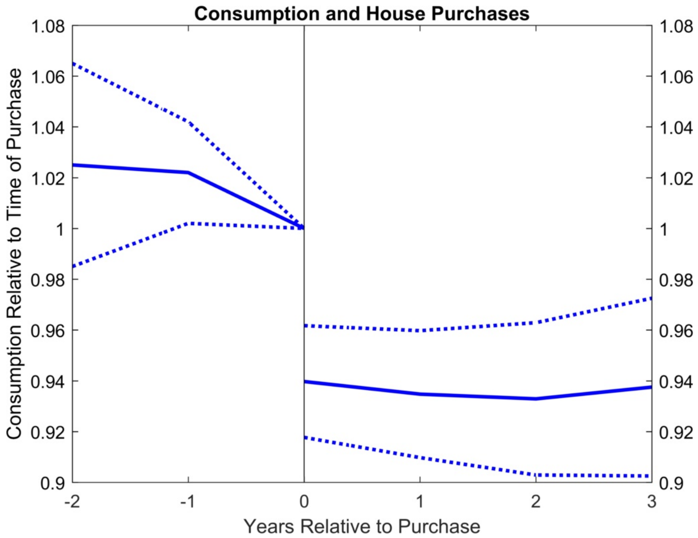

Note. The sold line represents the average level of consumption conditional on a household buying additional housing, relative to the level just prior to the purchase. The dashed lines represent the 90 per cent confidence interval.

Prior to purchasing additional housing stock, households anticipating the impending increase in debt and corresponding fall in non-durable consumption increase their savings rate to better smooth consumption in future periods. This anticipatory saving effect is stronger the higher the probability that households will find it optimal to buy a house in the near future.

Thus when households receive a small increase in liquid wealth there are three effects on household consumption. The wealth and liquidity effects both lead to an increase in the rate of consumption. However due to the increase in the probability that a household will choose to buy a house in future periods, the anticipatory saving effect may cause households to decrease their consumption. For the vast majority of households, the first two effects outweigh the third as only a relatively small share of households adjust their housing stock annually.

The same logic holds for households that are close to selling a portion of their housing stock. When a household decides to reduce its stock of housing, it uses the proceeds from the sale to pay back its mortgage debt, thus reducing the probability that it will become credit constrained in the future. This increase in liquid assets will lead to the household increasing its non-durable consumption after reducing its stock of housing. Anticipating this increase in consumption, households will choose to increase their consumption as their portfolio of assets approaches the optimal downsizing point. Since a small increase in liquid assets moves a household further away from this point, this will result in the household lowering their level of consumption as the probability that they will sell a portion of their housing stock decreases.

While this anticipatory saving effect reduces households' marginal propensity to consume, whether they adjust their stock of housing higher or lower, the size of the effect is asymmetric. This is because the transaction costs associated with adjusting the stock of housing reduces a household's wealth, regardless of whether they are buying or selling. This decrease in wealth ensures that the decrease in consumption associated with the purchase of additional housing stock is larger than the increase in consumption associated with selling the same amount. Thus the anticipatory saving effect is stronger for households who expect to buy housing compared with those who expect to sell.

This effect is illustrated in Figure 1.3 which shows how a household's consumption varies with respect to liquid assets for a fixed level of housing and income. If the households stock of liquid assets becomes too high (low) it buys (sells) housing stock which results in a fall (rise) in non-durable consumption post-adjustment. The household anticipates the possibility of this adjustment occurring and begins to smooth consumption towards the post-adjustment level. This results in consumption falling (rising) as liquid wealth rises (falls) close to the optimal adjustment point which in turn implies that the marginal propensity to consume is negative.

In practice, this phenomenon is caused by the combination of the cost of adjusting housing stock and the interest rate spread between borrowers and savers. When

Figure 1.3: Consumption as a Function of Liquid Assets

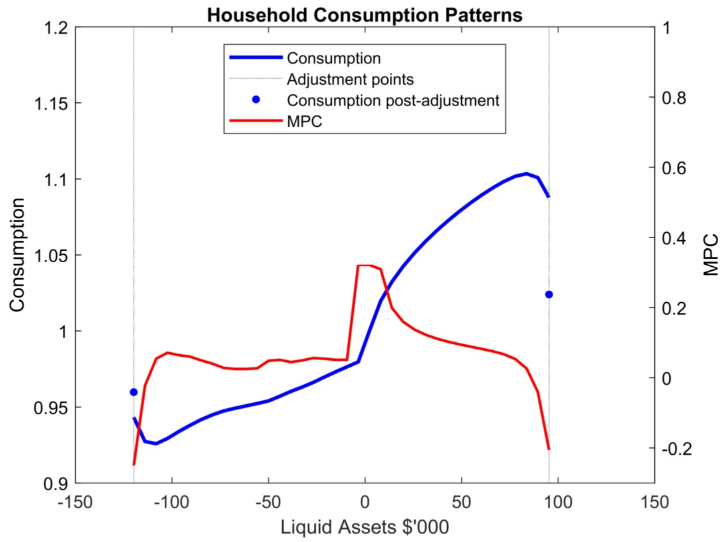

Note. This figure shows consumption as a function of the stock of liquid assets holding income and the stock of housing constant.

renters decide to enter the housing market they change from

savers to borrowers. This increases the interest rate they face when their liquid asset position shifts and causes them to decrease their level of non-durable consumption. If the stock of housing can be frictionlessly adjusted, then households are able to continuously adjust their ownership of housing and any debt as needed. A household's ability to slowly change the mix of their asset portfolio eliminates the large jumps in debt that are necessary in the presence of adjustment costs and allows consumption of both housing and non-durable goods to increase as household wealth and income increase. Alternatively, if the financial friction is removed then the interest rate households face remains unchanged, even after they take out a mortgage reducing the change in consumption.

Why would households choose to lower their consumption as they anticipate adjusting their housing stock and deviate from the optimal consumption path imposed by the short-run Euler equation? Households could alternatively choose to buy additional housing only when they have sufficient liquid wealth to avoid going into debt and risk becoming credit constrained. The answer is that households are better off financing their house purchases with debt because the alternative entails either (i) paying the transaction cost more often as households have less scope to expand their stock of housing; or (ii) remaining with an inefficiently small amount of housing as households save sufficient liquid wealth to be able buy more; or (iii) having to reside in inferior rental stock for longer which is subject to mark-ups by real estate firms.

1.4.6 Robustness

To test the robustness of this result I re-estimate the model over a range of externally calibrated parameters. Figure 1.4 shows how the main result, the negative marginal propensity to consume for households close to their adjustment point, varies with the externally calibrated parameters. I measure this result by calculating the mean marginal propensity to consume for households that have a high likelihood of adjusting their stock of housing in the next year.

Figure 1.4: Mean MPC of Households Close to Adjustment Point

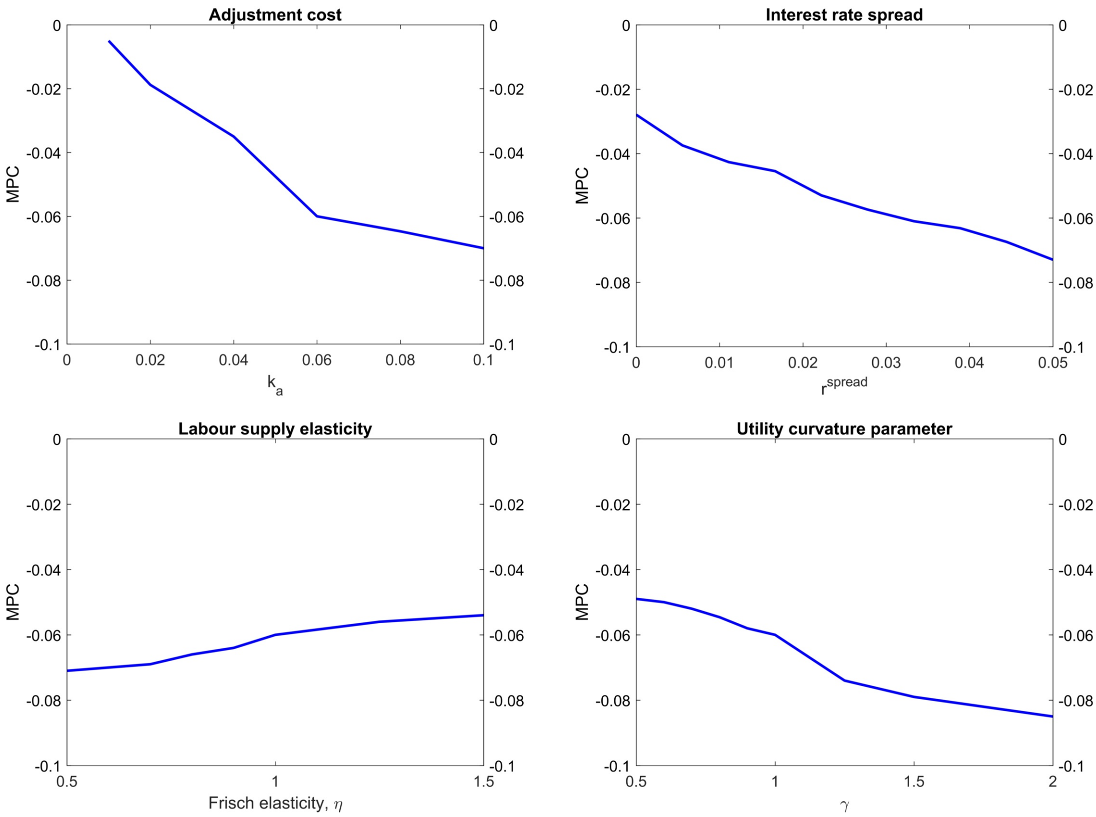

Note. Across the range of different calibrations I define households being close to the adjustment point if they have a greater than 90 per cent probability of adjusting their stock of housing in the next year.

The existence of households with a negative marginal propensity to consume is relatively robust over a range of calibrations. However, the importance of adjustment costs and financial frictions are highlighted by these results. As these parameters are lowered, the proportion of households with negative marginal propensities to consume tends towards zero. As the adjustment cost decreases households are able adjust their stock of housing more frequently, as their wealth increases, without going into debt. The reduction of leveraged purchases decreases the proportion of households that will decrease their consumption when they buy housing stock.

In Appendix 1.C I explore two different model specifications. The first allows for the debt constraint to apply only when households adjust their stock of housing. At all other times households are able to keep the stock of debt constant, even if a fall in house prices lowers the value of their housing stock. This breaks the tight link between house prices and household debt as discussed by Iacoviello and Pavan (2013). The second model relaxes the assumption that household labour supply is identically supplied across households. Instead I assume that the disutility of labour supply is separable from consumption and thus varies with household wealth. This assumption decreases the prevalence of households with negative marginal propensities to consume, as households have a second margin which they can adjust as they anticipate buying or selling housing stock. However while the proportion of households with a negative marginal propensity to consume is smaller, the phenomenon still exists across a range of reasonable calibrations.

1.5 Empirical Results

In this section I compare the predictions from the general equilibrium model against empirical results derived from household surveys. To compute the key cross-section moments, I require household-level panel data on income, consumption and household debt. There are a variety of techniques in the literature used to calculate parameters such as households' marginal propensity to consume, the most commonly used are the US PSID and the US Consumer Expenditure Survey. However since the Consumer Expenditure Survey does not follow households who move address, it excludes those who buy or sell their home - the main population of interest. I thus use the PSID to verify the predictions.

1.5.1 Measuring Marginal Propensities to Consume

The first exercise I conduct is to compare the predicted marginal propensity to consume of the model with the data. The model generates three broad predictions concerning the dynamics of non-durable consumption: (i) that households close to their borrowing constraint have a high marginal propensity to consume; (ii) that households who have close to zero liquid wealth have a high marginal propensity to consume; and (iii) that households who are close to adjusting their stock of housing have a low, or even negative, marginal propensity to consume. As there is already considerable empirical evidence supporting the first two predictions as outlined in my literature review, I will focus on the more novel third prediction.

To estimate marginal propensities to consume I follow the identification strategy of Blundell, Pistaferri, and Preston 2008 and refined by Kaplan and Violante (2010). Blundell, Pistaferri, and Preston 2008 show that by assuming household income is governed by a process with a permanent and an i.i.d. component, then given appropriate theoretical restrictions, an estimator for households' marginal propensities to consume can be calculated with panel data on consumption and income.[12]

The estimator for the transmission of transitory income shocks to consumption is given by

$$\widehat{MPC} = \frac{\operatorname{cov}\left(\Delta c_{it}, \Delta y_{i,t+1}\right)}{\operatorname{var}\left(\Delta y_{i,t}, \Delta y_{i,t+1}\right)} \tag{1.5.1}$$

which can be estimated using an instrumental variable regression, where $\Delta c_{i,t}$ is regressed on $\Delta y_{i,t}$, instrumented by $\Delta y_{i,t+1}$ with panel data of at least three periods.

Kaplan and Violante (2010) show that this estimator is highly robust at identifying the true marginal propensity to consume in a wide variety of models, including models in which there is a forecastable component of future income. Berger et al. (2015)

12 Two assumptions are required for the estimator to be consistent. The first is that households have no information about future shocks, the second is that consumption is memory-less and does not vary in response to the lags of the transitory shock.

further show that this method remains robust in the presence of a housing market with transaction costs and the option to rent. I replicate the exercises conducted by Kaplan and Violante (2010) and show that the true marginal propensity to consume is reliably measured in my model of jump-diffusion income processes.

1.5.1.1 PSID Data

While the PSID started collecting data in 1968, it only started surveying non-food consumption in 1999. Although it is possible to impute aggregate consumption by inverting the demand for food using data from the Consumer Expenditure Survey (Blundell et al., 2004), this measure can have high levels of measurement error as shown by Attanasio and Weber (1995). Accordingly, I calculate a broader measure of consumption which is measured bi-annually from 1999 to 2015. In this sub-sample, the PSID surveys over 70 per cent of all consumption items available in the Consumer Expenditure Survey and includes expenditure on food, utilities, gasoline, car repairs, public transportation, childcare, education and medical services. Following Kaplan and Violante (2010), I first purge the data of non-model features by regressing log consumption and log income on year & cohort dummies and a range of demographic variables and using the residuals to calculate households' marginal propensities to consume.

1.5.1.2 Measuring Expected Changes in Debt

To verify the model's predictions on the heterogeneity of households' marginal propensities to consume, I need to identify households who plan on adjusting their housing stock in the near future. Using actual changes in household balance sheets may bias our results due to unexpected shocks which affect both a household's consumption and their decision to take on debt.[13]

To control for this potential bias I use households' surveyed expectation that they will move in the next couple of years. Specifically the PSID asks households:

A51. Do you think you might move in the next couple of years?

13 For example an unexpected persistent, positive income shock may lead households to increase both consumption and their willingness to hold debt.

If a household replies in the affirmative, this question is followed with:

> A52. Would you say you definitely will move, probably will move, or are you more uncertain?

This variable, when lagged, is uncorrelated to contemporaneous shocks and highly correlated with future decisions to buy new houses as shown by Table 1.5.

Table 1.5: PSID Surveyed Expectations

	N	Share of households that:			
		subsequently move	... into a larger house	... into a similar sized house	... into a smaller house
Do not expect to move	47,305	0.17	0.80	0.04	0.16
Expectation - uncertain	4,988	0.34	0.84	0.03	0.13
Expectation - probably	8,758	0.50	0.86	0.02	0.12
Expectation - definitely	13,123	0.72	0.87	0.03	0.10

1.5.1.3 Results

In order to compute the marginal propensities to consume, the households must be grouped into buckets. I pursue three different strategies to identify households who anticipate buying or selling housing stock. First, I group households by their self-reported expectations in the previous survey. The second approach is to estimate a logit model of the probability a household will move, regressed on lags of household expectations, demographic variables, family income and employment. The equation provides us with a continuous measure of the likelihood that a household will move. Households are then grouped into quartiles based on this estimated likelihood of moving. Although the majority of households who adjust their stock of housing increase their stock, both of these strategies combine households regardless of whether they are net buyers or net sellers. To separate out the minority of households who move into smaller houses, my third approach groups households by whether they actually did move in the subsequent period, and whether the house they moved into had a smaller, unchanged or larger number of rooms.

Figure 1.5 shows the estimated marginal propensities to consumes of households according to their expectations about the likelihood of moving in the next couple of years.

Figure 1.5: Marginal Propensities to Consume and Household Expectations

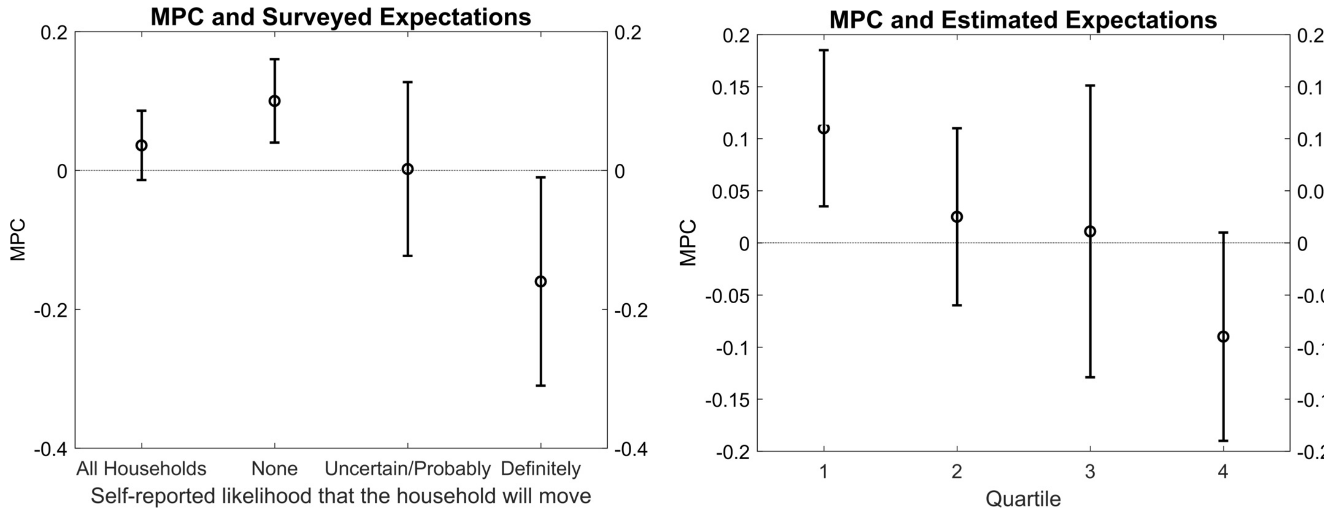

Note. This figure shows the estimated marginal propensity to consume with households grouped by their surveyed expectation of moving in the next two years (left) or by the likelihood that they will move estimated using a logit model regressing households' decisions to move on lagged expectations and demographic variables.

These results show that households with a higher expectation of moving house consistently have a lower marginal propensity to consume. Indeed for households that report that they "Definitely expect to move" in subsequent periods or which are in the top quartile of the estimated equation, the estimator is statistically below zero as predicted by the model. Furthermore when households are grouped by their actual change in housing tenure, I show that this effect is stronger for those households who move into larger houses.

This result can be explained by the covariances which comprise the Blundell, Pistaferri, and Preston estimator. The correlation between Δy_{t+1} and Δc_t is negative for the population as a whole, but is positive for the subset of households who report that they have a high probability of moving. This change in sign occurs because households who expect to move, but subsequently receive a negative income shock, are more likely to delay or cancel their planned purchase of housing. By choosing not to adjust their stock of housing they remain free from mortgage debt and subject to a lower interest rate, which results in a higher level of consumption.

Figure 1.6: MPC and Changes in Household Tenure

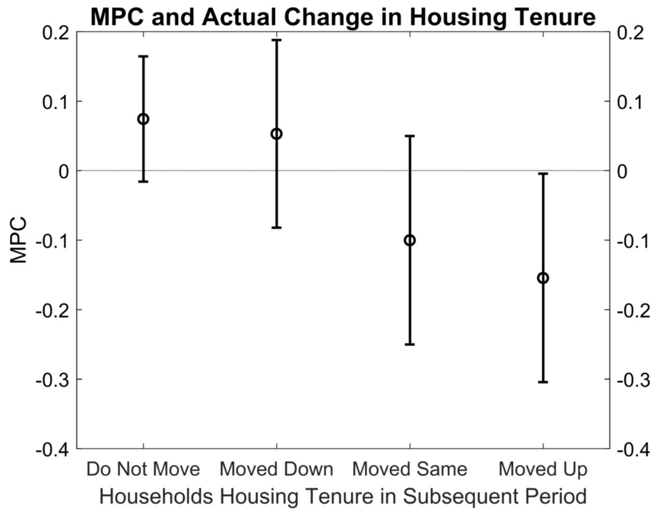

Note. This figure shows the estimated marginal propensity to consume with households grouped by whether they actually moved in the subsequent two years and whether they moved into a house with less, the same or more rooms.

While the standard errors around the estimates produced with this method are large, the covariance between anticipated increases in debt and households' marginal propensities to consume, and the negative result for those who are highly likely to buy additional housing stock, are consistent with the results of the model.

The identification of households with a negative marginal propensity to consume is a novel result. The closest related work is Martin (2003) who shows a similar effect in a partial equilibrium model of the housing market with transaction costs. There are several possible explanations for why this result has not occurred in the previous literature. First, the share of households who adjust their stock of housing in any given period is quite small. As empirical studies usually compare households across quantiles, unless the analysis deliberately isolates such households it is likely that they will be subsumed by the majority of households that do not adjust their stock of housing and have a conventionally non-negative marginal propensity to consume. Secondly, because the optimal adjustment point varies with income, even analyses using administrative data such as Fagereng et al. (2016), which is able to partition households by their balance sheets in finer gradients, may fail to find this effect. Finally, while the finding that some households have a negative marginal propensity to consume is new to the literature, these results are consistent with the broader conclusions on the negative correlation between households' marginal propensity to consume and their ownership of liquid assets.

1.5.2 Consumption Growth Regressions

The second exercise I conduct is to estimate how anticipated changes in debt affect the level of consumption of a household. To estimate this effect I adapt the model used by Dynan (2012) to estimate how changes in debt affect the growth in consumption,

$$\Delta \log c_{i,t} = \beta_0 + \beta_r r_{t-1} + \beta_y \Delta \log y_{i,t} + \beta_{Debt} \Delta D_{i,t} + \beta_W \Delta \log NW_{i,t} + \boldsymbol{\beta} \boldsymbol{X}_{i,t} + \varepsilon_{i,t}, \quad (1.5.2)$$

where $c_{i,t}$ is the log of real consumption, r_t is the real interest rate, $y_{i,t}$ is the household's real income, $D_{i,t}$ is the debt-to-asset ratio of each household and $\boldsymbol{X}_{i,t}$ is a set of demographic variables. Note that since the PSID is a bi-annual survey, I define Δ such that $\Delta x_t = x_t - x_{t-2}$. I use this model to examine (i) how anticipated changes in household debt and housing stock affect non-durable consumption and (ii) how consumption responds to changes in households' expectations of when they will adjust their stock of housing.

In order to identify the impact of expected debt on household consumption, I instrument the debt-to-asset ratio with households' lagged expectation of the probability that they will move along with lags for income, consumption and demographic variables. In Appendix 1.D I extend Hansen and Singleton (1982) to show that if the set of instruments used to identify the endogenous variable contain only information from time $t-2$ then an instrumental variable regression will produce a consistent estimator of the expected change in debt, excluding any unanticipated changes. This result is driven by the assumption of rational expectations, in which i.i.d deviations from prior expectations are orthogonal to information available at time $t-2$. By using the household's self reported expectation of moving as an instrument for changes in debt, I produce a consistent estimator of the effect of an expected change in debt on household consumption. It should be noted that while this estimator is consistent, it is only unbiased when the marginal effect of an expected increase in debt is equal to that of an unexpected change. However, simulating the model under a wide range of parameters suggests that this bias will be quantitatively small given our sample size.

Table 1.6: Regression Results

	(1)	(2)	(3)	(4) 1999-2007	(5) OLS
ΔDebt-Asset$_{i,t}$	-0.043 *	-0.043 *	-0.046 *	-0.051 *	0.011
ΔFamily Size$_{i,t}$	0.015	0.015	0.012	0.020	0.079***
r_{t-2}	-0.09***	-0.09***	-0.09***	0.00	-0.04*
Δlog$y_{i,t}$		-0.005	-0.005	-0.005	0.013***
Δlog$W_{i,t}$			-0.002	-0.005*	0.000
Observations	40,965	38,921	37,560	22,601	37,560

Source: Authors regressions of nested specifications of the model $\Delta \log c_{i,t} = \beta_0 + \beta_r r_{t-1} + \beta_y \Delta \log y_{i,t} + \beta_{Debt} \Delta D_{i,t} + \beta_W \Delta \log NW_{i,t} + \boldsymbol{\beta X}_{i,t} + \varepsilon_{i,t}$. The dependent variable is the change in log of real consumption. All regressions instrument contemporaneous variables using lags of endogenous variables and households expectation of moving, unless stated otherwise. Sample includes households with consecutive, complete set of interviews from 1999 to 2015. Also included, but not reported, are controls for the age and education of the household head.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

These results show that anticipated increases in household debt have a significant impact on household consumption across a range of specifications, even when the increases are anticipated years in advance.

To compare these estimates to the predictions of the model I calculate the mean change in non-durable consumption that occurs when households adjust their stock of housing, instrumenting again with household expectations, in Table 1.7. My estimates suggest that households that buy their first house decrease their consumption by approximately 10 per cent, relative to households who choose to keep renting. This figure increases to 12 per cent if we only include households that take out a mortgage in the process. By contrast, households who buy additional housing stock (as proxied by an increase in the number of rooms when homeowners move) decrease their non-durable consumption by approximately 7 per cent. The results show that even households who anticipate their adjustment of housing stock decrease their non-durable consumption in response to the purchase.

These estimates are somewhat larger then the simulated outcome of the model, which predicts a fall in consumption of approximately 4 per cent upon upgrading housing stock. Part of this difference is due to unanticipated income shocks that exist within the model. A household that is subject to a positive income shock will simultaneously increase their consumption of both housing and non-durable goods. This will bias the simulated model results towards zero.

Table 1.7: Alternate Regressions

	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)
ΔDebt-Asset$_{i,t}$	-0.04 *	-0.05 *						
ΔMortgage$_{i,t}$			-0.13***	-0.11***				
ΔFirst Home Owner$_{i,t}$					-0.11***	-0.10***		
ΔBuy additional stock$_{i,t}$							-0.06*	-0.07*
ΔFamily Size$_{i,t}$	0.02	0.01	0.02	0.02	0.02 *	0.02 *	0.01	0.01
r_{t-2}	-0.09***	-0.09***	-0.05***	-0.05***	-0.05***	-0.06***	-0.01***	-0.01***
$\Delta \log y_{i,t}$	-0.01	-0.01	0.03 *	0.03 *	0.03 *	0.03 *	0.02	0.020
$\Delta \log W_{i,t}$		-0.00		-0.03		-0.03		-0.00
Observations	37,560	37,560	37,560	37,560	37,560	37,560	37,560	37,560

Source: Authors regressions of the baseline model with alternative measures of debt and home ownership.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

These results, using a more comprehensive measure of consumption, confirm the findings from Martin (2003). Namely, that households which choose to purchase additional housing stock decrease their non-durable consumption. This result accords with the predictions of the model and highlights the importance of transaction costs and financial frictions when modeling a household's consumption dynamics.

Table 1.8: Changes in Household Expectations

		Period t			
		No expectation	Uncertain	Probably	Definitely
	No expectation	-	-0.011	-0.011	-0.073^{***}
Period $t-2$	Uncertain	-0.015	-0.031	-0.010	-0.063^{***}
	Probably	0.007	-0.017	-0.027^{*}	-0.048^{***}
	Definitely	0.020^{*}	-0.065	-0.029	-0.057^{***}

Estimated impact of a change in household expectations on household consumption.
Source: Authors regressions of nested specifications of the model $\Delta \log c_{i,t} = \beta_0 + \beta_r r_{t-1} + \beta_y \Delta \log y_{i,t} + \boldsymbol{\beta_{Exp} \Delta E_{i,t}} + \beta_W \Delta \log NW_{i,t} + \boldsymbol{\beta X}_{i,t} + \varepsilon_{i,t}$, where $\boldsymbol{\Delta E_{i,t}}$ is a vector of dummy variables which represent all potential transitions in household expectations across two periods.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Finally, I also find evidence that consumption of non-durable goods decreases as household expectations of moving rise, even when their stock of housing remains constant. Table 1.8 shows the estimated elasticity of changes in household expectations on consumption. I find that households who definitely expect to move have lower levels of consumption, and that the fall is larger for households that report a bigger change in expectations. By contrast there is evidence that households who report a lower probability that they will move have a higher level of consumption. These results are consistent with the predictions from my model in which households decrease their level of non-durable consumption in anticipation of making a leveraged purchase of housing in the near future.

These results are in line with Benmelech et al. (2017), who find that non-durable consumption falls after the purchase of a house. However my results indicate that the decrease is significantly larger. This difference is driven by the use of household expectations as an instrumental variable. The ordinary least squares regression results are much more similar in magnitude. This implies that the instrumental variable approach, which remove the impact of contemporaneous shocks that may impact both non-durable consumption and the decision to purchase housing, is important to identifying the true marginal effect.

1.6 Fiscal Transfers and the Housing Market

Subsidies for home ownership are a common feature of housing markets (Andrews et al., 2011). These policies range from subsidised access to credit, tax preferences for household debt and outright transfers to home buyers. While many of these subsidies are structural in nature, a common component of the fiscal response to the global financial crisis was to expand these programs to help stabilise housing markets and stimulate demand within the economy (Table 1.9).

The results from the previous sections raise the question of whether a transfer payment or tax credit to proto home-owners may decrease aggregate consumption in general equilibrium and thus be contractionary in nature. This is an important policy question as tax credits to households saving for their first deposit were enacted during the financial crisis in a number of countries. While these programs were primarily designed to help boost the housing and credit markets, it was also assumed that they would help boost the broader economy and increase output (Dynan et al., 2013).

Table 1.9: Selected Fiscal Transfers Targeted at First Home Buyers

Country	Size (% of median home price)	Dates	Notes
US	$7,500-8,000 (3.5%)	April 2008 - June 2010	Initially an interest free loan, it was converted into an outright grant for first home buyers.
UK	Discount of up to $140,000 for public housing tenants to buy their home (35-70% depending on length of tenancy)	2012 - Present	"Right to Buy" discounts were introduced in 1980, but greatly expanded in 2012.
Australia	$11,200 (4 %)	October 2008 - September 2009	This is in addition to $6,000 permanently available.
Canada	$700 (0.3%)	January 2009 - Present	

I consider two policy scenarios. The first is an unconditional transfer of money to renters who are on the margin of purchasing a house. The second is a conditional tax credit which is available to all renters who purchase additional housing stock within a one year window.

1.6.1 Unconditional Transfers

To estimate the impact of these programs, I model the impact of an $8,000 debt-funded transfer to renters who are sub-marginal first home buyers. I model this transfer as a one-off increase in liquid wealth to households who are currently renting and are on the margin of becoming a home owner. The transfer is funded by the government through an increase in debt, while tax rates remain unchanged. I then solve for the impact of this transfer in equilibrium allowing prices to adjust to clear markets.

I solve for the equilibrium path for the model as it returns to steady state after the initial transfer. Figure 1.4 shows the impact from this policy on the housing market and the broader economy. Unsurprisingly, the transfer increases both prices and the stock of housing owned by households, as the transfer to renters enables them to purchase more housing sooner than they otherwise would have. This leads to an

Figure 1.7: Impact of a Unconditional Transfer

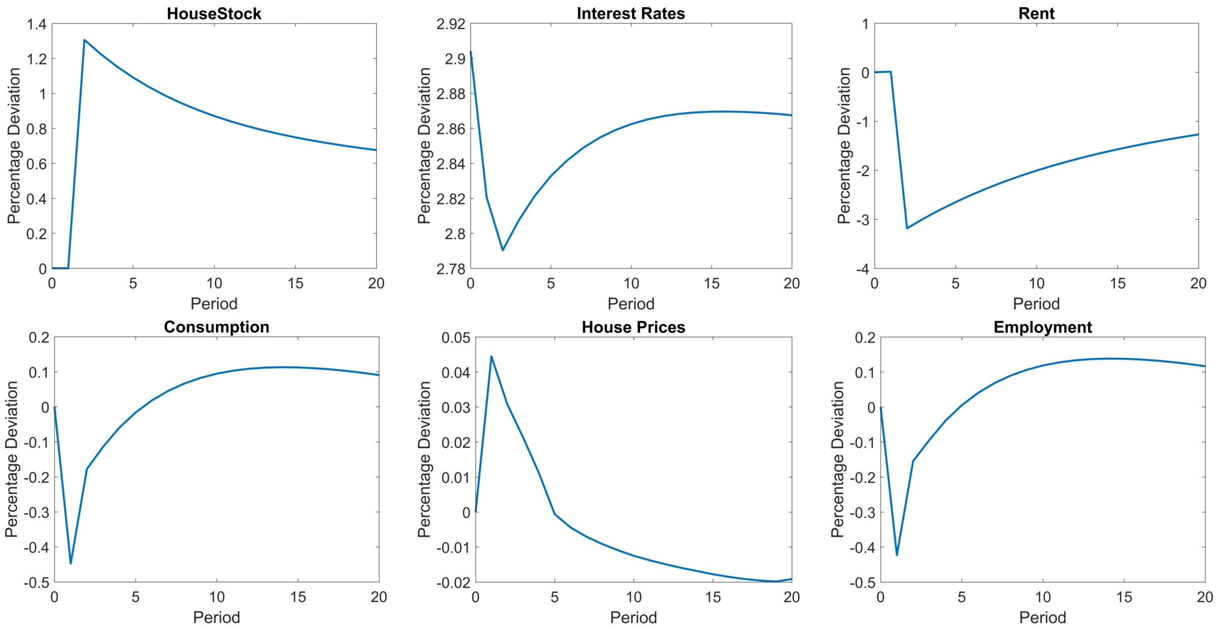

Note. These impulse response functions show the impact of an unconditional $8,000 transfer to renters

increase in household leverage within the model which results in households decreasing their consumption.

1.6.2 Conditional Transfers

An alternative policy design is to offer tax credits when renters purchase a house. I model this scenario by announcing a policy that will give renters an $8,000 payment if they purchase a house in the next four quarters. This policy differs from a straight fiscal transfer in two ways. First, by applying the policy over a longer period, it allows more renters to take advantage of the credit magnifying its impact on the goods and housing market. The second is that the tax credit provides the opportunity for renters to adjust their behavior in anticipation of taking advantage of the credit before it expires.

If the tax policy is open to all renters who buy a house within a set period, in this case one year, the general equilibrium impact of the policy has two distinct phases. From the time the tax credit is announced, renters decrease their consumption of both non-durable goods and rental stock in order to increase their stock of savings to a level high enough to enable them to buy a house while the credit is available. This

Figure 1.8: Impact of a Conditional Transfer

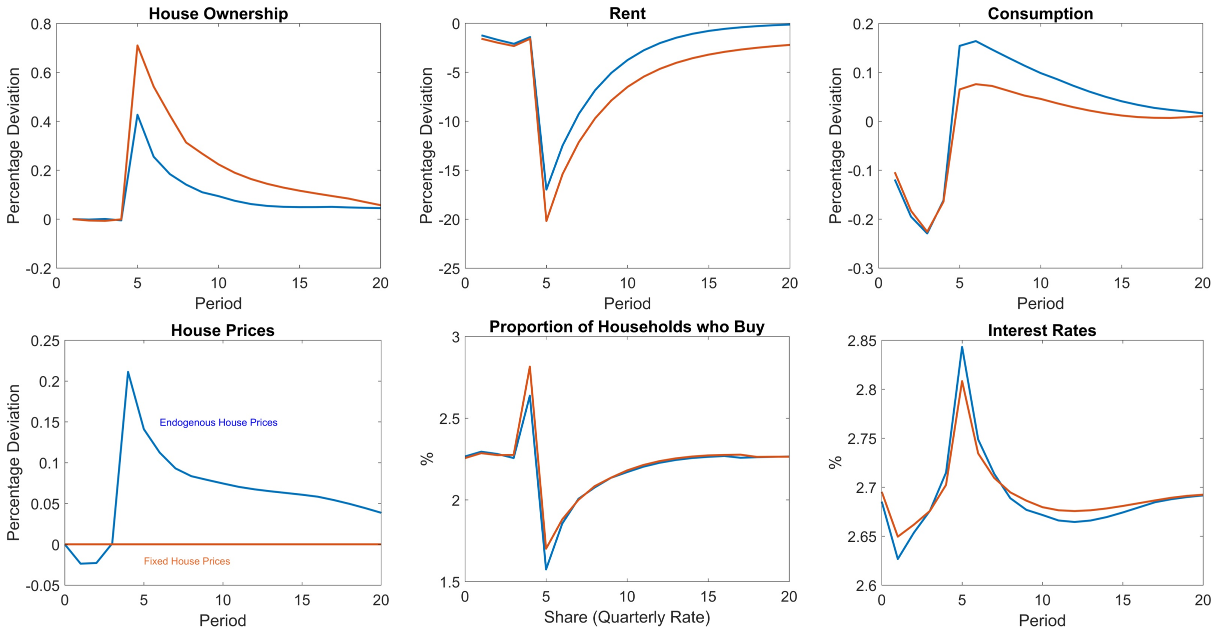

Note. These impulse response functions show the impact of an \$8,000 transfer to renters conditional on buying housing stock within the first year. The results are presented under the baseline assumption that the price of housing adjusts to clear the market, and the alternative assumption that the price is held constant.

temporary incentive to purchase a house thus results in a fall in consumption while households try to increase their stock of savings. When the tax credit is due to expire, renters' incentive to buy a house rather than delay and continue saving spikes, and demand for housing increases dramatically. This leads to an increase in purchases, house prices and interest rates on debt.

While home owners who are not able to access the tax credit directly, they are affected indirectly by the change in prices within the economy. The decline in house prices and wages and the increase in interest rates while the tax credit is available causes home owners to lower their non-durable consumption. Anticipating the increase in house prices that occurs just before the tax credit expires, home owners are incentivised to bring forward (delay) any increase (decrease) in their stock of housing, however this effect is quantitatively small. After the tax credit expires home owners consumption increases as the rise in house prices increases their wealth and the lower interest rate decreases their incentive to save.

Although liquid assets are also in zero net-supply, the variance of the distribution varies within the economy.[14] The tax credit encourages renters to take on high levels of leverage in order to purchase housing which increases the variance of the distribution of liquid wealth.

After the tax credit has expired the economy slowly returns to equilibrium. The overall result is to effect a shift in housing from the real estate sector to owners, due to the incentive provided by the temporary tax credit which pulls forward housing purchases from future periods. The incentive to pull forward the purchase of housing has the opposite effect on non-durable consumption, which is pushed back as households prioritise their investment in housing. Thus despite the stated intention of policy makers to stimulate the broader economy, I find that this tax credit has a contractionary impact on aggregate consumption and output.

Previous empirical studies of these programs (Dynan et al., 2013) found that the tax credit affected the price of housing, however they did not find a significant positive impact on the broader economy. These results suggest that a resolution to this puzzle is the tax credit merely encourages households to substitute consumption across time and goods, pulling forward the purchase of housing and pushing back the consumption of non-durable goods. This finding is consistent with Mian and Sufi (2012), who evaluated other temporary tax credits such as the "cash for clunkers" program, which offered tax credits for owners trading in fuel inefficient cars, and show that it had a limited effect on overall demand.

1.7 Conclusion

This chapter extends the two-asset Bewley model to allow for endogenous borrowing constraints determined by households' ownership of illiquid assets. I establish that the presence of financial frictions and transaction costs in the housing market generate important distortions in households' consumption of non-durable goods. In particular, I show that these frictions create an anticipatory saving effect. As the probability that a household will choose to adjust their stock of housing rises, their level of consumption will fall. This occurs because households anticipate that they

14 There is a small degree of variation in supply due to changes in the governments fiscal position.

will face higher interest rates and be credit constrained after making a debt-financed purchase of housing.

My results suggest that this effect is quantitatively meaningful. For households close to their optimal point of adjustment, an increase in liquid assets leads to a decrease in consumption as the anticipatory saving effect outweighs the wealth and liquidity effects. I provide empirical evidence for these results using micro-data from the PSID to show that a household's marginal propensity to consume is negatively correlated with their expectation of moving. I find that households which definitely expect that they will move have a negative marginal propensity to consume.

This result has important implications for the use of fiscal incentives in the housing market as a counter-cyclical policy tool. My general equilibrium model shows that tax credits for first home buyers decrease aggregate consumption due to an increase in mortgage debt and a larger incentive to save. These tax credits do boost prices and sales in the housing market, though much of this is the result of demand being pulled forward. While my analysis indicates that such policies lower output and employment, the financial system in my model is relatively simple. In the context of a bursting house price bubble and a financial crisis, a policy that helps to stabilise house prices may be welfare-improving despite the decrease in consumption due to higher household leverage.

More broadly, these results suggest that the considerable evidence of heterogeneity in households' marginal propensities to consume has significant consequences for the optimal design of macroeconomic policy. Beyond the implications for fiscal transfers my work on the relationship between debt and consumption raises new questions, for example how this mechanism may affect the transmission of monetary policy, which I leave as an avenue for future work.

Appendix

1.A Details on Model Solution

1.A.1 Model Approximation

I present the households' Hamilton-Jacobi-Bellman (HJB) equation and Kolmogorov forward equation for the evolution of the distribution of households μ over their holdings of liquid assets, b, housing stock, a, and income, z. We first solve for the HJB equation assuming that households choose not to adjust their stock of housing. The stationary version of the households' HJB equation is given by

$$\begin{aligned}
(\rho+\lambda) V^s(a,b,z) = & \max_{c,l,c^r,\alpha} u(c,h,l) \\
& +V_b^s\left[zw(1-\tau)l + r(b)b - c - c^{rent}p^{rent} + T\right] \qquad (1.A.1)\\
& +V_a^s(\delta^a a) + V_z^s(-\beta z) \\
& +\lambda \int_{-\infty}^{\infty} \left(V^s(a,b,x) - V^s(a,b,z)\right)\phi(x)\, dx,
\end{aligned}$$

where the value function is determined by the current flow of utility and the expected change in state variables.

The optimal drift in liquid assets is equal to the level of household savings which is governed by the budget constraint. Since we assume that households do not buy or sell housing stock the optimal drift in housing stock is driven solely by depreciation. Finally, the expected change in household income is determined by the mean reversion and the expectation of idiosyncratic "jumps" which cause log-income to jump to a new level drawn from normal distribution of density ϕ with variance σ^2.

We assume that if a household chooses to adjust their stock of housing that they will choose the optimal portfolio of assets subject to their budget constraint

$$V^{a}\left(a,b,z\right)=\max V^{s}\left(a',b',z\right) \tag{1.A.2}$$
$$\text{s.t.}$$
$$p^{a}a+b=p^{a}a'+b'-k\left(a,a'\right).$$

The household re-allocates their wealth across the two assets subject to paying the cost of adjustment, $k\left(a,a'\right)$.

The households' final value function is given by

$$V\left(a,b,z\right)=\max\left\{V^{a}\left(a,b,z,\right),V^{s}\left(a,b,z,\right)\right\}, \tag{1.A.3}$$

such that households choose to "stay" or "adjust" depending on which value function is larger. The adjustment function can then be defined as

$$I^{adjust}=\begin{cases}1 & \text{if } V\left(a,b,z\right)=V^{a}\left(a,b,z\right)\\ 0 & \text{if } V\left(a,b,z\right)=V^{s}\left(a,b,z\right),\end{cases} \tag{1.A.4}$$

and, conditional on paying the adjustment cost, the optimal portfolio, $a^{*}\left(a,b,z\right)$ and $b^{*}\left(a,b,z\right)$, is solved by maximising 1.A.2.

The households' optimal choice variables $\left(c_{t},c_{t}^{rent},l_{t}\right)$ are determined by this value function and the inverse of the utility function

$$c_{i,j,k}=\left(\frac{\partial u}{\partial c}\right)^{-1}\left[\left(\frac{\partial V\left(a,b,z,\right)}{\partial b}\right)\right] \tag{1.A.5}$$
$$c_{i,j,k}^{rent}=\left(\frac{\partial u}{\partial h}\right)^{-1}\left[\frac{p_{t}^{r}}{\theta^{r}}\left(\frac{\partial V\left(a,b,z,\right)}{\partial b}\right)\right] \tag{1.A.6}$$
$$l_{i,j,k}=\left(\frac{\partial u}{\partial l}\right)^{-1}\left[w_{t}\left(\frac{\partial V\left(a,b,z,\right)}{\partial b}\right)\right] \tag{1.A.7}$$

The evolution of the joint distribution of households is described by a Kolmogorov forward equation. I denote the density of households over the joint-distribution $g\left(a,b,z,t\right)$ corresponding to the density function $\mu_{t}\left(a,b,z\right)$. Furthermore I denote the optimal drift function for liquid assets as $s\left(a,b,z\right)$. Then the stationary density must satisfy the Kolmogorov forward equations

$$
\begin{aligned}
0 = & - \partial_b \left(s\left(a,b,z\right) g\left(a,b,z\right) \right) \\
& - \partial_a \left(\delta^a g\left(a,b,z\right) \right) + \int I^{adjust} \delta\left(a-a^*\right) \delta\left(b-b^*\right) \, d\mu_t \\
& - \partial_z \left(-\beta z g\left(a,b,z\right) \right) - \alpha g\left(a,b,z\right) + \alpha \phi\left(z\right) \int_{-\infty}^{\infty} g\left(a,b,x\right) \, dx \\
& - \lambda g\left(a,b,z\right) + \lambda \delta\left(a-a_0\right) \delta\left(b-b_0\right) \delta\left(z-z_0\right),
\end{aligned}
\tag{1.A.8}
$$

where δ is the Dirac delta function and (a_0, b_0, z_0) are the birth level of assets and income.

This equation equates inflows and outflows of households across each point in the joint-distribution. The first line measures the outflow as households increase (decrease) their stock of liquid assets through saving (dis-saving). The second line relates to changes in households' stock of housing via depreciation and the decision to buy or sell. The third line describes how exogenous changes in household productivity impact the distribution. Finally, the stochastic death and the birth of households, with zero wealth is described in the last line.

1.A.2 Steady State

I solve (1.A.1) and (1.A.8) expanding on the finite difference method outlined by Achdou et al. (2017). To solve the HJB function for households who choose not to adjust their stock of housing, I approximate the function $V^s(a,b,z)$ on $I \times J \times K$ discrete equi-spaced points. I denote the grid points $b_i, i = 1, ...I$, $a_j, j = 1, ..., J$, $z_k, k = 1, ...K$, and define the discretised value function for households that do not adjust their stock of housing

$$v_{i,j,k} = V^s\left(a_j, b_i, z_k\right). \tag{1.A.9}$$

I approximate the derivatives for b with either a forward or backward difference approximation

$$
\begin{aligned}
\frac{\partial v_{i,j,k}}{\partial b} &\approx v_b^F = \frac{v_{i+1,j,k} - v_{i,j,k}}{\Delta b_i^+} \\
\frac{\partial v_{i,j,k}}{\partial b} &\approx v_b^B = \frac{v_{i,j,k} - v_{i-1,j,k}}{\Delta b_i^-},
\end{aligned}
\tag{1.A.10}
$$

and use an upwind method to choose when to use each approximation. I use the forward (backward) approximation whenever the drift of b is positive (negative). Since the drift terms for housing and income are always negative (due to depreciation and mean reversion) I use the backward difference approximation universally for these variables. As outlined by Achdou et al. (2017), this approximation allows us to solve for the discretised value function using a process of iteration.

$$
\begin{aligned}
\frac{v_{i,j,k}^{n+1} - v_{i,j,k}^{n}}{\Delta^{step}} + (\rho + \lambda)\, v_{i,j,k}^{n+1} \quad = \quad & u\left(c_{i,j,k}^n, h_{i,j,k}^n, l_{i,j,k}^n\right) \\
& + v_b^n \left[z_t w_t (1-\tau)\, l_{i,j,k}^n + r(b_i)\, b_i - c_{i,j,k}^n - c_{i,j,k}^{rent,n} p_t^{rent} (1+\tau_t) \right] \\
& + v_a^n (\delta^a a_j) + v_z^n (-\beta z_k) + \lambda \int_{-\infty}^{\infty} \left(v_{i,j,x}^n - v_{i,j,k}^n\right) \phi(x)\, dx,
\end{aligned}
\tag{1.A.11}
$$

which can in turn be written as

$$
\left[\left(\frac{1}{\Delta^{step}} + \rho + \lambda\right) \boldsymbol{I} - \boldsymbol{A}\right] v_{i,j,k}^{n+1} = u\left(c_{i,j,k}^n, h_{i,j,k}^n, l_{i,j,k}^n\right) + \frac{1}{\Delta^{step}} v_{i,j,k}^n, \tag{1.A.12}
$$

where $\left(c_{i,j,k}^n, h_{i,j,k}^n, l_{i,j,k}^n\right)$ are the choice variables optimised by the households under the value function $v_{i,j,k}^n$ and $\boldsymbol{A^n}$ is a $(I \times J \times K)^2$ square matrix which summarises the exogenous productivity process, depreciation of housing stock and optimal household saving.

The HJB equation is then solved for with the following algorithm:

1. Start with a guess for the value function $v_{i,j,k}^0$.
2. Compute the derivatives of $v_{i,j,k}^n$ using (1.A.10) and the upwind approximation.
3. Compute the optimal household choices $\left(c_{i,j,k}^n, h_{i,j,k}^n, l_{i,j,k}^n\right)$ using the first order conditions X and the approximated derivatives of $v_{i,j,k}^n$.
4. Find $v_{i,j,k}^{n+\frac{1}{2}}$ from (1.A.12).

5. Compute the value function for households conditional of moving with

$$\left(v^*_{i,j,k}\right)^{n+\frac{1}{2}} = \max_{a',b'} v^{n+\frac{1}{2}}_{i,j,k} \text{ s.t. } p^a a + b = p^a a' + b' - k\left(a, a'\right). \tag{1.A.13}$$

6. Given $v^{n+\frac{1}{2}}_{i,j,k}$ and $\left(v^*_{i,j,k}\right)^{n+\frac{1}{2}}$ calculate

$$v^{n+1}_{i,j,k} = \max\left\{v^{n+\frac{1}{2}}_{i,j,k}, \left(v^*_{i,j,k}\right)^{n+\frac{1}{2}}\right\} \tag{1.A.14}$$

7. Repeat steps 2 through 6 until the difference between $v^{n+1}_{i,j,k}$ and $v^{n}_{i,j,k}$ is below a given threshold.

Once the discretised value function has been calculated we can solve the Kolmogorov Forward equation. In the absence of adjustments in the stock of housing, the joint distribution can be conveniently be calculated by the linear system

$$0 = \boldsymbol{A}' g, \tag{1.A.15}$$

where $\boldsymbol{A}$ is the transpose of the transition matrix in (1.A.12) as shown by Achdou et al. (2017), with a minor adjustment for the stochastic birth and death process.[15]

To account for adjustments in the stock of housing we need to modify this transition matrix to

$$\boldsymbol{C}_{m,n} = \begin{cases} 0, & \text{if } I^{adjust}\left(n\right) = 1 \\ \boldsymbol{A}_{m,n} + \sum_{I^{adjust}(n)=1} \boldsymbol{A}_{m,n} & \text{if } I^{adjust}\left(n\right) = 0, \quad m'\left(m\right) = n \\ \boldsymbol{A}_{m,n} & \text{if } I^{adjust}\left(n\right) = 0, \quad m'\left(m\right) \neq n, \end{cases} \tag{1.A.16}$$

where m' is the optimal adjustment target conditional on a household deciding to change its stock of housing. I^{adjust} is a vector which indicates if it is optimal for a household to adjust.

The first row indicates that any point n that lies in the adjustment region will have a density of 0. Instead, the second row says that any households who would otherwise transition into the adjustment region instead move to the optimal adjustment target.

15 Specifically, I add, λ to each element of the row associated with the birth position (a_0, b_0, z_0) and subject λ from each diagonal element of $\boldsymbol{A}$.

Finally, any point that is not in the adjustment region, nor the optimal target for a point in the region, remains the same as the original transition matrix.[16] The density is thus given by

$$0 = \boldsymbol{C}' g.$$

where $g(a,b,z)$ is normalised such that

$$1 = \int_0^{a_J} \int_{b_1}^{b_I} \int_{z_1}^{z_K} g(a,b,z)\; da\, db\, dz. \tag{1.A.17}$$

Once we have solved for the density of households, $g(a,b,z)$, we can easily solve for the aggregate demand for liquid, B_t, and illiquid, A_t, assets

$$\begin{aligned} B_t &= \int_0^{a_J} \int_{b_1}^{b_I} \int_{z_1}^{z_K} b g(a,b,z)\; da\, db\, dz \\ A_t &= \int_0^{a_J} \int_{b_1}^{b_I} \int_{z_1}^{z_K} a g(a,b,z)\; da\, db\, dz. \end{aligned} \tag{1.A.18}$$

1.A.3 Transition Dynamics

Outside of the steady state, when changes in prices, wages or interest rates vary over the time, the time-dependent HJB equation is required

$$\begin{aligned} (\rho+\lambda) V^s(a,b,z,t) &= \max_{c,l,c^r,\alpha} u(c,h,l) \\ &\quad + V_b^s \left[z_t w_t (1-\tau) l + r_t(b) b - c - c^{rent} p_t^{rent} + T_t \right] \\ &\quad + V_a^s (\delta^a a) + V_z^s(-\beta z) \\ &\quad + \lambda \int_{-\infty}^{\infty} \left(V^s(a,b,x) - V^s(a,b,z) \right) \phi(x)\; dx \\ &\quad + V_t^s(a,b,z,t), \end{aligned} \tag{1.A.19}$$

where the last term accounts for changes in the value function over time. Solving this function over the same discrete grid as in the steady state results in the approximation

16 To ensure the matrix $\boldsymbol{C}$ is not singular we also set $\boldsymbol{C}_{m,m} = -\varepsilon$ for some small $\varepsilon > 0$, for all $I^{adjust}(m) = 1$. This ensures that all points in the adjustment region have a defined density of zero.

$$\begin{aligned}
&(\rho+\lambda)\, v_{i,j,k}^{n} = \\
&u\left(c_{i,j,k}^{n}, h_{i,j,k}^{n}, l_{i,j,k}^{n}\right) \\
&+v_{b}^{n}\left[z_{t} w_{t}(1-\tau)\, l_{i,j,k}^{n}+r\left(b_{i}\right) b_{i}-c_{i,j,k}^{n}-c_{i,j,k}^{rent,n} p_{t}^{rent}+T_{t}\right] \\
&+v_{a}^{n}\left(\delta^{a} a_{j}\right)+v_{z}^{n}\left(-\beta z_{k}\right)+\lambda \int_{-\infty}^{\infty}\left(v_{i,j,x}^{n}-v_{i,j,k}^{n}\right) \phi(x)\, d x \\
&+\frac{v_{i,j,k}^{n+1}-v_{i,j,k}^{n}}{\Delta t},
\end{aligned} \tag{1.A.20}$$

where $v_{i,j,k}^{n}$ is now short-hand for $v\left(a_{j}, b_{i}, z_{k}, t^{n}\right)$ and the terminal condition $v_{i,j,k}^{N}$ is equal to the steady state solution $v_{i,j,k}$. This can be written as

$$\left(\rho+\lambda-\left[\boldsymbol{A}^{n+1}\right]+\frac{1}{\Delta t}\right) v_{i,j,k}^{n}=u\left(c_{i,j,k}^{n+1}, h_{i,j,k}^{n+1}, l_{i,j,k}^{n+1}\right)+\frac{1}{\Delta t}\left(v_{i,j,k}^{n+1}\right), \tag{1.A.21}$$

which yields a solution for $v_{i,j,k}^{n}$ given $v_{i,j,k}^{n+1}$.

Similarly the joint distribution may vary over time, thus outside of steady state the time-dependent Kolmogorov equation is

$$\begin{aligned}
\partial_{t} g(a, b, z, t)= & -\partial_{b}(s(a, b, z)\, g(a, b, z)) \\
& -\partial_{a}\left(\delta^{a} g(a, b, z)\right)+\int I^{adjust} \delta\left(a-a^{*}\right) \delta\left(b-b^{*}\right) d \mu_{t} \\
& -\partial_{z}(-\beta z g(a, b, z))-\alpha g(a, b, z)+\alpha \phi(z) \int_{-\infty}^{\infty} g(a, b, x)\, d x \\
& -\lambda g(a, b, z)+\lambda \delta\left(a-a_{0}\right) \delta\left(b-b_{0}\right) \delta\left(z-z_{0}\right).
\end{aligned} \tag{1.A.22}$$

Solving this equation over the same discrete grid requires a two-step n

$$g^{n+1}=\left(\boldsymbol{I}-\Delta t\left(\boldsymbol{C}^{\boldsymbol{n}}\right)^{\prime}\right)^{-1} g^{n}, \tag{1.A.23}$$

where $\boldsymbol{C}^{n}$ is the transition matrix from (1.A.21) adjusted for households who choose to adjust their stock of housing.

Once we have solved for the discreteised steady state value function, $v(a, b, z)$, and density $g(a, b, z)$ it is relatively straightforward to solve for transition dynamics in response to a “MIT” shock - an unanticipated shock followed by a deterministic transition back to the (potentially new) steady state. To solve for the transition path we use the following algorithm.

1. Guess a path for wages, w_t, interest rates, r_t^b, and house prices, p_t^a.
2. Given the path of wages, interest rates and prices solve the HJB backwards in time, using $v_{i,j,k}$ as the terminal condition.
3. Using this time path for the value function, calculate the optimal rules for household consumption and housing stock adjustment for each time period.
4. Given the optimal saving and adjustment rules, solve the time-dependent Kolmogorov forward equation, using the steady state as the initial condition g^0.
5. Given the evolution of the density of households, calculate the aggregate household supply of labour and demand for liquid and illiquid assets over time.
6. Update the path for wages, interest rates and house prices such that
   a) $w_t' = w_t - \xi_w \left(N_t + N_t^c - \int zl\left(a,b,z,t\right)\, d\mu_t\right)$, such that the wage falls (rises) when there is excess labour supply (demand).
   b) $\left(r_t^b\right)' = r_t^b - \xi_r \left(\int b\, d\mu_t - p_t^A C_t^{rent} + B_t^{gov}\right)$, such that the interest rate falls (rises) when there is an excess supply (demand) of liquid assets.
   c) $\left(p_t^a\right)' = p_t^a - \xi_a \left(IA_t - \delta A_t + \int \dot{a}_t + \dot{c}_t^{rent}\, d\mu_t\right)$, such that the price of houses falls (rises) when there is an excess supply (demand) of housing.
7. Repeat steps 2 through 6 until the sum of the changes in the three time path of prices $\left(w_t, r_t^b, p_t^a\right)$ converges.

1.B Marginal Propensities to Consume

In this continuous time model I follow Achdou et al. (2017) in their definition of the marginal propensity to consume.

Definition 1. The marginal propensity to consume over a period τ for an individual with the state vector (a, b, z) is given by

$$MPC_\tau^x\left(a,b,z\right) = \frac{C_\tau\left(a,b+x,z\right) - C_\tau\left(a,b,z\right)}{x}, \tag{1.B.1}$$

where $C_\tau(a,b,z)$ is the sum of expected consumption of an individual household over the next τ periods.

$$C_\tau(a,b,z) = E\left[\int_0^\tau c(a_t,b_t,z_t)\,dt|a_0 = a, b_0 = b, z_0 = z\right] \tag{1.B.2}$$

This conditional expectation can be calculated using the Feynman-Kac formula. This formula links the conditional expectations of a stochastic process and the solution to partial differential equations.

1.C Alternate Model Specifications

1.C.1 Partial Borrowing Constraints

In the baseline version of the model I assumed that mortgage debt taken on by households is taken on a rolling basis. This implies that fluctuations in the price of housing will change the endogenous borrowing constraint faced by households. However, as discussed by Iacoviello and Pavan (2013), the assumption that the collateral constraint binds at all times is at odds with the reality that loan-to-value constraints only bind for agents that refinance. In practice, falls in the price of housing do not tighten the credit constraints of the majority who choose not to refinance. To better match the lagged correlation between measures of aggregate debt and house prices I relax this assumption to see how it affects the baseline results.

Concretely, I assume that the endogenous credit constraint applies directly to households that choose to adjust their stock of housing. For households who do not choose to adjust their stock of housing, I assume that stock of debt can be maintained at its current level even if the value of their housing stock falls due to a decline in house prices. This introduces an asymmetry into the relationship between house prices and the debt constraint. If house prices fall then households are not forced to immediately reduce their mortgage, but rather can maintain the size of the debt making repayments that cover the interest only. If house prices rise then households have the option of refinancing their mortgage and increasing their stock of debt. I assume that a household can costlessly refinance its current stock of housing. Thus the borrowing constraint becomes

$$\begin{aligned} b_t &\geq -M a_t p_t^a && \text{if } I_{\alpha_t} = 1 \\ b_t &\geq \min\left\{-M a_t p_t^a, b_t\right\} && \text{if } I_{\alpha_t} = 0. \end{aligned} \tag{1.C.1}$$

This changes both the steady state of the model, and its response to shocks around the equilibrium. While the price of housing is constant in the steady state, this alternative borrowing constraint loosens restrictions on households for whom the depreciation of their stock of housing would otherwise cause the debt constraint to bind. In equilibrium this relaxation allows households to take on greater levels of debt due to the reduced probability of being credit constrained. Accordingly, the equilibrium interest rate is lower.

The impact of this assumption on the estimated impact of the tax credit is marginal as house prices generally increase in response to the policy.

1.C.2 Heterogeneous Labour Supply

Previous research has indicated that households vary labour supply in response to changes in their balance sheet (Attanasio et al., 2012). In the baseline model, labour supply is invariant to a households' asset portfolio. Here I relax this assumption by considering a utility function in which the disutility from supplying labour is separable from consumption

$$u_t = \frac{\left(c_t^{1-\zeta} h_t^{\zeta}\right)^{1-\gamma}}{1-\gamma} - \phi z_t \frac{l_t^{1+\frac{1}{\eta}}}{1+\frac{1}{\eta}}. \tag{1.C.2}$$

Under this alternative function for flow utility, a household's optimal supply of labour will set by

$$z_t w_t \frac{\partial u_t}{\partial c_t} = -\frac{\partial u_t}{\partial l_t} \tag{1.C.3}$$

which in turn implies that each household's optimal labour supply $l_t = l_t\left(a, b, z\right)$ will vary according to the household's asset portfolio, though not with the productivity shock. To reflect the change in household preferences I re-calibrate the parameters within the utility function to match their original targets.

When I assume that labour supply varies with a households' assets, I create a second margin that households can adjust to respond to the change in debt that occurs when they adjust their stock of housing. Specifically, households increase their supply of labour in anticipation of, and after, buying additional stock. The reasoning behind this choice is the same as that of non-durable consumption, namely that households increase their labour supply to reduce the probability they will become bound by their credit constraint and reduce the time spent paying the higher interest rate on their mortgage. The same logic applies in reverse to households who anticipate decreasing their stock of housing.

Despite this second margin which households can use to offset fluctuations in the level of their non-durable consumption, the effect on marginal propensities to consume is quantitatively small given reasonable calibrations for the Frisch elasticity of supply.

1.D Rational Expectations and Instrumental Variables

Proposition 2. *Given a lagged instrument, an IV regression will produce a consistent estimator of the marginal effect of the expected change in the endogenous variable, ignoring the marginal effect of contemporaneous changes.*

Proof. Consider an endogenous variable which can be decomposed into an expected component based on information available at time $t-2$ and an unexpected component, ξ_t

$$X_t = E_{t-2}X_t + \xi_t \tag{1.D.1}$$

Such that

$$E_{t-2}(\xi_t) = 0 \tag{1.D.2}$$

Assume that in the true data generating process each component may have a different impact on the dependent variable, Y_t, such that

$$Y_t = \beta_1 E_{t-2}X_t + \beta_2 \xi_t + u_t \tag{1.D.3}$$

Now assume there is an instrument available, Z_{t-2}, that I wish to use to identify the marginal effect of the expected component, β_1, that is

Informative

$$E\left(Z_{t-2}E_{t-2}X_t\right) \neq 0 \tag{1.D.4}$$

Valid

$$E\left(Z_{t-2}u_t\right) = 0 \tag{1.D.5}$$

Only comprises information known at time $t-2$ such that

$$E\left(Z_{t-2}\xi_t\right) = 0 \tag{1.D.6}$$

Consider the regression of Y_t on X_t where X_t is instrumented with Z_{t-2}

$$\hat{\beta}_{IV} = \frac{\frac{1}{N}\sum_{i=1}^{N}\left(Z_{t-2}\right)\left(Y_t\right)}{\frac{1}{N}\sum_{i=1}^{N}\left(Z_{t-2}\right)\left(X_t\right)} \tag{1.D.7}$$

Therefore

$$\begin{aligned} \mathrm{p}\underset{N\to\infty}{lim}\hat{\beta}_{IV} &= \frac{\frac{1}{N}\sum_{i=1}^{N}\left(Z_{t-2}\right)\left(Y_t\right)}{\frac{1}{N}\sum_{i=1}^{N}\left(Z_{t-2}\right)\left(X_t\right)} \\ &= \frac{p\underset{N\to\infty}{lim}\beta_1\ \frac{1}{N}\sum_{i=1}^{N}Z_{t-2}E_{t-2}X + \beta_2\frac{1}{N}\sum_{i=1}^{N}Z_{t-2}\xi_t + \frac{1}{N}\sum_{i=1}^{N}Z_{t-2}u_t}{\mathrm{p}\underset{N\to\infty}{lim}\frac{1}{N}\sum_{i=1}^{N}\left(Z_{t-2}\right)\left(E_{t-2}X_t\right) + \ \mathrm{p}\underset{N\to\infty}{lim}\frac{1}{N}\sum_{i=1}^{N}\left(Z_{t-2}\right)\left(\xi_t\right)} \end{aligned} \tag{1.D.8}$$

By the law of large numbers

$$\begin{aligned} \mathrm{p}\underset{N\to\infty}{lim}\frac{1}{N}\sum_{i=1}^{N}\left(Z_{t-2}\right)E_{t-2}X &= E\left(\left(Z_{t-2}\right)E_{t-2}X\right) \neq 0 \\ \mathrm{p}\underset{N\to\infty}{lim}\frac{1}{N}\sum_{i=1}^{N}\left(Z_{t-2}\right)\left(\xi_t\right) &= E\left(\left(Z_{t-2}\right)\left(\xi_t\right)\right) = 0 \\ \mathrm{p}\underset{N\to\infty}{lim}\frac{1}{N}\sum_{i=1}^{N}\left(Z_{t-2}\right)u_t &= E\left(\left(Z_{t-2}\right)u_t\right) = 0 \end{aligned}$$

Therefore

$$\mathrm{p}\underset{N\to\infty}{lim}\hat{\beta}_{IV} = \beta_1 \tag{1.D.9}$$

Thus assuming that the instrument is valid and informative the IV estimator will be a consistent estimator of the marginal effect of the expected change in the independent variable. □

1.E Robustness of Empirical Analysis

In Table 1.10, I consider several alternative specifications for the regression of consumption on the change in household expectations. In Table 1.10, I show my results are robust over several alternative specifications for the regression of consumption on the change in household expectations. Finally, in Table 1.11, I consider several alternative specifications for the regression of consumption on the change in household expectations.

Table 1.10: Alternate Regressions

	(6) Baseline	(7) Exclude food stamps	(8) Pre-tax earnings	(9) Continuously married only	(10) Male heads only	(11) Only mortgage debt	(12) Include state unemployment
ΔDebt-Asset$_{i,t}$	-0.05 *	-0.07 **	-0.04	-0.08 *	-0.07 *	-0.11 **	-0.03 *
ΔFamily Size$_{i,t}$	0.01	0.02	0.01	0.01	0.01	0.01	0.01
r_{t-2}	-0.09***	-0.11***	-0.07*	-0.08***	-0.13***	-0.10**	-0.04*
$\Delta \log y_{i,t}$	-0.01	0.02	-0.01	-0.02	-0.03	-0.00	0.01
$\Delta \log W_{i,t}$	-0.00	0.01	-0.00	-0.00	0.01	0.02	-0.01
ΔState Unemployment$_{i,t}$							-0.05
Observations	37,560	37,560	37,560	21,242	26,185	37,560	37,560

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 1.11: Changes in Household Expectations

		Baseline	Exclude food stamps	Pre-tax earnings	Continuously married only
Change in household expectations	No expectation to Definitely	-0.073^{***}	-0.088^{***}	-0.069^{***}	-0.053^{*}
	Uncertain to Definitely	-0.063^{***}	-0.060^{**}	-0.055^{***}	0.013
	Probably to Definitely	-0.048^{***}	-0.052^{**}	-0.072^{***}	-0.022
	Definitely to Definitely	-0.057^{***}	-0.046^{*}	-0.007	-0.07

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

1.E.1 Durable good consumption

In Table 1.12, I consider the inclusion of durable good consumption. Similar to Benmelech et al. (2017) I find that consumption of durable goods, in particular house furnishings, increases when a household purchases a house. However I find that total consumption still falls when durable goods are included in the consumption regression. This difference may be due to omitted variable bias such as contemporaneous shocks to financial wealth or income which are controlled for under the instrumental variable method. I also consider the possibility that the future expectations regarding the decision to purchase a house are correlated with changes in household size. To test if this correlation impacts the results I re-run the regressions on sub-samples of the population limited to those who do not report a change in family size or households over the age of 40 for whom the average fertility rate is significantly lower. Table 1.13 shows that these alternative samples do not materially change the main results of the chapter.

Table 1.12: Alternate Regressions - Durable Goods

		All Buyers			First-time buyers			Repeat Buyers	
	Non-durable	Durable	Total	Non-durable	Durable	Total	Non-durable	Durable	Total
$\text{Buy}_{i,t}$	-0.11 ***	0.48 ***	-0.08 ***	-0.07 *	1.922 ***	-0.01	-0.62 *	-0.02	-0.02
$\Delta\text{Family Size}_{i,t}$	0.02	0.01	0.01	0.02	0.01	0.01	0.01	0.01	0.01
r_{t-2}	-0.07***	-0.01***	-0.08*	-0.08***	-0.03***	-0.10**	-0.04*	-0.01***	-0.08*
Observations	37,560	31,560	31,560	37,560	31,560	31,560	37,560	31,560	31,560

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 1.13: Alternate Regressions - Household Size

	All Households		Age > 40		Family size unchanged	
	Non-durable	Total	Non-durable	Total	Non-durable	Total
$\text{Buy}_{i,t}$	-0.11 ***	-0.08 ***	-0.15 ***	-0.14 ***	-0.12 ***	-0.10 ***
$\Delta\text{Family Size}_{i,t}$	0.02	0.01	0.01	0.01	0.01	0.00
r_{t-2}	-0.07***	-0.08*	-0.08***	-0.10**	-0.04*	-0.08*
Observations	37,560	31,560	27,668	17,261	33,168	28,810

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

2 Optimal Monetary Policy Rules under Downward Nominal Wage Rigidity[1]

Abstract

I estimate optimal monetary policy in a non-linear New Keynesian model in which nominal wages are downwardly rigid. In this environment a higher rate of inflation gives workers more flexibility when setting real wages, which comes at the cost of greater price dispersion in the goods market. After outlining a numerical algorithm to solve the model with aggregate uncertainty I use micro-data on the distribution of workers' change in nominal wages to calibrate the wage rigidity. I find that downward nominal wage rigidities bend the Phillips curve constraining the inflation rate from falling in times of low demand. This indicates that an inflation rate that is only moderately below its target can mask large falls in the output gap. Finally, I show that the optimal monetary policy rule places a high weight on stabilising wage inflation, relative to a symmetrical wage friction, and that an asymmetric Taylor rule can improve household welfare.

JEL Classification: E31, E43, E58, J30.

Keywords: Optimal Monetary Policy, Downward Nominal Wage Rigidity.

1 I would like to thank seminar participants at the University of Oxford, and the European Meeting of Econometric Society, Theories and Methods in Macroeconomics Annual Conference, and International Symposium on Money Banking and Finance for useful comments. I would like to acknowledge the use of the University of Oxford Advanced Research Computing (ARC) facility in carrying out this work (http://dx.doi.org/10.5281/zenodo.22558).

2.1 Introduction

Nominal price and wage rigidities are an important transmission mechanism of monetary policy (Woodford, 2003). While these nominal rigidities have been modeled in a variety of ways, the most common methods are symmetric in nature, restricting nominal prices from rising or falling in equal measure. Assuming that nominal rigidities are symmetric makes general equilibrium models significantly more tractable. However, this approach ignores the strong empirical evidence that nominal wage rigidities are asymmetric; they are downwardly rigid, but upwardly flexible. Individual-level data shows that workers' nominal wages rarely fall, and that a large share of workers instead have their nominal wages frozen. This chapter will ask two questions: how do downwardly rigid nominal wages affect the economy, and what is optimal monetary policy in the presence of these asymmetric rigidities?

This chapter makes three contributions to the literature. First, I develop a general equilibrium model which micro-founds downwardly rigid nominal wages using a time-dependent wage friction following the framework of Calvo (1983). Variation in workers' wages is induced by an idiosyncratic shock to the disutility of labour. The combination of nominal price and wage rigidities generates a meaningful welfare trade-off when choosing the optimal monetary policy rule. A social planner must balance the welfare costs of price dispersion in the goods market with the cost of distorted real wages. I use micro-data on workers' change in wages to calibrate newly introduced parameters in the model which govern how tightly the nominal wage rigidity binds and the standard deviation of workers' idiosyncratic preference shock.

The second contribution is an algorithm for numerically solving both the equilibrium and the non-linear dynamics of the model. I develop a method for computing equilibrium in heterogeneous agent models with aggregate uncertainty using impulse response functions as a numerical derivative to solve for the solution. This process is analogous to that of Boppart, Krusell, and Mitman (2017), who use a set of impulse response functions to calculate a linearised solution. However I extend this approach by showing how second order approximations can be calculated using a grid of impulse response functions. This higher order approximation allows the method to incorporate the non-linearities caused by the asymmetric wage rigidity. I find that downwardly rigid nominal wages cause output to be inefficiently low. Intuitively this is because the rigidity prevents workers from lowering their nominal wage when it is otherwise desirable to do so. Inefficiently high wages decrease firms' demand for labour, lowering employment and welfare within the model.

Finally, by simulating the model under different monetary policy rules, I solve for optimal policy in the presence of the downward nominal wage rigidity.

I find two main results concerning the problem of optimal monetary policy. The first is that central banks should place a higher weight on wage inflation, relative to price inflation, when setting monetary policy in comparison to when wages are flexible or modeled as symmetrically rigid. This is because the welfare cost of wage inflation volatility outweighs that of price inflation. Secondly, I find the central bank should react more aggressively when inflation is below its target, compared with when it is above target.

Previous studies analysing monetary policy where aggregate wages are downwardly rigid generally use one of two methods. The first approach assumes that all workers are identical and so have the same nominal wage.[2] Under this assumption the downward nominal wage constraint will not bind any workers if the wage inflation rate is positive, and will constrain all workers when it is negative. However, this approach ignores the widespread empirical evidence on downward nominal wage rigidities, which finds that a significant share of workers' wages are constrained by the nominal rigidities even when the inflation rate is high.[3]

The second approach to modeling downwardly rigid wages relaxes this assumption and instead assumes that an idiosyncratic shock which causes individual workers to choose different nominal wages. This approach is more complex to solve as the distribution of wages, rather than the mean level, is a state variable. However these models are better able to match the key patterns in empirical data on workers' nominal wage changes. Due to this complexity, previous studies that have taken this approach have

2 Examples of this approach include Carlsson and Westermark (2008); Fahr and Smets (2010); Kim and Ruge-Murcia (2009); Schmitt-Grohé and Uribe (2012, 2013) and Abo-Zaid (2013b).

3 Individual-level studies of nominal wage rigidities include Altonji and Devereux (2000); Babecký et al. (2010); Card and Hyslop (1997); Elsby et al. (2013); Kahn (1997) and Lebow et al. (2003).

only solved for optimal monetary policy for the model's steady state, ignoring the dynamic properties of the model. This thesis will build off the framework of Daly and Hobijn (2014) and solve for the non-linear dynamics and optimal monetary policy of a model which includes an idiosyncratic shock to workers. It is thus the first paper, to the author's knowledge, to analyse optimal monetary policy rules when heterogeneous workers are subject to downwardly rigid nominal wages in a dynamic setting.

Relative to the previous literature on nominal rigidities, in particular Erceg et al. (2000) who analyse optimal monetary policy under symmetric price and wage rigidities, I find that optimal monetary policy places a greater weight on wage inflation and the output gap when wages are downwardly rigid. This is because the distortion caused by downwardly rigid wages is significantly more costly than that of a symmetric nominal wage rigidity.

In the next section I will outline the literature on downwardly rigid nominal wages and optimal monetary policy in the presence of nominal wage rigidities. Section 2.4 will describe the model and how I micro-found the downward nominal wage rigidity. Section 2.5 describes how I solve for the model's steady state using value function iteration. Section 2.6 outlines optimal monetary policy under the presence of downward nominal wage rigidities. Section 2.7 concludes.

2.2 Literature Review

This thesis contributes to the literature on both downward nominal wage rigidity and the formulation of optimal monetary policy. I will discuss the previous literature on each topic in turn.

The presence of downward nominal wage rigidities have been a long noted feature of the labour market dating back to Keynes (1936) who wrote that "[w]hilst workers will usually resist a reduction of money-wages [nominal wages], it is not their practice to withdraw their labour whenever there is a rise in the price of wage-goods [a reduction in real wages]". The reason behind this nominal rigidity is commonly ascribed to a combination of workers' money illusion, the confusion of real and nominal prices, and a perception that wage cuts are inherently unfair (Kahneman et al., 1991). The seminal investigation into why wages are downwardly sticky is Bewley (1999) who conducted extensive interviews with a range of executives, labour leaders and professional recruiters. Bewley reported that executives were averse to cutting the nominal wages of current employees, or even new hires, during economic downturns when demand for their products fell sharply. Cuts to nominal wages were thought to be highly damaging to worker morale and degrading to the firms' productivity, which relied on a high degree of cooperation with the workforce.

Empirical studies of downward nominal wage rigidities have focused on individual-level data, tracking the distribution of changes in nominal wages across a sample of workers. These empirical studies report three stylised facts about the distribution of changes in the nominal wage. The first is that there is a spike at zero, indicating that a large share of workers report an unchanged nominal wage over a 12 month period. This can be seen in Figure 2.1 which shows the reported change in hourly nominal wages in 2012 according to the Current Population Survey.

Figure 2.1: Percentage Change in Hourly Nominal Wages Over 1980-2015

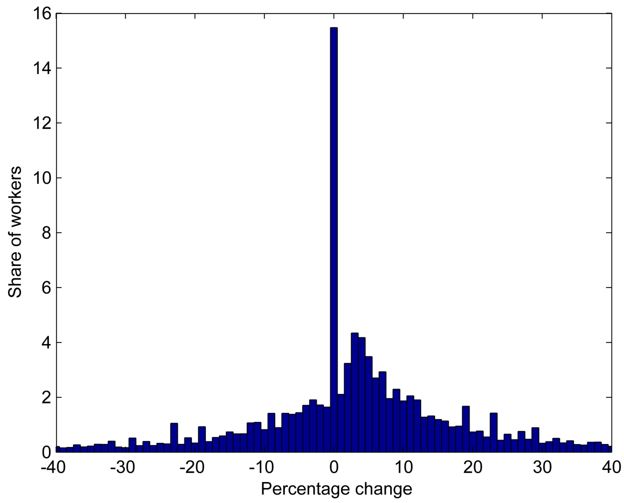

All workers (hourly and salary, job-stayers and job-switchers)
Source: Current Population Survey, Daly and Hobijn (2014)

The second stylised fact is that very few workers report a decrease in their nominal wage. This suggests that an aversion to lowering nominal wages is causing workers who might otherwise have been expected to receive a pay cut to receive a wage freeze instead. The third stylised fact is that the share of workers reporting an unchanged nominal wage is counter-cyclical and lower in times of high inflation.

Table 2.1 reports a range of empirical studies which attempt to measure the share of workers with an unchanged nominal wage over a 12 month period across countries and time periods. Several papers have questioned whether this spike may be caused by measurement error or noise. For example, respondents may round their wage to the nearest whole figure when surveyed. This will cause small changes in nominal wages to be misreported as a pay freeze. After controlling for these rounding problems and measurement errors, Lebow et al. (2003) calculate that almost 40 per cent of the spike at zero is due to rounding, while Smith (2000) argues that eliminating measurement error could cut the spike by half. Even after taking this into account, however, the empirical evidence still suggests that there is an out-sized proportion of workers whose nominal wages are unchanged.

Table 2.1: Empirical Results on Downward Nominal Wage Rigidities

Paper	Size of spike at 0 per cent	Notes
Card and Hyslop (1997)	13 - 17 %	CPS data, 1981-93, excludes workers who change industries
Card and Hyslop (1997)	8 - 11 %	PSID data, 1976-88, paid hourly, excludes workers who change jobs
Altonji and Devereux (2000)	6 - 22 %	PSID 1971-92, paid hourly, excludes workers who change jobs
Kahn (1997)	8 %	PSID data, 1970-1988, paid hourly
Lebow et al. (2003)	18 %	BLS data, job-level data (as opposed to individual).
Babecký et al. (2010)	10 %	ECB survey covering 15 countries. Across the countries the spike ranged from 3 to 27 %.
Elsby et al. (2013)	14 - 19 %	CPS data, 1982-2012, hourly workers
Daly and Hobijn (2014)	9 - 16 %	CPS data, 1984 - 2012, includes workers who change jobs

Sources: Current Population Survey, Panel Study of Income Dynamics, Bureau of Labor Statistics

Despite the weight of empirical evidence there are relatively few theoretical papers that explicitly model the impact of downward nominal wage rigidities on aggregate wages and employment. The theoretical literature generally takes one of two approaches when modeling the nominal rigidity.

The first approach assumes that all households are identical, varying only by the type of labour they supply. In these models all households have the same nominal wage and choose to supply the same amount of labour in each period. When nominal wages are constrained from falling, the real wage will be inefficiently high. The high real wage will cause firms' demand for labour to decrease, lowering both employment and welfare in the model.

The assumption of homogeneous households makes models more tractable. However this approach ignores key stylised facts about the labour market. It assumes the downward nominal wage rigidity is not binding when wage inflation is high, or indeed positive. This implies that the nominal wage constraint will have no impact on the model's steady state.[4] However empirical data suggests that there is always a significant share of workers who are constrained by downward nominal wage rigidity even in periods of high inflation (Altonji and Devereux, 2000). By assuming that the nominal wage rigidity only constrains workers in times of low inflation this method underestimates the impact that downward nominal wage rigidities will have on aggregate wage growth and unemployment.

The second approach to modeling downwardly rigid wages relaxes this assumption. Instead of assuming all households are identical, they are differentiated by an idiosyncratic utility shock (Daly and Hobijn, 2014) or productivity shock (Fagan and Messina, 2009; Benigno and Ricci, 2011). This idiosyncratic shock causes workers to optimally set different nominal wages generating a non-degenerate distribution of wages within the economy. Moreover, it causes a fraction of individual workers to optimally choose to lower their nominal wage, even when aggregate wage inflation is positive. Models with differentiated households are more complex to solve as the distribution of wages becomes a state variable, rather than just the aggregate wage level. However, these models are better able to match the key patterns highlighted in the data. Under this assumption there is always a share of workers who are subject to the downward nominal wage constraint even when aggregate wage inflation is high. In addition this approach matches the cyclical trend of the share of workers constrained by the downward nominal wage rigidities increasing during downturns and periods of

4 Assuming a non-negative rate of inflation in equilibrium.

low inflation.

Relaxing this assumption, however, comes at a cost. In general these models can not be solved analytically, instead requiring computationally expensive numerical techniques. Because of this complexity, most papers which take this approach only solve for the model's steady state which limits the types of questions the models can be used to answer.

Examples of the first approach to modeling optimal monetary policy with downward nominal wage rigidities include Schmitt-Grohé and Uribe (2012) and Schmitt-Grohé and Uribe (2013). In these papers Schmitt-Grohé and Uribe use a somewhat ad-hoc approach to modeling downwardly rigid nominal wages. They constrain aggregate nominal wages, W_t, such that

$$W_t \geq \gamma W_{t-1}$$

where γ governs the strength of the downward nominal wage rigidity. The strength of the nominal wage rigidity can either be exogenous, as in Schmitt-Grohé and Uribe (2012), where γ is a parameter which is calibrated, or an endogenous function as in Schmitt-Grohé and Uribe (2013) where it varies with the state of the labour market.

The second approach to modeling downward nominal wage rigidities, using differentiated households with non-identical wages, was pioneered by Benigno and Ricci (2011) who modeled a continuum of households who choose their wages subject to a downwardly rigid nominal wage constraint and both idiosyncratic sectoral shocks and aggregate shocks. They then solve for the model's equilibrium and analyse the long-run relationship between inflation and output. They find that the long-run Phillips curve, relating the output gap to wage inflation, is convex, being vertical at high inflation rates and flatter when wage growth is subdued. They then discuss how a higher rate of inflation can be used to unwind this nominal friction and increase welfare in equilibrium.

Daly and Hobijn (2014) extend the approach of Benigno and Ricci (2011) to solve for the model's short-run Phillips Curve. Their model contains a continuum of differentiated households that are subject to both an idiosyncratic preference shock and the downwardly rigid nominal wage constraint. They use value function iteration to solve for the dynamics of the model when subject to aggregate demand shocks and aggregate supply shocks. They find that the wage constraint bends the short-run Phillips Curve placing a floor on wage inflation even in response to large adverse demand (or expansionary supply) shocks.

The seminal work on optimal monetary policy in the presence of rigid wages is Erceg et al. (2000). They extend the basic New Keynesian framework and solve it with both wage and price frictions, modeled according to Calvo (1983), and analyse optimal monetary policy in the model. They find that the optimality of strict inflation targeting (also know as the "divine coincidence") breaks down and that the first best allocation can no longer be replicated using monetary policy. Optimal policy can be implemented by a policy rule which strictly targets a weighted average of price and wage inflation, with the optimal weighting determined by the relative stickiness of the respective prices within the model. The more rigid prices are, relative to wages, the more weight is placed on price inflation in the optimal weighted average. In addition they find that targeting the output gap can be a reasonable second best policy which is robust to errors in estimating the relative rigidities.

Schmitt-Grohé and Uribe (2013) model downwardly rigid nominal wages as an occasionally binding constraint, restraining aggregate nominal wages from falling. They then embed this friction within a small open economy to analyse the welfare implications of currency pegs in the context of emerging economies. They find that this friction can create an asymmetry within the business cycle, such that recessions cause labour to be demand determined due to the inability of wages to fall, while the labour market clears during booms. They report that optimal exchange rate policy can alleviate unemployment and call for large devaluations during external crises. However, they abstract from the problem of domestic monetary policy by treating the real interest rate as purely exogenous.

2.3 The Model

This section extends the model of Daly and Hobijn (2014) by incorporating monopolistic competition and a Calvo-style pricing friction among the firms. This creates a welfare cost to increasing the rate of inflation creating a meaningful trade-off when setting monetary policy.

2.3.1 Households

The economy is populated by a representative household comprised of a unit measure of workers who share consumption among themselves, but each supply a differentiated type of labour, $L_{i,t}$, in exchange for a nominal wage, $W_{i,t}$. Workers are differentiated by an idiosyncratic, and uninsurable, shock to the disutility derived from working $Z_{i,t}$.[5] Household utility in period t is derived from consumption, C_t, which is uniform across workers. The households' disutility from working is given by aggregating across the individual workers' disutility which is given by $Z_{i,t}L_{i,t}^{\frac{\gamma+1}{\gamma}}$. Households seek to maximise the present value of their expected utility in period t, U_t, as shown below

$$U_t = \sum_{s=0}^{\infty} \beta^s E_t e^{-\sum_{u=0}^{s-1} D_u} \left\{ \ln C_s - \frac{\gamma}{\gamma+1} \int_0^1 Z_{i,s} L_{i,s}^{\frac{\gamma+1}{\gamma}} \, di \right\}, \tag{2.3.1}$$

where the parameter γ is the Frisch elasticity of labour supply, and β is the steady state discount rate. D_t is an exogenous shock in period t to the households' discount factor. The idiosyncratic preference shock is drawn from a truncated log-normal distribution such that $\ln Z \sim N\left(-\frac{\sigma_z^2}{2}, \sigma_z\right)$, where σ_z is the standard deviation of the shock. The first term in this sum denotes the utility gained from consumption using a logarithmic utility function. The second term gives the aggregate disutility across workers.

Since the utility function is separable with respect to consumption and leisure, consumption will be shared equally among all members of the household.[6] However, the optimal level of labour to supply will vary according to each worker's idiosyncratic preference shock. The household will seek to maximise the present value of utility by choosing average consumption, C_t, the workers' wages and labour supply

5 The idiosyncratic preference shock has zero persistence. However given the shock is calibrated to match to the distribution of annual changes in nominal wages, adding persistence would not materially change the mechanics of the model.

6 This is because the marginal utility from consumption is invariant to both the workers' labour supply and their preference shock. Thus given the diminishing returns to consumption, the household will maximise its welfare by sharing consumption equally between members.

$\{W_{i,t}, L_{i,t}\}_{i=0,t=0}^{1,\infty}$. The representative household solves this maximisation problem subject to four constraints.

The first is the budget constraint, which governs the evolution of household assets in the form of a risk-free bond, where the amount of bonds held at the end of period t is denoted B_t.

$$B_t = (1 + i_{t-1}) B_{t-1} + \int_0^1 W_{i,t} L_{i,t} \, di - P_t C_t - \tau_t + \Pi_t \tag{2.3.2}$$

where P_t is the aggregate price level, i_t is the nominal interest rate, τ_t are lump sum taxes, and Π_t is any profits from the firms in period t.

The second constraint is the firms' demand for each type of labour. The demand for each type of labour i is governed by aggregate employment, L_t, the aggregate real wage w_t, and the real wage for the type of labour $w_{i,t}$,

$$L_{i,t} = \left(\frac{w_{i,t}}{w_t} \right)^{-\eta} L_t. \tag{2.3.3}$$

The parameter η is the elasticity of demand for labour and measures how substitutable the different types of labour are with each other. It should be noted that the real wage is detrended by both the level of productivity, A_t, and the price level, P_t such that

$$w_{i,t} = \frac{W_{i,t}}{A_t P_t}. \tag{2.3.4}$$

The third constraint is the AR(1) process for the exogenous shock to the households' stochastic discount rate.

$$D_t = \rho_D D_{t-1} + \varepsilon_t^D \tag{2.3.5}$$

Where ρ_D is the degree of persistence in aggregate demand shocks and ε_t^D is an independently normally distributed random variable, $\varepsilon_t^D \sim N(0, \sigma_D)$.

The final constraint is that each worker is subject is a stochastic downward nominal wage rigidity. This is modeled using the time-dependent framework from Calvo (1983). I assume that each period a randomly selected fraction of workers, governed by the parameter $\lambda \in [0, 1)$, are unable to decrease their nominal wage and thus they are subject to the constraint $W_{i,t} \geq W_{i,t-1}$. The remaining $(1-\lambda)$ share of workers are free to set their nominal wage optimally, but do so in the knowledge that they may be subject to the downward nominal wage rigidity in future periods.

Each household takes the path for aggregate nominal wage index, W_t, total employment L_t, productivity a_t, nominal interest rates i_t and the price level P_t as given.

2.3.1.1 Optimal Consumption Decision

Solving the first order conditions with respect to bonds, B_t and consumption, C_t, yields the standard inter-temporal Euler equation

$$\frac{1}{C_t} = \beta e^{-D_t} \left(1 + r_t\right) \frac{1}{E_t C_{t+1}}, \tag{2.3.6}$$

where consumption in period t is governed by the real interest rate r_t and expectations of consumption in period $t+1$. In turn the real interest rate is governed by the Fisher equation

$$1 + r_t = \frac{1 + i_t}{1 + E_t \pi_{t+1}}. \tag{2.3.7}$$

2.3.1.2 Optimal Wage Decision

I will solve the optimal wage setting problem first in the simplified case of flexible wages, and then in the more general case where downward nominal wage rigidities apply. Each household sets wages to maximise the expected present value of their utility, thus while they internalise the wage-setting decision of other workers within the household, they do not internalise the decisions of other households, and so they treat aggregate real wages, w_t, aggregate employment, L_t, and price dispersion, s_t, as exogenous.

As I show in Appendix 2.A, the wage optimising problem can be simplified to the household choosing the path of wages, $w_{i,t}$, to maximise the expression

$$E_t \sum_{t=0}^{\infty} \beta^t e^{-\sum_{s=0}^{t-1} D_s} \left(\Omega \left(Z_{i,t}, w_{i,t}\right)\right) \tag{2.3.8}$$

where $\Omega\left(Z_{it}, w_{it}\right)$ is per period pay-off to the household with respect to worker i's choice of wage $w_{i,t}$. It can be written as function of the worker specific variables, the disutility shock Z_{it} and real wage w_{it}, and aggregate variables

$$\Omega\left(Z_{i,t}, w_{i,t}\right) = s_t \left(\frac{w_{i,t}}{w_t}\right)^{1-\eta} - \frac{\gamma}{\gamma+1} Z_{i,t} \left[\left(\frac{w_{i,t}}{w_t}\right)^{-\eta} L_t\right]^{\frac{\gamma+1}{\gamma}}, \tag{2.3.9}$$

where w_t, L_t and s_t are the aggregate variables for the real wage, total amount of labour and price dispersion among firms respectively.

Since each worker is subject to downward nominal wage rigidities with constant probability λ, the optimal wage setting is dynamic. The wage set in the current period may constrain the worker's choice set in the subsequent periods and thus will affect the household's expected pay-off. To solve the dynamic optimisation problem given in equation 2.3.8 we can write the problem in the form of a Bellman equation, where the state variable is the real wage the worker starts with at the beginning of the period. This is the real wage which the worker will be constrained by if it is subject to the downward nominal wage rigidity.

$$\begin{aligned}
&V_t\left(w\right) = \\
&\left(1-\lambda\right) \int_0^\infty \max_{w_{i,t}\geq 0} \left\{\Omega\left(Z_{i,t}, w_{i,t}\right) + \beta e^{-D_t} V_{t+1}\left(\frac{w_{i,t}}{\left(1+\pi_{t+1}\right)\left(1+a_{t+1}\right)}\right)\right\} dF\left(Z\right) \\
&+ \lambda \int_0^\infty \max_{w_{i,t}\geq w} \left\{\Omega\left(Z_{i,t}, w_{i,t}\right) + \beta e^{-D_t} V_{t+1}\left(\frac{w_{i,t}}{\left(1+\pi_{t+1}\right)\left(1+a_{t+1}\right)}\right)\right\} dF\left(Z\right).
\end{aligned} \tag{2.3.10}$$

The first term in this expression pertains to workers who are not subject to the downwardly rigid nominal wages and thus are able to choose any non-negative wage. They will set this wage in order to maximise the sum of their expected utility in the current period and the discounted value function for the subsequent period. Although the workers have complete wage flexibility in the current period they know they may be bound by the downward nominal wage rigidity in the next period. Thus choosing a high real wage may constrain their ability to set a lower wage in the subsequent period. The second term pertains to the workers that are subject to the downward nominal wage rigidity. Similar to the unconstrained workers they will seek to optimise the sum of their expected utility in the current period with the discounted value from the next period. However they will do so under the constraint that they cannot lower their wage, $w_{i,t}$, below the wage they started the period with w.

Note that even though it is nominal wages that are downwardly inflexible, the value function is written in terms of the detrended real wage. This is because the real wage is stationary, unlike the nominal wage. This means we must deflate the real wage from the current period by inflation and productivity growth, to obtain the binding real wage in the next period. This is because even if the rigidity forces workers to keep the same nominal wage, their real detrended wage will decrease as prices rise and their labour productivity increases. Thus the constraint written in terms of the real wage $w_{i,t} \geq \frac{w_{i,t-1}}{(1+\pi_t)(1+a_t)}$ is the equivalent of the nominal constraint $W_{i,t} \geq W_{i,t-1}$

Under Flexible Wages

If there are no nominal wage rigidities this problem can be solved analytically. In this case the Bellman equation collapses to

$$V_t \quad (w) = \int_0^\infty \max_{w_{it}\geq 0} \left\{ \Omega \left(Z_{it}, w_{it}^{flex} \right) \right\} dF\left(Z_{it}\right) \tag{2.3.11}$$

which will be invariant with respect to w. Since wages are perfectly flexible they can be re-optimised in every period. Thus workers will set their wage to optimise their current period pay-off without having to weigh the consequences for future periods. The wage function can thus be derived by maximising the per period pay-off given in equation 2.3.9. The optimal wage is a function of the aggregate variables, which are assumed to be exogenous by the household, and the individual worker's idiosyncratic preference shock Z_{it} as shown below

$$w_{i,t}^{flex} = \left(\frac{\eta}{\eta - 1} \right)^{\frac{\gamma}{\gamma+\eta}} Z_{i,t}^{\frac{\gamma+\eta}{\gamma}} L_t^{\frac{\gamma+1}{\gamma+\eta}} s_t^{-\frac{\gamma}{\gamma+\eta}} w_t. \tag{2.3.12}$$

With Downwardly Rigid Wages

In the presence of downward nominal wage rigidity, workers who are not subject to the nominal wage constraint will set their wage according to the function

$$w_t\left(Z_{i,t}\right) = \underset{w}{\operatorname{argmax}}\left\{\Omega\left(Z_{it}, w\right) + \beta V_{t+1}\left(\frac{w}{\left(1+\pi_{t+1}\right)\left(1+a_{t+1}\right)}\right)\right\} \tag{2.3.13}$$

If workers are unaffected by the downward nominal wage constraint, all workers who draw the same preference shock, $Z_{i,t}$, choose the same wage.

This wage function will differ from the flexible wage schedule even for workers who are not subject to downward wage rigidities. This is because workers will set their wage lower than they otherwise would if wages were perfectly flexible to offset the possibility that they will be stuck at an inefficiently high nominal wage in the next period.

Workers who are subject to the nominal wage constraint fall into two categories. The first group consists of workers who draw a preference shock such that $w\left(Z_{it}\right) \geq w$. This group of workers will have a preferred wage, $w\left(Z_{i,t}\right)$, that is higher than their wage at the beginning of the period, w. Thus the downward nominal wage rigidities will not be a binding constraint. The second group are workers who draw a preference shock $w\left(Z_{it}\right) < w$. For these workers the constraint is binding. They would ideally like to lower their nominal wage, but are prevented by the downward wage rigidity. Instead these workers will choose to keep their nominal wage constant, keeping their real wage fixed at w, as shown in Appendix 2.A.

2.3.2 Final Goods Firms

Final goods firms produce goods for consumption, Y_t, by combining intermediate goods, $Y_{j,t}$, using a Dixit-Stiglitz aggregator. These firms are perfectly competitive and thus have zero profit. Their prices, P_t, are perfectly flexible and they choose the amount of each intermediate good, $Y_{j,t}$, to maximise profit in the current period as shown below.

$$\max_{Y_{j,t}} P_t Y_t - \int_0^1 P_{j,t} Y_{j,t}\, dj, \tag{2.3.14}$$

subject to the Dixit-Stiglitz aggregator which transforms intermediate goods into final goods according to

$$Y_t = \left[\int_0^1 Y_{j,t}^{\frac{\epsilon-1}{\epsilon}} \, dj \right]^{\frac{\epsilon}{\epsilon-1}}. \tag{2.3.15}$$

Solving for the first order conditions of this optimisation problem, shown in detail in Appendix 2.A, yields the demand curve

$$Y_{j,t} = \left(\frac{P_{j,t}}{P_t} \right)^{-\epsilon} Y_t, \tag{2.3.16}$$

while the zero profit condition pins down the aggregate price index

$$P_t = \left[\int_0^1 P_{j,t}^{1-\epsilon} \right]^{\frac{1}{1-\epsilon}}. \tag{2.3.17}$$

2.3.3 Intermediate Goods Firms

There is a continuum of firms who produce intermediate goods in this economy, denoted by the subscript j, each of which produce a differentiated good in a monopolistically competitive market. Firms produce goods using each type of labour supplied by the workers that comprise the household and are subject to pricing frictions a la Calvo (1983). Each period a randomly selected fraction, θ, of firms are allowed to re-set their price. Those that are not able to flexibly adjust their price instead partially index it by the central bank's target rate of inflation, $\overline{\pi^p}$. The firms also receive a subsidy which unwinds the mark-up caused by their market power. The firms thus optimise the present discounted value of profits by optimally choosing the amount of output to produce $Y_{j,t}$, how much of each type of labour, $L_{j,t}$, to hire and, if they are able to, the optimal price to set $P_{j,t}^*$. When setting their price they maximise the discounted stream of expected profits

$$\max_{\{P_{j,t}^*, Y_{j,t}\}} \sum_{k=0}^{\infty} \beta^k \frac{\lambda_{t+k}}{\lambda_t} \theta^k \left[\frac{P_{j,t}^* \left(1 + \overline{\pi^p}\right)^{\mu k}}{P_{t+k}} Y_{j,t+k} - \frac{\epsilon - 1}{\epsilon} MC_{j,t+k}^r Y_{j,t+k} \right], \tag{2.3.18}$$

where $MC_{j,t}^r$ is the real marginal cost of firm j in period t, ϵ is the elasticity of substitution between goods and μ is the degree of price indexation to the central bank's target rate of inflation. Since the firms are owned by the households the stochastic discount factor of the firms is governed by the households' marginal utility of consumption λ_t. The stochastic discount factor between periods t and $t+k$ is given by the term $\beta^k \frac{\lambda_{t+k}}{\lambda_t}$. Each firm is provided a subsidy of $\frac{\epsilon-1}{\epsilon}$ which unwinds the mark-up caused by the firms' market power.

The discounted expected stream of profits are maximised subject to three constraints. The first is the firms' linear production function

$$Y_t = A_t L_t, \tag{2.3.19}$$

where A_t is the level of labour productivity and L_t represents an aggregate index of the different types of labour provided by the household which is given by

$$L_t = \left[\int_0^1 L_{i,t}^{\frac{\eta-1}{\eta}}\, di\right]^{\frac{\eta}{\eta-1}}, \tag{2.3.20}$$

where η is the Frisch elasticity of demand for individual workers. The firm will choose each type of labour, $L_{i,t}$, to minimise its wage bill subject to the aggregate amount of labour required. This cost minimisation results in the labour demand function

$$L_{i,t} = \left(\frac{W_t}{W_{i,t}}\right)^{\eta} L_t, \tag{2.3.21}$$

where the nominal wage aggregate W_t is given by

$$W_t = \left[\int_0^1 W_{i,t}^{1-\eta}\, di\right]^{\frac{1}{1-\eta}}. \tag{2.3.22}$$

Growth in labour productivity is given by the exogenous process

$$(a_t - \overline{a}) = \rho_a \left(a_{t-1} - \overline{a}\right) + \varepsilon_t^a, \tag{2.3.23}$$

where $\overline{a}$ is the steady state rate of growth in the economy, ρ_a is the degree of persistence in shocks to the labour productivity growth rate and ε_t^a is an independently normally distributed random variable, $\varepsilon_t^a \sim N\left(0, \sigma_a\right)$.

The second constraint is the demand for the firm's intermediate good.

$$Y_{j,t} = \left[\frac{P_{j,t}}{P_t}\right]^{-\epsilon} Y_t. \tag{2.3.24}$$

Finally, a firm's marginal cost is governed by the aggregate wages set by the household, which are assumed to be exogenous by the firms

$$MC^r_{j,t} = \frac{W_t}{P_t A_t} = w_t. \tag{2.3.25}$$

Solving this constrained optimisation problem, as outlined in Appendix 2.A, gives the solution for the firms' optimal price. Note that since all firms are identical we can drop the subscript j. The solution for the optimal price to choose, P^*_t, is given by

$$P^*_t = \frac{K^r_t}{J^r_t}, \tag{2.3.26}$$

where K^r_t and J^r_t are recursive variables defined by

$$K^r_t = MC^r_t \left(P_t\right)^{\epsilon} + \beta\theta \left(1+\overline{\pi^p}\right)^{-\mu\epsilon} E_t \left[K^r_{t+1}\right] \tag{2.3.27}$$

$$J^r_t = P_t^{\epsilon-1} + \beta\theta \left(1+\overline{\pi^p}\right)^{\mu(1-\epsilon)} E_t \left[J^r_{t+1}\right]. \tag{2.3.28}$$

The aggregate price level evolves according to the price that firms with price flexibility set in period t and the prices that the remainder of firms had in the previous period, indexed by the central bank's inflation target.

$$P_t^{1-\epsilon} = \theta \left(1+\overline{\pi}\right)^{\mu(1-\epsilon)} P_{t-1}^{1-\epsilon} + \left(1-\theta\right) P_t^{*1-\epsilon}. \tag{2.3.29}$$

Once the price level has been pinned down we can solve for the inflation rate which is given by $\pi_t = \frac{P_t}{P_{t-1}} - 1$. Since both prices and wages are subject to nominal constraints we separately track wage inflation, π^w_t, using the change in the real wage, price inflation and rate of productivity growth according to

$$1+\pi^w_t = \frac{W_t}{W_{t-1}} = \frac{w_t}{w_{t-1}} \left(1+a_t\right)\left(1+\pi_t\right). \tag{2.3.30}$$

2.3.4 Monetary Policy

Monetary policy is set by a central bank which uses the nominal interest rate to stabilise price inflation and the output gap. In the baseline model the central bank follows a simple Taylor Rule, setting interest rates based on deviations in inflation from its target, $\overline{\pi}$, and in movements in the output gap y_t

$$1+i_t = \frac{\left(1+\overline{\pi^p}\right)\left(1+\overline{a}\right)}{\beta}\left(\frac{y_t}{\overline{y}}\right)^{\phi_y}\left(\frac{1+\pi_t^p}{1+\overline{\pi^p}}\right)^{\phi_\pi}. \tag{2.3.31}$$

The output gap is given by

$$y_t = \frac{Y_t}{A_t}. \tag{2.3.32}$$

where $\overline{y}$ is the steady state outcome associated with an output gap of zero. The parameters ϕ_y and ϕ_π are the non-negative Taylor coefficients on the output gap and inflation respectively and determine how aggressively the central bank changes the nominal interest rate in response to deviations from the model's steady state.

2.3.5 Market Clearing Conditions

In equilibrium the goods market must clear such that aggregate consumption is equal to the firms' output

$$C_t = Y_t. \tag{2.3.33}$$

Finally, we need a market clearing condition for each of the intermediate goods. Since each firm will in general have different prices due to the Calvo price setting framework, when aggregating output across firms we need to calculate the degree of price dispersion. This price dispersion creates a wedge between the aggregate labour employed and the total output produced by firms as described in detail in Appendix 2.A.

$$Y_t = \frac{1}{s_t}A_t L_t, \tag{2.3.34}$$

where s_t is a measure of price dispersion within the economy. This measure of price dispersion evolves according to

$$s_t = (1-\theta)\left(\frac{P_t^*}{P}\right)^{-\epsilon} + \theta\left(1+\overline{\pi^p}\right)^{-\mu\epsilon} s_{t-1} \tag{2.3.35}$$

2.4 Equilibrium

The equilibrium for this economy is a set of the endogenous variables $\{L_t, Y_t, C_t, \pi_t^p, \pi_t^w, w_t, i_t, r_t, a_t, D_t, s_t, b_t, P_t^*, K_t^r, J_t^r\}$ and the endogenous functions $\{w(Z_{it}), V_t(w)\}$ which satisfies equations 2.3.2 to 2.3.35 for all periods t given the exogenous variables . I will consider the equilibrium of the model under the special case where wages are fully flexible, $\lambda = 0$, and the more general case in which the downward nominal wage constraint applies.

2.4.1 Flexible Wages

Under the special case of flexible wages the equilibrium for the model can be solved analytically. Since the workers' wages are perfectly flexible the households set wages to maximise utility in the current period only, since they have no impact on future periods. Given the workers' optimal wage setting function under flexible wages is given by equation 2.3.12 we can integrate over the distribution of possible shocks to find the aggregate real wage. This gives the flexible-wage labour supply curve which relates the amount of aggregate level labour supplied as a function of the real wage

$$L_t = \left(\frac{\eta-1}{\eta}\right)^{\frac{\gamma}{\gamma+1}} s_t^{\frac{\gamma}{\gamma+1}} w_t^{\frac{\eta}{\gamma+1}} \left[\frac{1}{Z_t}\right]^{\frac{\gamma}{1+\gamma}}, \qquad (2.4.1)$$

where the aggregate level of the disutility from working is given by

$$Z_t = \left[\int_0^1 \left(Z_{it}^{\frac{\gamma(1-\eta)}{\gamma+\eta}}\right) di\right]^{-\frac{\eta+\gamma}{\gamma(\eta-1)}} = e^{-\frac{1}{2}\frac{\eta(1+\gamma)}{\gamma+\eta}\sigma_z^2}. \qquad (2.4.2)$$

This labour supply curve completes the flexible wage model and allows us to solve the steady state analytically.

2.4.2 Downwardly Rigid Nominal Wages

In the more general case with downward nominal wage rigidity the equilibrium will depend on the distribution of the real wages within the economy. Thus to solve the model we need to track the distribution of real wages. The distribution of real detrended wages at time t, is denoted $G_t(w)$, and has a corresponding density function

$g_t(w)$. To see how the distribution is defined within the model consider the share of workers who have a real wage lower than w' in a given period. This set is comprised of three types of workers.

The first is those that are not subject to the downward nominal wage rigidity and optimally choose a wage lower than w'. This occurs because the idiosyncratic shock to their disutility of labour preferences, $z_{i,t}$ is lower than a cut-off $\overline{z}$, such that

$$z_{i,t} < \overline{z} = Z_t\left(w'\right), \tag{2.4.3}$$

where $Z_t(w)$ is the inverse of the optimal wage function in period t. Given the distribution of the preference shock $F(.)$, this group will comprise $(1-\lambda)F(Z_t(w'))$ of all workers.

The second type of workers are those that are subject to the downwardly rigid nominal wage constraint, but are not bound by it. These workers start period t with a wage w which is lower than w' and choose to increase their wage such that it remains below w'. This group of workers who are subject to the wage constraint, but not bound by it, and have wages lower than w' will thus have a preference shock $z_{i,t}$ such that

$$Z_t(w) \leq z_{i,t} \leq Z_t\left(w'\right). \tag{2.4.4}$$

Finally, the third category of workers with a wage lower than w' are those who are subject to the downward nominal wage constraint and are bound by it. These workers will begin the period with a detrended real wage that is lower than w' and draw a shock that causes them to prefer an even lower real wage, $z_{it} < Z(w) \leq Z_t(w')$. However, due to the nominal wage rigidity they are forced to maintain their previous nominal wage. The optimality of choosing the same nominal wage when the downward wage constraint binds is shown in Appendix 2.A.

The proportion of workers who start period t with a real wage equal to or lower than w is given by $G_{t-1}(w(1+\pi_t)(1+a_t))$. Note that it is necessary to inflate the real wage by inflation and the rate of productivity growth to translate the wage distribution between periods.

Thus the combined share of workers in these last two groups is given by

$$\lambda F(Z_t(w))\, G_{t-1}(w(1+\pi_t)(1+a_t)) \tag{2.4.5}$$

so combining these three groups of workers gives the evolution of distribution of the real wage

$$G_t(w) = (1-\lambda) F(Z_t(w)) + \lambda F(Z_t(w)) G_{t-1}(w(1+\pi_t)(1+a_t)). \tag{2.4.6}$$

Note that as inflation or productivity growth increases $G_{t-1}(w(1+\pi_t)(1+a_t))$ approaches 1. This demonstrates how a higher rate of inflation or productivity growth unwinds the labour market friction, moving the wage distribution to the flexible wages case where $G_t(w) = F(Z_t(w))$.

Given the distribution of wages and the optimal wage setting function we can calculate the equilibrium level of labour supplied as shown in Appendix 2.A.

$$L_t = \left(\frac{\eta-1}{\eta}\right)^{\frac{\gamma}{\gamma+1}} s_t^{-\frac{\gamma}{\gamma+\eta}} w_t^{\frac{\eta}{\gamma+1}} \times \left(\frac{1}{Z_t^*}\right)^{\frac{\gamma}{1+\gamma}}, \tag{2.4.7}$$

where Z_t^* is a distorted measure of the aggregate disutility of labour given by

$$\begin{aligned} Z_t^* \quad = \quad & \left\{(1-\lambda)\int_0^\infty \left(\frac{\hat{w}_t(Z)}{w_t(Z)}\right)^{\eta-1} Z^{\frac{\gamma(1-\eta)}{\gamma+\eta}}\, dF(Z)\right. \\ & +\lambda\int_0^\infty \left(\frac{\hat{w}_t(Z)}{w_t(Z)}\right)^{\eta-1} Z^{\frac{\gamma(1-\eta)}{\gamma+\eta}} G_{t-1}(w_t(Z)(1+\pi_t)(1+a_t))\, dF(Z) \\ & +\lambda\int_0^\infty Z^{\frac{\gamma(1-\eta)}{\gamma+\eta}} \left[\int_{w_t(Z)}^\infty \left(\frac{\hat{w}_t(Z)}{w_t(Z)}\right)^{\eta-1} (1+\pi_t)(1+a_t) \ldots\right. \\ & \left.\left. \times g(w_t(Z)(1+\pi_t)(1+a_t))\, dw\right] dF(Z)\right\}^{-\frac{\gamma+\eta}{(\eta-1)\gamma}}. \end{aligned} \tag{2.4.8}$$

Each of the terms in this expression corresponds with the three types of workers whose disutility we are aggregating. Critical to the analysis is the third term which pertains to workers for which the wage rigidity is a binding constraint. The inability to lower their nominal wage, in order to satisfy their relatively low draw for the disutility shock, means that they provide an inefficiently low level of labour creating a distortion in the economy.

The distribution of the changes in workers' nominal wages between the periods in the model's equilibrium is shown in Figure 2.2. Under downward nominal wage rigidities the model is able to replicate the two main stylised facts from the empirical data, namely the large spike of workers whose nominal wage is unchanged between periods and the relatively small share of workers who receive a cut in their nominal wage.

Figure 2.2: Distribution of Changes in Workers' Nominal Wages Over One Quarter

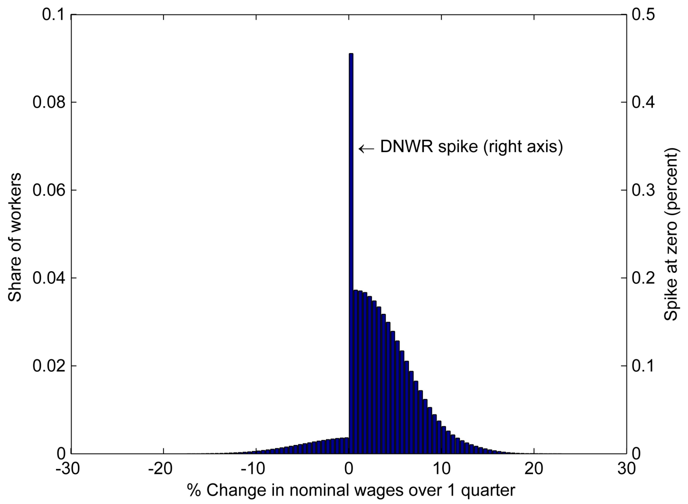

2.5 Numerical Results

Since this non-linear approach to nominal rigidities involves tracking the distribution of real wages the model's steady state cannot be solved using standard approximation techniques. Instead it is solved using the following algorithm.

2.5.1 Steady State

First I assume that in steady state the inflation rate is equal to the central bank's inflation target. Following Ascari and Sbordone (2014), this allows us to pin down the steady state variables associated with the firm's pricing problem and aggregate wage growth using equations 2.3.26 to 2.3.30.

I then guess a value for the equilibrium level of output, and iterate over the households' value function for its wage setting problem, $V(w)$, until it converges to a stable function. Given the households' value function $V(w)$ we can solve for the optimal wage setting function $w(z)$ using equation 2.3.13. Using the inverse of this function, $Z(w)$, I solve for the distribution of real wages, $G(w)$, by iterating over equation 2.4.6 until convergence.

Finally, I use the optimal wage function and the distribution of real wages to solve for the level of output at which the firms' marginal cost is equal to the real wage using equation 2.4.7. I use this value to update our original guess for the level of output and continue to iterate over the process until the level of output converges to its equilibrium value.

2.5.2 Dynamics

Once we have solved for the steady state of the model we can solve its response to the exogenous shocks. There are two sources for stochastic dynamics in this model: a shock to the growth rate of productivity in the model and a shock to the households' discount rate. These exogenous processes can be thought of as shocks to aggregate supply and aggregate demand respectively.

We can numerically solve for the rational expectations solution for the model's dynamics by using the extended path method developed in Fair and Taylor (1980). This solves for the rational expectations solution of the model by assuming that all future stochastic shocks are equal to their expected value which is zero. We approximate the functions $V(w)$, $w(Z)$ and $Z(w)$ with polynomials as outlined in Appendix B.

This rational expectations algorithm solves for the path for all state variables over T periods, assuming that there is no aggregate uncertainty. It assumes that the model is in steady state in the periods $t = 0$ and $t = T + 1$. To solve the model I use a combination of forward and backwards iteration. First I assume that all variables are at their steady state across all periods with the exception of the exogenous shock processes D_t and a_t. I then use backwards iteration to solve for the households' value function, $V_t(w)$, and associated wage setting function, $w_t(Z_{i,t})$, taking the path of aggregate wages and price dispersion as constant. This allows us to calculate the paths for inflation, output and the real and nominal interest rates in response to the shocks.

The second stage then uses forward iteration to solve for the distribution of real wages, $G_t(w)$, and aggregate level of price dispersion, s_t, among the firms based on these updated values for inflation and output. Using the newly updated paths for the state variables we then repeat the process until the model converges to a stable solution.

2.5.3 Linearisation

Most previous approaches to linearisation rely on taking derivatives analytically in order to construct a Taylor approximation of the model. By contrast this method will use deterministic impulse response functions to calculate numerical derivatives for the endogenous variables with respect to the aggregate shocks.

This approach to simulating the model with aggregate uncertainty has two steps. The first is to calculate numerical derivatives of endogenous variables with respect to the exogenous shocks. This is done by calculating the deterministic transition path to a one-off shock that occurs when the economy is in it's steady state and tracing out the path of endogenous variables as it returns to the models steady state. Consider an impulse response function as a series of variables $\{f_{y,S,t}, f_{y,S,t+1} f_{y,S,t+2}, ..., f_{y,S,t+T}\}$ which correspond to the response of an aggregate variable, y, to a set of one-off shocks, S, at time t as shown in Figure 2.3.

Figure 2.3: Impulse Response Functions

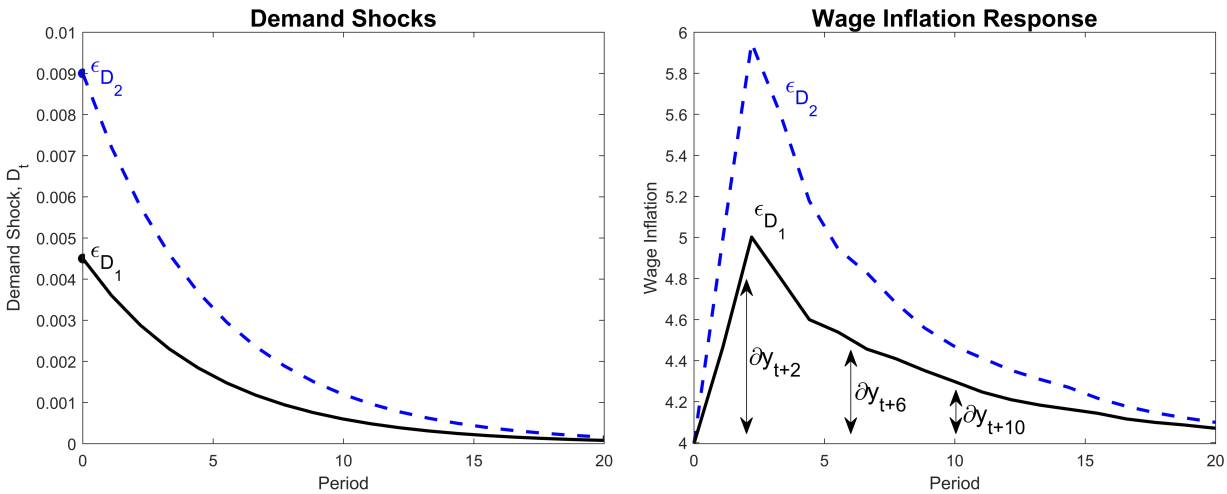

Note. These figures show two demand shocks ϵ_{D_1} and ϵ_{D_1}, and the respective responses of the endogenous variable wage inflation. To calculate the numerical derivatives we use the impulse functions functions which trace out the path of ∂y_t in response to the respective shocks at time t.

For a set S that comprises of a single small shock ϵ_D, this sequence can thus be thought of as a set of numerical derivatives showing how the endogenous variable responds over time to the demand shock that occurs at time t. Thus

$$\left\{\frac{\partial y_t}{\partial e_{D,t}}, \frac{\partial y_{t+1}}{\partial e_{D,t}}, \frac{\partial y_{t+2}}{\partial e_{D,t}}, ..., \frac{\partial y_{t+T}}{\partial e_{D,t}}\right\} = \left\{\frac{f_{y,\{\epsilon_D\},t}}{e_D - 0}, \frac{f_{y,\{\epsilon_D\},t+1}}{e_D - 0}, \frac{f_{y,\{\epsilon_D\},t+2}}{e_D - 0}, ..., \frac{f_{y,\{\epsilon_D\},t+T}}{e_D - 0}\right\}. \tag{2.5.1}$$

Similarly, by calculating the impulse response functions across a set of multiple sizes or combinations of shocks this process can be used to approximate higher order numerical derivatives as described in Appendix 2.B.

With these numerical derivatives in hand a model with aggregate shocks can be simulated by drawing a sequence of shocks

$$\{\boldsymbol{e}_1, \boldsymbol{e}_2, \boldsymbol{e}_3, ...\}$$

and then calculating the path of endogenous variables with the derivatives. For a linear approximation the simulation of each endogenous variable, y_t, is generated by

$$y_t = \overline{y} + \nabla_{y,0}\boldsymbol{e}_t + \nabla_{y,-1}\boldsymbol{e}_{t-1} + \nabla_{y,t-2}\boldsymbol{e}_{t-2} + ... + +\nabla_{y,-T}\boldsymbol{e}_{t-T} \tag{2.5.2}$$

where $\nabla_{y,j}$ is the vector of partial derivatives for the endogenous variable y with time lag j.

In order to capture the non-linearities in aggregate variables within this model, I use the second order approximation

$$y_t = \overline{y} + \sum_{j=t}^{t-20} \nabla_j \boldsymbol{e}_j + \left\{\sum_{j=t}^{t-20} \boldsymbol{e}_j H_j \boldsymbol{e}_j\right\}, \tag{2.5.3}$$

where H is the Hessain matrix of second derivatives.[7]

2.5.4 Calibration

Table 2.2 reports the baseline calibration for the model. The discount rate, β, is set such that the steady state real interest rate is 2 per cent. The parameters governing the firms' pricing friction are set to match those estimated by previous studies, in this case Cogley and Sbordone (2008). The targeted level of inflation is assumed to

7 The Hessian matrix is comprised of second derivatives $H = \begin{bmatrix} \frac{\partial^2 y_t}{\partial a^2} & \frac{\partial^2 y_t}{\partial a \partial D} \\ \frac{\partial^2 y_t}{\partial D \partial a} & \frac{\partial^2 y_t}{\partial D^2} \end{bmatrix}$

be 2 per cent, following the Federal Reserve's stated target (Bernanke, 2010), and labour productivity growth is calibrated to 2 per cent. The parameters within the central bank's monetary policy rule, ϕ_π and ϕ_y, govern the response of interest rates to changes in the output gap and rate of inflation. For the baseline model these are set at $\phi_\pi = 0.5$ and $\phi_y = 0.5$ following Rudebusch (2009).

Table 2.2: Baseline Calibration

Parameter	Value	Interpretation	Reference
γ	0.5	Frisch labour supply elasticity	Daly and Hobijn (2014)
η	2	Labour demand elasticity for individual workers	Daly and Hobijn (2014)
λ	0.8	Share of workers subject to downward nominal wage constraint	Internally calibrated
σ_z	0.6	Standard deviation of the idiosyncratic shock to workers' disutility of labour	Internally calibrated
θ	0.7	Share of firms subject to nominal price rigidity	Cogley and Sbordone (2008)
χ	0.05	Indexation of firms' price to inflation target	Cogley and Sbordone (2008)
ϵ	13	Price elasticity of good	Mark up of 11 per cent
ϕ_π	0.5	Baseline Taylor Rule Parameter on Inflation	Rudebusch (2009)
ϕ_y	0.5	Baseline Taylor Rule Parameter on Inflation	Rudebusch (2009)
$\overline{\pi}$	0.005	Steady state rate of inflation	2 per cent annualised inflation
β	0.995	Discount factor	2 per cent real interest rate
$\overline{a}$	0.005	Steady state rate of productivity growth	2 per cent annualised labor productivity growth
σ_D	0.008	Standard deviation of the shock to worker's discount factor	Estimated
σ_a	0.007	Standard deviation of the shock to the rate of productivity growth	Estimated
ρ_D	0.8	Persistence of the shock to worker's discount factor	Estimated
ρ_a	0.7	Persistence of the shock to the rate of productivity growth	Estimated

2.5.5 Estimation

Two sets of parameters are estimated using a simulated method of moments. The first are the parameters governing the labour market which are set to match the distribution of workers' change in nominal wages over four quarters. The distance between the distributions is minimised by setting the rate at which workers are subject to downward nominal wage rigidity, λ, at 0.8 and the volatility of the idiosyncratic shock to the disutility of labour, σ_z, at 0.6. Under this calibration 13 per cent of workers have no change in their nominal wage and only a quarter experience a fall in nominal wage over four quarters (see Figure 2.4).

Figure 2.4: Percentage Change in Hourly Nominal Wages over 4 Quarters vs the Model

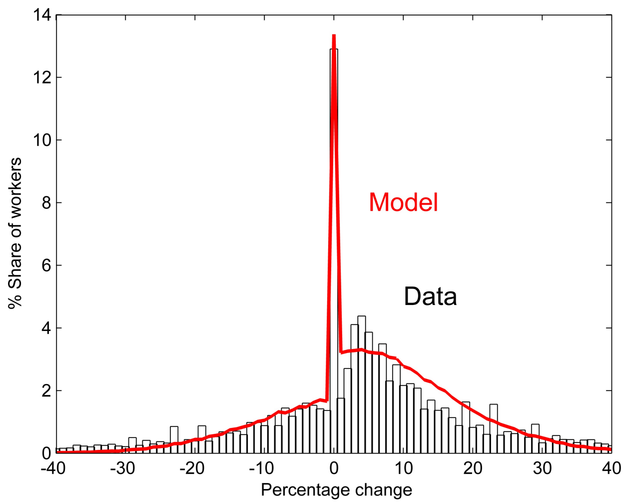

Data covers all workers from 1980-2015 (hourly and salary, job-stayers and job-switchers)
Source: Current Population Survey

To estimate the parameters governing the exogenous shock processes I set them to match the moments of the aggregate data, namely the variances and co-variances of inflation and the output gap. I use a generalised method of moments to match the variances and co-variances of inflation, wage inflation and the output gap. For the set of parameters $\Sigma = \left[\rho^a, \rho^D, \varepsilon^a, \varepsilon^D\right]$, I estimate the moments $m = \left[\text{var}\left(\pi_t\right), \text{var}\left(y_t\right), \text{var}\left(\pi_t^w\right), \text{cov}\left(\pi_t, y_t\right)\right]$ and choose the calibration that minimises the distance

$$\hat{\Sigma} = \arg\min_{\Sigma} \left[\left(m - \hat{m}\left(\Sigma\right)\right)' W \left(m - \hat{m}\left(\Sigma\right)\right)\right] \tag{2.5.4}$$

where the weight matrix, W, is a diagonal matrix consisting of the variance of moments which ensures that the estimators are efficient.

2.5.6 Impulse Response Functions

To demonstrate the impact of the asymmetric downward nominal wage constraint I solve the model's response to a demand shock for both the baseline calibration of the model and when nominal wages are fully flexible. Figure 2.5 shows the impulse response functions from a range of positive and negative shocks to the household's discount rate in the flexible wage model. As expected positive shocks to aggregate demand increase price and wage inflation and output, while a negative shock causes these variables to fall. With flexible wages the responses are relatively symmetric; negative and positive shocks have largely the same magnitude.

Figure 2.5: Impulse Response Functions to a Demand Shock with Flexible Wages

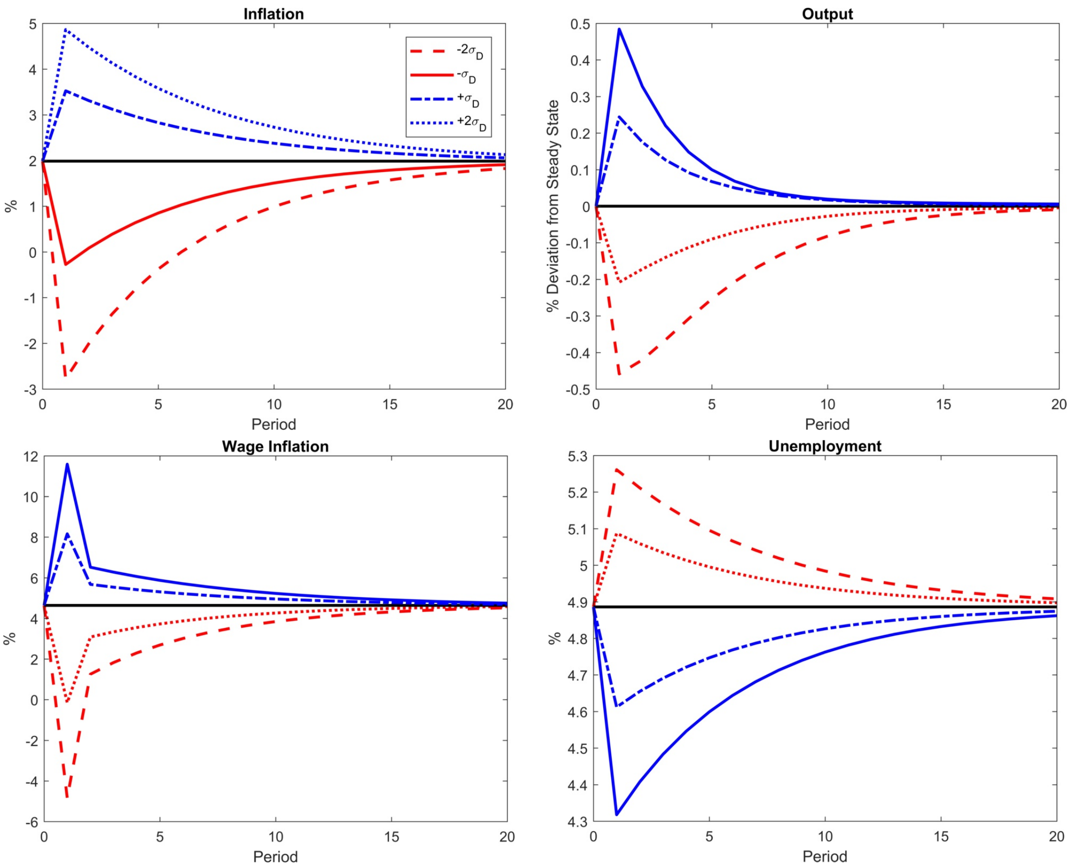

Figure 2.6 shows the impulse response functions when wages are downwardly rigid. The responses of price and wage inflation are notably skewed, with the magnitude of

the effect on inflation from a negative demand shock being significantly smaller than a positive shock of the same size. For example, while a large positive demand shock may increase wage inflation by approximately 3.2 percentage points, a negative shock of similar magnitude will only lower it by 2.6 percentage points. In particular, wage deflation only occurs in response to very large negative demand shocks, if ever, within the model.

Figure 2.6: Impulse Response Functions to a Demand Shock with DNWR

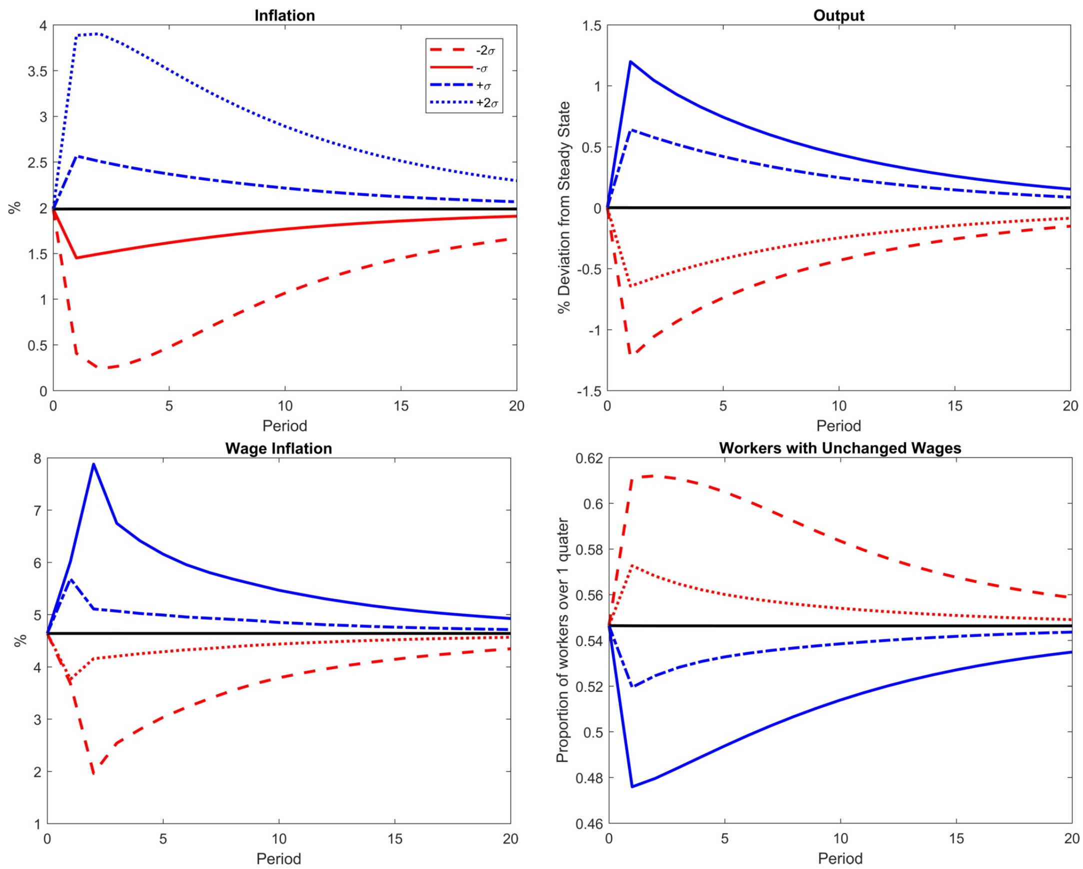

This is a result of the downward nominal wage constraint which limits the ability of households to lower their nominal wage in response to lower demand for their labour. This friction constrains wage inflation from falling below its equilibrium, but does not hamper its ability to rise above it. The inability to lower nominal, and consequently real, wages feeds directly into firms' labour costs, which in turn leads to a smaller decrease in price inflation. Thus the downward nominal wage rigidity limits the degree to which price deflation can occur either due to positive supply, or adverse demand shocks. However, since the nominal wage rigidity only constrains marginal costs from falling, price inflation is able to increase to a much greater degree. The response of output is also asymmetric, though to a smaller degree. Negative shocks have a slightly larger effect on output, compared with positive shocks, as a higher share of workers are constrained by the nominal wage rigidity when inflation is low. With more workers subject to the downward nominal wage constraint, real wages will be held inefficiently high and demand for their labour will decrease, lowering output. This asymmetry also occurs in response to productivity shocks as shown in Figure 2.7.

Figure 2.7: Impulse Response Functions to a Productivity Shock with DNWR

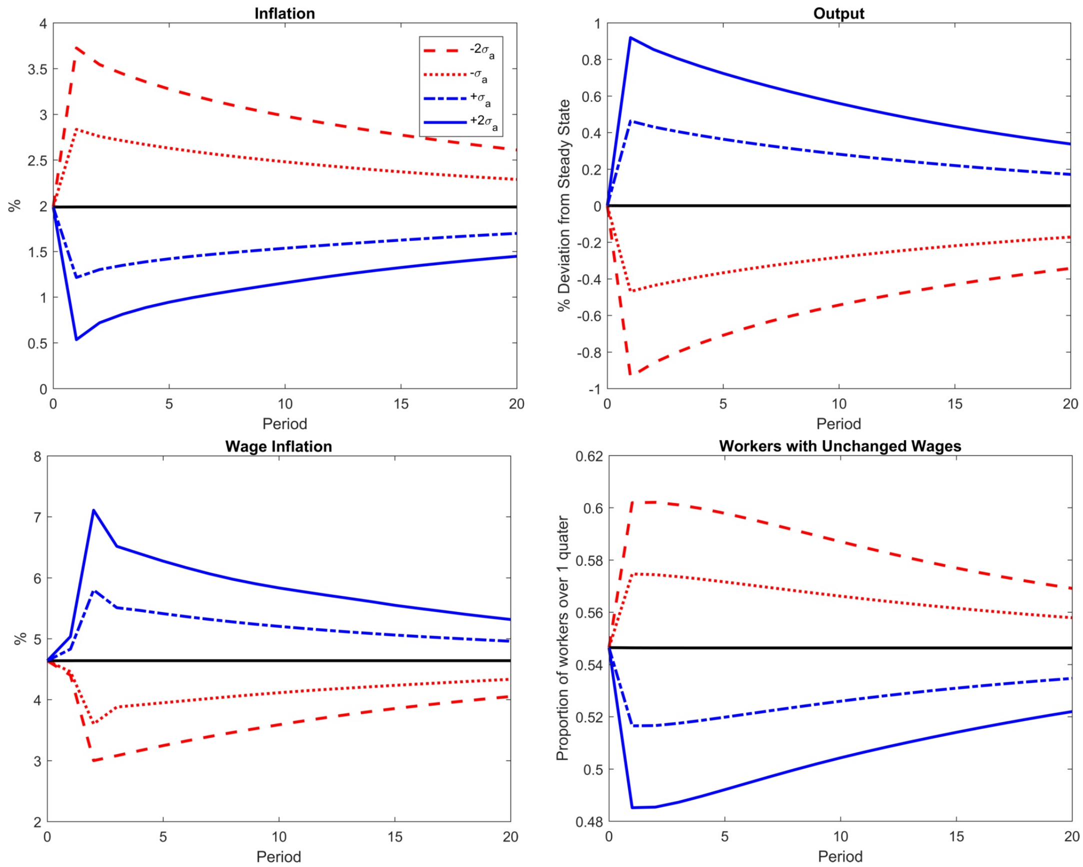

The short run Phillips Curve can be traced using the demand shock's impulse response functions. Figure 2.8 shows the path of the output gap and wage inflation in response to four aggregate demand shocks of different sizes and signs. For small perturbations around the equilibrium the Phillips Curve is generally linear. However for larger shocks the Phillips Curve is clearly convex. Large decreases in the output gap cause only moderate falls in wage inflation.

Figure 2.8: Short-run Phillips curve

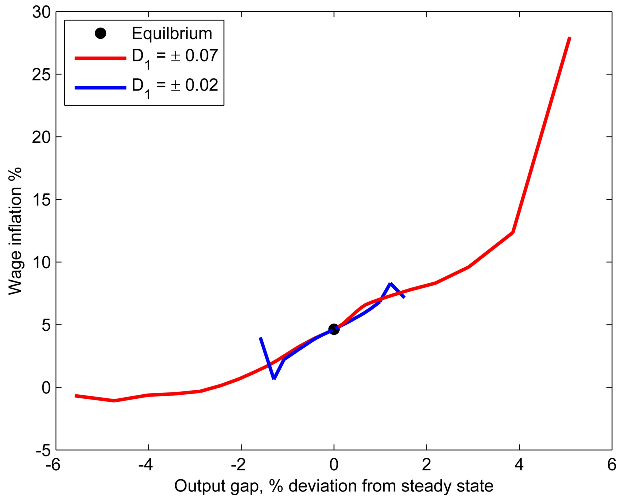

This asymmetry provides one plausible explanation for the post-crisis experience of the developed economies and the US in particular. As Daly, Hobijn, and Lucking (2012) note, the downturn saw a sharp decrease in output and employment, but wage growth was only moderately weaker. By contrast the recovery saw a period of relatively strong growth in jobs but inflation remained subdued. Seen through the lens of this model the Great Recession can be represented as a large shock to aggregate demand. As shown in the impulse response functions, such a shock has a relatively large impact on the output gap and employment as is expected. However the impact on inflation is significantly smaller than might be expected from a linear extrapolation of the slope of the Phillips Curve at the equilibrium. This demonstrates how downward nominal wage rigidity can bend the Phillips Curve, thus providing a possible explanation for the post crisis macro-economic experience.

2.6 Optimal Monetary Policy Rule

To analyse the optimal monetary policy rule within the model I simulate it under a range of policy regimes. I first consider optimal monetary policy when wages are fully flexible where I replicate the canonical result of divine coincidence from Blanchard and Galí (2007). I then simulate optimal monetary policy over a range of policy rules when wages are downwardly rigid and compare the results to Erceg et al. (2000)'s model of symmetrically rigid wage and price frictions. Finally, I consider whether a constrained monetary policy rule would benefit from an asymmetric approach to fluctuations below vs above the model's steady state.

2.6.1 Flexible Wages

I first consider optimal monetary policy where wages are perfectly flexible, but prices remain rigid.[8] The flexible wage model model closely resembles that of Blanchard and Galí (2007), the main difference being the addition of the idiosyncratic preference shock among workers. Blanchard and Galí show that, in a model which only contains a Calvo-style pricing friction and shocks that only affect household preferences and productivity level, monetary policy can replicate the first best outcome.[9]

Optimal monetary policy in this model consists of a strict inflation target. By perfectly stabilising prices this ensures that the output gap remains at zero and price dispersion is minimised. While there remains a significant degree of wage dispersion, this cannot be mitigated with monetary policy as it is driven by the real differences in worker preferences. Welfare loss is calculated as a fraction of Pareto-optimal consumption.

Table 2.3: Optimal Monetary Policy Under Flexible Wages

	Baseline	High ϕ_y	Strict price inflation target
ϕ_π	1.3	1.3	∞
ϕ_y	1.0	1.5	0
Welfare	-1.1	-1.7	0
Variance of output	3×10^{-5}	5×10^{-5}	0
Variance of price inflation	5×10^{-6}	7×10^{-6}	0

As Table 2.3 shows, in the absence of any nominal wage rigidities this model replicates that result. The optimal monetary policy for the central bank is to impose a strict inflation target. Placing any positive weight on output within the policy rule only increases the variance of the output gap and decreases household utility.

8 In the flexible wage case the wedge between price and wage inflation is exogenously fixed by the rate of productivity growth.

9 A subsidy is also required to offset the mark up caused by the monopolistically competitive firms.

2.6.2 Downwardly Rigid Nominal Wages

I now consider optimal monetary policy in the model with downward nominal wage rigidity. A useful baseline for this analysis is Erceg et al. (2000) who analyse optimal monetary policy in a simple New Keynesian model with symmetric nominal frictions a la Calvo (1983) on both prices and wages within the model. They report three relevant results. The first best outcome cannot be replicated in a model where both price and wage nominal rigidities exist. Monetary policy will not be able to fully unwind the effects of these nominal rigidities as is the case when there is only a single friction. The second is that optimal policy can be implemented by a policy rule which strictly targets a weighted average of price and wage inflation, with the weights determined by the relative stickiness of the two. If wages are more rigid than prices then the optimal average inflation target will place a higher rate on wage inflation than on price inflation. Intuitively they find that the welfare costs of inflation volatility increases with the degree of price rigidity in the sector. The third finding is that a Taylor rule that strictly targets the output gap is not in general optimal policy, however it can achieve an outcome that is fairly close to the optimal outcome.

Since the general model includes both price and wage nominal rigidities I augment the baseline Taylor rule to include both price and wage inflation as shown in equation 2.6.1.

$$1 + i_t = \frac{(1+\overline{\pi^p})(1+\overline{a})}{\beta} \left(\frac{y_t}{\overline{y}}\right)^{\phi_y} \left(\frac{1+\pi_t^p}{1+\overline{\pi^p}}\right)^{\phi_{\pi^p}} \left(\frac{1+\pi_t^w}{1+\overline{\pi^w}}\right)^{\phi_w} \tag{2.6.1}$$

I first consider the monetary policy rules that strictly target a single variable. The results are summarised in Table 2.4.

Table 2.4: Optimal Monetary Policy with Downwardly Rigid Wages

	Baseline	Strict price inflation target	Strict wage inflation target	Strict output gap target
Welfare	-1.5	-2.4	-1.1	-1.0
Variance of price inflation	2×10^{-5}	0	2×10^{-5}	2×10^{-5}
Variance of wage inflation	5×10^{-5}	9×10^{-5}	0	1×10^{-6}
Variance of output	3×10^{-5}	7×10^{-4}	3×10^{-5}	0

I find that household utility is maximised when the central bank strictly targets either the output gap or wage inflation. Targeting price inflation is inferior with household welfare being substantially lower. I also compare the welfare effects from strictly targeting a weighted average of price and wage inflation as defined below

$$1 + \pi_t^c = \xi \left(1 + \pi_t^p\right) + \left(1 - \xi\right)\left(1 + \pi_t^w\right), \tag{2.6.2}$$

where ξ is the weight placed on price inflation relative to wage inflation. I find that welfare is maximised by setting $\xi = 0.15$ as shown in Figure 2.9.

Figure 2.9: Strict Monetary Policy Targets

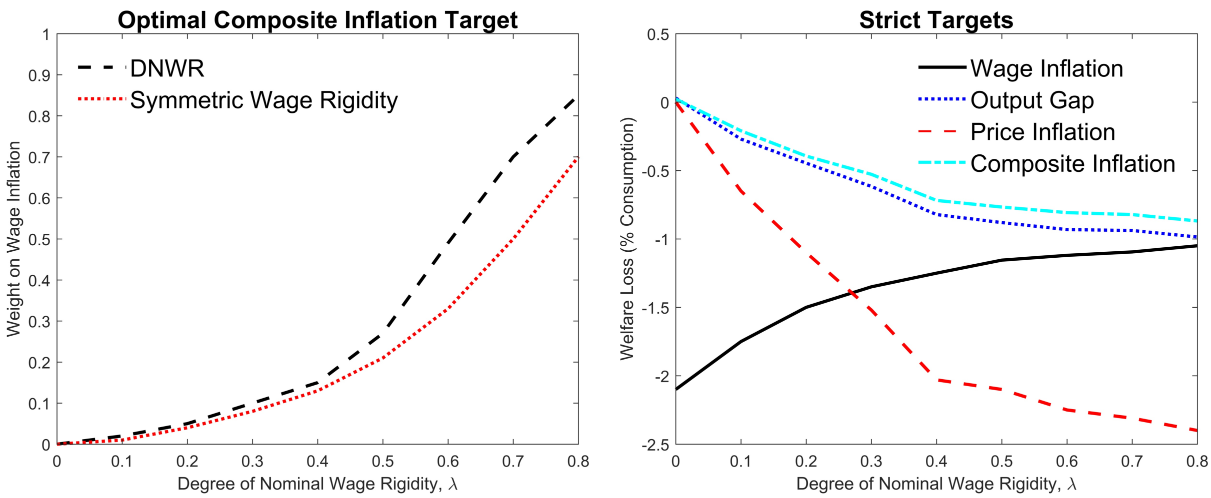

These results imply that the welfare costs from deviations in wage and price inflation are not equal. The high optimal weight on wage inflation implies that deviations in wage inflation have a far greater welfare cost compared with the price inflation volatility. Thus a policy that stabilises wage inflation will be preferable to one that stabilises price inflation. By contrast calibrating a model with symmetric price and wage rigidities as described by Erceg et al. (2000), when calibrated to the same data, places an equal weight on price and wage inflation.

2.6.3 Asymmetric policy rules

When I consider a range of simple monetary policy rules such as 2.6.1 I find a similar result, with the highest weight being placed on stabilising wage inflation. However given the asymmetric nature of this economy, I also consider if an asymmetric monetary policy rule is superior to a conventionally symmetric Taylor rule. I thus simulate a set of asymmetric policy rules in which central banks can act more aggressively to deviations above or below the model's steady state. Concretely I simulate a series of policy rules given by

$$1 + i_t = \frac{(1+\overline{\pi^p})(1+\overline{a})}{\beta}\left(\frac{y_t}{\overline{y}}\right)^{\phi_y+\eta_y}\left(\frac{1+\pi_t^p}{1+\overline{\pi^p}}\right)^{\phi_{\pi^p}+\eta_p}\left(\frac{1+\pi_t^w}{1+\overline{\pi^w}}\right)^{\phi_w+\eta_w} \tag{2.6.3}$$

where

$$\eta_y = \begin{cases} \overline{\eta} & \text{if } y_t - \overline{y} < 0 \\ -\overline{\eta} & \text{if } y_t - \overline{y} \geq 0 \end{cases} \tag{2.6.4}$$

and similarly, for η_p, η_y. To simulate this non-linear policy I extend my linearisation algorithm to calculate separate numerical derivatives for both positive and negative shocks. This greater degrees of freedom allow the approximation to capture the dynamics of the non-linear monetary policy rule. This framework allows the central bank to be more or less reactive to deviations on either side of the target variable's steady state, while maintaining the same baseline level of of hawkishness. That is if a central bank wants to act more actively on one side of the model's steady state it must be be equally passive on the opposite side.[10] I find that an asymmetric response to deviations in the output gap yield no increase in welfare. Indeed as the asymmetry increases welfare actually falls. An asymmetric response to price inflation yields small improvements in household welfare. The largest degree of asymmetry, and those that yield the largest gains in welfare come from responding asymmetrically to wage inflation. This is because wage inflation is the target variable which exhibits the largest degree of asymmetry due to the non-linear nature of the nominal wage rigidity. As such by responding aggressively to small deviations in wage inflation below target, the central bank is better able to adjust interest rates to stabilise the economy.

10 This restriction forces the central bank to reveal its preference over true asymmetry, rather than the optimality that a strict target brings.

Figure 2.10: Optimal Policy Rule Asymmetry

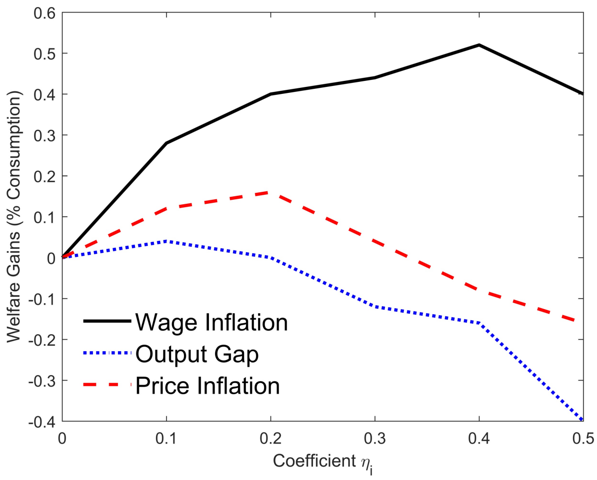

Note: Effect on welfare from different degrees of asymmetry.

2.7 Conclusion

In this chapter I have made three main contributions. First I have solved a non-linear New Keynesian model in which workers are subject to both downward nominal wage rigidities and an idiosyncratic preference shock. Extending previous research on downward nominal wage rigidities, we have expanded the model to include nominal price rigidities creating a meaningful trade-off when setting optimal monetary policy. I then solve this model for its non-linear transitional dynamics in response to demand and productivity shocks, tracking the distribution of wages and the share of workers unable to lower their nominal wages over the equilibrium path.

The second contribution was to use this model to understand how downward nominal wage rigidities skew the distribution of aggregate variables such as wage and price inflation. I show that both measures of inflation are skewed since wage rigidities constrain their ability to fall. This bends the short-run Phillips Curve in the model and helps explain recent trends in economic data which have combined sharp rises in unemployment with only moderate falls in inflation and wage growth.

Finally, I have analysed optimal monetary policy rules in the presence of downward nominal wage rigidities. I found that in comparison to the previous literature on optimal monetary policy under price and wage rigidities, the welfare gains from stabilising wage inflation are relatively high when wages are downwardly rigid. This is because the cost associated with an increase in the share of workers unable to lower their wage is significantly higher than that of price dispersion. In addition, we find that below target inflation can often mask large falls in the output gap. Accordingly a price inflation targeting central bank which responds more aggressively to below average inflation compared to above target inflation could improve welfare within the model.

This model features, with the exception of the labour market, a fairly simple New Keynesian framework. This is done to maintain tractability and to better pin down the effects of the nominal wage rigidity. However, this does suggest several enticing extensions that could be pursued by future research. These include augmenting the model with a role for capital, a constraint on the Taylor rule at the zero lower bound and the inclusion of real rigidities such as a shock to firms' and workers' mark-ups.

Despite the model's simplicity these results provide several insights into the conduct of monetary policy. This model suggests there is some merit placing more weight than the previous literature has suggested on wage inflation rather than price inflation within the monetary policy rule. Finally, despite the simplified nature of this model, these results suggest that the costs that result from downwardly rigid nominal wages can be significant. While their inclusion with a general equilibrium model can be computationally expensive, the data both from micro-level surveys and the macro-level Phillips Curve, suggest they are a key part of understanding the business cycle and how monetary policy can help stabilise it.

Appendix

2.A Derivation of the Model

In this Appendix we derive the model described in Section 2.3 in greater detail.

2.A.1 Households

Households seek to maximise the present value of their utility

$$U_t = \sum_{s=0}^{\infty} \beta^s E_t e^{-\sum_{u=0}^{s-1} D_u} \left\{ \ln C_s - \frac{\gamma}{\gamma+1} \int_0^1 Z_{i,s} L_{i,s}^{\frac{\gamma+1}{\gamma}} \, di \right\} \tag{2.A.1}$$

Subject to the budget constraint

$$B_t = (1 + i_{t-1}) B_{t-1} + \int_0^1 W_{i,t} L_{i,t} \, di - P_t C_t - \tau_t \tag{2.A.2}$$

and the demand for its labour

$$L_{it} = \left(\frac{w_{i,t}}{w_t} \right)^{-\eta} L_t \tag{2.A.3}$$

The Lagrangian for this problem is thus

$$L_t = \sum_{s=0}^{\infty} \beta^s E_t e^{-\sum_{u=0}^{s-1} D_u} \left\{ \begin{array}{c} \ln C_s - \frac{\gamma}{\gamma+1} \int_0^1 Z_{i,s} L_{i,s}^{\frac{\gamma+1}{\gamma}} \, di \\ +\lambda_s \left(-b_t + \frac{1+i_{t-1}}{(1+\pi_t)(1+a_t)} b_{t-1} + \int_0^1 w_{i,t} L_{i,t} \, di - c_t \right) \end{array} \right\} \tag{2.A.4}$$

Where the budget constraint has been detrended by the price and productivity level such that $c_t = \frac{C_t}{A_t}$

2.A.1.1 Optimal Consumption Decision

Solving the first order conditions with respect to output yields

$$\begin{aligned} 0 &= \frac{\partial L}{\partial y_t} = \beta^t e^{-\sum_{s=0}^{t-1} D_s} \left[\frac{1}{c_t} - \lambda_t \right] \\ \therefore \lambda_t &= \frac{1}{c_t} \end{aligned}$$

and with respect to bonds

$$\begin{aligned} 0 &= \frac{\partial L}{\partial b_t} = \beta^t e^{-\sum_{s=0}^{t-1} D_s} \left[-\lambda_t + \beta e^{D_t} \frac{(1+i_t)}{(1+\pi_{t+1})(1+a_{t+1})} \lambda_{t+1} \right] \\ \therefore \lambda_t &= \beta e^{D_t} \frac{(1+i_t)}{(1+\pi_{t+1})(1+a_{t+1})} \lambda_{t+1} \end{aligned}$$

Combining these conditions yields the Euler equation

$$\frac{1}{C_t} = \beta e^{-D_t} (1+r_t) \frac{1}{E_t C_{t+1}} \tag{2.A.5}$$

2.A.1.2 Optimal Wage Decision

When considering the optimal wage setting problem we can drop all terms that do not include the wage and labour supply for member i.

$$L_t^{wage} = \sum_{s=0}^{\infty} \beta^s E_t e^{-\sum_{u=0}^{s-1} D_u} \left\{ -\frac{\gamma}{\gamma+1} \int_0^1 Z_{i,s} L_{i,s}^{\frac{\gamma+1}{\gamma}} \, di + \lambda_s (w_{it} L_{it}) \right\}$$

First we substitute out the amount of labour supplied with the labour demand function which gives the expression.

$$L_t^{wage} = \sum_{s=0}^{\infty} \beta^s E_t e^{-\sum_{u=0}^{s-1} D_u} \left\{ -\frac{\gamma}{\gamma+1} L_s^{\frac{\gamma+1}{\gamma}} Z_{i,s} \left(\frac{w_{i,s}}{w_s} \right)^{-\eta \frac{\gamma+1}{\gamma}} + \lambda_s L_s \left(\frac{w_{i,s}}{w_s} \right)^{1-\eta} \right\}$$

Substituting out for the Lagrange multiplier from the optimal consumption decision yields

$$L^{wage} = \sum_{s=0}^{\infty} \beta^s E_t e^{-\sum_{u=0}^{s-1} D_u} \left\{ -\frac{\gamma}{\gamma+1} L_s^{\frac{\gamma+1}{\gamma}} Z_{i,s} \left(\frac{w_{i,s}}{w_s} \right)^{-\eta \frac{\gamma+1}{\gamma}} + \frac{L_s}{c_s} \left(\frac{w_{i,s}}{w_s} \right)^{1-\eta} \right\}$$

Finally, we can use the aggregate resource constraint to simplify the labour to consumption ratio such that

$$L^{wage} = \sum_{s=0}^{\infty} \beta^s E_t e^{-\sum_{u=0}^{s-1} D_u} \left\{ -\frac{\gamma}{\gamma+1} L_s^{\frac{\gamma+1}{\gamma}} Z_{i,s} \left(\frac{w_{i,s}}{w_t}\right)^{-\eta\frac{\gamma+1}{\gamma}} + s_s \left(\frac{w_{i,s}}{w_s}\right)^{1-\eta} \right\}$$

Which gives equation 2.3.9, which is the contribution to the representative household's utility from worker i's wage related terms.

$$\Omega\left(Z_{it}, w_{it}\right) = s_t \left(\frac{w_{i,t}}{w_t}\right)^{1-\eta} - \frac{\gamma}{\gamma+1} Z_{it} \left[w_{it}^{-\eta} w_t^{\eta} L_t\right]^{\frac{\gamma+1}{\gamma}}$$

Where w_t, L_t and s_t are the aggregate variables for the real wage, total amount of labour and price dispersion among firms respectively. It is this expression we will determine the optimal wage for each worker i.

With flexible wages

When there is no nominal wage rigidity, wages are set by maximising this function only. The first order condition with respect to $w_{i,t}$ is

$$0 = \frac{\partial \Omega\left(Z_{it}, w_{it}\right)}{\partial w_{i,t}} = (1-\eta)\, s_t w_{it}^{-\eta} w_t^{\eta-1} + \eta Z_{it} w_{it}^{-\eta\frac{\gamma+1}{\gamma}-1} \left[w_t^{\eta} L_t\right]^{\frac{\gamma+1}{\gamma}}$$

Which gives the wage setting function

$$w_{i,t}^{flex}\left(Z_{i,t}\right) = \left(\frac{\eta}{\eta-1}\right)^{\frac{\gamma}{\gamma+\eta}} Z_{i,t}^{\frac{\gamma+\eta}{\gamma}} L_t^{\frac{\gamma+1}{\gamma+\eta}} s_t^{-\frac{\gamma}{\gamma+\eta}} w_t \tag{2.A.6}$$

With downwardly rigid wages In the presence of downward nominal wage rigidity the optimal wage problem becomes a dynamic one the wage chosen in the current period will affect the expected utility in the second period. The optimal wage function is thus given by maximising the sum of the current period pay-off and the expected value function in the subsequent period

$$w\left(Z_{i,t}\right) = \underset{w}{\operatorname{argmax}} \left\{ \Omega\left(Z_{it}, w\right) + \beta V_{t+1}\left(\frac{w}{\left(1+\pi_{t+1}\right)\left(1+a_{t+1}\right)}\right) \right\} \tag{2.A.7}$$

We will now show that $w\left(Z_{i,t}\right) \leq w^{flex}\left(Z_{i,t}\right)$ for all possible preference shocks $Z_{i,t}$.

To simplify this proof we transform the decision variable such that $x_{i,t} = \left(w_{i,t}\right)^{1-\eta}$, thus the transformed Bellman equation is

$$V_t(x) =$$
$$(1-\lambda)\int_0^{\infty}\max_{x_{it}\geq 0}\left\{\Omega(Z_{it},x_{i,t})+\beta V_{t+1}\left(x_{i,t}\left[(1+\pi_{t+1})(1+a_{t+1})\right]^{\eta-1}\right)\right\}dF(Z) \tag{2.A.8}$$
$$+\lambda\int_0^{\infty}\max_{x_{it}\leq x}\left\{\Omega(Z_{it},x_{i,t})+\beta V_{t+1}\left(x_{i,t}\left[(1+\pi_{t+1})(1+a_{t+1})\right]^{\eta-1}\right)\right\}dF(Z)$$

Where

$$\Omega(Z_{i,t},x_{i,t}) = s_t\left(\frac{1}{w_t}\right)^{1-\eta}x_{i,t}-\frac{\gamma}{\gamma+1}Z_{it}\left[(x_{i,t})^{\frac{-\eta}{1-\eta}}w_t^{\eta}L_t\right]^{\frac{\gamma+1}{\gamma}}$$

First we note from the Bellman equation that $V'(x)\geq 0$, since a higher x unwinds the constrain in the second term of the Bellman equation. This equivalent to a lower real wage unwinding the constrain placed by the downward nominal wage rigidity.

Second we note that $\Omega(Z_{i,t},x_{i,t})$ is concave with respect to $x_{i,t}$ since

$$\begin{aligned}\frac{\partial}{\partial x_{i,t}}\Omega(Z_{i,t},x_{i,t}) &= s_t\left(\frac{1}{w_t}\right)^{1-\eta}x_{i,t}-\frac{\gamma}{\gamma+1}Z_{it}\left[(x_{i,t})^{\frac{-\eta}{1-\eta}}w_t^{\eta}L_t\right]^{\frac{\gamma+1}{\gamma}}\\ &= s_t\left(\frac{1}{w_t}\right)^{1-\eta}+\frac{\eta}{1-\eta}Z_{it}\left[w_t^{\eta}L_t\right]^{\frac{\gamma+1}{\gamma}}(x_{i,t})^{\frac{\eta}{\eta-1}\frac{\gamma+1}{\gamma}-1}\end{aligned}$$

And

$$\begin{aligned}\frac{\partial^2}{\partial x_{i,t}^2}\Omega(Z_{i,t},x_{i,t}) &= \left(\frac{\eta}{\eta-1}\frac{\gamma+1}{\gamma}-1\right)\frac{\eta}{1-\eta}Z_{it}\left[w_t^{\eta}L_t\right]^{\frac{\gamma+1}{\gamma}}(x_{i,t})^{\frac{\eta}{\eta-1}\frac{\gamma+1}{\gamma}-2}\\ &= \frac{\eta}{1-\eta}\left(\frac{\eta}{\eta-1}\frac{\gamma+1}{\gamma}+1\right)Z_{it}\left[w_t^{\eta}L_t\right]^{\frac{\gamma+1}{\gamma}}(x_{i,t})^{\frac{\eta}{\eta-1}\frac{\gamma+1}{\gamma}-2}\end{aligned}$$

Thus since $\eta>1$, $\gamma>0$, and the aggregate variables are all positive this second derivative will be negative.

From the flexible wage optimality problem we know that

$$\frac{\partial}{\partial x_{i,t}}\Omega\left(Z_{i,t},x_{i,t}^{flex}\right)=0$$

Under the downward nominal wage rigidity this condition is

$$\frac{\partial}{\partial x_{i,t}}\Omega\left(Z_{i,t},x_{i,t}\right)+\beta\left[\left(1+\pi_{t+1}\right)\left(1+a_{t+1}\right)\right]^{\eta-1}V'_{t+1}\left(x_{i,t}\left[\left(1+\pi_{t+1}\right)\left(1+a_{t+1}\right)\right]^{\eta-1}\right)=0$$

Since we know that $\beta\left[\left(1+\pi_{t+1}\right)\left(1+a_{t+1}\right)\right]^{\eta-1}V'_{t+1}\left(x_{i,t}\left[\left(1+\pi_{t+1}\right)\left(1+a_{t+1}\right)\right]^{\eta-1}\right)\geq 0$ this implies that $\frac{\partial}{\partial x_{i,t}}\Omega\left(Z_{i,t},x_{i,t}\right)\leq 0$.

Given the two conditions

$$\begin{aligned}\frac{\partial}{\partial x_{i,t}}\Omega\left(Z_{i,t},x_{i,t}\right) &\leq 0\\ \frac{\partial}{\partial x_{i,t}}\Omega\left(Z_{i,t},x_{i,t}^{flex}\right) &= 0\end{aligned}$$

And the fact that the second derivative, $\frac{\partial^2}{\partial x_{i,t}^2}\Omega$, is negative, this implies that

$$x_{i,t}^{flex}\leq x_{i,t}$$

And thus

$$w^{flex}\left(Z_{i,t}\right)\geq w\left(Z_{i,t}\right)\quad\forall\ Z_{i,t}$$

Thus in the presence of downward nominal wage rigidities workers will always choose to set a lower wage compared with what they otherwise would if wages were perfectly flexible.

Finally, we will show that if a worker is subject to and bound by the downward nominal constraint that they will choose to keep their nominal wage constant. Consider the case where the constraint is just binding, such that $w_t\left(Z\right)=w$. Then all shocks less than Z will also be binding, since otherwise the worker can increase their utility by lowering their wage.

Given that that constraint will be binding for all preference shocks shocks less than Z, any worker subject to the downward nominal wage rigidity at any wage $w>w_t\left(Z\right)$ will keep their wage fixed.

2.A.2 Final Goods Firms

Final goods firms set their prices and output to maximise profit in the current period as shown below.

$$\max P_t Y_t - \int_0^1 P_{j,t} Y_{j,t}\, dj$$

Subject to the Dixit-Stiglitz aggregator which transforms intermediate goods into final goods according to

$$Y_t = \left[\int_0^1 Y_{j,t}^{\frac{\epsilon-1}{\epsilon}}\, dj\right]^{\frac{\epsilon}{\epsilon-1}}$$

Solving for the first order conditions of the is optimisation process with respect to output yields

$$0 = P_t + \lambda_t$$

And with respect to the intermediate good input

$$0 \quad = \quad -P_{j,t} - \lambda_t \left[\int_0^1 Y_{j,t}^{\frac{\epsilon-1}{\epsilon}}\, dj\right]^{\frac{\epsilon}{\epsilon-1}-1} Y_{j,t}^{\frac{\epsilon-1}{\epsilon}-1}$$

Combining these conditions yields the demand curve for intermediate goods

$$Y_{j,t} \quad = \quad \left(\frac{P_{j,t}}{P_t}\right)^{-\epsilon} Y_t$$

Finally we can impose the zero profit condition to solve for the aggregate price index

$$0 = P_t Y_t - \int_0^1 P_{j,t} Y_{j,t}\, dj$$

While yields

$$P_t = \left[\int_0^1 P_{j,t}^{1-\epsilon}\right]^{\frac{1}{1-\epsilon}}$$

2.A.3 Intermediate Goods Firms

Intermediate firms set their optimal price to maximise the discounted stream of expected profits

$$\max_{\{P^*_{j,t},Y_{j,t}\}} \sum_{k=0}^{\infty} \beta^k \frac{\lambda_{t+k}}{\lambda_t} \theta^k \left[\frac{P^*_{j,t} \left(\overline{\pi^p}\right)^{\mu k}}{P_{t+k}} Y_{j,t+k} - \frac{\epsilon-1}{\epsilon} MC^r_{j,t+k} Y_{j,t+k} \right]$$

Profits are maximised subject to three constraints. The first is the firms' linear production function

$$Y_{t,j} = A_t L^d_{t,j}$$

The second constraint is the demand for the firm's output good.

$$Y_{j,t} = \left[\frac{P^*_{j,t}}{P_t}\right]^{-\epsilon} Y_t$$

Finally, a firm's marginal cost is governed by the aggregate wages set by the household, which are assumed to be exogenous by the firms

$$MC^r_{j,t} \quad = \quad w_t$$

Note that since all firms are identical we can drop the subscript j. Substituting in the constraints yields

$$\max_{\{P^*_{j,t},Y_{j,t}\}} \sum_{k=0}^{\infty} \beta^k \frac{\lambda_{t+k}}{\lambda_t} \theta^k \left[\left[\frac{P^*_{j,t} \left(1+\overline{\pi^p}\right)^{\mu k}}{P_{t+k}} \right]^{1-\epsilon} Y_{t+k} - \frac{\epsilon-1}{\epsilon} w_t \left[\frac{P^*_{j,t} \left(1+\overline{\pi^p}\right)^{\mu k}}{P_{t+k}} \right]^{-\epsilon} Y_{t+k} \right] \tag{2.A.9}$$

The first order conditions with respect to the to the optimal price, P^*_t, is

$$0 = \sum_{k=0}^{\infty} \beta^k \frac{\lambda_{t+k}}{\lambda_t} \theta^k \left[(1-\epsilon) \frac{1}{P^*_t} \left[\frac{P^*_t \left(1+\overline{\pi^p}\right)^{\mu k}}{P_{t+k}} \right]^{1-\epsilon} Y_{t+k} + (\epsilon-1)\, w_{t+k} \frac{1}{P^*_t} \left[\frac{P^*_t \left(1+\overline{\pi^p}\right)^{\mu k}}{P_{t+k}} \right]^{-\epsilon} Y_{t+k} \right]$$

Which can be solved to show that the optimal price is

$$P^*_t = \frac{\sum_{k=0}^{\infty} \beta^k \theta^k \left[\frac{(1+\overline{\pi^p})^{\mu k}}{P_{t+k}}\right]^{-\epsilon} w_{t+k}}{\sum_{k=0}^{\infty} \beta^k \theta^k \left[\frac{(1+\overline{\pi^p})^{\mu k}}{P_{t+k}}\right]^{1-\epsilon}}$$

This solution can written in terms of two recursive variables K^r_t and J^r_t where

$$P^*_t = \frac{K^r_t}{J^r_t} \tag{2.A.10}$$

$$\begin{aligned} K_t^r &= MC_t^r (P_t)^{\epsilon} + \beta\theta (1+\overline{\pi^p})^{-\mu\epsilon} E_t \left[K_{t+1}^r\right] \\ J_t^r &= P_t^{\epsilon-1} + \beta\theta (1+\overline{\pi^p})^{\mu(1-\epsilon)} E_t \left[J_{t+1}^r\right] \end{aligned}$$

The aggregate price level is defined by

$$P_t^{1-\epsilon} = \int_0^1 P_{j,t}^{1-\epsilon}\, dj$$

The continuum of firms can be divided into two groups, those who reset their optimal price and those who index it from the previous period.

$$P_t^{1-\epsilon} = \int_\theta^1 (P_t^*)^{1-\epsilon}\, dj + \int_0^\theta (1+\overline{\pi^p})\, P_{t-1}^{1-\epsilon}\, dj$$

Which can be aggregate to give the evolution of the price level.

$$P_t^{1-\epsilon} = \theta (1+\overline{\pi})^{1-\epsilon} P_{t-1}^{1-\epsilon} + (1-\theta) P_t^{*1-\epsilon}$$

Similarly the measure of price dispersion can found by integrating over the production of the intermediate firm

$$L_{j,t} = \frac{1}{A_t} Y_{j,t} \tag{2.A.11}$$

$$\begin{aligned} L_t &= \frac{Y_t}{A_t} \left[\int_0^1 \left(\frac{P_{j,t}}{P_t}\right)^{-\epsilon} di\right] \\ L_t &= \frac{Y_t}{A_t} s_t \end{aligned}$$

Where $s_t = \int_0^1 \left(\frac{P_{j,t}}{P_t}\right)^{-\epsilon} di$

This measure of price dispersion can again be divided into two groups, those who reset their optimal price and those who index it from the previous period.

$$P_t^{-\epsilon} = \int_\theta^1 (P_t^*)^{-\epsilon}\, dj + \int_0^\theta (1+\overline{\pi^p})\, P_{t-1}^{-\epsilon}\, dj$$

Which can be aggregated to give the evolution of the price level.

$$s_t = (1-\theta) \left(\frac{P_t^*}{P}\right)^{-\epsilon} + \theta (1+\overline{\pi^p})^{-\chi\epsilon} s_{t-1}$$

2.A.4 Equilibrium

Differentiation of equation 2.4.6 yields a formula for the evolution of the wage density function.

$$\begin{aligned} g_t(w) &= (1-\lambda) Z_t'(w) f(Z_t(w)) \\ &+\lambda G_{t-1}(w(1+\pi_t)(1+a_t)) Z_t'(w) f(Z_t(w)) \\ &+\lambda(1+\pi_t)(1+a_t) g_{t-1}(w(1+\pi_t)(1+a_t)) F(Z_t(w)) \end{aligned} \tag{2.A.12}$$

Each of these lines corresponds to the three types of worker: those for who free to set their own wage, those that are subject to the wage rigidity but are not bound by it, and those that are both subject to and bound by the wage rigidity.

In equilibrium the real detrended wages of the workers must aggregate to the variable w_t.

$$w_t = \left[\int_0^1 \left(\frac{1}{w_{it}}\right)^{\eta-1} di\right]^{\frac{1}{1-\eta}}$$

We can expand the integral over the workers to the three categories of workers with respect to the downward nominal wage rigidity. This gives

$$\begin{aligned} \therefore 1 &= \left(\frac{\eta}{\eta-1}\right)^{\frac{\gamma}{\gamma+\eta}} s_t^{-\frac{\gamma}{\gamma+\eta}} L_t^{\frac{\gamma+1}{\gamma+\eta}} w_t^{\frac{-\gamma}{\gamma+\eta}} \\ &\times\left\{(1-\lambda)\int_0^\infty \left(\frac{\hat{w}_t(Z)}{w_t(Z)}\right)^{\eta-1} Z^{\frac{\gamma(1-\eta)}{\gamma+\eta}} dF(Z)\right. \\ &+\lambda\int_0^\infty \left(\frac{\hat{w}_t(Z)}{w_t(Z)}\right)^{\eta-1} Z^{\frac{\gamma(1-\eta)}{\gamma+\eta}} G_t(w_t(Z)(1+\pi_t)(1+a_t))\, dF(Z) \\ &\left.+\lambda\int_0^\infty \left[\int_{w_t(Z)}^\infty \left(\frac{\hat{w}_t(Z)}{w_t(Z)}\right)^{\eta-1} Z^{\frac{\gamma(1-\eta)}{\gamma+\eta}} G_t(w_t(Z)(1+\pi_t)(1+a_t))\right] dF(Z)\right\}^{\frac{1}{1-\eta}} \end{aligned}$$

We can now define

$$
\begin{aligned}
Z_t^* \;=\; & \left\{ (1-\lambda) \int_0^{\infty} \left(\frac{\hat{w}_t(Z)}{w_t(Z)} \right)^{\eta-1} Z^{\frac{\gamma(1-\eta)}{\gamma+\eta}} \, dF(Z) \right. \\
& + \lambda \int_0^{\infty} \left(\frac{\hat{w}_t(Z)}{w_t(Z)} \right)^{\eta-1} Z^{\frac{\gamma(1-\eta)}{\gamma+\eta}} G_t \left(w_t(Z) (1+\pi_t)(1+a_t) \right) dF(Z) \\
& \left. + \lambda \int_0^{\infty} \left[\int_{w_t(Z)}^{\infty} \left(\frac{\hat{w}_t(Z)}{w_t(Z)} \right)^{\eta-1} Z^{\frac{\gamma(1-\eta)}{\gamma+\eta}} G_t \left(w_t(Z)(1+\pi_t)(1+a_t) \right) \right] dF(Z) \right\}^{-\frac{\gamma+\eta}{(\eta-1)\gamma}}
\end{aligned} \tag{2.A.13}
$$

A distorted measure of the aggregate disutility of labour among workers which takes into account the distortion of wage and labour choices caused by the downward nominal wage rigidity. This equilibrium condition can be written as

$$
L_t = \left(\frac{\eta-1}{\eta} \right)^{\frac{\gamma}{\gamma+1}} s_t^{-\frac{\gamma}{\gamma+\eta}} w_t^{\frac{\eta}{\gamma+1}} \times \left(\frac{1}{Z_t^*} \right)^{\frac{\gamma}{1+\gamma}}
$$

2.B Linearisation Approach

My approach to simulating the model with aggregate uncertainty has two steps. The first is to calculate numerical derivatives of endogenous variables with respect to the exogenous shocks by calculating the deterministic transition path to a one-off shock that occurs when the economy is in it's steady state and tracing out the path of endogenous variables as it returns to the models steady state. Consider two impulse response functions pertaining to two different, but small, shocks such that $\epsilon_{D_1} = \frac{1}{2}\epsilon_{D_2}$. These impulse response functions are calculated with the following algorithm.

1. Guess a path for the distribution real wages, w_t.
2. Given the path of wages, interest rates and prices solve the Value function backwards in time, using $v(w)$, the steady state solution, as the terminal condition.
3. Using this time path for the value function, calculate the optimal rules for household labour supply and wage choice for each time period.
4. Given the optimal wage choice, solve the forward for the distribution of real wages, using the steady state as the initial condition g^0.
5. Repeat steps 2 through 4 until the time path of prices for the distribution of wages converges.

6. Stop when the solution has converged, when the difference in the path for real wages reaches an arbitrarily small point.

This process produces impulse response functions $\{f_{y,i,t}, f_{y,i,t+1} f_{y,i,t+2}, ..., f_{y,i,t+T}\}$ for each aggregate variables, y, and aggregate shocks, i, to a one-off shock at time t. This sequence can thus be thought of as a set of numerical derivatives showing how the endogenous variable responds over time to the aggregate shock. Thus

$$\{f_{y,i,t}, f_{y,i,t+1} f_{y,i,t+2}, ..., f_{y,i,t+T}\} = \left\{ \frac{\partial y_t}{\partial e_{i,t}}, \frac{\partial y_{t+1}}{\partial e_{i,t}}, \frac{\partial y_{t+2}}{\partial e_{i,t}}, \frac{\partial y_{t+3}}{\partial e_{i,t}}, ... \right\}.$$

By calculating the impulse response functions across multiple sizes of initial shocks, say ϵ_{D_1} and ϵ_{D_2} (see Figure 2.11), we can then calculate second order derivatives using

$$\frac{\partial^2 y_t}{\partial e_{D,t}^2} = \frac{\frac{\partial y_t}{\partial e_{D,2,t}} - \frac{\partial y_t}{\partial e_{D,1,t}}}{e_{D_2,t} - e_{D_1,t}} = \frac{f_{y,\epsilon_{D_2},t} - f_{y,\epsilon_{D_1},t}}{e_{D_1,t}}.$$

Similarly, cross-derivatives can be calculated by comparing how the how the impulse response functions differ when two shocks occur simultaneously

$$\frac{\partial}{\partial e_{a,t}} \frac{\partial y_t}{\partial e_{D,t}} = \frac{\frac{\partial y_t}{\partial e_{\{D_1,a_1\},t}} - \frac{\partial y_t}{\partial e_{\{D_1,a_0\},t}}}{e_{a,t} - 0} = \frac{f_{y,\epsilon_{\{D_1,a_1\},t},t} - f_{y,\{D_1,a_0\},t,t}}{e_{a,t}}.$$

Figure 2.11: Impulse Response Functions

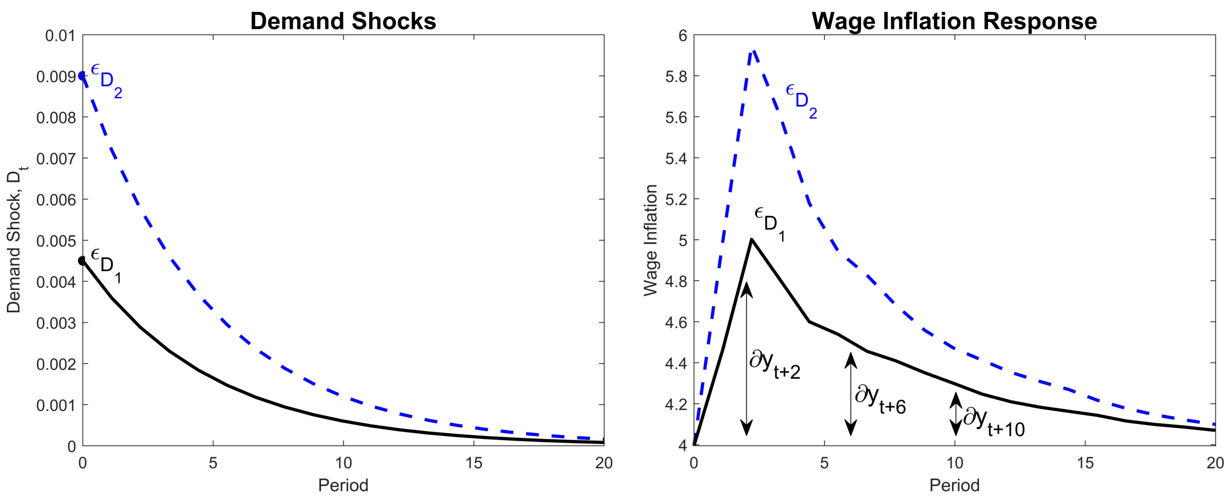

Once the derivatives have been calculated we can simulated a draw of shocks for a large number of periods and then use the numerical derivatives to approximate the impact on outcome for the aggregate variables.[11] Given the asymmetric nature of the downward nominal wage rigidity ...

Given the sequence of shocks

$$\{\boldsymbol{e}_1, \boldsymbol{e}_2, \boldsymbol{e}_3, ...\}$$

In order to capture the non-linearities in aggregate variables, I use the second order approximation

$$y_t = \overline{y} + \sum_{j=t}^{t-20} \nabla_j \boldsymbol{e}_j + \left\{ \sum_{j=t}^{t-20} \boldsymbol{e}_j H_j \boldsymbol{e}_j \right\},$$

where H is the Hessain matrix of second derivatives.

$$H_j = \begin{bmatrix} \frac{\partial^2 y_t}{\partial a^2} & \frac{\partial^2 y_t}{\partial a \partial D} \\ \frac{\partial^2 y_t}{\partial D \partial a} & \frac{\partial^2 y_t}{\partial D^2} \end{bmatrix}$$

Figure 2.12: Impulse Response Functions and Approximations

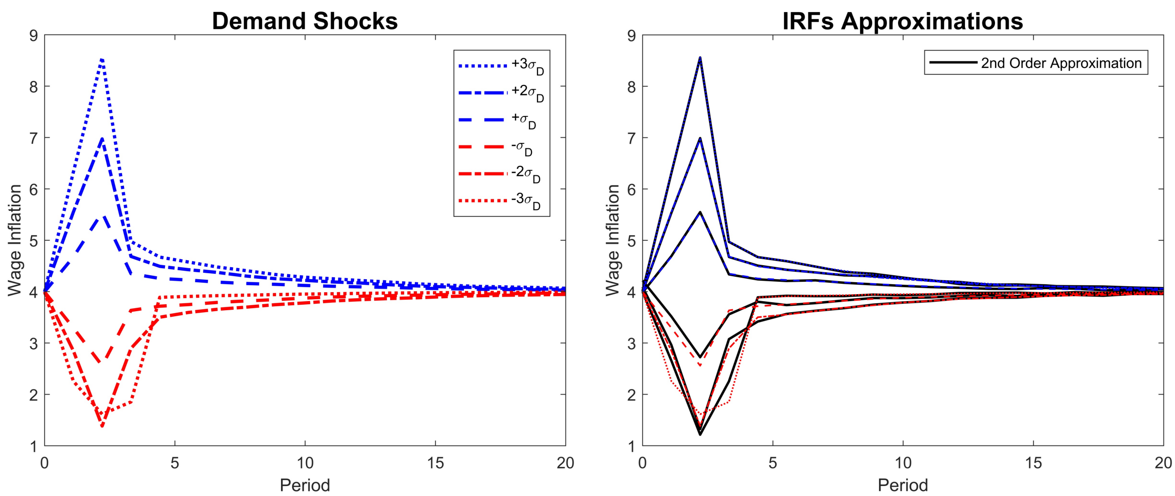

11 I simulate 10,500 periods. We initialise the model at steady state and then dropping the first 500 periods from the sample.

3 Downward Nominal Wage Rigidity and the Zero Lower Bound[1]

Abstract

I estimate the optimal rate of inflation in a non-linear New Keynesian model in which downwardly rigid wages and a zero lower bound on nominal interest rates are jointly modeled. In this economy a higher rate of inflation gives workers more flexibility when setting real wages and the central bank more room to vary interest rates, at the cost of greater price dispersion in the goods market. I discuss how downwardly rigid wages can amplify or mitigate the welfare loss caused by the zero lower bound and how this varies with the structure of the model. Finally, I find that the optimal rate of inflation is increased by the presence of both nominal frictions, when modeled either separately or in tandem, and is 3 per cent in the baseline calibration.

JEL Classification: E31, E43, E58, J30.

Keywords: Optimal Inflation Rate, Downward Nominal Wage Rigidity, Zero Lower Bound.

1 I would like to thank seminar participants at the University of Oxford, and the European Economic Association Annual Congress, Money Macro and Finance Annual Conference for useful comments. I would like to acknowledge the use of the University of Oxford Advanced Research Computing (ARC) facility in carrying out this work (http://dx.doi.org/10.5281/zenodo.22558).

3.1 Introduction

Downwardly rigid wages and the zero lower bound on nominal interest rates are the two of the most commonly cited reasons for a positive inflation target Coibion et al. (2012). While this view is supported by a broad range of literature, the two frictions are commonly modeled independently. This chapter will develop a model which contains both nominal rigidities in a framework with heterogeneous workers and answer two questions: how do downwardly rigid nominal wages interact with a zero lower bound on nominal interest rates, and how do the two frictions impact the optimal rate of inflation. I find that both frictions justify a positive rate of inflation individually and, in contrast to previous literature, a higher optimal rate when modeled jointly.

Previous studies analysing monetary policy where aggregate wages are downwardly rigid and there is a zero lower bound on nominal interest rates have generally assumed that all workers have an identical nominal wage. Under this framework the downward nominal wage constraint binds zero workers if the wage inflation rate is positive, and all workers if it would otherwise be negative. However, this approach ignores the widespread empirical evidence on downward nominal wage rigidities and implies that wage rigidities will only constrain the economy when wage inflation hits 0 per cent - an exceptionally rare event even during recessions (see Figure 3.1). Examples of this approach include Schmitt-Grohé and Uribe (2012) who model both downwardly rigid wages and an effective lower bound on nominal interest rates to generate a model with multiple equilibrium and Coibion et al. (2012) who find that a constraint on negative wage inflation will stabilise firms' marginal costs and reduce frequency of episodes at the zero lower bound. More recently the conventional wisdom of the effect of these frictions has been challenged by Amano, Gnocchi, et al. (2017) who find that in a representative agent model with nominal price stickiness, downwardly rigid wages can increase welfare by reducing the both the probability and the severity of episodes at the zero lower bound. The consequence of this result is that when downward nominal wage rigidities are included in the model of the zero lower bound they can lead to a lower optimal rate of inflation.

The main contribution of this chapter is show that that both of these findings break down in a heterogeneous model of downward nominal wage rigidities, and that both the zero lower bound and asymmetric wage rigidities lead to an increase in the optimal rate of inflation.

Figure 3.1: Nominal Wage Growth 1985-2018

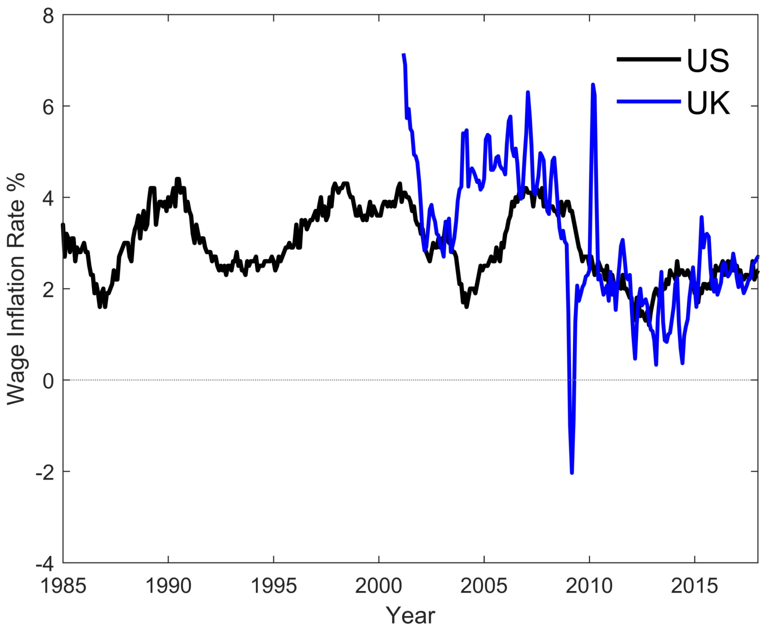

Note: US series is annual growth in hourly earnings of production and non-supervisory employees, total private sector, monthly. UK series is the nominal wage component of annual growth in average weekly earnings, including bonuses.

This thesis will build off the framework of Daly and Hobijn (2014) and Dordal-i Carreras et al. (2016) and solve the optimal rate of inflation of the model which includes an idiosyncratic preference shock to workers, a zero lower bound on nominal interest rates and infrequent, but long-lived, zero-bound episodes.

This chapter makes several further contributions to the literature. First, we examine the margins through which the two nominal rigidities interact. I show that episodes at the zero lower bound decrease real wage flexibility, as inflation falls when interest rates are constrained. Furthermore I find that downwardly rigid wages change the probability of nominal interest rates being constrained at the lower bound - however the direction of this effect will vary depending on the coefficients in the Taylor Rule. Finally, I demonstrate that downwardly rigid wages will lower real interest rates when nominal rates are fixed at zero, as the wage friction prevents the inflation rate from falling in times of low demand. Since these margins affect the optimal rate of inflation in opposite directions, it is not clear ex-ante which effect dominates. However in the baseline version of my model I find that when both frictions are modeled jointly the optimal rate of inflation increases.

The key driver behind my results is how wage heterogeneity breaks down the insurance-like role of downwardly rigid wages. One of the key benefits of downwardly rigid wages is their ability to prop up the inflation rate in times of low demand. This helps avoid interest rates hitting the zero lower bound and lowers the real interest rate when they do so. When all households have the same nominal wage this increase in inflation will only occur when the wage inflation rate approaches 0 per cent i.e. in periods of low demand. Thus the wage rigidities only restrain workers' choices when the benefit of doing so is high, making it an effective form of insurance against a deflation spiral.

By contrast when households are subject to an idiosyncratic preference shock and the downward nominal wage rigidity applies stochastically, wages will be constrained even when wage inflation is greater then zero, and a degree of wage flexibility will remain even in times of very low inflation. This makes the downward nominal wage rigidity a much less effective form of insurance against the possibility of the zero lower bound. In this environment the costs of downwardly rigid wages are higher and less counter-cyclical which is why a higher rate of inflation, which unwinds the friction, becomes optimal.

When I simulate the model I estimate the optimal rate of inflation to be 3 per cent, higher than the majority of inflation targets among the OECD countries. In addition we show that that the optimal rate of inflation varies inversely with the trend rate of productivity growth in the economy. This result highlights a new mechanism through which secular stagnation, the widespread decline in productivity and growth rates, might lead to a higher optimal rate of inflation.

In the next section I will describe how this chapter relates to the literature. Section 3.3 outlines the model and how I micro-founded the downwardly nominal wage rigidity. Section 3.4 describes how I solve for the model's steady state using value function iteration. Section 3.5 discusses the various channels through which the two frictions interact. Finally, I analyse the optimal rate of inflation under the presence of both rigidities in section 3.7 before concluding.

3.2 Literature Review

This thesis contributes to the literature on the intersection of the topics of downward nominal wage rigidity and the optimal rate of inflation. We will discuss the previous literature on each topic in turn. The existence of downward nominal wage rigidity has long been used to justify a positive long-term inflation target. Tobin (1972) informally outlined the case for a positive rate of inflation in order to “grease the wheels” of the labour market. A positive inflation rate allows real wages to decrease even if nominal wages are unable to fall. Thus by increasing the rate of inflation the downward nominal wage rigidity is unwound and its effect on the labour market is minimised. While Tobin presented an informal argument in favour of positive rates of inflation, several papers have since used formal models to attempt to estimate the impact of downwardly rigid wages on optimal monetary policy.

However this intuition does not always bear out in formal models. Coibion et al. (2012) was one of the first papers to formally model the combination of downwardly rigid wages with a zero lower bound. They find wage rigidities significantly lowers the optimal inflation rate. This is because downward wage rigidity lowers the volatility of marginal costs and hence of inflation. In addition the wage rigidities moderate declines in inflation decreasing the likelihood of nominal interest rates being constrained. In a similar vein, Amano et al. (2017) considers the interaction of downward nominal wage rigidity and the zero lower bound on nominal interest rates in a representative agent framework with sticky prices. They find that combining the zero lower bound with downward nominal wage rigidities reduces both the frequency and duration of liquidity traps and the welfare cost of these periods. This is because the downwardly rigid wages act to moderate the fall in inflation during large downturns when the zero lower bound binds. In this environment downward nominal wage rigidity actually lowers the optimal rate of inflation.

Kim and Ruge-Murcia (2009) set up a simple model in which changes in nominal wages are subject to an asymmetric adjustment cost in which decreases in nominal wages cost more than increases. They then compare this model to one with quadratic adjustment costs akin to Rotemberg (1982). They consider a symmetric equilibrium in which all households supply the same labour, at the same wage. Estimating the model using perturbation methods, they solve the dynamics of the model in the presence of a productivity and a preference shock. In order to avoid the costs of wage deflation, they find a positive rate of inflation, around 0.4 per cent, is optimal.

Similarly, Abo-Zaid (2013b) constructs a model containing downward nominal wage rigidity using an asymmetric wage adjustment cost and augment the model with a search and matching framework. Abo-Zaid maintains the assumption that all households and firms are identical and thus agree to the same nominal wage. Thus in steady state, the optimal rate of inflation is zero, or negative when money demand is included. To generate a positive optimal rate of inflation Abo-Zaid simulates the model under the presence of an aggregate productivity shock to find an optimal rate of inflation of 1 per cent. While this method remains tractable the assumption that downward wage rigidities do not affect the economy when inflation is greater than or equal to its steady state level, it is at odds with the empirical data and likely biases the optimal rate of inflation and the welfare cost of these nominal wage rigidities downwards.

Schmitt-Grohé and Uribe (2012) uses the same framework to examine the impact of downward nominal wage rigidities and how they interact with a monetary policy rule which is subject to the zero lower bound. However, they assume that nominal wages can fall when the unemployment rate is elevated. They develop a model which can generate a liquidity trap combined with a jobless recovery. This result is due to the existence of two equilibria, the first at full employment with inflation at its target level and the second with high unemployment, deflation and interest rates trapped at the zero lower bound. They show that confidence shocks can shift the economy between the two equilibria and consider how monetary policy can be used to return the economy to the full employment equilibrium. Unconventionally, they argue that high interest rates can boost inflation expectations and thereby increase growth contra to the standard intuition. Although Schmitt-Grohé and Uribe (2012) consider optimal monetary policy under downwardly rigid wages, they do so only in the context of escaping the liquidity trap and switching between equilibria.. They do not consider how monetary policy should respond to small deviations around the equilibria.

Fagan and Messina (2009) take a different approach to solve for the optimal rate of inflation. They assume that workers are subject to idiosyncratic productivity shocks and that nominal wage changes are subject to an asymmetric adjustment cost, with decreases in nominal wages costing more than increases. They then solve for the steady state of a general equilibrium model which also includes capital and symmetric pricing frictions. They report that the optimal inflation rate is positive, but varies widely across countries according to the different characteristics of their labour markets.

3.3 The Model

This section extends the model of Chapter 2 by incorporating a zero-lower bound on interest rates, a mark-up shock in the monopolistically competitive markets, and a long-term risk premium shock a la Dordal-i Carreras et al. (2016) which can generate infrequent, but long-lived episodes at the zero lower bound. This extension keeps the model as parsimonious as possible while creating a meaningful policy trade-off when determining the optimal rate of inflation in the presence of downward nominal wage rigidities and the zero lower bound.

3.3.1 Households

The economy is populated by a representative household. Each household is comprised of a unit measure of workers who share consumption among themselves, but each supply a differentiated type of labour, $L_{i,t}$, in exchange for a nominal wage, $W_{i,t}$. Workers are differentiated by an idiosyncratic, and uninsurable, shock to the disutility derived from working $Z_{i,t}$. Household utility in period t is derived from consumption, C_t, which is uniform across workers. The households' disutility from working is given by aggregating across the individual workers' disutility which is given by $Z_{i,t}L_{i,t}^{\frac{\gamma+1}{\gamma}}$. Households seek to maximise the present value of their expected utility in period t, U_t, as shown below

$$U_t = \sum_{s=0}^{\infty} \beta^s E_t e^{-\sum_{u=0}^{s-1} D_u} \left\{ \ln C_s - \frac{\gamma}{\gamma+1} \int_0^1 Z_{i,s} L_{i,s}^{\frac{\gamma+1}{\gamma}} \, di \right\} \tag{3.3.1}$$

Where the parameter γ is the Frisch elasticity of labour supply, and β is the steady state discount rate. D_t is an exogenous shock in period t to the households' discount factor. The idiosyncratic preference shock is drawn from a truncated log-normal distribution such that $\ln Z \sim N\left(-\frac{\sigma_z^2}{2}, \sigma_z\right)$, where σ_z is the standard deviation of the shock. The first term in this sum denotes the utility gained from consumption using a logarithmic utility function. The second term gives the aggregate disutility across workers.

Since the utility function is separable with respect to consumption and leisure, consumption will be shared equally among all members of the household.[2] However, the optimal level of labour to supply will vary according to each worker's idiosyncratic disutility shock. The household will seek to maximise the present value of utility by choosing average consumption, C_t, and the workers' wages and labour supply $\{W_{i,t}, L_{i,t}\}_{i=0,t=0}^{1,\infty}$. The representative household solves this maximisation problem subject to four constraints.

The first is the budget constraint, which governs the evolution of household assets in the form of a risk-free bond, where the amount of bonds held at the end of period t is denoted B_t

$$B_t = (1 + i_{t-1}) Q_t B_{t-1} + \int_0^1 W_{i,t} L_{i,t} \, di - P_t C_t - \tau_t + \Pi_t, \tag{3.3.2}$$

where P_t is the aggregate price level, i_t is the nominal interest rate, τ_t are lump sum taxes, and Π_t is any profits from the firms in period t. Q_t is a Markov-switching risk premium shock which generates infrequent, but long-lived, periods of low interest rates.

$$Q_t = \begin{cases} -I_Q & \text{for } T_Q \text{periods with probability } P_Q \\ 0 & \text{otherwise} \end{cases} \tag{3.3.3}$$

2 This is because the marginal utility from consumption is invariant to both the workers' labour supply and their preference shock. Thus given the diminishing returns to consumption, the household will maximise its welfare by sharing consumption equally between members.

The second constraint is the firms' demand for each type of labour. The demand for each type of labour i is governed by aggregate employment, L_t, the aggregate real wage w_t, and the real wage for the type of labour $w_{i,t}$,

$$L_{i,t} = \left(\frac{w_{i,t}}{w_t}\right)^{-\eta} L_t \tag{3.3.4}$$

The parameter η is the elasticity of demand for labour and measures how substitutable the different types of labour are from each other. It should be noted that the real wage is detrended by both the level of productivity, A_t, and the price level, P_t such that

$$w_{i,t} = \frac{W_{i,t}}{A_t P_t} \tag{3.3.5}$$

The third constraint is the AR(1) process for the exogenous shock to the households' stochastic discount rate.

$$D_t = \rho^D D_{t-1} + \varepsilon_t^D \tag{3.3.6}$$

Where ρ_D is the degree of persistence in aggregate demand shocks and ε_t^D is an independently normally distributed random variable, $\varepsilon_t^D \sim N(0, \sigma_D)$.

The final constraint is that each worker is subject is a stochastic downward nominal wage rigidity. This is modeled using the time-dependent framework from Calvo (1983). We assume that each period a randomly selected fraction of workers, governed by the parameter $\lambda \in [0, 1)$, are unable to decrease their nominal wage and thus they are subject to the constraint $W_{i,t} \geq W_{i,t-1}$. The remaining $(1 - \lambda)$ share of workers are free to set their nominal wage optimally, but do so in the knowledge that they may be subject to the downward nominal wage rigidity in future periods.

Each household takes the path for aggregate nominal wage index, W_t, total employment L_t, productivity a_t, nominal interest rates i_t and the price level P_t as given.

3.3.1.1 Optimal Consumption Decision

Solving the first order conditions with respect to bonds, B_t and consumption, C_t, yields the standard inter-temporal Euler equation

$$\frac{1}{C_t} = \beta e^{-D_t+Q_t}(1+r_t)\frac{1}{E_tC_{t+1}} \tag{3.3.7}$$

where consumption in period t is governed by the real interest rate r_t and expectations of consumption in period $t+1$. In turn the real interest rate is governed by the Fisher equation

$$1+r_t = \frac{1+i_t}{1+E_t\pi_{t+1}} \tag{3.3.8}$$

3.3.1.2 Optimal Wage Decision

The solution for the households' optimal wage setting problem follows the framework of Chapter 2. We will solve the optimal wage setting problem first in the simplified case of flexible wages, and then in the more general case where downward nominal wage rigidities and symmetric nominal wage rigidities apply.

Each household sets wages to maximise the expected present value of their utility, thus while they internalise the wage-setting decision of other workers within the household, they do not internalise the decisions of other households, and so treat aggregate real wages, w_t, aggregate employment, L_t, and price dispersion, s_t, as exogenous.

As we show in Appendix 3.A, the wage optimising problem can be simplified to the household choosing the path of wages, $w_{i,t}$, to maximise the expression

$$E_t\sum_{t=0}^{\infty}\beta^t e^{-\sum_{s=0}^{t-1}D_s}\left(\Omega\left(Z_{i,t},w_{i,t}\right)\right), \tag{3.3.9}$$

where $\Omega\left(Z_{it},w_{it}\right)$ is the per period pay-off to the household with respect to the worker i's choice of wage $w_{i,t}$. It can be written as a function of the worker specific variables, the disutility shock Z_{it}, real wage w_{it}, and aggregate variables

$$\Omega\left(Z_{i,t},w_{i,t}\right) = s_t\left(\frac{w_{i,t}}{w_t}\right)^{1-\eta} - \frac{\gamma}{\gamma+1}Z_{i,t}\left[\left(\frac{w_{i,t}}{w_t}\right)^{-\eta}L_t\right]^{\frac{\gamma+1}{\gamma}}, \tag{3.3.10}$$

where w_t, L_t and s_t are the aggregate variables for the real wage, total amount of labour and price dispersion among firms respectively.

Under flexible wages If there are no nominal wage rigidities this problem can be solved analytically. In this case the Bellman equation collapses to

$$V_t \quad (w) = \int_0^\infty \max_{w_{it}\geq 0} \left\{ \Omega\left(Z_{it}, w_{it}^{flex}\right)\right\} dF\left(Z_{it}\right) \tag{3.3.11}$$

Which will be invariant with respect to w. Since wages are perfectly flexible they can be re-optimised in every period. Thus workers will set their wage to optimise their current period pay-off without having to weigh the consequences for future periods. The wage function can thus be derived by maximising the per period pay-off given in equation 3.3.10. The optimal wage is a function of the aggregate variables, which are assumed to be exogenous by the household, and the individual worker's idiosyncratic preference shock Z_{it} as shown below

$$w_{i,t}^{flex} = \left(\frac{\eta}{\eta-1}\right)^{\frac{\gamma}{\gamma+\eta}} Z_{i,t}^{\frac{\gamma+\eta}{\gamma}} L_t^{\frac{\gamma+1}{\gamma+\eta}} s_t^{-\frac{\gamma}{\gamma+\eta}} w_t. \tag{3.3.12}$$

With downwardly rigid wages In the presence of downward nominal wage rigidity, workers who are not subject to the nominal wage constraint will set their wage according to the function

$$w_t\left(Z_{i,t}\right) = \underset{w}{\operatorname{argmax}} \left\{ \Omega\left(Z_{it}, w\right) + \beta V_{t+1}\left(\frac{w}{\left(1+\pi_{t+1}\right)\left(1+a_{t+1}\right)}\right)\right\}. \tag{3.3.13}$$

If workers are unaffected by the downward nominal wage constraint, all workers who choose the same wage will have drawn the same preference shock, $Z_{i,t}$.

Workers who are subject to the nominal wage constraint fall into two categories. The first group consists of workers who draw a preference shock such that $w\left(Z_{it}\right) \geq w$. This group of workers will have a preferred wage, $w\left(Z_{i,t}\right)$, that is higher than their wage at the beginning of the period, w. Thus the downward nominal wage rigidities will not be a binding constraint. The second group are workers who draw a preference shock $w\left(Z_{it}\right) < w$. For these workers the constraint is binding. They would ideally like to lower their nominal wage, but are prevented by the downward nominal wage rigidity. Instead these workers will choose to keep their nominal wage constant, keeping their real wage fixed at w, as shown in Appendix 3.A.

With symmetrically rigid wages In the presence of a symmetric rigidity workers who are subject to the constraint will have their nominal wage fixed a la Calvo (1983). Workers who are not subject to the nominal wage constraint will set their wage according to the value function

$$\begin{aligned} V_t^{sym}(w) = \\ (1-\lambda)\int_0^\infty \max_{w_{i,t}\geq 0}\left\{\Omega\left(Z_{i,t},w_{i,t}\right)+\beta e^{-D_t}V_{t+1}^{sym}\left(\frac{w_{i,t}}{\left(1+\pi_{t+1}\right)\left(1+a_{t+1}\right)}\right)\right\}dF\left(Z\right) \\ +\lambda\int_0^\infty\left\{\Omega\left(Z_{i,t},w\right)+\beta e^{-D_t}V_{t+1}^{sym}\left(\frac{w}{\left(1+\pi_{t+1}\right)\left(1+a_{t+1}\right)}\right)\right\}dF\left(Z\right) \end{aligned} \tag{3.3.14}$$

If workers are unaffected by the symmetric wage constraint, all workers who choose the same wage will have drawn the same preference shock, $Z_{i,t}$ according to the wage setting function

$$w_t^{sym}\left(Z_{i,t}\right)=\underset{w}{\operatorname{argmax}}\left\{\Omega\left(Z_{it},w\right)+\beta V_{t+1}^{sym}\left(\frac{w}{\left(1+\pi_{t+1}\right)\left(1+a_{t+1}\right)}\right)\right\} \tag{3.3.15}$$

3.3.2 Final Goods Firms

Final goods firms produce goods for consumption, Y_t, by combining intermediate goods, $Y_{j,t}$, using a Dixit-Stiglitz aggregator. These firms are perfectly competitive and thus have zero profit. Their prices, P_t, are perfectly flexible and they choose the amount of each intermediate good, $Y_{j,t}$, to buy in the current period as shown below

$$\max_{Y_{j,t}} P_t Y_t - \int_0^1 P_{j,t}Y_{j,t}\,dj \tag{3.3.16}$$

subject to the Dixit-Stiglitz aggregator which transforms intermediate goods into final goods according to

$$Y_t=\left[\int_0^1 Y_{j,t}^{\frac{\epsilon-1}{\epsilon}}\,dj\right]^{\frac{\epsilon}{\epsilon-1}}. \tag{3.3.17}$$

Solving for the first order conditions of this optimisation problem, shown in detail in Appendix 3.A, yields the standard demand curve

$$Y_{j,t} = \left(\frac{P_{j,t}}{P_t}\right)^{-\epsilon} Y_t \tag{3.3.18}$$

while the zero profit condition pins down the aggregate price index

$$P_t = \left[\int_0^1 P_{j,t}^{1-\epsilon}\right]^{\frac{1}{1-\epsilon}}. \tag{3.3.19}$$

3.3.3 Intermediate Goods Firms

There is a continuum of firms who produce intermediate goods in this economy, denoted by the subscript j, each of which produce a differentiated good in a monopolistically competitive market. Firms produce goods using each type of labour supplied by the workers that comprise the households and are subject to pricing frictions a la Calvo (1983). Each period a randomly selected fraction, θ, of firms are allowed to reset their price. Those that are not able to flexibly adjust their price instead partially index it by the central bank's target rate of inflation, $\overline{\pi^p}$. The firms also receive a subsidy which unwinds the mark-up caused by their market power. The firms thus optimise the present discounted value of profits by optimally choosing the amount of output to produce $Y_{j,t}$, how much of each type of labour, $L_{j,t}$, to hire and, if they are able to, the optimal price to set $P^*_{j,t}$. When setting their price they maximise the discounted stream of expected profits

$$\max_{\{P^*_{j,t}, Y_{j,t}\}} \sum_{k=0}^{\infty} \beta^k \frac{\lambda_{t+k}}{\lambda_t} \theta^k \left[\frac{P^*_{j,t}\,(1+\overline{\pi^p})^{\mu k}}{P_{t+k}} Y_{j,t+k} - \frac{\epsilon-1}{\epsilon} MC^r_{j,t+k} Y_{j,t+k}\right], \tag{3.3.20}$$

where $MC^r_{j,t}$ is the real marginal cost of firm j in period t, ϵ is the elasticity of substitution between goods and μ is the degree of price indexation to the central bank's target rate of inflation. Since the firms are owned by the households the stochastic discount factor of the firms is governed by the households' marginal utility of consumption λ_t. The stochastic discount factor between periods t and $t+k$ is given by the term $\beta^k \frac{\lambda_{t+k}}{\lambda_t}$. Each firm is provided a subsidy of $\frac{\epsilon-1}{\epsilon}$ which unwinds the mark-up caused by the firms' market power.

The discounted expected stream of profits are maximised subject to three constraints. The first is the firms' linear production function

$$Y_t = A_t L_t, \tag{3.3.21}$$

where A_t is the level of labour productivity and L_t represents an aggregate index of the different types of labour provided by the household which is given by

$$L_t = \left[\int_0^1 L_{i,t}^{\frac{\eta-1}{\eta}} \, di \right]^{\frac{\eta}{\eta-1}}, \tag{3.3.22}$$

where η is the Frisch elasticity of demand for individual workers. The firm will choose each type of labour, $L_{i,t}$, to minimise its wage bill subject to the aggregate amount of labour required. This cost minimisation results in the labour demand function

$$L_{i,t} = \left(\frac{W_t}{W_{i,t}} \right)^{\eta} L_t, \tag{3.3.23}$$

where the nominal wage aggregate W_t is given by

$$W_t = \left[\int_0^1 W_{i,t}^{1-\eta} \, di \right]^{\frac{1}{1-\eta}}. \tag{3.3.24}$$

Growth in labour productivity is given by the exogenous process

$$(a_t - \overline{a}) = \rho^a \left(a_{t-1} - \overline{a} \right) + \varepsilon_t^a, \tag{3.3.25}$$

where $\overline{a}$ is the steady state rate of growth in the economy, ρ_a is the degree of persistence in shocks to the productivity growth rate and ε_t^a is an independently normally distributed random variable, $\varepsilon_t^a \sim N(0, \sigma_a)$.

The second constraint is the demand for the firm's intermediate good

$$Y_{j,t} = \left[\frac{P_{j,t}}{P_t} \right]^{-\epsilon} Y_t. \tag{3.3.26}$$

Finally, a firm's marginal cost is governed by the aggregate wages set by the household, which are assumed to be exogenous by the firms

$$MC_{j,t}^r = \frac{W_t}{P_t A_t} = w_t. \tag{3.3.27}$$

Solving this constrained optimisation problem, as outlined in Appendix 3.A, gives the solution for the firms' optimal price. Note that since all firms are identical we can drop the subscript j. The solution for the optimal price to choose, P_t^*, is given by

$$P_t^* = \frac{K_t^r}{J_t^r} + \epsilon_t^K \tag{3.3.28}$$

Where K_t^r and J_t^r are recursive variables defined by

$$K_t^r = MC_t^r (P_t)^{\epsilon} + \beta\theta (1+\overline{\pi^p})^{-\mu\epsilon} E_t \left[K_{t+1}^r\right] \tag{3.3.29}$$

$$J_t^r = P_t^{\epsilon-1} + \beta\theta (1+\overline{\pi^p})^{\mu(1-\epsilon)} E_t \left[J_{t+1}^r\right] \tag{3.3.30}$$

The aggregate price level evolves according to the price that firms with price flexibility set in period t and the prices that the remainder of firms had in the previous period, indexed by the central bank's inflation target.

$$P_t^{1-\epsilon} = \theta (1+\overline{\pi})^{\mu(1-\epsilon)} P_{t-1}^{1-\epsilon} + (1-\theta) P_t^{*1-\epsilon}. \tag{3.3.31}$$

Once the price level has been pinned down we can solve for the inflation rate which is given by $\pi_t = \frac{P_t}{P_{t-1}} - 1$. Since both prices and wages are subject to nominal constraints we separately track wage inflation, π_t^w, using the change in the real wage, price inflation and rate of productivity growth according to

$$1+\pi_t^w = \frac{W_t}{W_{t-1}} = \frac{w_t}{w_{t-1}} (1+a_t)(1+\pi_t). \tag{3.3.32}$$

3.3.4 Monetary Policy

Monetary policy is set by a central bank which uses the nominal interest rate to stabilise the economy. In the baseline model the central bank follows a simple Taylor Rule, setting interest rates based on deviations in inflation from its target, $\overline{\pi}$, and in movements in the output gap y_t subject to a zero lower bound on nominal interest rates

$$i_t = \max\left\{\frac{(1+\overline{\pi^p})(1+\overline{a})}{\beta}\left(\frac{y_t}{\overline{y}}\right)^{\phi_y}\left(\frac{1+\pi_t^p}{1+\overline{\pi^p}}\right)^{\phi_\pi}\left(1+\epsilon_t^M\right) - 1, 0\right\} \tag{3.3.33}$$

The output gap is given by

$$y_t = \frac{Y_t}{A_t} \tag{3.3.34}$$

where $\overline{y}$ is the steady state outcome associated with an output gap of zero. The parameters ϕ_y and ϕ_π are the non-negative Taylor coefficients on the output gap and inflation respectively and determine how aggressively the central bank changes the nominal interest rate in response to deviations from the model's steady state.

3.3.5 Market Clearing Conditions

In equilibrium the goods market must clear such that aggregate consumption is equal to the firms' output

$$C_t = Y_t \tag{3.3.35}$$

Finally, we need a market clearing condition for each of the intermediate goods. Since each firm will in general have different prices due to the Calvo price setting framework, when aggregating output across firms we need to calculate the degree of price dispersion. This price dispersion creates a wedge between the aggregate labour employed and the total output produced by firms as described in detail in Appendix 3.A.

$$Y_t = \frac{1}{s_t} A_t L_t \tag{3.3.36}$$

Where s_t is a measure of price dispersion within the economy. This measure of price dispersion evolves according to

$$s_t = (1-\theta)\left(\frac{P_t^*}{P}\right)^{-\epsilon} + \theta\left(1+\overline{\pi^p}\right)^{-\mu\epsilon} s_{t-1} \tag{3.3.37}$$

3.4 Equilibrium

The equilibrium for this economy is a set of the endogenous variables $\{L_t, Y_t, C_t, \pi_t^p, \pi_t^w, w_t, i_t, r_t, a_t, D_t, s_t, b_t, P_t^*, K_t^r, J_t^r\}$ and the endogenous functions $\{w\left(Z_{it}\right), V_t\left(w\right)\}$ which satisfies equations 3.3.2 to 3.3.37 for all periods t given the exogenous variables

$\left\{\varepsilon_t^a, \varepsilon_t^D, \epsilon_t^K\right\}$. I will consider the equilibrium of the model under three assumptions about nominal wage rigidity: perfect flexibility, symmetric nominal rigidity a la Calvo (1983) and a downward nominal wage rigidity.

3.4.1 Flexible Wages

Under the special case of flexible wages the equilibrium for the model can be solved analytically. Since the workers' wages are perfectly flexible the households set wages to maximise utility in the current period only, since they have no impact on future periods. Given the workers' optimal wage setting function under flexible wages given by equation 3.3.12 we can integrate over the distribution of possible shocks to find the aggregate real wage. This gives the flexible-wage labour supply curve which relates the amount of aggregate labour supplied as a function of the real wage

$$L_t = \left(\frac{\eta-1}{\eta}\right)^{\frac{\gamma}{\gamma+1}} s_t^{\frac{\gamma}{\gamma+1}} w_t^{\frac{\eta}{\gamma+1}} \left[\frac{1}{Z_t}\right]^{\frac{\gamma}{1+\gamma}}, \tag{3.4.1}$$

where the aggregate level of the disutility from working is given by

$$Z_t = \left[\int_0^1 \left(Z_{it}^{\frac{\gamma(1-\eta)}{\gamma+\eta}}\right) di\right]^{-\frac{\eta+\gamma}{\gamma(\eta-1)}} = e^{-\frac{1}{2}\frac{\eta(1+\gamma)}{\gamma+\eta}\sigma_z^2}. \tag{3.4.2}$$

This labour supply curve completes the flexible wage model and allows us to solve the steady state analytically.

3.4.2 Downwardly Rigid Nominal Wages

In the case with downward nominal wage rigidity the equilibrium will depend on the distribution of the real wages within the economy. Thus to solve the model we need to track the distribution of real wages. The distribution of real detrended wages at time t, is denoted $G_t(w)$, and has a corresponding density function $g_t(w)$. To see how the distribution is defined within the model consider the share of workers who have a real wage lower than w' in a given period. This set is comprised of three types of workers.

The first are those that are not subject to the downward nominal wage rigidity and optimally choose a wage lower than w'. This occurs because their idiosyncratic shock to their disutility of labour preferences, $z_{i,t}$ is lower than a cutoff $\overline{z}$, such that

$$z_{i,t} < \overline{z} = Z_t\left(w'\right), \tag{3.4.3}$$

where $Z_t(w)$ is the inverse of the optimal wage function in period t. Given the distribution of the preference shock $F(.)$, this group will comprise $(1-\lambda) F(Z_t(w'))$ of all workers.

The second type of workers are those that are subject to the downwardly rigid nominal wages constraint, but are not bound by it. These workers start period t with a wage w which is lower than w' and choose to increase their wage such that it remains below w'. This group of workers who are subject to the wage constraint, but not bound by it, and have wages lower than w' will thus have a preference shock $z_{i,t}$ such that

$$Z_t(w) \leq z_{i,t} \leq Z_t\left(w'\right). \tag{3.4.4}$$

Finally, the third category of workers with a wage lower than w' are those who are subject to the downward nominal wage constraint and are bound by it. These workers will begin the period with a detrended real wage that is lower than w' and draw a shock that causes them to prefer an even lower real wage, $z_{it} < Z(w) \leq Z_t(w')$. However, due to the nominal wage rigidity they are forced to maintain their previous nominal wage. The optimality of choosing the same nominal wage when the downward wage constraint binds is shown in Appendix 3.A.

The proportion of workers who start period t with real wage equal to or lower than w is given by $G_{t-1}(w(1+\pi_t)(1+a_t))$. Note that it is necessary to inflate the real wage by inflation and the rate of productivity growth to translate the wage distribution between periods.

Thus the combined share of workers in these last two groups is given by

$$\lambda F(Z_t(w))\, G_{t-1}(w(1+\pi_t)(1+a_t)). \tag{3.4.5}$$

Combining these three groups of workers gives the evolution of distribution of the real wage.

$$G_t(w) = (1-\lambda) F(Z_t(w)) + \lambda F(Z_t(w))\, G_{t-1}(w(1+\pi_t)(1+a_t)) \tag{3.4.6}$$

Note that as inflation or productivity growth increases $G_{t-1}(w(1+\pi_t)(1+a_t))$ approaches 1. This demonstrates how a higher rate of inflation or productivity growth unwinds the labour market friction, moving the wage distribution to the flexible wages case where $G_t(w) = F(Z_t(w))$.

Given the distribution of wages and the optimal wage setting function we can calculate the equilibrium level of labour supplied as shown in Appendix 3.A.

$$L_t = \left(\frac{\eta-1}{\eta}\right)^{\frac{\gamma}{\gamma+1}} s_t^{-\frac{\gamma}{\gamma+\eta}} w_t^{\frac{\eta}{\gamma+1}} \times \left(\frac{1}{Z_t^*}\right)^{\frac{\gamma}{1+\gamma}} \tag{3.4.7}$$

Where Z_t^* is a distorted measure of the aggregate disutility of labour given by

$$\begin{aligned} Z_t^* \quad = \quad & \Bigg\{(1-\lambda)\int_0^\infty \left(\frac{\hat{w}_t(Z)}{w_t(Z)}\right)^{\eta-1} Z^{\frac{\gamma(1-\eta)}{\gamma+\eta}}\, dF(Z) \\ & +\lambda\int_0^\infty \left(\frac{\hat{w}_t(Z)}{w_t(Z)}\right)^{\eta-1} Z^{\frac{\gamma(1-\eta)}{\gamma+\eta}} G_{t-1}(w_t(Z)(1+\pi_t)(1+a_t))\, dF(Z) \\ & +\lambda\int_0^\infty Z^{\frac{\gamma(1-\eta)}{\gamma+\eta}} \left[\int_{w_t(Z)}^\infty \left(\frac{\hat{w}_t(Z)}{w_t(Z)}\right)^{\eta-1} (1+\pi_t)(1+a_t)\ ... \right. \\ & \left. \times g(w_t(Z)(1+\pi_t)(1+a_t))\, dw\right] dF(Z)\Bigg\}^{-\frac{\gamma+\eta}{(\eta-1)\gamma}} \end{aligned} \tag{3.4.8}$$

Each of the terms in this expression corresponds with the three types of workers whose disutility we are aggregating. Critical to the analysis is the third term which pertains to workers for which the wage rigidity is a binding constraint. The inability to lower their nominal wage, in order to satisfy their relatively low draw for the disutility shock, means that they provide an inefficiently low level of labour creating a distortion in the economy.

3.4.3 Symmetrically Rigid Nominal Wages

In the second counter factual we assume that nominal wages are symmetrically rigid. The distribution of real detrended wages at time t, is denoted $G_t^{sym}(w)$, and has a corresponding density function $g_t^{sym}(w)$. To see how the distribution is defined within the model consider the share of workers who have a real wage lower than w' in a given period. This set is comprised of two types of workers; those who are able to freely set their wage and those whose nominal wage is fixed at its previous value.

The proportion of workers who start period t with real wage equal to or lower than w is given by $G_{t-1}^{sym}(w(1+\pi_t)(1+a_t))$. Thus the distribution of real wages evolves according to

$$G_t^{sym}(w) = (1-\lambda) F^{sym}(w) + \lambda G_{t-1}^{sym}(w(1+\pi_t)(1+a_t)), \tag{3.4.9}$$

where $F^{sym}(w)$ is the distribution of real wages for workers who are unconstrained by the nominal wage rigidity.

Given the distribution of wages and the optimal wage setting function we can calculate the equilibrium level of labour supplied as shown in Appendix 3.A.

$$L_t = \left(\frac{\eta-1}{\eta}\right)^{\frac{\gamma}{\gamma+1}} s_t^{-\frac{\gamma}{\gamma+\eta}} w_t^{\frac{\eta}{\gamma+1}} \times \left(\frac{1}{Z_t^{sym*}}\right)^{\frac{\gamma}{1+\gamma}}, \tag{3.4.10}$$

where Z_t^* is a distorted measure of the aggregate disutility of labour given by

$$\begin{aligned} Z_t^{sym*} \quad = \quad & \Bigg\{(1-\lambda)\int_0^\infty \left(\frac{\hat{w}_t(Z)}{w_t^{sym}(Z)}\right)^{\eta-1} Z^{\frac{\gamma(1-\eta)}{\gamma+\eta}}\, dF(Z) \\ & +\lambda \int_0^\infty \int_0^\infty \left(\frac{\hat{w}_t(Z)}{w_t^{sym}(Z)}\right)^{\eta-1} (1+\pi_t)(1+a_t)^{\eta-1} \\ & \times g_{t-1}^{sym}(w_t(Z)(1+\pi_t)(1+a_t))\, dw\, Z^{\frac{\gamma(1-\eta)}{\gamma+\eta}}\, dF(Z)\Bigg\}^{-\frac{\gamma+\eta}{(\eta-1)\gamma}} \end{aligned} \tag{3.4.11}$$

3.5 Numerical Results

Since this non-linear approach to nominal rigidities involves tracking the distribution of real wages the model's steady state cannot be solved using standard approximation techniques. Instead it is solved numerically using the algorithm outlined in Chapter 2. However instead of approximating the model's numerical derivatives around the steady state of the model, I compute the marginal effect of each aggregate shock for each state the of risk premium shock. This allows us to control for the possibility that aggregate shocks have different marginal effects when the economy is in a liquidity trap compared with when the economy is in its steady state.

Finally, to incorporate the zero-lower bound within the numerical approximation I augment the algorithm from Chapter 2 as follows. First I draw a set of residuals for the mark up, productivity, risk premium and preference shocks and simulate the model in the absence of a strict lower bound on nominal interest rates. I then use the numerical derivatives to calculate the path of endogenous variables over the simulated period. To account for the zero lower bound, I then update the path of the monetary policy shock via the formula

$$\epsilon_{t,j+1}^{M} = \begin{cases} \epsilon_{t,j}^{M} - \varphi_m\left(i_t\right) & \text{if } i_t < 0 \\ \epsilon_{t,j}^{M} - \varphi_m\left(i_t\right) & \text{if } i_t \geq 0 \end{cases} \tag{3.5.1}$$

where $\varphi_m > 0$ determines the rate at which the path of the shock adjusts. This process stops when the condition

$$\max\left|\epsilon_t^M i_t\right| < 10^{-3} \tag{3.5.2}$$

is met. This process ensures that contractionary monetary policy shocks are enacted such that nominal interest rates are subject to the zero lower bound in the numerical approximation.

3.5.1 Calibration Strategy

The parameters of the model are calibrated using three approaches. First, structural parameters that are commonly used in the literature are set to match key stylised facts or previously estimated results. The second calibrates the parameters related to the labour market friction to match the distribution of individual wage changes measured using matched-sample data from the Current Population Survey. Finally, I estimate the parameters governing the exogenous shock processes using a generalised method of moments.

Table 3.1 outlines the first set of parameters and the source of their calibration. The discount rate, β, is assumed to be 2 per cent. The parameters governing the firms' pricing friction are set to match previous studies of the labour market, in this case Smets and Wouters (2003). The targeted level of inflation is assumed to be 2 per cent, following the Federal Reserve's stated target (Bernanke, 2010), and labour productivity growth is calibrated to 2 per cent. The parameters within the central bank's monetary policy rule, ϕ_π and ϕ, governing the response of interest rates to changes in the output gap and rate of inflation are set at $\phi_\pi = \phi_y = 0.5$ following Rudebusch (2009). The parameters govern the firms' price rigidities are calibrated according to the Phillips Curve as estimated by Cogley and Sbordone (2008).

Table 3.1: Baseline Parameters

Parameter	Value	Interpretation	Reference
γ	0.5	Frisch labour supply elasticity	Daly and Hobijn (2014)
η	2	Labour demand elasticity for individual workers	Daly and Hobijn (2014)
θ	0.7	Share of firms subject to nominal price rigidity	Cogley and Sbordone (2008)
χ	0.05	Indexation of firms' price to inflation target	Cogley and Sbordone (2008)
ϵ	13	Price elasticity of good	Cogley and Sbordone (2008)
ϕ_π	0.5	Baseline Taylor Rule Parameter on Inflation	Rudebusch (2009)
ϕ_y	0.5	Baseline Taylor Rule Parameter on Inflation	Rudebusch (2009)
$\overline{\pi}$	0.005	Steady state rate of inflation	2 per cent annualised inflation
β	0.995	Discount factor	2 per cent discounting annually
$\overline{a}$	0.005	Steady state rate of productivity growth	2 per cent annualised labor productivity growth

Matching the distribution of annual changes in nominal wages, I estimate the rate at which workers are subject to downward nominal wage rigidity, λ, at 0.8 and the volatility of the idiosyncratic shock to the disutility of labour, σ_z, at 0.6. Under this calibration 13 per cent of workers have no change in their nominal wage and only a quarter experience a fall in nominal wage over four quarters. In the case of the symmetric wage friction the estimated Calvo wage friction, λ^{sym}, is 0.6 and the volatility of the idiosyncratic shock to the disutility of labour, σ_z^{sym}, is 0.6.

The third set of parameters are estimated to match moments in the data. To es-

Figure 3.2: Models vs Data: Percentage Change in Hourly Nominal Wages

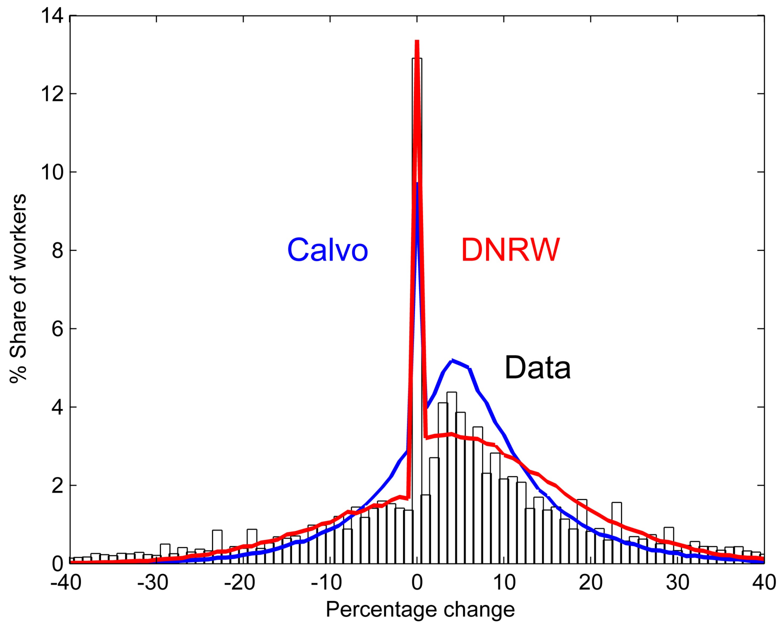

All workers (hourly and salary, job-stayers and job-switchers)
Source: Current Population Survey

timate the persistence parameters and variances of the exogenous shock processes I use a generalised method of moments to match the variances and co-variances of inflation, wage inflation and the output gap. For each set of parameters $\Sigma = \left[\rho^a, \rho^D, \rho^k, \varepsilon^a, \varepsilon^D, \varepsilon^k\right]$, I estimate the moments

$$m = \left[\text{var}\left(\pi_t\right), \text{var}\left(\pi_t^w\right), \text{var}\left(y_t\right), \text{cov}\left(\pi_t, \pi_t^w\right), \text{cov}\left(y_t, \pi_t^w\right), \text{cov}\left(\pi_t, y_t\right)\right] \tag{3.5.3}$$

and choose the calibration that minimises the distance

$$\hat{\Sigma} = \arg\min_{\Sigma} \left[\left(m - \hat{m}\left(\Sigma\right)\right)' W \left(m - \hat{m}\left(\Sigma\right)\right)\right] \tag{3.5.4}$$

where the weight matrix, W, is a diagonal matrix consisting of the variance of the moments which ensures that the estimator is efficient.

The parameters to govern the risk premium shock are set to match the unconditional frequency and length of periods at the zero lower bound. Using data from all advanced economies since 1950, I estimate an average frequency of being at the zero lower bound to be 7 per cent and an average duration of episodes to be 3 years. I match these moments by setting the size of the risk premium shock, I_Q, to 0.08, the length of risk premium shock, T_Q, to 12 periods and the probability of the risk premium shock to 2 per cent. Under this calibration the presence of the risk premium shock is highly correlated with episodes at the zero lower bound.

Table 3.2: Estimated Parameters

Parameter	Value	Interpretation
λ	0.8	Share of workers subject to DNWR constraint
σ_z	0.6	Std of the idiosyncratic shock to workers' disutility of labour
σ^D	0.15	Std of the shock to worker's discount factor
σ^a	0.05	Std of the shock to the rate of productivity growth
σ^k	0.20	Std of the shock to firms' mark-ups
ρ^D	0.5	Persistence of the shock to worker's discount factor
ρ^a	0.4	Persistence of the shock to firms' mark-ups
ρ^k	0.7	Persistence of the shock to the rate of productivity growth
I_Q	0.08	Size of risk premium shock
T_Q	12	Length of risk premium shock
P_Q	0.02	Probability of risk premium shock

3.6 Interactions between Downwardly Rigid Wages and the Zero Lower Bound

In order to better understand how these two nominal rigidities will effect the optimal rate of inflation within the economy, I will first discuss how they will impact each other in partial equilibrium. There are two channels through which wage rigidities interact with a zero lower bound on nominal interest rates. An extensive margin where nominal wage rigidities change the probability that a zero lower bound episode will occur, and an intensive margin in which they effect the severity of the episodes when they do.

3.6.1 Extensive Margin

The first is an extensive margin whereby wage rigidities moderate the degree of deflation in the economy and thus affect the probability that the monetary policy rule will require negative interest rates. While previous literature has shown that this channel lowers the probability of episodes at the zero lower bound, I find that this effect is contingent on the calibration of the monetary policy rule. This is because while downwardly rigid wages limit the fall in inflation, they exacerbate the fall in the output gap. Thus the overall impact on the nominal interest rates will be determined by the relative weights that central bank places on inflation and the output gap respectively. A central bank that places a larger weight on deviations of inflation from target will respond more aggressively to a negative demand shock under a Calvo-style wage friction or flexible wages and thus be more likely to hit the zero lower bound. Alternatively if the central bank places more weight on deviations in the output gap, then nominal interest rates will fall by more if wages are downwardly rigid. The direction of this effect is shown in Figure 3.3 which shows the difference in interest rate movements from a unconstrained Taylor rule between an economy where wages are downwardly rigid relative to one in which wages are flexible. A positive value (dark blue) indicates that the response of interest rates to a negative demand shock is smaller when wages are downwardly rigid compared with fully flexible wages, thus decreasing the probability of hitting the zero lower bound. Conversely a negative (light blue) value indicates a parameterisation under which interest rates are more volatile in the presence of downwardly rigid wages and thus more likely to result in interest rates being constrained.

Under the baseline calibration (where $\phi_\pi = \phi_y = 0.5$) the assumption of downward nominal wage rigidity dampens the response of interest rates to a negative demand shock thus reducing the probability that the zero lower bound will be binding relative to both the flexible wage and Calvo friction counter-factual. However, as Figure 3.4 the effect is quantitatively small. In response to a 4-quarter risk premium shock interest rates as set by a unconstrained Taylor rule are only on average only 0.2 per cent lower compared with both counter-factual wage assumptions.

Figure 3.3: Difference in Nominal Interest Rates, between DNWR and Flexible Wages

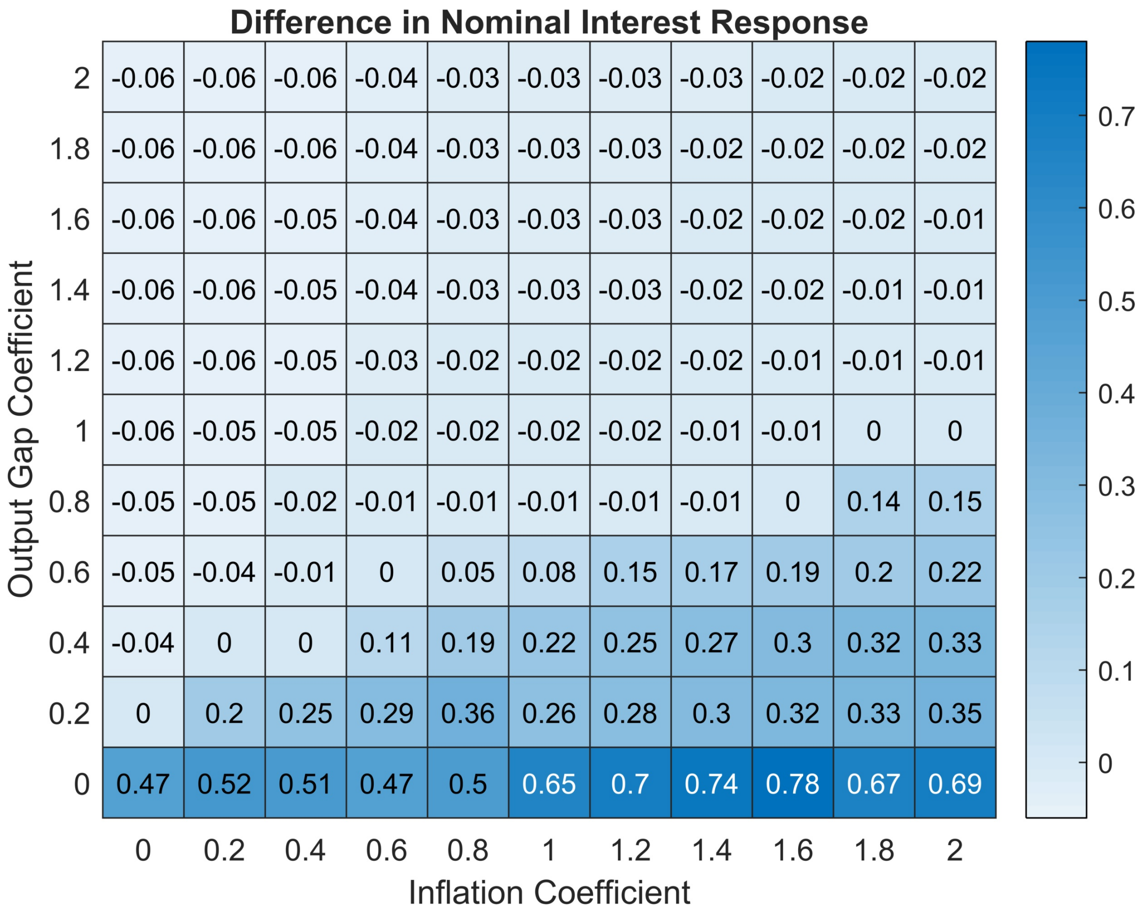

Notes: This heat map shows the difference in nominal interest rate impulse response functions to a 4 quarter risk premium shock, between an environment with downwardly rigid wages and flexible wages.

I also plot the impulse response functions for two alternative Taylor rules, one with a higher weight on deviations on price inflation from target and one which responds more to the output gap (Figure 3.5). Under the inflation-focused Taylor rule the assumption that wages are downwardly rigid generates interest rates that are 2 per cent higher than the case of flexible wages when subject to a 4 quarter risk premium shock. Under this parameterisation the assumption that wages are downwardly rigid will significantly decrease the likelihood of interest rates being constrained by the zero lower bound.

Figure 3.4: Nominal Interest Rates IRF, Baseline Calibration

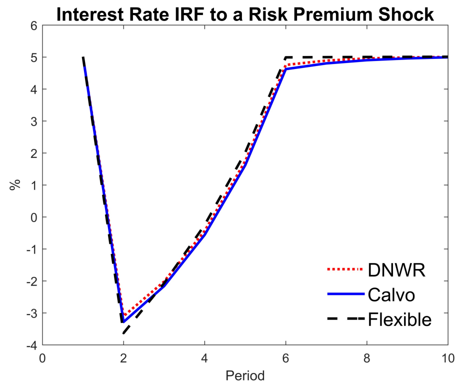

Notes: Response to a 4 quarter risk premium shock. Baseline Calibration.

Figure 3.5: Nominal Interest Rates IRF, Alternate Calibrations

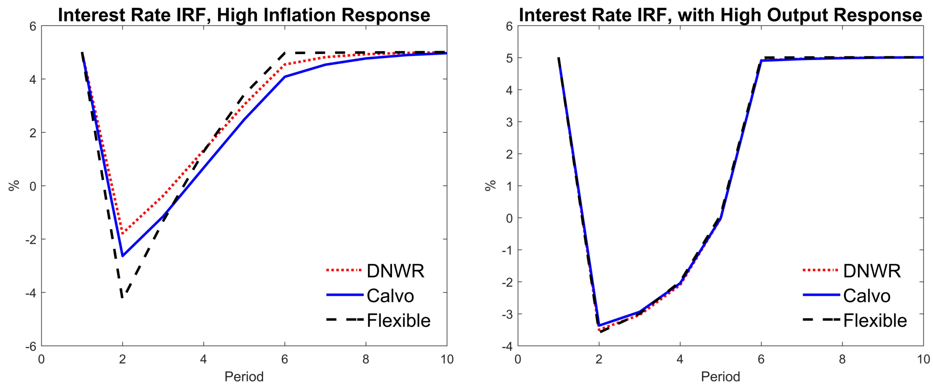

Notes: Response to a 4 quarter risk premium shock. The high inflation response corresponds to $\phi_\pi = 1.5$ and$\phi_y = 0.1$. The high output gap response corresponds to $\phi_\pi = 0.1$ and$\phi_y = 1.5$.

3.6.2 Intensive Margin

The second impact is on the intensive margin. As downwardly rigid nominal wages limit the degree to which real wages can fall, they will generate a higher inflation rate during periods at the zero lower bound. When nominal rates are fixed, this higher rate of inflation leads to a lower real interest rate, which mitigates the impact of the zero lower bound as shown in Figure 3.6.

However when interest rates are constrained at the zero lower bound this effectively tightens monetary policy, relative to an unconstrained Taylor rule, further lowering inflation and the output gap. A lower rate of inflation decreases the workers' real wage flexibility when their nominal wages are unable to fall, forcing wages to remain inefficiently high and reducing the demand for labour. This impact is shown in Figure

Figure 3.6: Impulse Response Functions with a ZLB

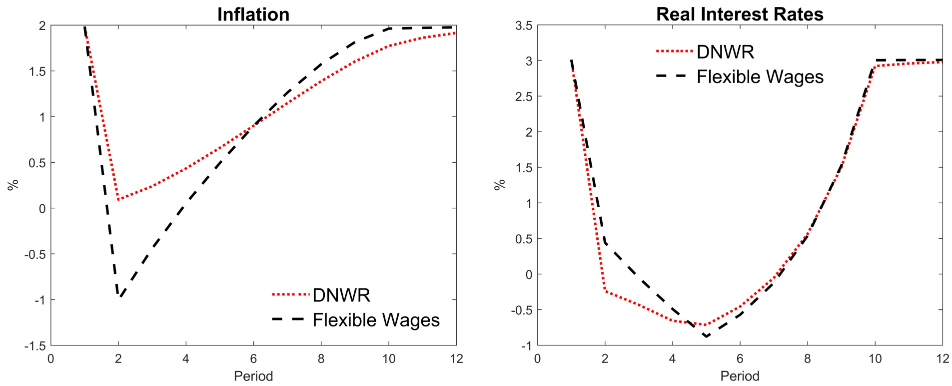

Notes: Response to a 8 quarter risk premium shock.

3.7 which shows the impulse response functions of a negative demand shock under two scenarios where the zero lower bound on nominal interest rates does and does not apply. The higher nominal interest rates caused by the zero lower bound cause larger falls in inflation, the output gap and consequently household welfare.

Figure 3.7: Impact of the Zero Lower Bound when Wages are Downwardly Rigid

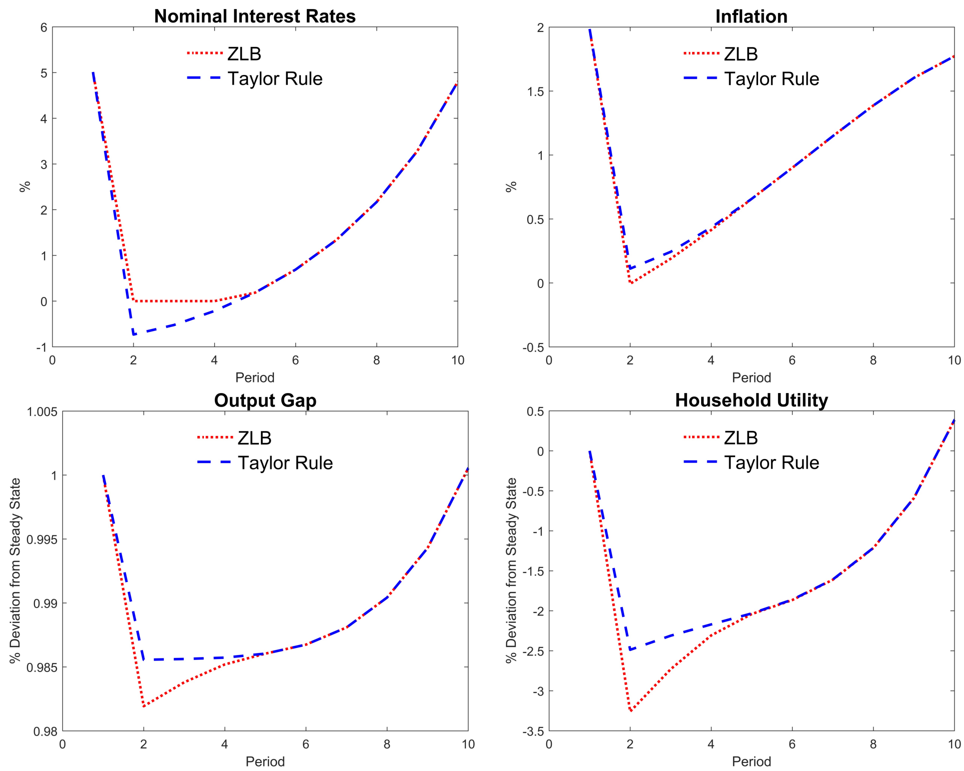

Notes: Response to a 4 quarter risk premium shock.

Both of these effects have been modeled by prior literature in representative agent frameworks (Amano et al., 2017; Coibion et al., 2012; Fahr and Smets, 2010). However the costs of downwardly rigid wages are significantly higher in a framework with heterogeneous households with differentiated nominal wages. When households are assumed to have an identical nominal wage these asymmetric rigidities act as an effective form of insurance: propping up inflation only when it is most beneficial, during periods of low demand in which there is a higher probability that nominal interest rates will be constrained by their lower bound, and allowing workers the flexibility to choose their optimal real wage at all other times. By contrast in an environment with heterogeneous workers the wage rigidities affect their ability to choose their optimal real wage when inflation is at, or even above, its inflation target (see Figure 3.8). Thus when workers are assumed to have differentiated wages, downward nominal wage rigidities are less effective as a form of insurance against the zero lower bound and impose larger welfare costs when the economy is not in a liquidity trap.

Figure 3.8: Representative vs Heterogeneous Environments

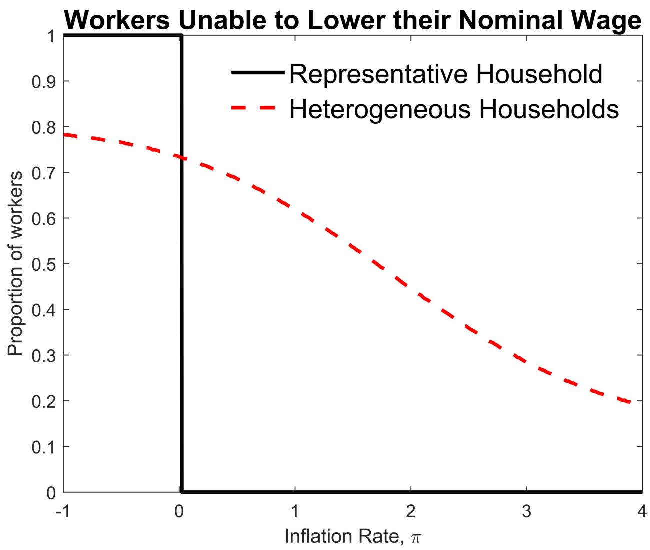

Notes: I have used the representative agent framework of Amano et al. (2017), in which identical households must choose a non-negative level of wage inflation such that $\pi_w \geq 0$.

3.7 Optimal Inflation Rate

In the absence of either wage rigidities or the zero lower bound the optimal rate of inflation is 0 per cent. A stable price level minimizes the level of price dispersion in the goods market, thus maximising household welfare. However under the presence of either of the constraints the optimal rate of inflation will be positive. The optimal rate of inflation in the presence of wage rigidities and the zero lower bound is summarised in Table 3.3.

Table 3.3: Optimal Rate of Inflation

Wage Friction	Taylor Rule	Zero Lower Bound
Downwardly Rigid	2.5%	3%
Flexible	0%	0.9%
Symmetrically Rigid	−1%	0.5%

3.7.1 Flexible Wages

In the absence of the nominal wage rigidities the results are in line with the previous literature. Under the assumption of flexible wages and an unconstrained interest rate rule the optimal rate of inflation is 0 per cent in order to minimise price dispersion. When nominal interests are constrained by the zero lower bound the optimal rate of inflation is approximately 1 per cent which decreases the unconditional probability of interest rates being constrained and the loss of output which occurs when they are.

Figure 3.9: Optimal Rate of Inflation - Flexible Nominal Wages

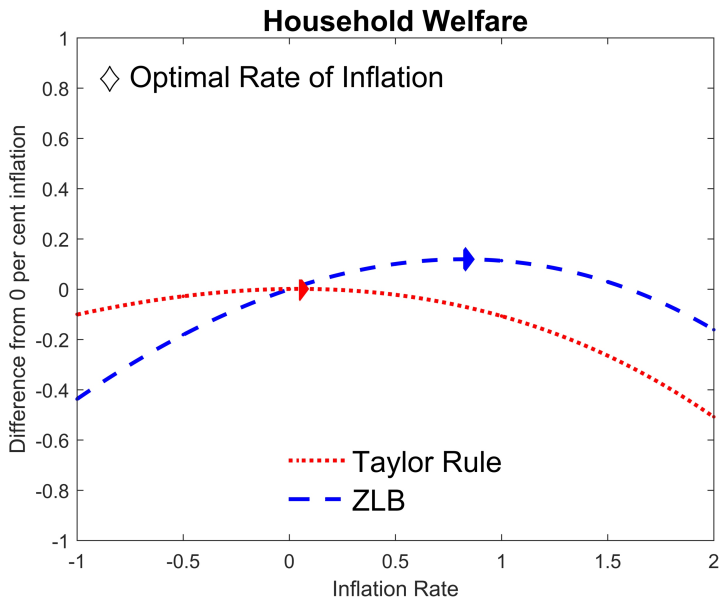

3.7.2 Downwardly Rigid Wages

When wages are assumed to be downwardly rigid, and nominal interest rates are unconstrained, the optimal rate of inflation is 2.5 per cent, increasing to 3 per cent if interest rates are subject to the zero lower bound. The assumption of downwardly rigid nominal wages increases the optimal rate of inflation relative to both counterfactuals in the presence of a zero lower bound. These results, while intuitive, differ from previous results in the literature such as Amano et al. (2016) and Coibion et al. (2012), who find that the assumption of downwardly rigid wages decreases the optimal rate of inflation due to the reduced probability of hitting the zero lower bound. This counter intuitive result is likely due to two factors. The first is the calibration of the Taylor Rule. Both Amano et al. (2016) and Coibion et al. (2012) place a relatively low weight on the output gap in the Taylor Rule which, as discussed in the previous section, will cause the assumption of downwardly rigid wages to decrease interest rate volatility. The second explanation is their assumption of homogeneous agents. The assumption that all agents have the same nominal wage will mean that the downwardly rigid wage frictions only constrain workers' in times of low inflation. By contrast my assumption that workers have different preferences over their labour supply will cause the downward nominal wage rigidity to distort the economy even when inflation is at or above its steady-state level. Thus in our model a higher rate of inflation will increase real wage flexibility in all periods, while under the alternative approach it only increases real wage flexibility when the wage inflation rate approaches zero - a time when wage flexibility actually reduces welfare.

3.7.3 Symmetric Wage Rigidity

In the case of a symmetric Calvo-style nominal wage friction the optimal rate of inflation is negative at -1 per cent. This is because a negative rate of inflation helps mitigate dispersion in wages which have a trend growth rate equal to the sum of productivity growth and the rate of inflation. An inflation rate of -1 per cent maximises household welfare by balancing the costs of dispersion in both prices and real wages. This result is in line with previous models which examine optimal monetary policy

Figure 3.10: Optimal Rate of Inflation - Downwardly Rigid Nominal Wages

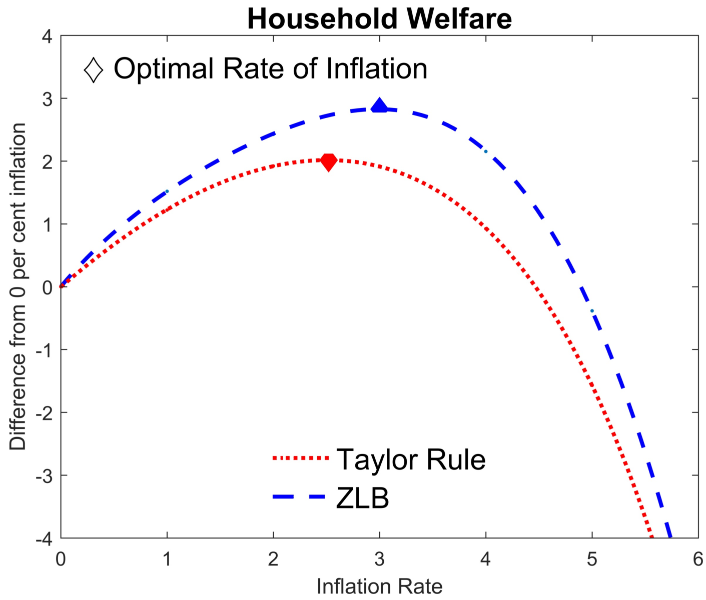

with the combination of trend productivity growth and nominal wage rigidities (see Amano et al. (2007)).

When the zero lower bound constraint is added to the model the optimal rate of inflation increases to 0.5 per cent to offset the disutility caused by high real interest rates imposed when nominal interest rates are constrained. When the zero lower bound binds interest rates the level of output is reduced, decreasing household welfare. However the impact of this fall in output is partially offset by a drop in the rate of inflation which decreases both price and real wage dispersion.

Figure 3.11: Optimal Rate of Inflation - Symmetrically Rigid Nominal Wages

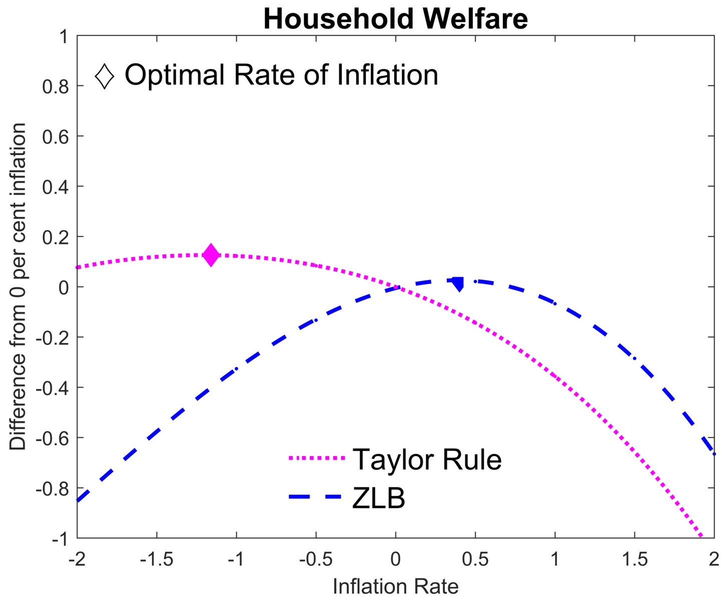

3.7.4 Trend productivity Growth

The issue of downward nominal wage rigidity is also relevant to the literature on secular stagnation. This literature examines the impact of the decline in real interest rates, the resulting increase in probability that the zero lower bound will constrain nominal rates, and the implications this has for optimal monetary policy (Eggertsson et al., 2017; Summers, 2014).Summers (2014) emphasises a range of channels which may be the cause of the decline in real interest rates, including a decrease in the trend growth rate of productivity. Eggertsson (2014), Gordon (2012) and Eichengreen (2015) detail the evidence that productivity growth rates have declined and how this has driven a fall in the steady state real, and consequently nominal, interest rates. This channel is present in this model as is shown by the equation for the long run real interest rate

$$r = \frac{1+a}{\beta} - 1. \tag{3.7.1}$$

However the ability for both price inflation and productivity growth to unwind the downward nominal wage friction suggests another channel through which declines in productivity growth might increase the optimal rate of inflation. When workers are constrained by the downward nominal wage rigidity they can achieve more real wage flexibility, either with higher inflation or higher productivity growth. If a secular decline in productivity growth occurs then this implies that the optimal rate of inflation will increase, irrespective of whether the zero lower bound exists.

Figure 3.12: Optimal Rate of Inflation and Trend Productivity Growth

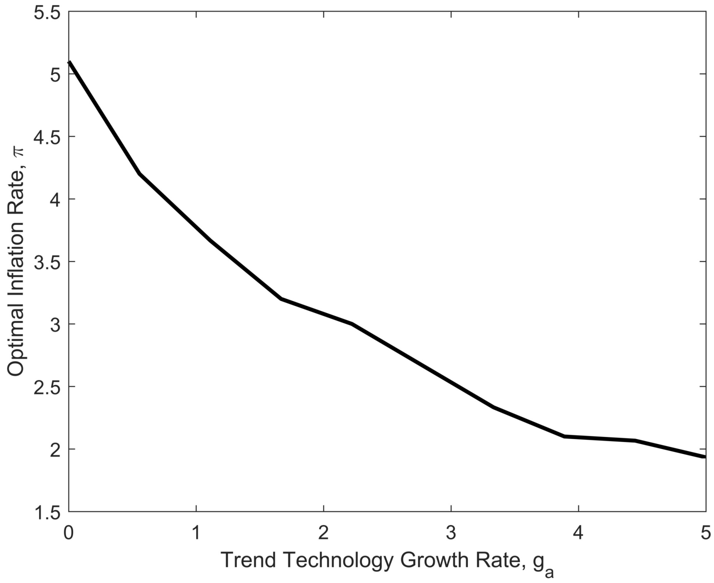

3.8 Conclusion

Downward nominal wage rigidities and the zero lower bound are both commonly cited motivations for a positive rate of inflation. This chapter shows that while each rigidity in isolation increases the optimal rate of inflation, modeling the two rigidities in tandem can amplify or mitigate the impact of the other on household welfare. When the zero lower bound binds, the lower rate of inflation that ensues decreases the flexibility of real wages, further distorting outcomes in the labour market. I find that in a heterogeneous agent model in which workers are subject to an idiosyncratic preference shock, that each friction leads to an increase in the optimal rate of inflation. This stands in contrast to previous results that find that downwardly rigid wages decrease the optimal rate of inflation in the presence of a zero lower bound on inflation. The relaxation of the representative agent assumption is key to these results. When all agents have the same nominal wage a downward nominal wage rigidity will constrain workers' choices if, and only if, demand in the economy is low. In this environment downward nominal wage rigidity is an effective form of insurance against the deflation spirals caused by a liquidity trap. However when workers' wages are differentiated, then these rigidities will constrain workers' choices at all times increasing their welfare cost and reducing their ability to moderate deflation.

The results described in this chapter have a number of potential implications. First, downward nominal wage rigidity may be a part of the explanation behind the "missing disinflation" puzzle of why advanced economies did not experience disinflation in line with what might have been expected by the large output gaps witnessed during the Great Recession. That is, downwardly rigid wages allow prices to remain relatively stable even in the face of persistently high output gaps. These results suggest that the combination of downwardly rigid wages and the zero lower bound on interest rates are reason to consider raising inflation targets in developed economies, if not in the long term then at least temporarily when the output gap is high.

This model is, aside from the two non-linear frictions, relatively parsimonious. A richer framework which includes the use of unconventional monetary policy may mitigate the loss of welfare caused by constrained monetary policy. However these results suggest that these non-linear frictions play a critical part in our understanding of how the rate of inflation within the economy relates to labour market outcomes and overall welfare within the economy.

Appendix

3.A Derivation of the Model

In this Appendix I derive the model described in Section 2.3 in greater detail.

3.A.1 Households

Households seek to maximise the present value of their utility

$$U_t = \sum_{s=0}^{\infty} \beta^s E_t e^{-\sum_{u=0}^{s-1} D_u} \left\{ \ln C_s - \frac{\gamma}{\gamma+1} \int_0^1 Z_{i,s} L_{i,s}^{\frac{\gamma+1}{\gamma}} \, di \right\}$$

subject to the budget constraint

$$B_t = (1 + i_{t-1}) \, Q_t B_{t-1} + \int_0^1 W_{i,t} L_{i,t} \, di - P_t C_t - \tau_t$$

and the demand for its labour

$$L_{it} = \left(\frac{w_{i,t}}{w_t} \right)^{-\eta} L_t.$$

The Lagrangian for this problem is thus

$$L_t = \sum_{s=0}^{\infty} \beta^s E_t e^{-\sum_{u=0}^{s-1} D_u} \left\{ \begin{array}{c} \ln C_s - \frac{\gamma}{\gamma+1} \int_0^1 Z_{i,s} L_{i,s}^{\frac{\gamma+1}{\gamma}} \, di \\ +\lambda_s \left(-b_t + \frac{1+i_{t-1}}{(1+\pi_t)(1+a_t)} b_{t-1} + \int_0^1 w_{i,t} L_{i,t} \, di - c_t \right) \end{array} \right\}$$

where the budget constraint has been detrended by the price and productivity level such that $c_t = \frac{C_t}{A_t}$.

3.A.1.1 Optimal Consumption Decision

Solving the first order conditions with respect to output yields

$$\begin{aligned} 0 &= \frac{\partial L}{\partial y_t} = \beta^t e^{-\sum_{s=0}^{t-1} D_s} \left[\frac{1}{c_t} - \lambda_t \right] \\ \therefore \lambda_t &= \frac{1}{c_t} \end{aligned}$$

and with respect to bonds

$$\begin{aligned} 0 &= \frac{\partial L}{\partial b_t} = \beta^t e^{-\sum_{s=0}^{t-1} D_s} \left[-\lambda_t + \beta e^{D_t} \frac{(1+i_t)}{(1+\pi_{t+1})(1+a_{t+1})} \lambda_{t+1} \right] \\ \therefore \lambda_t &= \beta e^{D_t} \frac{(1+i_t)}{(1+\pi_{t+1})(1+a_{t+1})} \lambda_{t+1}. \end{aligned}$$

Combining these conditions yields the Euler equation

$$\frac{1}{C_t} = \beta e^{-D_t} (1+r_t) \frac{1}{E_t C_{t+1}} \tag{3.A.1}$$

3.A.1.2 Optimal Wage Decision

When considering the optimal wage setting problem I can drop all terms that do not include the wage and labour supply for member i

$$L_t^{wage} = \sum_{s=0}^{\infty} \beta^s E_t e^{-\sum_{u=0}^{s-1} D_u} \left\{ -\frac{\gamma}{\gamma+1} \int_0^1 Z_{i,s} L_{i,s}^{\frac{\gamma+1}{\gamma}} \, di + \lambda_s (w_{it} L_{it}) \right\}.$$

First I substitute out the amount of labour supplied with the labour demand function which gives the expression.

$$L_t^{wage} = \sum_{s=0}^{\infty} \beta^s E_t e^{-\sum_{u=0}^{s-1} D_u} \left\{ -\frac{\gamma}{\gamma+1} L_s^{\frac{\gamma+1}{\gamma}} Z_{i,s} \left(\frac{w_{i,s}}{w_s} \right)^{-\eta \frac{\gamma+1}{\gamma}} + \lambda_s L_s \left(\frac{w_{i,s}}{w_s} \right)^{1-\eta} \right\}.$$

Substituting out for the Lagrange multiplier from the optimal consumption decision yields

$$L^{wage} = \sum_{s=0}^{\infty} \beta^s E_t e^{-\sum_{u=0}^{s-1} D_u} \left\{ -\frac{\gamma}{\gamma+1} L_s^{\frac{\gamma+1}{\gamma}} Z_{i,s} \left(\frac{w_{i,s}}{w_s} \right)^{-\eta \frac{\gamma+1}{\gamma}} + \frac{L_s}{c_s} \left(\frac{w_{i,s}}{w_s} \right)^{1-\eta} \right\}.$$

Finally, I can use the aggregate resource constraint to simplify the labour to consumption ratio such that

$$L^{wage} = \sum_{s=0}^{\infty} \beta^s E_t e^{-\sum_{u=0}^{s-1} D_u} \left\{ -\frac{\gamma}{\gamma+1} L_s^{\frac{\gamma+1}{\gamma}} Z_{i,s} \left(\frac{w_{i,s}}{w_t} \right)^{-\eta\frac{\gamma+1}{\gamma}} + s_s \left(\frac{w_{i,s}}{w_s} \right)^{1-\eta} \right\}.$$

Which gives equation 2.3.9, which is the contribution to the representative household's utility from worker i's wage related terms

$$\Omega\left(Z_{it}, w_{it}\right) = s_t \left(\frac{w_{i,t}}{w_t} \right)^{1-\eta} - \frac{\gamma}{\gamma+1} Z_{it} \left[w_{it}^{-\eta} w_t^{\eta} L_t \right]^{\frac{\gamma+1}{\gamma}}$$

where w_t, L_t and s_t are the aggregate variables for the real wage, total amount of labour and price dispersion among firms respectively. It is this expression I will determine the optimal wage for each worker i.

With flexible wages

When there is no nominal wage rigidity, wages are set by maximising this function only. The first order condition with respect to $w_{i,t}$ is

$$0 = \frac{\partial \Omega\left(Z_{it}, w_{it}\right)}{\partial w_{i,t}} = (1-\eta)\, s_t w_{it}^{-\eta} w_t^{\eta-1} + \eta Z_{it} w_{it}^{-\eta\frac{\gamma+1}{\gamma}-1} \left[w_t^{\eta} L_t \right]^{\frac{\gamma+1}{\gamma}}$$

which gives the wage setting function

$$w_{i,t}^{flex}\left(Z_{i,t}\right) = \left(\frac{\eta}{\eta-1} \right)^{\frac{\gamma}{\gamma+\eta}} Z_{i,t}^{\frac{\gamma+\eta}{\gamma}} L_t^{\frac{\gamma+1}{\gamma+\eta}} s_t^{-\frac{\gamma}{\gamma+\eta}} w_t.$$

With downwardly rigid wages

In the presence of downward nominal wage rigidity the optimal wage problem becomes a dynamic one. The wage chosen in the current period will affect the expected utility in the second period. The optimal wage function is thus given by maximising the sum of the current period pay-off and the expected value function in the subsequent period

$$w\left(Z_{i,t}\right) = \underset{w}{\operatorname{argmax}} \left\{ \Omega\left(Z_{it}, w\right) + \beta V_{t+1} \left(\frac{w}{\left(1+\pi_{t+1}\right)\left(1+a_{t+1}\right)} \right) \right\}.$$

I can now show that $w\left(Z_{i,t}\right) \leq w^{flex}\left(Z_{i,t}\right)$ for all possible preference shocks $Z_{i,t}$.

To simplify this proof I transform the decision variable such that $x_{i,t} = \left(w_{i,t}\right)^{1-\eta}$, thus the transformed Bellman equation is

$$V_t(x) =$$
$$(1-\lambda)\int_0^\infty \max_{x_{it}\geq 0}\left\{\Omega(Z_{it},x_{i,t}) + \beta V_{t+1}\left(x_{i,t}\left[(1+\pi_{t+1})(1+a_{t+1})\right]^{\eta-1}\right)\right\}dF(Z)$$
$$+\lambda\int_0^\infty \max_{x_{it}\leq x}\left\{\Omega(Z_{it},x_{i,t}) + \beta V_{t+1}\left(x_{i,t}\left[(1+\pi_{t+1})(1+a_{t+1})\right]^{\eta-1}\right)\right\}dF(Z)$$

where

$$\Omega(Z_{i,t},x_{i,t}) = s_t\left(\frac{1}{w_t}\right)^{1-\eta}x_{i,t} - \frac{\gamma}{\gamma+1}Z_{it}\left[(x_{i,t})^{\frac{-\eta}{1-\eta}}w_t^\eta L_t\right]^{\frac{\gamma+1}{\gamma}}.$$

First I note from the Bellman equation that $V'(x)\geq 0$, since a higher x unwinds the constraint in the second term of the Bellman equation. This is equivalent to a lower real wage unwinding the constraint placed by the downward nominal wage rigidity.

Second, note that $\Omega(Z_{i,t},x_{i,t})$ is concave with respect to $x_{i,t}$ since

$$\begin{aligned}\frac{\partial}{\partial x_{i,t}}\Omega(Z_{i,t},x_{i,t}) &= s_t\left(\frac{1}{w_t}\right)^{1-\eta}x_{i,t} - \frac{\gamma}{\gamma+1}Z_{it}\left[(x_{i,t})^{\frac{-\eta}{1-\eta}}w_t^\eta L_t\right]^{\frac{\gamma+1}{\gamma}} \\ &= s_t\left(\frac{1}{w_t}\right)^{1-\eta} + \frac{\eta}{1-\eta}Z_{it}\left[w_t^\eta L_t\right]^{\frac{\gamma+1}{\gamma}}(x_{i,t})^{\frac{\eta}{\eta-1}\frac{\gamma+1}{\gamma}-1}\end{aligned}$$

and

$$\begin{aligned}\frac{\partial^2}{\partial x_{i,t}^2}\Omega(Z_{i,t},x_{i,t}) &= \left(\frac{\eta}{\eta-1}\frac{\gamma+1}{\gamma}-1\right)\frac{\eta}{1-\eta}Z_{it}\left[w_t^\eta L_t\right]^{\frac{\gamma+1}{\gamma}}(x_{i,t})^{\frac{\eta}{\eta-1}\frac{\gamma+1}{\gamma}-2} \\ &= \frac{\eta}{1-\eta}\left(\frac{\eta}{\eta-1}\frac{\gamma+1}{\gamma}+1\right)Z_{it}\left[w_t^\eta L_t\right]^{\frac{\gamma+1}{\gamma}}(x_{i,t})^{\frac{\eta}{\eta-1}\frac{\gamma+1}{\gamma}-2}.\end{aligned}$$

Thus since $\eta > 1$, $\gamma > 0$, and the aggregate variables are all positive this second derivative will be negative.

From the flexible wage optimality problem

$$\frac{\partial}{\partial x_{i,t}}\Omega\left(Z_{i,t},x_{i,t}^{flex}\right) = 0.$$

Under the downward nominal wage rigidity this condition is

$$\frac{\partial}{\partial x_{i,t}}\Omega\left(Z_{i,t},x_{i,t}\right)+\beta\left[\left(1+\pi_{t+1}\right)\left(1+a_{t+1}\right)\right]^{\eta-1}V'_{t+1}\left(x_{i,t}\left[\left(1+\pi_{t+1}\right)\left(1+a_{t+1}\right)\right]^{\eta-1}\right)=0$$

since $\beta\left[\left(1+\pi_{t+1}\right)\left(1+a_{t+1}\right)\right]^{\eta-1}V'_{t+1}\left(x_{i,t}\left[\left(1+\pi_{t+1}\right)\left(1+a_{t+1}\right)\right]^{\eta-1}\right)\geq 0$ this implies that $\frac{\partial}{\partial x_{i,t}}\Omega\left(Z_{i,t},x_{i,t}\right)\leq 0$.

Given the two conditions

$$\begin{aligned}\frac{\partial}{\partial x_{i,t}}\Omega\left(Z_{i,t},x_{i,t}\right) &\leq 0\\ \frac{\partial}{\partial x_{i,t}}\Omega\left(Z_{i,t},x_{i,t}^{flex}\right) &= 0\end{aligned}$$

and the fact that the second derivative, $\frac{\partial^2}{\partial x_{i,t}^2}\Omega$, is negative, this implies that

$$x_{i,t}^{flex}\leq x_{i,t}$$

and thus

$$w^{flex}\left(Z_{i,t}\right)\geq w\left(Z_{i,t}\right)\quad\forall\; Z_{i,t}$$

Thus in the presence of downward nominal wage rigidities workers will always choose to set a lower wage compared with what they otherwise would if wages were perfectly flexible.

Finally, I will show that if a worker is subject and bound by the downward nominal constraint that they will choose to keep their nominal wage constant. Consider the case where the constraint is just binding, such that $w_t\left(Z\right)=w$. Then all shocks less then Z will also be binding, since otherwise the worker can increase their utility by lowering their wage.

Given that constraint will be binding for all preference shocks less than Z, any worker subject to the downward nominal wage rigidity at any wage $w>w_t\left(Z\right)$ will keep their wage fixed.

3.A.2 Final Goods Firms

Final goods firms set their prices and output to maximise profit in the current period as shown below.

$$\max P_t Y_t - \int_0^1 P_{j,t} Y_{j,t}\, dj$$

subject to the Dixit-Stiglitz aggregator which transforms intermediate goods into final goods according to

$$Y_t = \left[\int_0^1 Y_{j,t}^{\frac{\epsilon-1}{\epsilon}}\, dj \right]^{\frac{\epsilon}{\epsilon-1}}.$$

Solving for the first order conditions of the problem with respect to output yields

$$0 = P_t + \lambda_t$$

and with respect to the intermediate good input

$$0 \quad = \quad -P_{j,t} - \lambda_t \left[\int_0^1 Y_{j,t}^{\frac{\epsilon-1}{\epsilon}}\, dj \right]^{\frac{\epsilon}{\epsilon-1}-1} Y_{j,t}^{\frac{\epsilon-1}{\epsilon}-1}$$

combining these conditions yields the demand curve for intermediate goods

$$Y_{j,t} \quad = \quad \left(\frac{P_{j,t}}{P_t} \right)^{-\epsilon} Y_t.$$

Finally, I can impose the zero profit condition to solve for the aggregate price index

$$0 = P_t Y_t - \int_0^1 P_{j,t} Y_{j,t}\, dj$$

which yields

$$P_t = \left[\int_0^1 P_{j,t}^{1-\epsilon} \right]^{\frac{1}{1-\epsilon}}.$$

3.A.3 Intermediate Goods Firms

Intermediate firms set their optimal price to maximise the discounted stream of expected profits

$$\max_{\{P_{j,t}^*, Y_{j,t}\}} \sum_{k=0}^{\infty} \beta^k \frac{\lambda_{t+k}}{\lambda_t} \theta^k \left[\frac{P_{j,t}^* \left(\overline{\pi^p}\right)^{\mu k}}{P_{t+k}} Y_{j,t+k} - \frac{\epsilon - 1}{\epsilon} MC_{j,t+k}^r Y_{j,t+k} \right].$$

Profits are maximised subject to three constraints. The first is the firms' linear production function

$$Y_{t,j} = A_t L^d_{t,j}$$

The second constraint is the demand for the firm's output good.

$$Y_{j,t} = \left[\frac{P^*_{j,t}}{P_t}\right]^{-\epsilon} Y_t$$

Finally, a firm's marginal cost is governed by the aggregate wages set by the household, which are assumed to be exogenous by the firms

$$MC^r_{j,t} \quad = \quad w_t$$

Note that since all firms are identical we can drop the subscript j. Substituting in the constraints yields

$$\max_{\{P^*_{j,t}, Y_{j,t}\}} \sum_{k=0}^{\infty} \beta^k \frac{\lambda_{t+k}}{\lambda_t} \theta^k \left[\left[\frac{P^*_{j,t}\left(1+\overline{\pi^p}\right)^{\mu k}}{P_{t+k}}\right]^{1-\epsilon} Y_{t+k} - \frac{\epsilon-1}{\epsilon} w_t \left[\frac{P^*_{j,t}\left(1+\overline{\pi^p}\right)^{\mu k}}{P_{t+k}}\right]^{-\epsilon} Y_{t+k}\right]$$

The first order conditions with respect to the optimal price, P^*_t, is

$$0 = \sum_{k=0}^{\infty} \beta^k \frac{\lambda_{t+k}}{\lambda_t} \theta^k \left[(1-\epsilon)\frac{1}{P^*_t} \left[\frac{P^*_t\left(1+\overline{\pi^p}\right)^{\mu k}}{P_{t+k}}\right]^{1-\epsilon} Y_{t+k} + (\epsilon-1)\, w_{t+k} \frac{1}{P^*_t} \left[\frac{P^*_t\left(1+\overline{\pi^p}\right)^{\mu k}}{P_{t+k}}\right]^{-\epsilon} Y_{t+k}\right]$$

which can be solved to show that the optimal price is

$$P^*_t = \frac{\sum_{k=0}^{\infty} \beta^k \theta^k \left[\frac{\left(1+\overline{\pi^p}\right)^{\mu k}}{P_{t+k}}\right]^{-\epsilon} w_{t+k}}{\sum_{k=0}^{\infty} \beta^k \theta^k \left[\frac{\left(1+\overline{\pi^p}\right)^{\mu k}}{P_{t+k}}\right]^{1-\epsilon}}.$$

This solution can be written in terms of two recursive variables K^r_t and J^r_t where

$$P^*_t = \frac{K^r_t}{J^r_t}\left(1+\epsilon^K_t\right)$$

$$\begin{aligned} K^r_t &= MC^r_t \left(P_t\right)^{\epsilon} + \beta\theta\left(1+\overline{\pi^p}\right)^{-\mu\epsilon} E_t\left[K^r_{t+1}\right] \\ J^r_t &= P_t^{\epsilon-1} + \beta\theta\left(1+\overline{\pi^p}\right)^{\mu(1-\epsilon)} E_t\left[J^r_{t+1}\right] \end{aligned}$$

The aggregate price level is defined by

$$P_t^{1-\epsilon} = \int_0^1 P_{j,t}^{1-\epsilon}\, dj$$

The continuum of firms can be divided into two groups, those who reset their optimal price and those who index it from the previous period.

$$P_t^{1-\epsilon} = \int_\theta^1 \left(P_t^*\right)^{1-\epsilon}\, dj + \int_0^\theta \left(1+\overline{\pi^p}\right) P_{t-1}^{1-\epsilon}\, dj$$

which can be aggregated to give the evolution of the price level.

$$P_t^{1-\epsilon} = \theta \left(1+\overline{\pi}\right)^{1-\epsilon} P_{t-1}^{1-\epsilon} + (1-\theta)\, P_t^{*1-\epsilon}$$

Similarly the measure of price dispersion can be found by integrating over the production of the intermediate firm

$$L_{j,t} = \frac{1}{A_t} Y_{j,t}$$

$$\begin{aligned} L_t &= \frac{Y_t}{A_t}\left[\int_0^1 \left(\frac{P_{j,t}}{P_t}\right)^{-\epsilon} di\right] \\ L_t &= \frac{Y_t}{A_t} s_t \end{aligned}$$

where $s_t = \int_0^1 \left(\frac{P_{j,t}}{P_t}\right)^{-\epsilon} di$

This measure of price dispersion can again be divided into two groups, those who reset their optimal price and those who index it from the previous period.

$$P_t^{-\epsilon} = \int_\theta^1 \left(P_t^*\right)^{-\epsilon}\, dj + \int_0^\theta \left(1+\overline{\pi^p}\right) P_{t-1}^{-\epsilon}\, dj$$

which can be aggregated to give the evolution of the price level.

$$s_t = (1-\theta)\left(\frac{P_t^*}{P}\right)^{-\epsilon} + \theta \left(1+\overline{\pi^p}\right)^{-\chi\epsilon} s_{t-1}$$

3.A.4 Equilibrium

Differentiation of equation 2.4.6 yields a formula for the evolution of the wage density function.

$$\begin{aligned} g_t(w) &= (1-\lambda) Z_t'(w) f(Z_t(w)) \\ &\quad +\lambda G_{t-1}(w(1+\pi_t)(1+a_t)) Z_t'(w) f(Z_t(w)) \\ &\quad +\lambda(1+\pi_t)(1+a_t) g_{t-1}(w(1+\pi_t)(1+a_t)) F(Z_t(w)) \end{aligned}$$

Each of these lines corresponds to the three types of worker: those who are free to set their own wage, those that are subject to the wage rigidity but are not bound by it, and those that are both subject to and bound by the wage rigidity.

In equilibrium the real detrended wages of the workers must aggregate to the variable w_t.

$$w_t = \left[\int_0^1 \left(\frac{1}{w_{it}}\right)^{\eta-1} di\right]^{\frac{1}{1-\eta}}$$

We can expand the integral over the three categories of workers with respect to the downward nominal wage rigidity. This gives

$$\begin{aligned} \therefore 1 &= \left(\frac{\eta}{\eta-1}\right)^{\frac{\gamma}{\gamma+\eta}} s_t^{-\frac{\gamma}{\gamma+\eta}} L_t^{\frac{\gamma+1}{\gamma+\eta}} w_t^{\frac{-\gamma}{\gamma+\eta}} \\ &\quad \times \left\{ (1-\lambda) \int_0^\infty \left(\frac{\hat{w}_t(Z)}{w_t(Z)}\right)^{\eta-1} Z^{\frac{\gamma(1-\eta)}{\gamma+\eta}} dF(Z) \right. \\ &\quad +\lambda \int_0^\infty \left(\frac{\hat{w}_t(Z)}{w_t(Z)}\right)^{\eta-1} Z^{\frac{\gamma(1-\eta)}{\gamma+\eta}} G_t(w_t(Z)(1+\pi_t)(1+a_t))\, dF(Z) \\ &\quad \left. +\lambda \int_0^\infty \left[\int_{w_t(Z)}^\infty \left(\frac{\hat{w}_t(Z)}{w_t(Z)}\right)^{\eta-1} Z^{\frac{\gamma(1-\eta)}{\gamma+\eta}} G_t(w_t(Z)(1+\pi_t)(1+a_t))\right] dF(Z) \right\}^{\frac{1}{1-\eta}} \end{aligned}$$

We can now define

$$\begin{aligned} Z_t^* &= \Bigg\{ (1-\lambda) \int_0^\infty \left(\frac{\hat{w}_t(Z)}{w_t(Z)} \right)^{\eta-1} Z^{\frac{\gamma(1-\eta)}{\gamma+\eta}} \, dF(Z) \\ &\quad + \lambda \int_0^\infty \left(\frac{\hat{w}_t(Z)}{w_t(Z)} \right)^{\eta-1} Z^{\frac{\gamma(1-\eta)}{\gamma+\eta}} G_t\left(w_t(Z)(1+\pi_t)(1+a_t)\right) dF(Z) \\ &\quad + \lambda \int_0^\infty \left[\int_{w_t(Z)}^\infty \left(\frac{\hat{w}_t(Z)}{w_t(Z)} \right)^{\eta-1} Z^{\frac{\gamma(1-\eta)}{\gamma+\eta}} G_t\left(w_t(Z)(1+\pi_t)(1+a_t)\right) \right] dF(Z) \Bigg\}^{-\frac{\gamma+\eta}{(\eta-1)\gamma}} \end{aligned}$$

A distorted measure of the aggregate disutility of labour among workers which takes into account the distortion of wage and labour choices caused by the downward nominal wage rigidity. This equilibrium condition can thus be written as

$$L_t = \left(\frac{\eta-1}{\eta} \right)^{\frac{\gamma}{\gamma+1}} s_t^{-\frac{\gamma}{\gamma+\eta}} w_t^{\frac{\eta}{\gamma+1}} \times \left(\frac{1}{Z_t^*} \right)^{\frac{\gamma}{1+\gamma}} .$$

Bibliography

Abo-Zaid, S. (2013a). Optimal monetary policy and downward nominal wage rigidity in frictional labor markets. *Journal of Economic Dynamics and Control 37*(1), 345–364. (document)

Abo-Zaid, S. (2013b). Optimal monetary policy and downward nominal wage rigidity in frictional labor markets. *Journal of Economic Dynamics and Control 37*(1), 345–364. 2, 3.2

Achdou, Y., J. Han, J.-M. Lasry, P.-L. Lions, and B. Moll (2017). Income and wealth distribution in macroeconomics: A continuous-time approach. Technical report, National Bureau of Economic Research. 1.1, 1.2, 1.3, 1.A.2, 1.A.2, 1.A.2, 1.B

Ahn, S., G. Kaplan, B. Moll, T. Winberry, and C. Wolf (2017a). Micro heterogeneity and aggregate consumption dynamics. *NBER Macroeconomics Annual, forthcoming*. (document)

Ahn, S., G. Kaplan, B. Moll, T. Winberry, and C. Wolf (2017b). When inequality matters for macro and macro matters for inequality. In *NBER Macroeconomics Annual 2017, volume 32*. University of Chicago Press. 1.2

Aiyagari, S. R. (1994). Uninsured idiosyncratic risk and aggregate saving. *The Quarterly Journal of Economics 109*(3), 659–684. 1.2

Altonji, J. G. and P. J. Devereux (2000). Is there nominal wage rigidity? evidence from panel data. *Research in Labor Economics 19*, 383–431. 3, 2.1, 2.2

Amano, R., D. Brouillette, S. Gnocchi, N. Kyui, et al. (2016). the micro and macro of downward nominal wage rigidity. *Bank of Canada Review 2016*(Spring), 15–28. 3.7.2

Amano, R., K. Moran, S. Murchison, and A. Rennison (2007). Trend inflation, wage and price rigidities, and welfare. Technical report, Bank of Canada Working Paper. 3.7.3

Amano, R. A., S. Gnocchi, et al. (2017). *Downward Nominal Wage Rigidity Meets the Zero Lower Bound*. Bank of Canada Banque du Canada. 3.1, 3.2, 3.6.2, **??**

Andrews, D., A. C. Sanchez, and Å. Johansson (2011). Housing markets and structural policies in oecd countries. *OECD Economic Department Working Papers* (836), 0_1. 1.6

Ascari, G. and A. M. Sbordone (2014). The macroeconomics of trend inflation. *Journal of Economic Literature 52*(3), 679–739. 2.5.1

Attanasio, O. P., R. Bottazzi, H. W. Low, L. Nesheim, and M. Wakefield (2012). Modelling the demand for housing over the life cycle. *Review of Economic Dynamics 15*(1), 1–18. (document), 2, 1.2, 1.C.2

Attanasio, O. P. and G. Weber (1995). Is consumption growth consistent with intertemporal optimization? evidence from the consumer expenditure survey. *Journal of political Economy 103*(6), 1121–1157. 1.5.1.1

Auclert, A. (2017). Monetary policy and the redistribution channel. Technical report, National Bureau of Economic Research. (document), 1.2

Babeckỳ, J., P. Du Caju, T. Kosma, M. Lawless, J. Messina, and T. Rõõm (2010). Downward nominal and real wage rigidity: Survey evidence from european firms. *The Scandinavian Journal of Economics 112*(4), 884–910. 3, 2.1

Benigno, P. and L. A. Ricci (2011). The inflation-output trade-off with downward wage rigidities. *American Economic Review 101*(4), 1436–66. 2.2

Benmelech, E., A. Guren, and B. T. Melzer (2017). Making the house a home: The stimulative effect of home purchases on consumption and investment. Technical report, National Bureau of Economic Research. 1.2, 7, 1.5.2, 1.E.1

Berger, D., V. Guerrieri, G. Lorenzoni, and J. Vavra (2015). House prices and consumer spending. Technical report, National Bureau of Economic Research. 1.2, 1.5.1

Bernanke, B. S. (2010). Monetary policy and the housing bubble. In *speech at the annual meeting of the American Economic Association, Atlanta, Georgia*, Volume 3. 2.5.4, 3.5.1

Best, M. C. and H. J. Kleven (2013). Housing market responses to transaction taxes: Evidence from notches and stimulus in the uk. *London School of Economics*. 1.2, 7

Bewley, T. (1983). A difficulty with the optimum quantity of money. *Econometrica: Journal of the Econometric Society*, 1485–1504. 1.2

Bewley, T. F. (1999). *Why wages don't fall during a recession*. Harvard University Press. 2.2

Blanchard, O. and J. Galí (2007). Real wage rigidities and the new keynesian model. *Journal of Money, Credit and Banking 39*(s1), 35–65. 2.6, 2.6.1

Blundell, R., L. Pistaferri, and I. Preston (2004). Imputing consumption in the psid using food demand estimates from the cex. 1.5.1.1

Blundell, R., L. Pistaferri, and I. Preston (2008). Consumption inequality and partial insurance. *The American Economic Review 98*(5), 1887–1921. (document), 1.1, 1.2, 1.5.1, 1.5.1.3

Blundell, R. and I. Preston (1998). Consumption inequality and income uncertainty. *The Quarterly Journal of Economics 113*(2), 603–640. (document)

Boppart, T., P. Krusell, and K. Mitman (2017). Exploiting mit shocks in heterogeneous-agent economies: The impulse response as a numerical derivative. Technical report, National Bureau of Economic Research. 2.1

Broda, C. and J. A. Parker (2014). The economic stimulus payments of 2008 and the aggregate demand for consumption. *Journal of Monetary Economics 68*, S20–S36. (document), 2, 1.4.5, 1.4.5

Caballero, R. J. and E. Farhi (2014). On the role of safe asset shortages in secular stagnation. *Secular stagnation: Facts, causes and cures*, 111. 1.1

Calvo, G. A. (1983). Staggered prices in a utility-maximizing framework. *Journal of monetary Economics 12*(3), 383–398. 2.1, 2.2, 2.3, 2.3.1, 2.3.3, 2.6.1, 2.6.2, 3.3.1, 3.3.3

Campbell, J. Y. and J. F. Cocco (2007). How do house prices affect consumption? evidence from micro data. *Journal of monetary Economics 54*(3), 591–621. (document), 2, 1.2

Card, D. and D. Hyslop (1997). Does inflation "grease the wheels of the labor market"? pp. 71–122. 3, 2.1

Carlsson, M. and A. Westermark (2008). Monetary policy under downward nominal wage rigidity. *The BE Journal of Macroeconomics 8*(1). 2

Carroll, C. D. and W. E. Dunn (1997, July). Unemployment expectations, jumping (s,s) triggers, and household balance sheets. Working Paper 6081, National Bureau of Economic Research. 2, 1.2

Cogley, T. and A. M. Sbordone (2008). Trend inflation, indexation, and inflation persistence in the new keynesian phillips curve. *The American Economic Review 98*(5), 2101–2126. 2.5.4, 2.2, 3.5.1, 3.1

Coibion, O., Y. Gorodnichenko, and J. Wieland (2012). The optimal inflation rate in new keynesian models: should central banks raise their inflation targets in light of the zero lower bound? *The Review of Economic Studies*, rds013. 3.1, 3.2, 3.6.2, 3.7.2

Daly, M., B. Hobijn, and B. Lucking (2012). Why has wage growth stayed strong? *FRBSF Economic Letter 10*(2). 2.5.6

Daly, M. C. and B. Hobijn (2014). Downward nominal wage rigidities bend the phillips curve. *Journal of Money, Credit and Banking 46*(S2), 51–93. 2.1, **??**, 2.1, 2.2, 2.3, 2.2, 3.1, 3.1

Dordal-i Carreras, M., O. Coibion, Y. Gorodnichenko, and J. Wieland (2016). Infrequent but long-lived zero-bound episodes and the optimal rate of inflation. Technical report, National Bureau of Economic Research. 3.1, 3.3

Dynan, K. (2012). Is a household debt overhang holding back consumption? *Brookings Papers on Economic Activity 2012*(1), 299–362. 1.1, 1.5.2

Dynan, K., T. Gayer, and N. Plotkin (2013). An evaluation of federal and state homebuyer tax incentives. *Washington, DC: The Brookings Institution*. 1.2, 1.6, 1.6.2

Eggertsson, G. (2014). This version :. 3.7.4

Eggertsson, G. B., N. R. Mehrotra, and J. A. Robbins (2017, January). A model of secular stagnation: Theory and quantitative evaluation. Working Paper 23093, National Bureau of Economic Research. 3.7.4

Eichengreen, B. (2015). Secular stagnation: the long view. *American Economic Review 105*(5), 66–70. 3.7.4

Elsby, M. W., D. Shin, and G. Solon (2013). Wage adjustment in the great recession. 3, 2.1

Engelhardt, G. V. (1996). Consumption, down payments, and liquidity constraints. *Journal of Money, Credit and Banking 28*(2), 255–271. (document)

Erceg, C. J., D. W. Henderson, and A. T. Levin (2000). Optimal monetary policy with staggered wage and price contracts. *Journal of Monetary Economics 46*(2), 281–313. 2.1, 2.2, 2.6, 2.6.2, 2.6.2

Fagan, G. and J. Messina (2009). Downward wage rigidity and optimal steady-state inflation. 2.2, 3.2

Fagereng, A., M. B. Holm, and G. J. Natvik (2016). Mpc heterogeneity and household balance sheets. 2, 1.2, 1.4.5, 1.4.5, 1.5.1.3

Fahr, S. and F. Smets (2010). Downward wage rigidities and optimal monetary policy in a monetary union*. *The Scandinavian Journal of Economics 112*(4), 812–840. 2, 3.6.2

Fair, R. C. and J. B. Taylor (1980). Solution and maximum likelihood estimation of dynamic nonlinear rationalexpectations models. 2.5.2

Favilukis, J., S. C. Ludvigson, and S. Van Nieuwerburgh (2017). The macroeconomic effects of housing wealth, housing finance, and limited risk sharing in general equilibrium. *Journal of Political Economy 125*(1), 140–223. 1.3.2

Fernandez-Villaverde, J. and D. Krueger (2011). Consumption and saving over the life cycle: How important are consumer durables? *Macroeconomic dynamics 15*(5), 725–770. 1.2

Garriga, C. and A. Hedlund (2016). Mortgage debt, consumption, and illiquid housing markets in the great recession. Technical report, Working Paper. 1.2

Gordon, R. J. (2012). Is us economic growth over? faltering innovation confronts the six headwinds. Technical report, National Bureau of Economic Research. 3.7.4

Greenwood, J., Z. Hercowitz, and G. W. Huffman (1988). Investment, capacity utilization, and the real business cycle. *The American Economic Review*, 402–417. 1.3.1

Greer, J. M., J. A. Parker, and N. S. Souleles (2006). Household expenditure and the income tax rebates of 2001. *The American Economic Review 96*(5), 1589–1610. 1.2

Guerrieri, L. and M. Iacoviello (2017). Collateral constraints and macroeconomic asymmetries. *Journal of Monetary Economics 90*, 28–49. (document), 1.2

Guvenen, F., F. Karahan, S. Ozkan, and J. Song (2015). What do data on millions of us workers reveal about life-cycle earnings risk? Technical report, National Bureau of Economic Research. 1.4.2, 1.4.2

Hansen, L. P. and K. J. Singleton (1982). Generalized instrumental variables estimation of nonlinear rational expectations models. *Econometrica: Journal of the Econometric Society*, 1269–1286. 1.5.2

Huggett, M. (1993). The risk-free rate in heterogeneous-agent incomplete-insurance economies. *Journal of economic Dynamics and Control 17*(5), 953–969. 1.2

Iacoviello, M. and M. Pavan (2013). Housing and debt over the life cycle and over the business cycle. *Journal of Monetary Economics 60*(2), 221–238. 1.2, 1.4.6, 1.C.1

Justiniano, A., G. E. Primiceri, and A. Tambalotti (2015). Credit supply and the housing boom. Technical report, National Bureau of Economic Research. 1.2

Kahn, S. (1997). Evidence of nominal wage stickiness from microdata. *The American Economic Review*, 993–1008. 3, 2.1

Kahneman, D., J. L. Knetsch, and R. H. Thaler (1991). Anomalies: The endowment effect, loss aversion, and status quo bias. *The Journal of Economic Perspectives*, 193–206. 2.2

Kaplan, G., B. Moll, and G. L. Violante (2016). Monetary policy according to hank. Technical report, National Bureau of Economic Research. (document), 1.1, 1.3, 1.4.2

Kaplan, G. and G. L. Violante (2010). How much consumption insurance beyond self-insurance? *American Economic Journal: Macroeconomics 2*(4), 53–87. 1.1, 1.5.1, 1.5.1, 1.5.1.1

Kaplan, G. and G. L. Violante (2014). A model of the consumption response to fiscal stimulus payments. *Econometrica 82*(4), 1199–1239. 3, 1.1, 1.2, 1.4.5

Kaplan, G., G. L. Violante, and J. Weidner (2014). The wealthy hand-to-mouth. Technical report, National Bureau of Economic Research. (document), 2, 1.4.5

Keynes, J. M. (1936). *General theory of employment, interest and money.* Palgrave Macmillan. 2.2

Kim, J. and F. J. Ruge-Murcia (2009). How much inflation is necessary to grease the wheels? *Journal of Monetary Economics 56*(3), 365–377. 2, 3.2

Krusell, P. and A. A. Smith, Jr (1998). Income and wealth heterogeneity in the macroeconomy. *Journal of political Economy 106*(5), 867–896. (document), 1.2

Lebow, D. E., R. E. Saks, and B. A. Wilson (2003). Downward nominal wage rigidity: Evidence from the employment cost index. *Advances in Macroeconomics 3*(1). 3, 2.2, 2.1

Martin, R. F. (2003). Consumption, durable goods, and transaction costs. 1.2, 1.5.1.3, 1.5.2

McKay, A. and R. Reis (2016). The role of automatic stabilizers in the us business cycle. *Econometrica 84*(1), 141–194. (document)

Mian, A., K. Rao, and A. Sufi (2013). Household balance sheets, consumption, and the economic slump. *The Quarterly Journal of Economics 128*(4), 1687–1726. (document), 2

Mian, A. and A. Sufi (2012). The effects of fiscal stimulus: Evidence from the 2009 cash for clunkers program. *The Quarterly journal of economics 127*(3), 1107–1142. 1.6.2

Mian, A. and A. Sufi (2014). What explains the 2007–2009 drop in employment? *Econometrica 82*(6), 2197–2223. 2

Misra, K. and P. Surico (2014). Consumption, income changes, and heterogeneity: Evidence from two fiscal stimulus programs. *American Economic Journal: Macroeconomics 6*(4), 84–106. 1.1

Ngai, L. R., K. D. Sheedy, et al. (2016). The decision to move house and aggregate housing-market dynamics. Technical report. (document), 1.4.3

Parker, J. A., N. S. Souleles, D. S. Johnson, and R. McClelland (2013). Consumer spending and the economic stimulus payments of 2008. *The American Economic Review 103*(6), 2530–2553. 1.2

Ravn, M. O. and V. Sterk (2016). Macroeconomic fluctuations with hank and sam: An analytical approach. (document)

Reiter, M. (2009). Solving heterogeneous-agent models by projection and perturbation. *Journal of Economic Dynamics and Control 33*(3), 649–665. (document)

Rotemberg, J. J. (1982). Sticky prices in the united states. *The Journal of Political Economy*, 1187–1211. 3.2

Rudebusch, G. D. (2009). The fed's monetary policy response to the current crisis. *FRBSF economic Letter*. 2.5.4, 2.2, 3.5.1, 3.1

Saiz, A. (2010). The geographic determinants of housing supply. *The Quarterly Journal of Economics 125*(3), 1253–1296. 1.4.3

Schmitt-Grohé, S. and M. Uribe (2012). The making of a great contraction with a liquidity trap and a jobless recovery. 2, 2.2, 3.1, 3.2

Schmitt-Grohé, S. and M. Uribe (2013). Downward nominal wage rigidity, currency pegs, and involuntary unemployment. *Unpublished manuscript, Columbia University*. 2, 2.2

Smets, F. and R. Wouters (2003). An estimated dynamic stochastic general equilibrium model of the euro area. *Journal of the European economic association 1*(5), 1123–1175. 3.5.1

Smith, J. C. (2000). Nominal wage rigidity in the united kingdom. *The Economic Journal 110*(462), 176–195. 2.2

Summers, L. H. (2014). Reflections on the new secular stagnation hypothesis. *Secular stagnation: Facts, causes and cures*, 27–38. 3.7.4

Tobin, J. (1972). Inflation and unemployment. *American economic review 62*(1), 1–18. 3.2

Woodford, M. (2003). *Interest and Prices: Foundations of a Theory of Monetary Policy.* Connecticut, USA: Princeton University Press. 2.1

Yang, F. (2009). Consumption over the life cycle: How different is housing? *Review of Economic Dynamics 12*(3), 423–443. 1.2