Chapter 4

Heterogeneous Discounting and Unemployment

Abstract

We propose a novel channel of amplification in the labour market via heterogeneity in discounting of future values between workers and firms. If workers and firms discount the continuation of matches differently in response to aggregate shocks, we illustrate in a stylised environment that the degree of *dispersion* between the discount factors of different households introduces an additional amplification channel via its effect on the job creation margin. We characterise this mechanism in the context of a standard TANK model extended to allow for search frictions in the labour market. In the TANK model discount factor dispersion is driven by the relative response of profit income to labour income in response to shocks. We use the model to quantify the strength of this novel amplification channel, finding under a standard calibration that this channel amplifies both demand and supply shocks. Quantitatively we find that the amplification is significant for demand shocks due to the higher sensitivity of profit income, whilst irrelevant for the transmission of supply shocks owing to the well-known insensitivity of profits to changes in productivity in this environment.

Keywords: *Search & matching; TANK models; Household heterogeneity; Stochastic discount factors*

JEL Codes: *E24, E30, E32, E44, J23*

4.1 Introduction

In this paper we study of novel channel of amplification for the transmission of shocks to the labour market which operates through heterogeneity in stochastic discount factors across households. In the presence of labour market frictions employment matches have an asset value. In an environment where households are heterogeneous in their stochastic discount factors this implies that some workers will value the continuation of matches differently to the firms they are matched with. Firstly, in a stylised two-agent environment with search frictions and heterogeneity in stochastic discount factors, we illustrate that in the presence of SDF heterogeneity, the dispersion in SDFs matters for the response of the joint match surplus to shocks and ultimately for the transmission of these shocks to the labour market via the job creation channel. In principle this channel can either amplify or dampen the effect of aggregate shocks depending on the direction of dispersion, which itself is driven by the underlying differences in household income sources and their relative responses to shocks.

We embed this mechanism into an otherwise standard two-agent New Keynesian model ('TANK') based on Bilbiie (2008) extended to allow for search frictions in the labour market frictions in order investigate the determinants of discount factor dispersion. We use this environment for two main reasons. Firstly, by definition TANK models adopt a tractable two-agent structure of household heterogeneity. Secondly, the presence of the NK block allows the presence of firms with market power who earn profits in equilibrium which are typically only remunerated to one type of household. In equilibrium this implies that consumption heterogeneity is driven by differences in income sources which generates heterogeneity in SDFs. Moreover, the presence of nominal rigidities allows us to study the implications of discount factor dispersion for the transmission of demand shocks (as well as more commonly studied supply shocks). Through the lens of the resulting framework, we illustrate that fluctuations in discount factor dispersion is entirely driven by differences in the response of labour income relative to profit income to shocks.

Next, we utilise the model under a fairly standard calibration to quantify the direction and size of this novel amplification channel in response to both demand and supply shocks. In both cases we find that this channel amplifies the effects of aggregate shocks, though to very different degrees. In the first instance, we find that this channel delivers significant amplification in the case of a demand shock. This is for two key reasons. Firstly, profits in the NK model are highly sensitive to demand shocks. Secondly, as is well-known, profits are countercyclical with respect to demand shocks in the TANK model. This means that profit and labour income are pushed in opposite directions in the case of a demand shock. Taken together, this generates sizeable volatility in discount factor dispersion and quantitatively significant amplification of the demand shock. This contrasts starkly to the case of a supply (productivity) shock. In this case the model has the well-known feature that profits are relatively insensitive to productivity shocks, whilst labour and profit incomes co-move positively in this case. Overall the amplification generated under supply shocks is quantitatively negligible.

Finally, we investigate the sensitivity of these findings to several key parameters, notably those governing the SDF sensitivity relating to household preferences and worker bargaining power. Using non-recursive preferences allows us to examine separately the implications of increasing household intertemporal elasticity of substitution (IES) compared to the degree of risk aversion. Overall we illustrate that attitudes towards inter temporal substitution are more important for determining SDF volatility, and that lower levels of IES amplify the impact of heterogeneous discounting.

Related literature: This paper is chiefly related to several broad strands of the literature bridging household heterogeneity, labour market frictions, and the role of discounting for labour market fluctuations.

Firstly, the paper is related to the vast literature studying the role of household heterogeneity for the transmission of aggregate shocks. This covers both an earlier literature which utilised more tractable two-agent structures (as we do in this paper) such as contributions (e.g. Bilbiie 2008), and the more recent literature using richer models of household heterogeneity based on Kaplan et al. (2018) who embed a standard incomplete markets model based on Huggett (1993) into a New Keynesian framework. This literature has tended to emphasise the role of heterogeneity in household marginal propensities to consume (MPCs), both for the transmission of shocks and for the appropriate designs of monetary and fiscal policy.[1] Whilst in these environments there is heterogeneity in the stochastic discount

[1]Notable more recent examples using TANK models are: Galí et al. (2007), Nisticò (2016), Ascari et al. (2017), Broer et al. (2020), Cantore and Freund (2021), Bilbiie (2023), and Debortoli and Galí (2024). For HANK models see: McKay et al. (2016), Acharya and Dogra (2020), Auclert (2019), Auclert et al. (2020), Nuno and Thomas (2020), Werning (2015), Bayer et al. (2020), and Luetticke (2021).

factors between households, in the vast majority of this literature there is typically no role for those of constrained households as these households are off their Euler equation. Moreover, there are usually no frictions in the labour market which would otherwise give employment matches an asset value.

Secondly, there is a small but growing literature examining the interaction between household heterogeneity, incomplete markets and labour market frictions. A key contribution is the framework developed by Ravn and Sterk (2017, 2021), who embed search & matching frictions into a tractable incomplete markets model in Broer et al. (2020) and illustrate how the link between endogenous unemployment risk and precautionary savings can significantly amplify aggregate shocks. Broer et al. (2023) use a similar environment augmented to allow for a more empirically consistent job creation margin and use the environment to quantify the role of unemployment risk over the business cycle. Gornemann et al. (2021) build and solve a rich quantitative model featuring incomplete markets, search & matching frictions, and allowing for heterogeneity in preferences and worker skills. They use the model to study the distributional consequences of monetary policy. All these papers rely on the same mechanism for amplification - the interaction between aggregate demand, precautionary saving under incomplete markets and endogenous income risk from frictions in the labour market. In contrast, we emphasise a very different channel of amplification which does not rely on endogenous unemployment risk. We show that even when there is perfect risk-sharing within the family, household heterogeneity can still generate amplification in the labour market when there is heterogeneity in the discounting of match continuation values.

Finally, this paper also contributes to a small literature studying the relationship between discount rates and the labour market. This literature views job creation as another form of financial investment which should be subject to the same rate of discount as financial assets. The key contribution here is Hall (2017), who finds that feeding an estimated process for discount factors on US stock market data into a standard search & matching model leads to empirically plausible fluctuations in unemployment and overcomes the volatility puzzle. Kehoe et al. (2023) develop a standard search & matching model where household preferences are consistent with time-varying risk over the business cycle (via recursive preferences) and find quantitatively that the model is able to replicate US unemployment fluctuations without generating further empirical puzzles. Finally, a recent contribution by Martellini et al. (2021) builds a rich search framework with heterogeneous match quality, and instead find that shocks to financial discount rates on their own do not generate sufficient volatility in labour market transition probabilities to induce large unemployment fluctuations. Relative to these papers, our contribution is to highlight that heterogeneity in the rates of financial discounting, induced by heterogeneity in income sources and the ability to smooth consumption, endogenously generates additional volatility in the labour market via its effect on the job creation margin.

Layout. The rest of the paper is organised as follows. Section 4.2 outlines the core mechanism in a stylised environment. Section 4.3 embeds the mechanism into a standard TANK model extended to allow for search frictions in the labour market and outlines the determinants of heterogeneous discounting. Section 4.4 details the main results of the paper where we quantify the role of heterogeneous discounting in amplifying demand and supply shocks. Section 4.5 concludes.

4.2 A Stylised Example

In this section we illustrate how the presence of heterogeneity in discounting between workers and firms affects job creation incentives in an otherwise standard environment with search & matching frictions. We derive the implications of discount factor heterogeneity for the key equation in this model - to job creation condition.

Environment. Time is discrete and runs forever. The economy is populated by a continuum of infinitely-lived households and is subject to fluctuations due to aggregate shocks. Households are risk-averse, derive utility from consuming a final consumption good, and discount the future according to $\beta \in (0, 1)$. Households are heterogeneous in their ability to smooth consumption and are indexed by $i = \{C, U\}$ to be one of two types: (i) constrained households who only receive income from working, and (ii) unconstrained households who trade assets (including shares in firm profits). The composition of household types is fixed, where the fraction of constrained households is given by $\lambda \in (0, 1)$. We assume there is perfect risk-sharing within the two types of household, but not *across* households. The consumption of different households will therefore respond differently to aggregate shocks, generating dispersion in stochastic discount factors $\beta^i_{t,t+\tau}$ where:

$$\beta^i_{t,t+\tau} = \beta \frac{u'(C^i_{t+\tau})}{u'(C^i_t)}$$

where $u'(\cdot) \geq 0$ and $u''(\cdot) \leq 0$.

Households inelastically supply labour to the market but face search frictions such that not all household members successfully find jobs. If matched, worker-firm pairs produce output $A > 0$ and the worker is paid a wage w^i, where wages are bargained such that the worker receives fraction $\eta \in (0,1)$ of the joint surplus of the match. If unemployed workers instead receive flow value $b > 0$. Matches separate exogenously with probability $\rho \in (0,1)$. Firms post vacancies v at per period cost $\chi > 0$. Matches are formed according to a constant returns matching function $M(V, U)$, where U is the number of job searchers. Contact rates for workers and vacancies $\{p(\theta), q(\theta)\}$ depend on the ratio of vacancies-to-searchers, i.e. $\theta = V/U$, which is endogenously determined by job creation.

Match surplus. Workers from different households discount the continuation of a match at different rates in the presence of aggregate shocks (i.e. out of steady state). This implies heterogeneity in the match surplus and therefore in wages w^i across households. To see this, define the total surplus of a match with worker from household i as $\mathbb{S}_i$, and the worker and firm surpluses as $\{\mathbb{S}_i^w, \mathbb{S}_i^f\}$. Given the assumptions made above can write:

$$\mathbb{S}_{U,t} = A - b + (1-\rho)\mathbb{E}_t \beta_{t,t+1}^U \mathbb{S}_{U,t+1}$$

$$\mathbb{S}_{C,t} = A - b + (1-\rho)\mathbb{E}_t \left[\beta_{t,t+1}^C \mathbb{S}_{C,t+1}^w + \beta_{t,t+1}^U \mathbb{S}_{C,t+1}^f\right]$$

For matches where workers are members of unconstrained households, both members of the match discount its continuation using the same discount factor. Otherwise, the continuation value of the match is discounted differently by either side.

Job creation. The implications for the job creation will determine whether the presence of heterogeneous discounting amplifies or dampens the response of the labour market to aggregate shocks. Imposing the free entry condition implies:

$$\chi = q(\theta)(1-\eta) \int \mathbb{S}_{i,t} di = q(\theta)(1-\eta) \cdot [(1-\lambda)\mathbb{S}_{U,t} + \lambda \mathbb{S}_{C,t}]$$

where $q(\theta)$ is the job filling rate and $q'(\cdot) \leq 0$. Job creation is driven by fluctuations in the weighted average of the joint surpluses, which somewhat complicates further derivations. To get around this we make a simplifying assumption that for the purposes of wage bargaining all workers use the *average* discount factor, defined as:

$$\beta_{t,t+1} = \lambda \beta_{t,t+1}^C + (1-\lambda)\beta_{t,t+1}^U$$

which also ensure that all workers are paid a single wage which we can characterise analytically.[2]

Proposition 4.1. *The Nash wage is given by the usual weighted average between worker & firm reservation wages:*

$$w_t = (1-\eta)b + \eta\bar{w}_t$$

where the firm's reservation wage $\bar{w}_t$ is now given by:

$$\bar{w}_t = A_t + \mathbb{E}_t\beta^U_{t,t+1}\chi\theta_{t+1} + \underbrace{\lambda\left[\frac{\chi}{q(\theta_{t+1})}(1-\rho-p(\theta_{t+1}))(\beta^U_{t,t+1}-\beta^C_{t,t+1})\right]}_{\text{SDF heterogeneity}}$$

Proof. See Appendix 4.A. □

In this environment, heterogeneous discounting affects labour market volatility via its effect on the firm's reservation wage and therefore job creation (as illustrated in Corollary 4.2).

Corollary 4.2. *The job creation condition in this environment can be written as:*

$$\frac{\chi}{q(\theta_t)} = (1-\eta)(A-b) + \mathbb{E}_t\frac{\chi}{q(\theta_{t+1})}\cdot\left[(1-\rho)\beta^U_{t,t+1} - \eta\lambda(1-\rho-p(\theta_{t+1}))(\beta^U_{t,t+1}-\beta^C_{t,t+1})\right]$$

where in the absence of household heterogeneity (i.e. $\lambda = 0$ and $\beta^C_{t,t+1} = \beta^U_{t,t+1}$) this channel disappears.

Whether this novel channel amplifies or dampens the response of the labour market to aggregate shocks ultimately depends on how the dispersion in discount factors itself responds to the shock, or more specifically on the direction of the inequality:

$$\beta^U_{t,t+1} - \beta^C_{t,t+1} \lesseqgtr 0$$

Overall the modified job creation in Corollary 4.2 suggests that the quantitative importance of this amplification channel will depend on (i) the sensitivity of SDFs to shocks, (ii) the fraction of constrained households λ, and (iii) worker bargaining power η. Without further structure we cannot characterise what forces determine fluctuations in discount factor dispersion. In the following section, we embed this stylised environment into a TANK model with search frictions in order to explore the implications of this channel in an otherwise well-understood environment.

[2] One rationalization for this could be that workers are represented by a third party in the wage bargain (e.g. a union) who bargains with firms on behalf of all workers, rather than just the worker in the match, in order to prevent wage differentials.

4.3 Quantification: Insights from a TANK Model

In this section we quantify the contribution of heterogeneous discounting to labour market volatility. To do this, we first embed the environment outlined above into a relatively standard two-agent New Keynesian model ('TANK') based on Bilbiie (2008). We then briefly discuss what drives heterogeneity in discount factor fluctuations in the model, before outlining a standard calibration of the model. The main results from simulating the responses to demand and supply shocks are presented in the following section.

4.3.1 A Simple TANK Model

We utilise a TANK model to quantify the role of heterogeneous discounting for two main reasons: (i) we maintain the tractable two-agent approach to capturing heterogeneity in consumption smoothing, and (ii) allowing for a New Keynesian block allows firms to make excess profits from market power which are typically remunerated to unconstrained households such that households have different income sources. Moreover, the NK block allows us to study the implications of heterogeneous discounting for demand shocks (as well as supply shocks). Below we briefly outline the key additional features of the model relative to our previous stylised environment. Further details about the TANK model can be found in Appendix 4.B.

Households. We broadly retain the same assumptions as in Section 4.2. A fraction $1-\lambda$ of "unconstrained" (U) households smooth income using riskless one-period bonds B_t yielding the gross nominal return R_t, and shares in an equity fund which owns firm profits S_t which can be purchased at (real) price Q_t. Optimization yields Euler equations for these two assets. The remaining fraction λ of "constrained" (C) households consume their labour income in every period. Households consume a continuum of differentiated goods sold by monopolistically competitive firms indexed by $k \in [0,1]$ at price $P_{k,t}$. Total consumption of household i is given by $C_t^i \equiv \left(\int_0^1 {C_{k,t}^i}^{\frac{\nu-1}{\nu}} \right)^{\frac{\nu}{\nu-1}}$, where $\nu > 0$ is the elasticity of substitution across goods.[3]

[3]The overall demand for each good variety in the economy is standard:

$$C_{k,t} = \left(\frac{P_{k,t}}{P_t}\right)^{-\nu} C_t$$

Labour market. Again we retain the same assumptions about the labour market as in Section 4.2. It follows that aggregate employment denoted by N_t obeys the following law of motion:

$$N_t = (1-\rho)N_{t+1} + M_t \tag{4.1}$$

where M_t are newly formed matches.

Firms. There are a continuum of monopolistically competitive firms indexed by k who produce a differentiated good which they then sell to consumers at price $P_{k,t}$. All firms employ a fraction of total workers $N_{k,t}$ who produce output according to $Y_{k,t} = A_t N_{k,t}$ and are each paid a real wage $w_{k,t}$. In addition to per period vacancy costs $\chi > 0$, firms also face quadratic price adjustment costs governed by $\Psi > 0$.[4] Firm k chooses employment and vacancy postings to maximise present discounted profits subject to a law of motion for firm employment, the production function, and the standard demand curve for their output. Standard optimization techniques yield a standard job creation condition and a non-linear Phillips curve:

$$\frac{\chi}{q(\theta_t} = \varepsilon A_t - w_t + (1-\rho)\mathbb{E}_t\left[\beta^U_{t,t+1}\frac{\chi}{q(\theta_{t+1})}\right] \tag{4.2}$$

$$1 - \Psi(\pi_t - \pi)\pi_t + \mathbb{E}_t[\beta^U_{t,t+1}\Psi(\pi_{t+1} - \pi)\pi_{t+1}\frac{Y_{t+1}}{Y_t}] = \nu(1-\varepsilon_t) \tag{4.3}$$

where π_t is the (gross) inflation rate and ε_t are real marginal costs. The fact that firms are identical ensures that optimal choices are symmetric across firms, which allows us to drop the k subscripts.

Wages. The determination of wages is identical to that described in Section 4.2, where we maintain the assumptions ensure a single equilibrium wage across all workers.[5] The only difference is that now (real) marginal costs ε_t also enter the firm's reservation wage.

Closing the model. Markets for goods, bonds and shares must clear. Goods market clearing requires that the total supply of goods (net of the costs of posting

where $P_t \equiv \left(\int_0^1 P_{k,t}^{1-\nu}\right)^{\frac{1}{1-\nu}}$ is the relevant price index.

[4]We assume Rotemberg (1982) price rigidities in order to stay closer to the TANK/HANK literature.

[5]Throughout the paper we maintain Nash bargaining for simplicity and tractability, but is just one possible way in which wages could be determined in relation to the worker's surplus that would preserve pass-through from heterogeneity in discount factors to the job creation margin, which is key for our main channel to operate.

vacancies and changing prices) must equal the total demand of households, given by $C_t \equiv \int_0^1 C_{k,t} dk = (1-\lambda)C_t^U + \lambda C_t^C$. This implies the following goods market clearing condition:

$$C_t = Y_t \left(1 - \frac{\Psi}{2}(\pi_t - \pi)\right) - \chi V_t \tag{4.4}$$

Bonds are in zero net supply and the market for shares must also clear, i.e. $B_t^U = 0, S_t^U = \frac{1}{1-\lambda}$. Lump-sum taxes are levied on each type of household to finance their unemployment benefits. Overall this implies the following equilibrium consumption allocation across households:

$$C_t^U = w_t N_t + \frac{1}{1=\lambda} D_t,\ C_t^C = w_t N_t \tag{4.5}$$

In this environment, consumption differences across the two types of households are due to the additional financial income unconstrained households obtain from holding firm shares.[6] Finally, we assume the presence of a central bank who determine R_t via a standard Taylor-type rule.

4.3.2 What Drives Heterogeneity in Discount Factors?

Through the lens of the TANK model, heterogeneity in discount factors is driven by the relative responses of different sources of income to aggregate shocks. Given the consumption allocation in equation (4.5), it follows that fluctuations in $\beta_{t,t+1}^C$ are entirely determined by fluctuations in labour income whereas $\beta_{t,t+1}^U$ captures changes to *both* labour and profit income in response to shocks. Hence factor dispersion in the model is driven by the response of profits *relative* to labour income.[7] We prove this formally below.

Proposition 4.3. *Suppose the utility function is a standard CRRA: $u'(C) = C^{-\sigma}$ It can be shown that to a first-order approximation discount factor dispersion can be expressed as:*

$$\beta_{t,t+1}^U - \beta_{t,t+1}^C \approx \frac{\beta\sigma}{1 + (1-\lambda)\gamma WN} \left[\mathbb{E}_t \Delta(w_{t+1} N_{t+1)} - \gamma \mathbb{E}_t \Delta D_{t+1} \right] \tag{4.6}$$

where $\gamma \equiv \frac{wN}{D}$ is the steady state ratio of labour income to profits.

[6] The importance of how firm profits are distributed across households in models of limited asset market participation models was first highlighted in Bilbiie (2008).

[7] This reflects another dimension of the importance of firm profits in New Keynesian models with limited asset market participation, a point originally raised in Bilbiie (2008) relating to how firm profits are distributed. Here we take the simple case of no redistribution of firm profits.

Proof. See Appendix 4.A. □

In response to an aggregate shock, whether the response of labour or profit income dominates depends on their relative direction and volatility, the size of labour income relative to profits in steady state (γ), as well as other factors anticipated in Section 4.2 such as λ (which controls the share of profits unconstrained households hold) and the preference parameters characterising the sensitivity of SDFs (here captured by σ). In the next section we outline a standard calibration of the model which we can then use to characterise this amplification channel quantitatively.

4.3.3 Parameterization

We outline a standard parameterization of the model at a quarterly frequency. For several values we calibrate them externally by setting them equal to standard values used in the literature on business cycles and search & matching models. Additionally we determine some parameters relating to the labour market by imposing external restrictions on steady state values such that the model matches key features of the US labour market. The parameterization is summarised in Table 4.2.

Functional forms. We use a more general non-recursive household utility function popularised by Epstein and Zin (1989), in order to allow for a richer parameterization of household stochastic discount factors. The expected discounted lifetime utility of household i is given by J_t^i, which is defined as:

$$J_t^i = \left[(1-\beta){C_t^i}^{1-\frac{1}{\varphi}} + \beta\left(\mathbb{E}_t\left[{J_{t+1}^i}^{1-\sigma}\right]\right)^{\frac{1-\frac{1}{\varphi}}{1-\sigma}}\right]^{\frac{1}{1-\frac{1}{\varphi}}}$$

where $\sigma > 0$ measures the degree of household risk aversion, and $\varphi > 0$ captures the household's intertemporal elasticity of substitution (IES).[8] The corresponding expression for the stochastic discount factor of household i is given by:

$$\beta_{t,t+1}^i = \beta\left(\frac{C_{t+1}^i}{C_t^i}\right)^{-\frac{1}{\varphi}}\left(\frac{J_{t+1}^i}{\mathbb{E}_t\left[{J_{t+1}^i}^{1-\sigma}\right]^{\frac{1}{1-\sigma}}}\right)^{\frac{1}{\varphi}-\sigma}$$

For matching we assume a standard Cobb-Douglas specification:

$$M(V_t, U_t) = \bar{m} V_t^{1-\xi} U_t^{\xi}$$

[8]Imposing the restriction $\sigma = \frac{1}{\varphi}$ recovers the standard power utility (CRRA) specification.

where ξ is the elasticity of new matches with respect to the stock of job searchers and $\bar{m}$ is a match efficiency constant.

Externally calibrated parameters. We calibrate several model parameters using standard values. The household discount factor β is set to 0.9926 to target a 3% annual average real interest rate. The intertemporal elasticity of substitution (IES) φ is set to 1.5. We initially set $\sigma = 1/\varphi$ such that preferences are CRRA. The elasticity of substitution across varieties is set to $\nu = 6$ to generate a 20% steady state markup. The price adjustment costs parameter Ψ is chosen to match an average price duration of 1 year. We choose a value of the match elasticity $\xi = 0.5$ which is standard in the context of the survey by Petrongolo and Pissarides (2001) and then set $\eta = \xi$ such that the Hosios' (1990) condition is satisfied. We choose a Taylor rule coefficient that is standard. Finally, the fraction of hand-to-mouth households λ is set t0 21% based on Debortoli and Galí (2024).

Steady state targets. Several parameters relating to the labour market are set to match standard quarterly targets for the US labour market. Specifically we set the values of $\rho, \chi, \bar{m}, b$ to match: (i) a 5.5% average unemployment rate, (ii) an average job filling rate equal to 0.7 as in Den Haan et al. (2000), (iii) an average quarterly job finding rate equal to 0.45 as in Shimer (2012), and (iv) a 40% replacement ratio as in Shimer (2005). Note that under this calibration of the replacement rate the standard Shimer (2005) critique will apply.

Model solution. The model is solved using standard perturbation techniques around a steady state with zero inflation.[9] A full description of the non-linear equilibrium conditions used to solve the model is given in Table 4.1 in Appendix 4.B.

4.4 Main Results

In this section we present the main quantitative findings. We use the calibrated TANK model to quantify the contribution of heterogeneous discounting by simulating the response to two different types of aggregate shocks - a demand (monetary policy) shock and a supply (productivity) shock - and comparing against the counterfactual where there is no heterogeneity in discounting (i.e. $\lambda = 0$).

[9] The model is initially solved using first-order perturbation, but later solved using higher order perturbation to allow the degree of risk aversion to influence the stochastic solution to the model when we deviate from using a power utility specification.

4.4.1 Demand shock

We first study the response of the model to a 1% contractionary shock to the monetary policy rate in the model. The responses of key variables in the model are presented in Figure 4.5. Declining demand is associated with a fall in output and inflation as standard, driven by downward pressure on labour demand by firms leading to falling wages and vacancies, lower labour market tightness and higher unemployment. As is well-known, a decline in the policy rate is also associated with an *increase* in firm profits in the NK environment.[10] Profits and labour income therefore move in different directions, inducing dispersion in the discount factors across households. Moreover, we find that profits are significantly more sensitive to demand shocks than labour income in our model. This induces the dispersion in discount factors to be positive, i.e. $\beta^U_{t,t+1} - \beta^C_{t,t+1} > 0$. As discussed above, this amplifies the effect of a contractionary shock via its effect on job creation.

The size of this additional amplification channel can be seen be examining the counterfactual response when $\lambda = 0$ in Figure 4.1. Our first main result is that discount factor heterogeneity amplifies the effect of the demand shock quite substantially. The greater fall in vacancies induces an additional 20% increase in the unemployment response, and leads to output falling by around an additional 1%. Note that because the discount factor heterogeneity channel is inherently forward-looking, its effect is entirely concentrated in the impact of the shock. Heterogeneity in discounting does not affect the persistence of the economy's response to the shock.

4.4.2 Technology shock

Secondly we study the response of the model to a 1% positive (labour) productivity shock. The responses of key variables in the model are presented in Figure 4.6. Job creation responds positively to the shock because the value of a worker to a firm increases, leading to falling unemployment and lower labour market tightness as standard. However the rise in real wages acts to dampens firms' incentives to hire more workers. In this case profits and labour income now move in the *same* direction. As the consumption level of constrained households increases by more (relative to steady state), the model predicts that $\beta^U_{t,t+1} - \beta_{t,t+1} < 0$, which acts to stimulate job creation.

[10]For a discussion of this channel, see Broer et al. (2020).

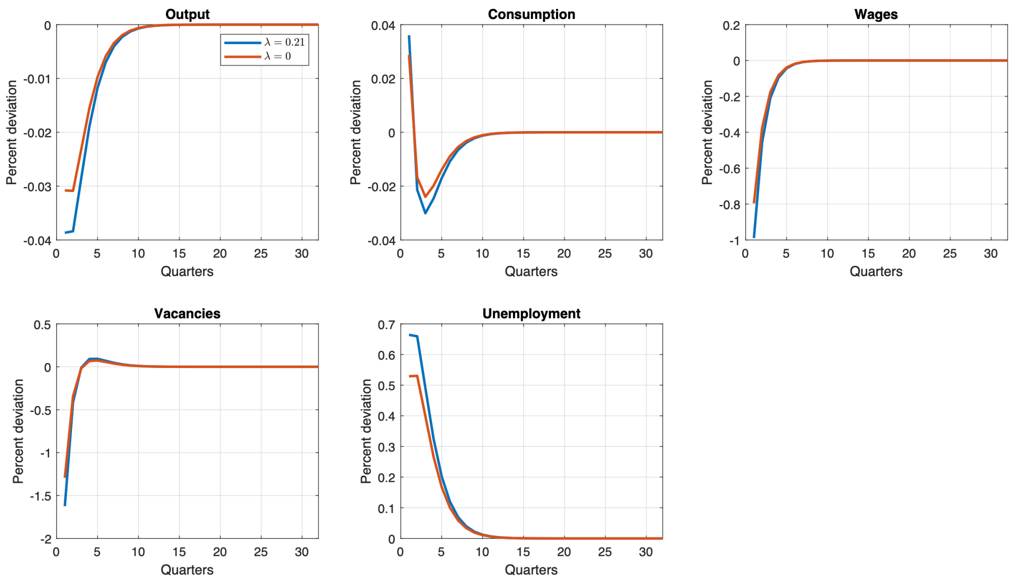

Figure 4.1. Amplification from SDF heterogeneity: Demand shock

Our second main result is that amplification of productivity shocks from our heterogeneous discounting channel is quantitatively insignificant. This is illustrated in Figure 4.2 where again we plot the responses relative to the counterfactual. This result is mainly driven by the fact that profits are not very sensitive to supply shocks, which is a well-known issue of our standard calibration approach (i.e. the unemployment volatility puzzle outlined in Shimer 2005). As a result, dispersion is SDFs is an order of magnitude smaller compared to the case of a monetary policy shock. Overall we see that the additional amplification from the heterogeneous discounting channel is vanishingly small.

Summary. We find that our heterogeneous discounting channel amplifies the response of the economy to standard demand and supply shocks. However, the quantitative significance of this channel is shown to depend on how firm profits respond to the aggregate shock relative to aggregate labour income, as well as their relative volatilities. A demand shock in the model moves labour and profit income in opposite directions and firm profits are roughly 4 times more volatile than labour income, so we find that heterogeneous discounting induces a fairly large amplification. In contrast, for a productivity shock labour income and profit income move in the same direction and have roughly equal volatilities in response to the shock, such that amplification in this case is quantitatively irrelevant.

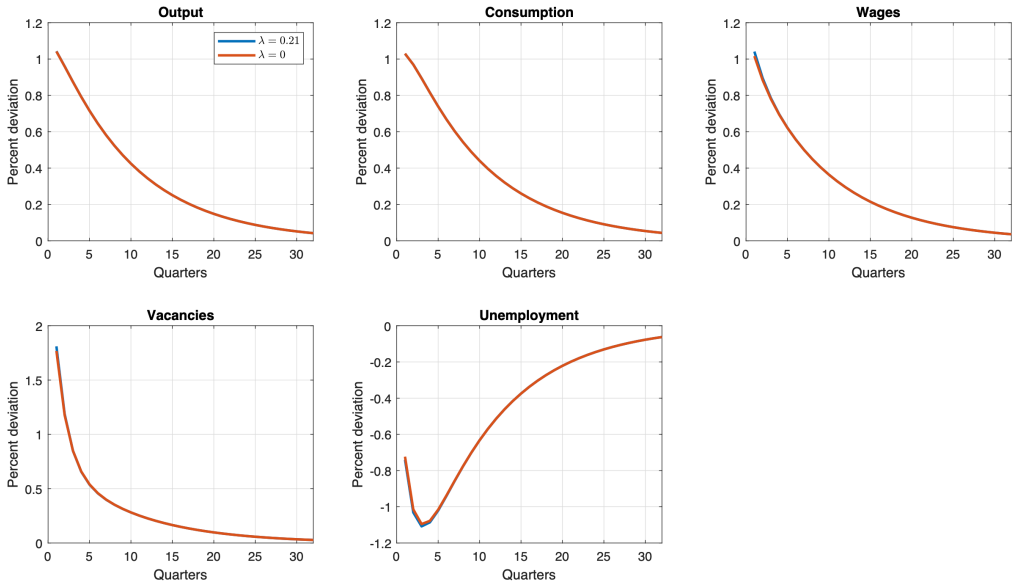

Figure 4.2. Amplification from SDF heterogeneity: Supply shock

4.4.3 Robustness

Finally, we investigate the sensitivity of amplification via this channel to changes in several key parameters outlined in Section 4.2, notably the parameters governing household preferences and the degree of worker bargaining power.

Household preferences. Epstein-Zin preferences feature two parameters: the IES parameter φ and degree of risk aversion σ. Under our baseline parameterization we set the former to a standard value and determine the latter by imposing the CRRA restriction, $\sigma = 1/\varphi$. We recompute the responses of the model to a contractionary demand shock under two alternative calibrations of these parameters. Firstly, we maintain the CRRA restriction but increase the value of φ from 1.5 to 3. Under the CRRA restriction this both simultaneously reduces household risk aversion and increases the appetite of households to engage in intertemporal substitution. Secondly, we move away from the CRRA restriction by fixing $\varphi = 1.5$ but increasing the degree of risk aversion σ to 5. This will allow us to identify whether risk aversion or intertemporal substitution appear to be more important for this amplification channel.

The IRFs are plotted in Figure 4.3. Firstly, we find that simultaneously increasing the degree of substitution elasticity whilst reducing risk aversion dampens ampli-

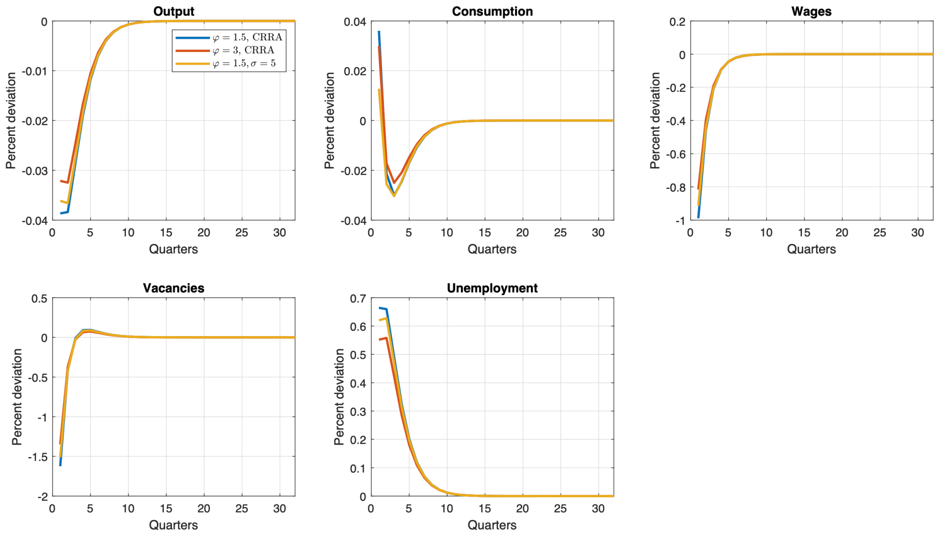

Figure 4.3. Robustness: Impact of IES φ and risk aversion σ.

fication. Quantitatively the effect on the amplification on impact is quite strong - with higher IES, the response now looks quantitatively similar to the baseline parameterization when we shut down heterogeneity ($\lambda = 0$) in Figure 2. Secondly, when we fix the IES parameter φ but increase the degree of risk aversion σ to a higher but plausible value (yellow line), we find that increasing household risk aversion without also changing attitudes towards intertemporal substitution dampens amplification. Taken together, these results suggest that household attitudes towards intertemporal substitution are more influential in determining volatility in stochastic discount factors (and therefore discount factor dispersion) than attitudes towards risk.

Bargaining power. Worker bargaining power η determines the influence of discount factor dispersion on job creation. We recompute the model responses for lower value of worker bargaining power $\eta = 0.4$. The results are plotted in Figure 4.4. Unsurprisingly, this significantly reduces amplification from the discount factor dispersion channel. Again quantitatively this is similar to maintaining the baseline calibration, but shutting down this channel altogether by eliminating household heterogeneity.

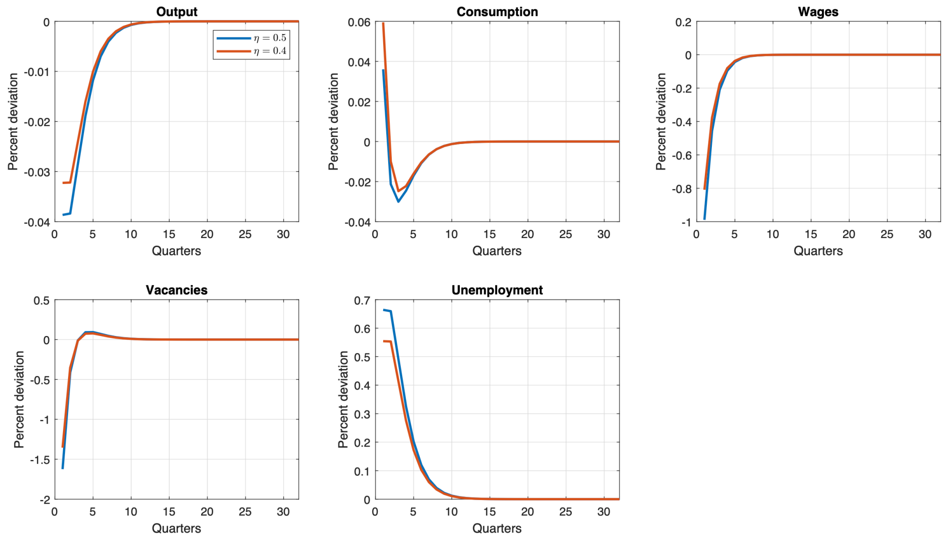

Figure 4.4. Robustness: Worker bargaining power.

4.5 Conclusion

This paper studies a novel channel of labour market amplification through heterogeneity in discounting. We show that when households are heterogeneous in their ability to smooth income and have different income sources, aggregate shocks generate heterogeneity in how households discount the continuation value of matches in the labour market, and that this matters for the transmission of aggregate shocks via the job creation margin. We illustrate this mechanism in the context of a standard TANK model extended to allow for search frictions in the labour market. We use a standard calibration of the model to quantify the effect of our heterogeneous discounting channel on the amplification of demand and supply shocks to the labour market. We find in the model that heterogeneous discounting amplifies both demand and supply shocks, though quantitatively the channel is only significant for demand shocks.

Chapter 4: Appendices

4.A Proofs

4.A.1 Proof of Proposition 4.1

Under Nash bargaining the wage satisfies the Nash sharing rule:

$$\mathbb{S}^w_t = \eta(\mathbb{S}^w_t + \mathbb{S}^f_t)$$

which states that the worker receive fraction η of the total match surplus, defined as the sum of the worker and firm surpluses. There in turn are defined as:

$$\mathbb{S}^w_t = w_t - b + \mathbb{E}_t\left[\beta_{t,t+1}(1-\rho-p(\theta_{t+1}))\mathbb{S}^w_{t+1}\right]$$

$$\mathbb{S}^f_t = A_t - w_t + (1-\rho)\mathbb{E}_t\left[\beta^U_{t,t+1}\mathbb{S}^f_{t+1}\right]$$

Substituting these expressions in the Nash sharing rule and rearranging gives:

$$w_t = (1-\eta)b + \eta A_t + \eta(1-\rho)\mathbb{E}_t\beta^U_{t,t+1}\mathbb{S}^f_{t+1} - (1-\eta)\mathbb{E}_t\beta_{t,t+1}(1-\rho-p(\theta_{t+1}))\mathbb{S}^w_{t+1}$$

Using $\mathbb{S}^w_{t+1} = \frac{\eta}{1-\eta}\mathbb{S}^f_{t+1}$:

$$w_t = (1-\eta)b + \eta\left[A_t + (1-\rho)\mathbb{E}_t\mathbb{S}^f_{t+1}(\beta^U_{t,t+1} - \beta_{t,t+1}) + \mathbb{E}_t\beta_{t,t+1}p(\theta_{t+1})\mathbb{S}^f_{t+1}\right]$$

then using the free entry condition $\frac{\chi}{q(\theta_t)} = \mathbb{S}^f_t$ yields:

$$w_t = (1-\eta)b + \eta\left[A_t + (1-\rho)\mathbb{E}_t\frac{\chi}{q(\theta_{t+1})}(\beta^U_{t,t+1} - \beta_{t,t+1}) + \mathbb{E}_t\beta_{t,t+1}\chi\theta_{t+1}\right]$$

Finally, using $\beta^U_{t,t+1} - \beta_{t,t+1} = \lambda(\beta^U_{t,t+1} - \beta^C_{t,t+1})$ and rearranging gives the final expression in Proposition 4.1:

$$w_t = (1-\eta)b + \eta\left[A_t + \mathbb{E}_t\left[\beta^U_{t,t+1}\chi\theta_{t+1}\right] + \lambda\mathbb{E}_t\left[\frac{\chi}{q(\theta_{t+1})}(1-\rho-p(\theta_{t+1}))(\beta^U_{t,t+1} - \beta^C_{t,t+1})\right]\right]$$

which has the standard interpretation of being a linear combination of the workers' and firms' reservation wages.

4.A.2 Proof of Proposition 4.3

The stochastic discount factor of some household i can be expressed as:

$$\beta^i_{t,t+1} = \beta\Big(\frac{u'(C^i_{t+1})}{u'(C^i_t)}\Big)$$

Assuming a standard CRRA utility function with parameter $\sigma > 0$ and taking a first-order Taylor expansion around the steady state gives:

$$\beta^i_{t,t+1} \approx \beta - \frac{\sigma\beta}{C^i}(C^i_{t+1} - C^i_t)$$

Using this, we can approximate discount factor dispersion by:

$$\beta^U_{t,t+1} - \beta^C_{t,t+1} \approx -\frac{\sigma\beta}{C^U}(C^U_{t+1} - C^U_t) + \frac{\sigma\beta}{C^C}(C^C_{t+1} - C^C_t)$$

$$= \frac{\sigma\beta}{C^U C^C}\left[(C^U - C^C)\mathbb{E}_t\Delta W_{t+1}N_{t+1} - \frac{C^C}{1-\lambda}\mathbb{E}_t D_{t+1}\right]$$

where the second equality follows from substituting in the (linear) equilibrium expressions for each type of household's consumption. Substituting out for the steady state consumption levels and rearranging gives the expression in Proposition 4.3:

$$= \frac{\sigma\beta}{1 + \frac{(1-\lambda)(WN)^2}{D}}\left[\mathbb{E}_t\Delta W_{t+1}N_{t+1} - \frac{WN}{D}\mathbb{E}_t\Delta D_{t+1}\right]$$

$$= \frac{\sigma\beta}{1 + (1-\lambda)WN\gamma}\left[\mathbb{E}_t\Delta W_{t+1}N_{t+1} - \gamma\mathbb{E}_t\Delta D_{t+1}\right]$$

where the final equality follows from defining $\gamma \equiv \frac{WN}{D}$ as the ratio of steady state aggregate labour income to firm profits.

4.B Model Appendix

Budget constraints. The period budget constraint for the representative unconstrained household is therefore given by:

$$C_t^U + \frac{B_t}{P_t} + Q_t S_t \leq R_{t-1}\frac{B_{t-1}}{P_t} + \int_0^1 w_{k,t} N_{k,t}^U dk + u_t b + [Q_t + D_t] S_{t-1} + T_t^U$$

where T_t^U denotes the transfers received by the unconstrained households, $N_{k,t}$ is the employment rate of agents in the household at firm k who earn the real wage $w_{k,t} \equiv \frac{W_{k,t}}{P_t}$, u_t is the fraction of household members who are unemployed, and b denotes unemployment benefits received by each unemployed member of the household.[11] Constrained households are excluded from asset markets. They consume their income in every period in a 'hand-to-mouth' fashion:

$$C_t^C = \int_0^1 w_{k,t} N_{k,t}^C dk + u_t b + T_t^C$$

where T_t^C are transfers received by constrained households. Hand-to-mouth behaviour implies perfect pass-through of labour income fluctuations to consumption for these households.

Euler equations: Unconstrained households choose consumption and asset holdings to maximise utility subject to their period budget constraint. This gives rise to the standard Euler equations for bonds and shares:

$$\frac{1}{R_t} = \mathbb{E}_t\left[\beta_{t,t+1}^U \frac{1}{\pi_{t+1}}\right] \tag{4.7}$$

$$Q_t = \mathbb{E}_t\left[\beta_{t,t+1}^U (Q_{t+1} + D_{t+1})\right] \tag{4.8}$$

where $\pi_{t+1} = P_{t+1}/P_t$ is the gross rate of price inflation in the economy. Note that the stochastic dis- count factor which matters for pricing assets belongs to unconstrained households due to the fact that constrained households are excluded from financial markets.

Labour market. The labour force is assumed to be constant and normalized to 1 such that $0 \leq N_t \leq 1$. The relevant stock of job searchers at the beginning of the period is given by:

$$U_t = 1 - N_{t-1} + \rho N_{t-1} = 1 - (1-\rho) N_{t-1} \tag{4.9}$$

[11] Note that the real wage is assumed to be common across workers from different household types. This assumption will be discussed in detail below.

where the job filling and job finding rates are defined as functions of tightness as standard:

$$\theta_t = \frac{V_t}{U_t}, \;\; q(\theta_t) = \frac{M_t}{V_t}, \;\; p(\theta_t) = \theta_t q(\theta_t) \tag{4.10}$$

Firms. Overall, firm k will choose employment, vacancies and prices to maximise the present discounted value of their profits:

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^U_{t,t+1} \left[\frac{P_{k,t}}{P_t} Y_{k,t} - \chi V_{k,t} - w_{k,t} N_{k,t} - \frac{\Psi}{2} \left(\frac{P_{k,t}}{P_{k,t-1}} - \pi \right)^2 Y_{k,t} \right]$$

where the parameter Ψ controls the severity of the adjustment costs and π is the steady state gross inflation rate. Profits are discounted using the discount factor of unconstrained households, as only these households trade shares in firm profits.

$$N_{k,t} = (1-\rho) N_{k,t-1} + q(\theta_t) V_{k,t}$$

$$Y_{k,t} = \left(\frac{P_{k,t}}{P_t} \right)^{-\nu} Y_t$$

Standard optimization techniques and re-arrangement of first-order conditions yield the job creation condition and optimal pricing condition reported in the main body.

Value functions: The marginal value to the firm of a filled vacancies $\mathbb{S}^f_t$ is given by:

$$\mathbb{S}^f_t = \varepsilon_t A_t - w_t + (1-\rho)\mathbb{E}_t \left[\beta^U_{t,t+1} \mathbb{S}^f_{t+1} \right]$$

where the assumption of free entry implies that the firm's surplus from a match is equal to the expected cost of posted a vacancy. The worker's surplus is given by:

$$\mathbb{S}^w_t = w_t - b + \mathbb{E}_t \left[\beta_{t,t+1} (1 - \rho - p(\theta_{t+1})) \mathbb{S}^w_{t+1} \right]$$

where discounting is with respect to the union's discount factor, which captures heterogeneity in discounting across the two types of households.

Monetary policy. The return on bonds is determined by a monetary authority who follow a simple Taylor-type rule:

$$R_t = R_{ss} \pi_t^{\phi_\pi} \zeta_t \tag{4.11}$$

where ζ_t is the monetary policy (demand) shock we use in our simulations.

$J_t^i = \left[(1-\beta){C_t^i}^{1-\frac{1}{\varphi}} + \beta\left(\mathbb{E}_t\left[{J_{t+1}^i}^{1-\sigma}\right]\right)^{\frac{1-\frac{1}{\varphi}}{1-\sigma}}\right]^{\frac{1}{1-\frac{1}{\varphi}}}$	(Household values)
$\beta_{t,t+1}^i = \beta\left(\frac{C_{t+1}^i}{C_t^i}\right)^{-\frac{1}{\varphi}}\left(\frac{J_{t+1}^i}{\mathbb{E}_t\left[{J_{t+1}^i}^{1-\sigma}\right]^{\frac{1}{1-\sigma}}}\right)^{\frac{1}{\varphi}-\sigma}$	(Stochastic discount factors)
$\frac{1}{R_t} = \mathbb{E}_t\left[\beta_{t,t+1}^U \frac{1}{\pi_{t+1}}\right]$	(Bond Euler equation)
$Q_t = \mathbb{E}_t\left[\beta_{t,t+1}^U (Q_{t+1} + D_{t+1})\right]$	(Firm share Euler equation)
$N_t = (1-\rho)N_{t-1} + p(\theta_t)U_t$	(Employment law of motion)
$U_t = 1 - (1-\rho)N_{t-1}$	(Definition of job searchers)
$u_t = 1 - N_t$	(Unemployment rate)
$\theta_t = \frac{V_t}{U_t}$	(Labour market tightness)
$q(\theta_t) = \bar{m}\theta_t^{-\xi}$	(Job filling rate)
$p(\theta_t) = \theta_t q(\theta_t)$	(Job finding rate)
$Y_t = A_t N_t$	(Aggregate production)
$\frac{\chi}{q(\theta_t} = \varepsilon A_t - w_t + (1-\rho)\mathbb{E}_t\left[\beta_{t,t+1}^U \frac{\chi}{q(\theta_{t+1})}\right]$	(Job creation condition)
$1 - \Psi(\pi_t - \pi)\pi_t + \mathbb{E}_t[\beta_{t,t+1}^U \Psi(\pi_{t+1} - \pi)\pi_{t+1}\frac{Y_{t+1}}{Y_t}] = \nu(1-\varepsilon_t)$	(Optimal pricing condition)
$w_t = (1-\eta)b + \eta\left[\varepsilon_t A_t + \mathbb{E}_t\beta_{t,t+1}^U \chi\theta_{t+1} + \lambda\left[\frac{\chi}{q(\theta_{t+1})}(1-\rho-p(\theta_{t+1}))(\beta_{t,t+1}^U - \beta_{t,t+1}^C)\right]\right]$	(Nash wage)
$C_t^U = w_t N_t + \frac{1}{1=\lambda}D_t$	(Unconstrained consumption)
$C_t^C = w_t N_t$	(Constrained consumption)
$C_t = Y_t\left(1 - \frac{\Psi}{2}(\pi_t - \pi)\right) - \chi V_t$	(Market clearing)
$D_t = Y_t\left(1 - \frac{\Psi}{2}(\pi_t - \pi)^2\right) - w_t N_t - \chi V_t$	(Firm profits)
$R_t = R_{ss}\pi_t^{\varphi_\pi}\zeta_t$	(Taylor rule)

Table 4.1. Model equilibrium conditions

4.B.1 List of Model Equilibrium Conditions

A summary of the equilibrium conditions used to solve and simulate the model is presented in Table 4.1.

4.B.2 Parameterization

A summary of the parameter values which result from our calibration strategy outlined in Section 4.3 is provided in Table 4.2.

4.C Additional Figures

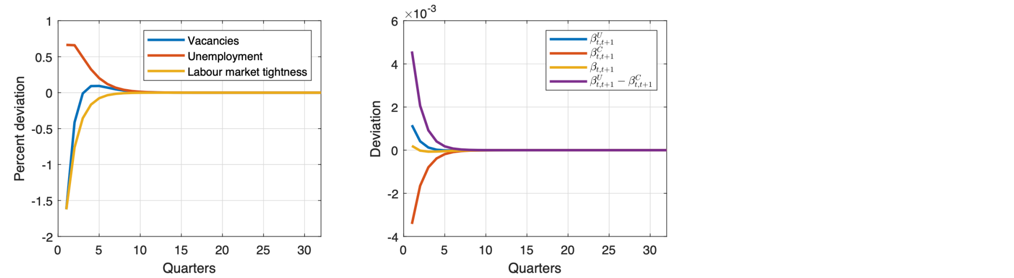

Figure 4.5. IRFs in response to a 1% contraction in monetary policy.

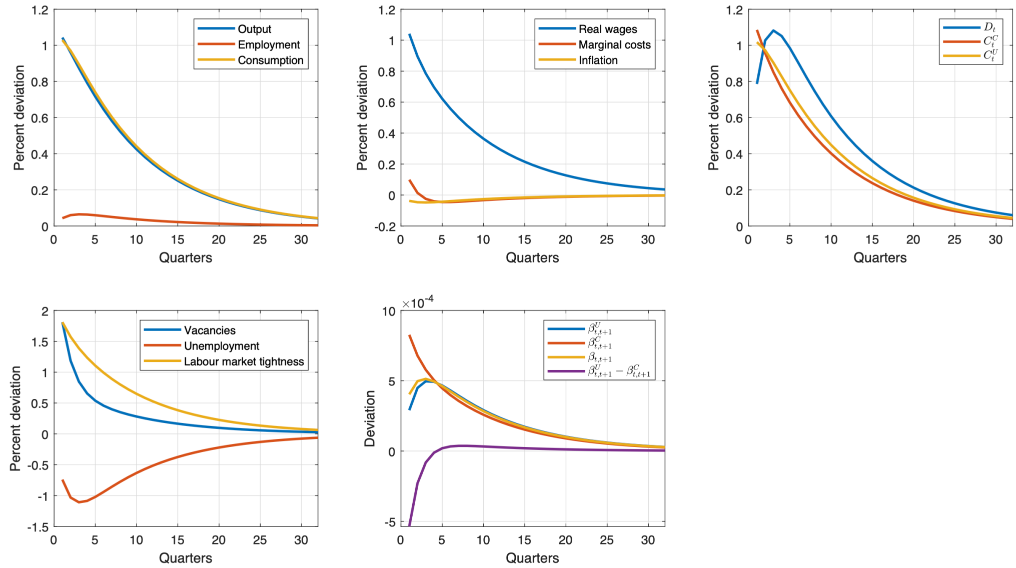

Figure 4.6. IRFs in response to a 1% increase in productivity.

Parameter	Description	Value	Source
Externally calibrated parameters:			
β	Discount factor	0.9926	3% avg.real interest rate
φ	IES	1.5	Standard
σ	Risk aversion	$1/\varphi$	CRRA utility
ν	Substitution elasticity	6	Standard
ξ	Match elasticity	0.5	Petrongolo and Pissarides (2001)
η	Workers' bargaining power	0.5	Hosios' condition
Ψ	Price adjustment cost	58.6969	avg. price duration of 1 year (Calvo)
ϕ_π	Taylor coefficient	2	Standard
A	Steady state productivity	1	Normalization
Internally calibrated parameters:			
ρ	Separation rate	0.0476	$u = 0.055$
χ	Vacancy posting costs	0.6553	$q(\theta) = 0.7$
$\bar{m}$	Match efficiency constant	0.5612	$p(\theta) = 0.45$
b	Unemployment benefits	0.3129	40% replacement rate
λ	Share of constrained households	0.21	Debortoli and Galí (2018)
Exogenous processes:			
ρ_a	Persistence of tech. shock	0.9	Standard
ρ_ξ	Persistence of mon. pol. shock	0.5	Standard

Table 4.2. Parameter Values