Fiscal Policy and Wealth Inequality Over the Business Cycle

A Heterogeneous Agent Real Business Cycle Model

Finn Robinson

Abstract

This paper extends the Krusell and Smith (1998) heterogeneous agent real business cycle model to examine how the wealth distribution responds to government consumption shocks over business cycle frequencies. The main contribution of the paper is that it is, to my knowledge, the first to model how the entire wealth distribution responds to policy shocks at this frequency. Until now, the literature has tended to use heterogeneity to relax the assumption of representative agents, and not to model the dynamics of wealth or income inequality in detail. The algorithm developed by Winberry (2016a) is used to solve the model. This algorithm is faster than other methods for solving heterogeneous agent models and approximates the entire wealth distribution. This enables me to construct three-dimensional plots of the evolution of the asset distribution in response to shocks, as well as obtain impulse response functions for the higher moments of the distribution. Other methods such as Krusell and Smith (1998) approximate the distribution using only the mean, which precludes them from offering such detailed descriptions of the dynamics of wealth and income inequality. The model finds that government consumption shocks result in increases in wealth inequality above steady-state, and the response of wealth inequality is twice as persistent as the government consumption shock. The stylised nature of the model might call into question the robustness of this key finding. I therefore analyse the same question using a mixed-frequency vector-autoregression and find that in the data, government consumption shocks also result in persistent increases in wealth inequality above steady-state.

Acknowledgments

I would first like to thank my supervisor Dr Dennis Wesselbaum of the Department of Economics at the University of Otago. He was always available to help when I ran into trouble or had silly questions to ask, but he also let me do my own work and figure out problems on my own.

I would also like to thank Dr Murat Ungor of the Department of Economics at the University of Otago for his helpful comments and reading suggestions. They certainly improved the quality of my thesis.

Contents

List of Tables

List of Figures

1. Introduction

The topics of wealth and income inequality have come to the forefront of academic and political debate in recent decades. Reports such as Collins and Hoxie (2017)'s 'Billionaire Bonanza: The Forbes 400 and the Rest of Us' highlight the extreme levels of income and wealth inequality that the world is currently experiencing. They found that in November of 2017 the wealthiest three US citizens, Bill Gates, Jeff Bezos and Warren Buffet, owned more wealth than the poorest half combined (or 160 million people).

Whilst it is clear that wealth and income inequality are becoming worse, the exact mechanisms driving this process are still debated in the literature. Thomas Piketty's book *Capital in the Twenty-First Century* (2014) provides a deep dive into the dynamics of capital accumulation, growth and inequality based on a richer data set than has previously been available. His dire warning for the world, which sparked immense media interest, is that when the rate of return to capital (r) is higher than the growth of output (g), capitalism 'automatically generates arbitrary and unsustainable inequalities' (p. 1). He writes that it seems likely the world will soon return to the state where $r > g$. At this point previously accumulated wealth grows faster than output and wages so that society will reach a state where a few wealthy 'rentiers' own everything and are 'dominant' over workers (p. 571). Piketty suggests that one of the key drivers of this trend is the rise of the 'supermanager,' those earning 'stratospheric' pay that can be fifty – 100 times the average income (p. 378, 417). Income and wealth are becoming more strongly related than was historically the case, which means that these 'supermanagers' will become the rentier class that accumulates everything.

Song et al. (2015) find evidence that contradicts this theory. Using US data they find that within-firm inequality has not actually increased between 1982 and 2012. Rather, the top firms are paying more than other firms. They describe this as 'stable inequality within firms' (p. 30). This is a direct contrast to Piketty who argues that it is inequality within firms that is causing rising inequality. Inequality is getting worse and is something that needs to be understood properly. Unfortunately, we do not yet understand it properly. This lack of understanding stems from a paucity of quality data on inequality, with wealth inequality data being particularly rare (Saez and Zucman, 2016). Another factor is that it is computationally difficult to incorporate inequality into theoretical models because the income and wealth distributions are infinite dimensional objects, making computation difficult. Despite these hurdles, the fact that two pieces of research can come to diametrically opposed conclusions about the dynamics behind inequality highlights the need for further exploration.

The novel contribution, and main focus, of this paper is to model the response of the wealth distribution to fiscal spending shocks over business cycle frequencies. As will be further developed in the literature review, the literature thus far has focussed on using heterogeneous agent models to relax the assumption of homogeneity. No paper that I am aware of specifically aims to model the response of the wealth distribution to government policy shocks over business cycle frequencies. The closest paper to mine is Ma (2017) which examines the response of private consumption to government consumption shocks. Households are divided into income quintiles. Although he reports how the means of the wealth distributions of each quintile change, he is limited by his choice of algorithm to only look at the mean of the distribution. My paper will be able to characterise the response of higher moments of the wealth distribution, enabling me to give a more detailed explanation for the response of wealth inequality to government spending shocks. In addition, my paper will use a mixed-frequency vector autoregression (MFVAR) to empirically characterise the response of wealth inequality to government spending shocks and test the robustness of the theoretical model. Using a VAR to examine how wealth inequality responds to government spending shocks is again something that I have not seen in the literature, and is another new contribution of this paper.

The new contributions of this paper will help expand our understanding of how wealth inequality responds to fiscal policy over business cycle frequencies. If policy makers are concerned about wealth inequality they must understand how their policies might impact the wealth distribution. In their seminal paper on introducing heterogeneity into dynamic stochastic general equilibrium (DSGE) models, Krusell and Smith (1998) note that it is possible to now use ‘equilibrium models to analyse the interrelation between business cycles, inequality and economic policy’ (p. 890). My paper therefore does just that, using an algorithm by Winberry (2016a) to extend the stylised Krusell and Smith real business cycle (RBC) model and incorporate government consumption. This makes it possible to plot impulse response functions (IRFs) for the moments of the wealth distribution and therefore gain an understanding of how government consumption shocks impact wealth inequality in the short run. To the best of my knowledge it is the first paper to use a heterogeneous agent RBC model to do this.

This paper introduces government consumption to the Krusell and Smith model by separating household consumption in the utility function into private and public consumption spending. Public consumption is a complement to private consumption in that increases in public consumption raise the marginal utility of private consumption. Heterogeneity is introduced by having two employment states, employed and unemployed, and a probability of households transitioning between states each period. This means that each household has its own employment history, and therefore its own level of accumulated assets. Government consumption spending enters the savings function of households due to its placement in the utility function of households. A shock to government consumption spending triggers a change in the savings behaviour of households, which then alters the asset (wealth) distribution.

As mentioned above, the computation of heterogeneous agent models is challenging. In the equilibrium of the Krusell and Smith (1998) model, agents' decisions are a function of the asset distribution which is an 'infinite dimensional object' because agents are modelled along a continuum (p. 869). Since their paper was published, there have been many attempts to write algorithms to solve heterogeneous agent models (Terry, 2017). This paper uses the Winberry algorithm because although it might not have the best prediction accuracy,[1] it is the fastest to complete simulations, and relatively easy to extend (Terry, 2017).

The main simulation results show how key macroeconomic variables (output, private consumption and investment, wages and the interest rate) and wealth inequality respond to a positive, temporary, one percent government consumption shock. Wealth in this model is the asset holdings that households accumulate over time. IRFs for the mean, variance and skewness of the asset distribution are reported, as well as a three-dimensional plot showing the evolution of the asset distribution in the first forty quarters after the shock.

In steady-state the mean of the employed asset distribution is higher than the unemployed distribution. The IRF results show that wealth inequality increases amongst both employed and unemployed households after the shock. The variances of the employed and unemployed asset distributions increase immediately after the government consumption shock. The mean of the employed asset distribution reaches a peak of 0.1018 percent above steady-state in the seventh quarter after the shock, with the variance peaking at 0.0078 percent above steady-state in the fourth quarter. The mean of the unemployed asset distribution reaches a slightly lower peak of

[1] Terry (2017) does note that all of the methods he surveyed gave approximately similar results, therefore even though the Krusell and Smith (1998) method is more accurate, the Winberry (2016a) algorithm is also accurate.

0.0935 percent above steady-state, whilst the variance has a much larger response compared to the employed variance. It reaches a peak of 0.0284 percent above steady-state, a response that is 3.6 times larger than the employed variance. The mechanism driving this is discussed in detail in section 5.3.1.

In the baseline simulation, the paper also considers the persistence of the impulse responses of the asset distribution. The simulated autocorrelation coefficients for the impulse responses of government consumption and the moments of the asset distribution show that the response of the asset distribution is twice as persistent as the government spending shock.

In the robustness section, different levels of complementarity between private and government consumption are explored. The simulated IRFs illustrate that the higher the complementarity between private and public consumption, the larger the response of the aggregate variables.[2] The response of wealth inequality to increased complementarity depends on whether households are employed or unemployed. For unemployed households, higher complementarity unambiguously increases wealth inequality. For employed households, the response is non-monotonic. The baseline level of complementarity results in lower wealth inequality amongst the employed households compared to both higher and lower levels of complementarity. As a second extension, changing the duration of unemployment spells is analysed. Doing this demonstrates how important it is to study the entire asset distribution when looking at the response of the economy to policy shocks. For the aggregate variables and the mean of the asset distribution, having unemployment spells last for one quarter or half a quarter makes very little difference. However, the increase in wealth inequality after the government consumption shock is much larger when unemployment spells are longer.

The Krusell and Smith model is very stylised, even with government consumption added, therefore I carry out a vector auto-regression (VAR) analysis using US data to check whether the main findings regarding wealth inequality and government consumption spending are present in the data. The data on wealth inequality from Saez and Zucman (2016) are available in annual frequency, whilst the data on the other macroeconomic variables are all at quarterly frequency. I therefore use a mixed-frequency VAR which allows for data of different frequencies to be used in the same regression. The varied data sampling (VARDAS) model developed by Qian (2010) is used to carry out the analysis.

[2] Output, consumption, investment, rental rate, wage, mean assets for employed and unemployed.

The results of the empirical analysis are consistent with the predictions of the theoretical model. In response to a government spending shock, wealth inequality experiences a persistent increase above steady-state. Investment and consumption are notable for having estimated impulse response functions that move in the opposite direction to the simulated responses. This results from the theoretical model being a stylised real business cycle model with inelastic labour supply and a simple method for combining public and private consumption. The important relationship between government spending shocks and wealth inequality is observed in the theoretical and empirical models.

The paper will proceed as follows: first the literature review will discuss how heterogeneity has been modelled, emphasising that heterogeneity and incomplete markets have been largely used to relax assumptions, with the dynamics of the wealth and income distributions receiving less attention. The review will then examine evidence on all the facets of my research topic, noting that no single paper addresses all facets simultaneously. Secondly, the extended Krusell and Smith model will be described and solved, closely following Winberry (2016b). Thirdly I will discuss the calibration of the model, which follows Krusell and Smith (1998) and Winberry (2016b) except for where government consumption enters the model. The next two sections will discuss the steady-state and dynamic results from simulating the model, with particular emphasis placed on the response of the wealth distribution to government consumption shocks. The question of how government consumption shocks impact the wealth distribution will then also be addressed using vector autoregression analysis to test the robustness of the theoretical model. The paper will then conclude.

2. Literature Review

The importance of modelling heterogeneity in macroeconomics has become evident in the last several decades. Over the past 30 years the top one percent of wealthy households in the US have held a third of total wealth in the economy, whilst over ten percent of households have little to no assets (de Nardi, 2015). Previously, theoretical work has used dynamic stochastic general equilibrium (DSGE) models with representative agents and complete markets. The reason for this was twofold: first, macroeconomists did not think that modelling heterogeneity was important for understanding the dynamics of the business cycle. Secondly, they did not have the computing power or technology to model heterogeneity (Heathcote, Storesletten and Violante, 2009). Recent research has shown that incorporating heterogeneity has changed our answers to many questions. This literature review will discuss the literature that is most relevant to modelling wealth inequality and government spending in a DSGE model.

2.1. Heterogeneity, Incomplete Markets and the Business Cycle

It is important to understand the distributional effects of business cycles, as well as the average effect (Mayer, Lopoo and Groves, 2016). Robert Lucas' 1985 book *Models of Business Cycles* is an example of how the average effect of policy hides heterogeneous effects for different types of people. Lucas finds that eliminating business cycles results in a tiny welfare gain for the representative consumer (Krusell et al., 2009; Heathcote, Storesletten and Violante, 2009). This was an important result because economists believe that business cycles reduce welfare and that governments should try and reduce or eliminate them using fiscal and/or monetary policy (Imrohoroglu, 2008). Two key assumptions that underpin Lucas' results are that agents are homogeneous, and that they can access complete capital markets. Imrohoroglu argues that it is feasible that for some consumers who have large savings, the costs of cycles would be minimal, whilst for consumers who are credit constrained and unable to insure against shocks, the impacts could be 'devastating' (p. 3).

One of the first papers to examine the effects of relaxing the assumptions imposed by Lucas was Imrohoroglu (1989). She compares two economies, one with a liquidity constraint and one without to determine whether it is appropriate to 'abstract from liquidity constraints' in models of business cycles (p. 1366). When consumers are restricted so that they can store wealth but not borrow, Imrohoroglu finds that the cost of business cycles for consumers is five times larger than Lucas originally calculated. Her results reject the hypothesis that one can make the simplifying assumption of perfect insurance. Note that this paper is still based on the assumption of homogeneous agents, therefore whilst Imrohoroglu shows that the impact of eliminating business cycles is large for the average consumer, nothing is said about how this impact changes along the income or wealth distributions.

More recent papers incorporate incomplete markets and heterogeneous consumers into their models and find that the welfare effects of eliminating business cycles can be quite significant, with this impact changing depending on the place of the consumer on the wealth or income distributions. Krusell et al. (2009) re-examine the welfare implications of removing business cycles, and they find that they are much larger than originally calculated by Lucas, with the magnitude of welfare gains depending on what type of consumer one is. For very poor consumers who are categorised as long-term unemployed, they find that removing the business cycle increases their consumption by thirty percent, whereas middle income (average) households experience only a very small increase in welfare. They describe the welfare effect as U-shaped. The poor gain from eliminating the business cycle because they have lower risk of losing their jobs, whilst the rich gain because precautionary savings in the economy falls which raises the interest rate and therefore their income from wealth (Krusell et al., 2009).

A different approach to resolving the question of the welfare costs of business cycles is implemented by Schulhofer-Wohl (2008). He distinguishes between two types of heterogeneity: heterogeneity in initial conditions, and in 'experiences' (p. 762). Heterogeneity in experiences is what we see in Krusell et al. (2009) where agents are, ex-ante, homogeneous, but experience uninsurable idiosyncratic shocks which lead to different consumption and income streams for households. If it were possible to fully insure against these shocks, then agents would remain homogeneous. Heterogeneity in initial conditions means that agents are ex-ante heterogeneous. Schulhofer-Wohl (2008) gives the examples of 'different preferences, endowments or technology' (p. 762). In his model, households have different levels of risk aversion. This heterogeneity opens up gains from trade because the least risk-averse households can sell insurance against risk to those households that have high levels of risk aversion.

The result from Schulhofer-Wohl's paper is that the least risk-averse agents gain from business cycles, whilst even infinitely risk averse households have small costs. Essentially, he shows that even with heterogeneity, if insurance markets are complete then ‘aggregate fluctuations in consumption are essentially irrelevant’ for all agents, not just the average household (p. 761). He replicates Lucas’ finding that the welfare costs of business cycles are tiny, and does so whilst relaxing the homogeneous agents assumption. It should be noted that his result requires that the complete markets assumption holds, which is a very strong assumption. As Imrohoroglu (1989) writes, it is almost ‘universally agreed’ that insurance markets are not complete, and agents cannot ‘perfectly’ insure against individual risks (p. 136).

Turning to empirical papers, de Giorgi and Gambetti (2015) use US consumer expenditure data to estimate the interaction between business cycles and the distribution of consumption. Using shocks to total factor productivity and economic policy uncertainty they show that the consumption responses on the upper end of the consumption distribution are much larger, whereas the response of people at the lower end of the consumption distribution is much less volatile. Barlevy and Tsiddon (2004) directly examine income inequality in their paper which looks at the connection between long run and cyclical variance in earnings inequality in the US. Their key finding is that recessions tend to ‘amplify long run trends’ (pg. 1). This means that if inequality is on an upward (downward) trend, then a recession will serve to increase (decrease) it.

Another paper by Hoover, Giedeman and Dibooglu (2009) investigates how economic growth, unemployment and immigration impact the different quintiles of the income distribution using US data. Their ‘primary objective’ is to see how income inequality interacts with these key business cycle variables (p. 280). They find using impulse response function analysis that in response to a positive shock to unemployment, income inequality increases, then returns to normal after four years. In response to a negative shock to unemployment, income inequality only falls for one year before returning to normal levels.

Morin (2013) looks at wage dispersion over the business cycle. Morin empirically and theoretically describes this relationship. The motivation for the paper is that although many believe that income inequality is countercyclical (that is, booms reduce inequality and busts increase it) the balance of empirical evidence suggests that it is in fact procyclical. Using US data, Morin shows that wage dispersion is positively correlated with GDP, negatively correlated with unemployment, and therefore procyclical. To rationalise this empirical finding, Morin constructs a DSGE model incorporating a labour market characterised by search and matching frictions in which monopsonistic firms set wages and compete for workers. This competition gives the (homogeneous) workers some market power. Wage dispersion is generated by the following process: in booms, vacancies are higher and workers receive more frequent and better job offers. The increased competition for workers reduces firms' market power in the labour market, so they post a higher proportion of higher paying vacancies. This then results in a rightwards shift in the upper tail of the wage distribution, increasing wage dispersion.

Morin's paper is similar to mine in that it looks at how the distribution of income changes in response to a shock over the business cycle. The key differences are firstly that in my model firms are competitive and households are ex-post heterogeneous, whilst in Morin's model firms are monopsonistic and households are homogeneous. Secondly, I am looking at how wealth inequality responds to government consumption shocks, whilst Morin explores how the income distribution is related to unemployment and output.

The papers in this sub-section demonstrate that to adequately understand the impacts of business cycles on the economy it is advisable to relax the assumptions of complete markets and homogeneous agents. As such, the model I use incorporates incomplete markets in the form of a borrowing constraint and ex-post heterogeneity to examine the impact of government spending shocks on macroeconomic variables and the asset (wealth) distribution.

2.2. Heterogeneity in DSGE Models

Having discussed the importance of modelling heterogeneity when looking at business cycles, this literature review will now turn to exactly how macroeconomists have been introducing heterogeneity into their models, and how these adjusted models have performed.

2.2.1. Winberry's Methods

I will begin with Winberry's works because the solution method of my model uses his algorithm. In two papers, Winberry uses this algorithm to solve two different models, one in which firms are heterogeneous, and one in which consumers are heterogeneous. Winberry (2016a) solves a stylised real business cycle model where the firms are heterogeneous. Firms have Cobb-Douglas production functions with an aggregate and idiosyncratic productivity shock in the function.[3] As a result of the idiosyncratic shock, firms will be distributed over both the idiosyncratic shock and the different levels of capital that result from their individual productivity path over time. Winberry (2016b) solves a model based on Krusell and Smith (1998) where households are either unemployed or employed, which is a two-state Markov process. This results in an endogenous distribution of assets across households. Given that my paper uses the Winberry algorithm to solve the Krusell and Smith model, this is the method of introducing heterogeneity that will be used here.

2.2.2. Other Approaches

de Nardi (2015) and Heathcote, Storesletten and Violante (2009) offer excellent surveys of the development of models with heterogeneity. They explain that the basic modelling foundation of this literature is the Bewley model. These are incomplete market models where households are ex-ante identical, but they all face the same stochastic process, for example shocks to income streams, which results in ex-post heterogeneity. This is the approach used by Heathcote (2005). In his model households receive idiosyncratic shocks to their labour efficiency, which cannot be insured against. This results in a distribution of asset holdings across households with different productivity histories. Kaplan, Moll and Violante (2016) have households face an idiosyncratic earnings process in their heterogeneous agent New Keynesian model.

de Nardi notes that the issue with the Bewley-type models is that they come nowhere near close to generating the income distribution that is observed in the data. The gini coefficient in the basic Bewley model is about half what is observed in the data (de Nardi, 2015). Gornemann, Kuester and Nakajima (2012)'s paper demonstrates that a lot of work is needed to make the model match the data. They construct a heterogeneous agent New Keynesian model with an employment status determined by a search and matching process. There is a probability that the household has a certain productivity level if employed, and this process generates an endogenous distribution of asset holdings. To generate the extreme inequality that is

[3] $y_{jt} = e^{z_t} e^{\epsilon_{jt}} k_{jt}^{\theta} n_{jt}^{v}$, $\theta + v < 1$. Where ϵ_{jt} is the idiosyncratic productivity shock faced by firm j and z_t is the aggregate productivity shock faced by all firms.

empirically observed, the highest productivity level is twenty-four times the second highest level. A Markov transition matrix then determines the probabilities of households moving to different productivity levels.

de Nardi discusses some other methods of introducing heterogeneity that have attempted to generate income and wealth distributions consistent with the data. One idea supported by microeconometric evidence is modelling heterogeneous preferences. This was explored by Krusell and Smith (1998) where they had an infinite horizon Bewley model that was augmented with stochastic risk aversion and discount factors for each dynasty (de Nardi, 2015). For overlapping generations models, intergenerational transmission of wealth is also considered, with parents investing in their children's human capital formation and giving monetary transfers. Neither of these methods generate the extreme wealth concentration that is empirically observed (de Nardi, 2015).

Modelling entrepreneurs has the potential to generate extreme wealth concentration. Entrepreneurs make up a small part of the population but have a large share of wealth (de Nardi, 2015). They are modelled as being borrowing constrained so that even though they are wealthy, they want to grow the size of their firms so that they can borrow more and earn higher returns from their capital. A final method that de Nardi discusses is introducing large earnings risk for high earners. As in Gornemann, Kuester and Nakajima (2012), the highest level of productivity is far higher than the second highest, and there is a probability of transitioning from the highest level to second highest level in the next period. High earning households thus have very high earnings risk, and accumulate massive savings to insure against this risk and smooth consumption.

The initial conditions of an economy might also be important in the modelling of heterogeneity (Heathcote, Storesletten and Violante, 2009). Heterogeneity in Schulhofer-Wohl (2008) for example is entirely driven by differences in risk tolerance by households. Keane and Wolpin (1997) state that around ninety percent of lifetime earnings dispersion between people can be explain by the initial conditions before they entered the labour market. Heathcote, Storesletten and Violante (2009) argue that it is important to have heterogeneity in initial market conditions, for example by augmenting an earnings process with an individual fixed effect.

Heathcote, Storesletten and Giovanni (2009) note that the current 'workhorse' for modelling heterogeneity is the standard incomplete markets model in which there are a large number of agents who have an idiosyncratic productivity level (pp. 2-3). This is what is used in Gornemann, Kuester and Nakajima (2012). Le Grand and Ragot (2017), Ma (2017) and Ruiz, Peralta-Alva and Puy (2017) are three recent working papers that also use idiosyncratic shocks to labour productivity as part of their methods for introducing heterogeneity into heterogeneous agent models. These papers are discussed in detail next because they give more attention to the response of the wealth and income distributions than the above papers that use heterogeneity to relax assumptions, rather than as a focus of the analysis.

2.2.3. Recent Heterogeneous Agent DSGE Model Literature

Many papers use heterogeneity to relax assumptions, rather than as a focus of the analysis. Ma (2017) gives more attention than most to the response of the wealth and income distributions. He first uses a VAR to look at the impact of government spending shocks on the consumption behaviour of the quintiles of the income distribution. He finds that the top quintile reduce consumption in response to an increase government spending, which is the real business cycle model prediction. The bottom quintile however, increase their consumption, which is more in line with IS-LM model predictions.

To provide an explanation for the heterogeneous response of the different quintiles, Ma constructs a stylised DSGE model with incomplete capital markets and households which face idiosyncratic shocks to their labour productivity. He reports impulse responses for the response of income quintiles to the government spending shock and finds that the results match the VAR. Overall, the increase in government spending reduces consumption inequality because it boosts the employment and therefore consumption of the poor, whilst the rich face a net loss of income due to the higher tax burden of funding government spending, which reduces their consumption. The after-tax wages for the quintiles are also reported and show that the bottom eighty percent of households experience higher incomes, whilst the top twenty percent experience lower incomes.

Ma (2017) is probably the closest paper to mine that I came across: he looks at the impact of a government spending shock on inequality at a business cycle frequency. Some key differences stand out however. Ma uses the method of Krusell and Smith (1998) to compute impulse response functions for the government spending shock. This approximates the distribution of household assets using the mean of the distribution. By comparison, I use the

Winberry (2016b) algorithm which approximates the entire distribution and allows me to calculate impulse response functions for more moments of the distribution (I calculate them for the mean, variance and skewness). In his model, government spending is interpreted as public capital and thus enters the production function. In my paper government consumption enters the utility function and is complementary to private consumption. Ma is also interested in income inequality and consumption inequality whereas I focus on wealth inequality.

Le Grand and Ragot (2017) is a paper also examining fiscal policy in a heterogeneous agent DGSE model, however they write their model as a constrained social planner problem. Their motivation for this is that whilst heterogeneous agent DSGE models are becoming more common in theoretical work, the computational difficulty they involve means not much is known about "optimal policies in this environment" (p. 2). Their paper focusses on various fiscal policies and the 'tradeoffs about redistribution, insurance and incentives (p.2). The paper uses another method for solving heterogeneous agent models called reduced heterogeneity. In the equilibrium of these models there are a finite number of levels of wealth, rather than a continuous wealth distribution. This makes the model much easier to solve and simulate as the equilibrium is a finite dimensional object.

The most thorough incorporation of heterogeneity into a DSGE model that I know of is by Ruiz, Peralta-Alva and Puy (2017). The authors are interested in trying to pin down the size of tax multipliers and their distributional impacts in the same model. Their main contribution to the literature is that they extend heterogeneity in ways that are otherwise 'absent from standard models' (p. 4). They have four methods for incorporating heterogeneity into their model. Firstly, they allow for the economy to produce three types of good, each using different amounts of labour. Second, each of these goods needs some amount of intermediate goods to be produced. Third, each goods sector employs workers with different skill levels. Finally, consumption spending by consumers is positively related to their position on the wealth distribution: the wealthier they are the larger is their consumption spending as a proportion of total spending.

2.3. Evaluating Heterogeneous Agent Models

There is a small literature that seeks to evaluate the accuracy of empirical and theoretical heterogeneous agent models. On the empirical side is Misra and Surico (2014)'s paper which compares the predictions of a homogeneous and heterogeneous agent model. Using 2001 and 2008 tax cuts in the US they compare the performance of the two models in predicting the response of consumption to the income change. The heterogeneous agent model allowed for unobserved heterogeneity in households that otherwise had the same characteristics. They find that the heterogeneous agent model was most consistent with the data because the homogeneous agent model was wrong in the tails of the consumption distribution. In both their model specifications, they reject the null of a homogeneous consumption response to the tax cuts, which further demonstrates that the heterogeneous agent model is more appropriate.

On the theoretical side, Cozzi (2014) uses Monte Carlo evidence to evaluate the performance of overlapping generation and infinitely lived heterogeneous agent models. For the overlapping generations model, Cozzi finds that when the welfare effect of a policy is greater than 0.2 percent, the heterogeneous agent model always gets the correct sign of the welfare effect, regardless of the sample size. As the sample size increases the range of the estimate falls, so the maximum error falls as the amount of information available increases. For the infinitely lived agent model which is what I am interested in for the purposes of this paper, Cozzi finds that the model always predicts the correct sign for welfare effects regardless of the sample size used, and the range of the estimates falls as the sample size rises. Therefore the maximum error that the heterogeneous agent model makes falls as the sample size increases. Overall it therefore appears that heterogeneous agent models perform very well, although Cozzi does note that if the welfare effect is small, robustness tests should be used to ensure that the sign of the predicted effect is correct.

2.4. Evidence

Whilst I am not aware of any papers that specifically aim to use a heterogeneous agent DSGE model to examine the response of wealth inequality to government spending shocks over business cycle frequencies, there are many papers that look at the separate aspects of my research topic. These include papers on fiscal policy over the business cycle, the effect of fiscal policy on inequality, and the relationship between inequality, redistributive policies, and economic growth. This literature review will now turn to these.

2.4.1. Government Spending Shocks

The effect of government spending has received extensive empirical treatment in the literature, with several papers explicitly addressing distributional issues and heterogeneity. One such paper by Mayer, Lopoo and Groves (2016) examines the effect of social spending by US state governments on the growth of family incomes over the income distribution. Social spending is measured as expenditure by the state government on primary and secondary schools, public welfare, health and hospitals. Increased social spending by the government had a positive and significant impact on income per capita at all points on the income distribution, however at lower levels it had a larger effect, and at the top of the distribution it had a smaller effect. The states whose social spending was the highest had the most growth and the lowest levels of inequality. A second paper that examines social security spending increases is Romer and Romer (2016). They find that an increase in government transfers in the US increases consumption significantly in the short term, but has an insignificant effect after half a year, and does not significantly impact output or employment. There are large differences in the response of consumption to social spending compared to taxes. To explain this, Romer and Romer cite evidence from US Federal Reserve discussions that the decision makers there view transfer increases as reasons to tighten monetary policy.

In the VAR literature Anderson, Inoue and Rossi (2016) examine the effect of government spending shocks on different types of consumers using a three-variable VAR similar to Ramey (2011). The real business cycle model prediction is that if the government increases its spending, consumers anticipate a future tax to pay for this, and therefore reduce their current consumption. Instead, the authors find a heterogeneous response: the wealthier consumers do exhibit the behaviour predicted by real business cycle models, reducing their consumption, whilst the poorer consumers actually increase their consumption in response to an increase in government spending. Overall this effect therefore reduces consumption inequality. The net effect of government spending shocks in this model will depend on the relative size and wealth of the rich and poor.

Other papers looking at government spending do not allow for heterogeneity, however their results are consistent with the papers examined so far, and look at a broader range of macroeconomic variables. Blanchard and Perotti (2002) find that a positive government spending shock increases output and private consumption, although it does crowd out private investment. Ramey (2011), using her new measure of defence news to measure government spending shocks, finds that a positive government spending shock increases output and consumption of services, but crowds out residential and non-residential investment, as well as reducing durable and non-durable consumption. Overall the literature seems to suggest that social spending increases raise consumption for families at the bottom of the income distribution, whilst positive government spending shocks increase output and reduce investment, and have an uncertain effect on consumption.

The papers covered in this sub-section do illustrate the impact of fiscal policy shocks on the economy at the same frequency as my paper, however they do not address how inequality responds. Anderson, Inoue and Rossi (2016) do allow for each quintile of the income distribution to have a separate response, however they do not show how the income distribution itself changes.

2.4.2. Inequality, Growth and Fiscal Policies

There is growing interest in the literature about the relationship between growth and inequality. By looking at the entire income distribution it is possible to determine the extent to which economic growth is benefitting all members of society, or just ‘certain subsets’ (Hoover, Giedeman and Dibooglu, 2009). In addition, authors have been looking into the effect of fiscal policies designed to reduce inequality, and whether these are conducive to economic growth. A paper by Biswas, Chakraborty and Hai (2017) examines how reducing income inequality using taxes would affect economic growth. Using US data they find that the effect of redistribution on economic growth depends on which sections of the income distribution are being brought closer by redistributive policies. Reducing income inequality between the low and middle income households is conducive to growth, but reducing inequality between middle and high income households reduces growth.

One issue with Biswas, Chakraborty and Hai (2017)’s paper is that they do not examine different types of taxes: they only look at income taxes. Garcia-Penalosa and Turnovsky (2007) use an AK model with heterogeneous households and a tax-funded investment subsidy to look at the relationship between growth, inequality and fiscal policy. They note that in AK models, the equilibrium results in a sub-optimal level of growth, so the government can improve overall welfare with an investment subsidy. However, to do this, they need to levy a tax. Since taxes have a negative impact on output, and have distributional consequences, this is a good framework to evaluate the relationship between growth, inequality and fiscal policy. The three taxes are a tax on capital, a tax on labour income, and a tax on consumption. They find that the capital tax is the best for reducing inequality in welfare between agents, whilst the consumption tax is best for encouraging growth. Both of these taxes improve growth and reduce inequality. The income tax is not a good option as it increases inequality and is not the best tax for promoting growth. These results suggest that growth and redistribution are not mutually exclusive depending on how they are both financed.

These papers have focussed on how one can deal with inequality without reducing growth. Other papers consider the growth consequences of worsening inequality. One such paper is by Alesina and Perotti (1996). They argue that inequality reduces growth because it has a negative effect on political stability, which reduces investment and therefore output. Using a sample of seventy countries from 1960-1985 they find that a one standard deviation increase in the income share of the middle class results in a one quarter standard deviation reduction in political instability, which then results in a one percentage point increase in the investment share of GDP. Increasing taxes on capitalists and investors might, *ceteris paribus*, reduce their incentives to invest, but by reducing inequality (and therefore political uncertainty) the government can create an environment that is more conducive to capital accumulation and investment, resulting in higher growth than before the redistribution.

The papers that I have discussed in this literature review illustrate that there is a gap in the literature regarding fiscal policy and wealth inequality. None of the empirical papers look at wealth inequality: they all look at income inequality and fiscal policy. This is likely due to the lack of data. The theoretical models do address wealth inequality in that an endogenous asset distribution is a consequence of having heterogeneous households in many of the models. That said, they do not look at how the distribution of wealth changes in response to the various shocks they introduce into their models. This paper will therefore try to fill in the gap by examining how government spending shocks impact the wealth distribution over business cycle frequencies.

2.5. Monetary Policy

Heterogeneous agent models have had important implications for monetary policy as well as fiscal policy, and the literature on the former is somewhat richer. Kaplan, Moll and Violante (2016) evaluate the consumption response to monetary policy of households using both a representative and heterogeneous agent DSGE model. This is motivated by their observation that in representative agent New Keynesian models, the consumption response of agents is driven by the Euler equation and the substitution of consumption between periods. The data, however, do not support this prediction: macroeconometric time series analysis has found that consumption is not actually very responsive to the interest rate (Kaplan, Moll and Violante, 2016). Yogo (2004) shows that the elasticity of intertemporal substitution is not statistically significantly different from zero. Canzeroni, Cumby and Diba (2007) find that the Euler equation rate is different to the money market rate. They argue that this 'comes as no surprise' because a 'sizeable literature' tries and fails to show that the data is consistent with models that rely on the link between the Euler equation and money market rate. [4]

Kaplan, Moll and Violante introduce heterogeneity to the New Keynesian model by having households face an idiosyncratic earnings process. They find that with the heterogeneous agent version of the New Keynesian model, the direct effect of an interest rate shock (that is, the Euler equation channel) is small, whilst the indirect effect makes up 90 percent of the first quarter consumption response. This indirect effect results from monetary policy changing labour demand, which then changes household income (Kaplan, Moll and Violante, 2016). Their model is more consistent with the data, with both finding that the Euler equation is not the most important element of the consumption response to monetary policy shocks.

Heterogeneous agent models not only raise questions about the transmission of monetary policy, but they also enrich our understanding of who is affected by policy changes, and by how much. Gornemann, Kuester and Nakajima (2012) construct a heterogeneous agent New Keynesian DSGE model in which households have different employment, skill and shareholding levels. They find that a contractionary monetary policy shock reduces labour income for all segments of society, but the income of the most asset rich houses rises because they receive an increase in the income they derive from those assets.

[4] Canzeroni et al. (2007) explain that these models equate the Euler equation interest rate and the money market rate that central bank policy targets. This then establishes a connection between the Hicksian demand of consumers, and the policy rate used by the central bank.

3. The Model

The model used in this paper is based heavily upon Winberry (2016b), which itself follows Krusell and Smith (1998). The model is at a quarterly frequency.

3.3. Households

In Krusell and Smith (1998) and Winberry (2016b) there is a continuum of j infinitely lived households indexed such that $j \in [0,1]$. Households maximise

$$E\sum_{t=0}^{\infty}\beta^t\frac{c_{jt}^{1-\sigma}-1}{1-\sigma}, \qquad (1)$$

where $\beta > 0$ is the subjective discount factor, c_{jt} is private consumption, $\frac{1}{\sigma}$ is the elasticity of intertemporal substitution.

3.3.1. Adding Government Spending

One way of introducing government spending into the economy is to assume that household utility depends both on private consumption and public (government) consumption (see for example Karras, 1994; McGrattan, 1994; Christiano and Eichenbaum, 1992; Feve, Matheron and Sahue, 2013; Finn, 1998; Ganelli and Tervala, 2009). The consumption component of the household utility function is then

$$c_{jt} = c_{jt}^{p} + \alpha^{g} G_t, \quad (2)$$

and the utility function becomes

$$U\left(c_{jt}^{p}, G_t\right) = \frac{\left(c_{jt}^{p} + \alpha^{g} G_t\right)^{1-\sigma} - 1}{1-\sigma}, \qquad (3)$$

where c_{jt}^{p} is private consumption and G_t is public consumption. Government consumption follows a simple AR(1) process

$$G_t = \rho^{g} G_{t-1} + \varepsilon_t^{g}, \quad where\ \varepsilon_t^{g} \sim i.i.d.N(0,1). \quad (4)$$

It is assumed that the government raises money for public consumption via a lump-sum tax. As Aschauer and Greenwood (1985) note, this then means the timing of the taxes becomes ‘irrelevant’ for how real variables react (p. 107).

The coefficient α^g on government consumption determines if and how government consumption impacts household utility. McGrattan (1994) notes that if the assumption $\alpha^g \neq 0$ is imposed, then government spending will impact household utility. Finn (1998) restricts α^g to $0 \leq \alpha^g \leq 1$. This formulation means that government spending is a substitute for private consumption, with the level of α^g determining how substitutable government consumption and private consumption are. The coefficient on government consumption does not have to be restricted as in Finn (1998). If it is negative, then this implies complementarity between private and public consumption (Karras, 1994). It is then feasible to observe a positive reaction of consumption to an innovation in government spending (Feve, Matheron and Sahue, 2013).

What exactly this government spending consists of matters for the calibration of the model: as Fiorito and Kollintzas (2004) note, how macroeconomic variables respond to government spending depends on what ‘categories’ of spending have changed (p. 1367). In Finn (1998), government activities are divided into government consumption, and government investment, with the government consumption component entering into the utility function as in this model. Bassetto (2005) models household preferences over consumption and government spending, where government spending is a public good that the government provides. An example he gives to motivate this is the government spending money on enforcing property rights which then facilitates private consumption. In this sense, government consumption is a complement to private consumption.

The question of complementarity between private and public consumption was comprehensively addressed in a paper by Fiorito and Kollintzas (2004). For twelve European countries they split government consumption between public goods (which includes defence, public order and justice) and merit goods (for example health and education) that can be provided privately. In contrast to the interpretation of Bassetto (2005), Fiorito and Kollintzas find that public good spending substitutes for private consumption, whilst merit good spending complements it. The positive impact of spending on merit goods always outweighs the negative (or insignificant) impact of spending on public goods in their data, which means that there is a positive elasticity of private consumption with respect to government spending. The authors suggest some reasons for the complementarity between merit good consumption and private consumption. First, government and private provision of merit goods could result in 'inefficiency' which then causes more private consumption than would otherwise be the case (p. 1394). Second, there could be positive externalities whereby public provision of schooling and health could induce more private consumption by households.

Government consumption in equation 2 is therefore to be interpreted along the lines of Fiorito and Kollintzas as spending on merit goods. An increase in government spending, for example on education, results in higher private consumption because households no longer have to pay as much to provide education themselves. Because government spending is financed through a lump sum tax, there are no negative feedback effects from having to raise distortionary taxes in order to pay for this extra consumption.

The way that government consumption works in this utility function is explained concisely by Karras (1994) and Ganelli and Tervala (2009). When α^g is negative, an increase in government consumption will raise the marginal utility of private consumption, thus encouraging more private consumption. When α^g is positive, higher government consumption reduces the marginal utility of consumption, leading to less private consumption. This can be seen by taking the derivative of the marginal utility of private consumption ($U_{c_{jt}^p}$) with respect to government spending

$$U_{c_{jt}^p}\left(c_{jt}^p, G_t\right) = (c_{jt}^p + \alpha^g G_t)^{-\sigma}, \quad (5)$$

$$U_{c_{jt}^p, G_t}\left(c_{jt}^p, G_t\right) = -\sigma(c_{jt}^p + \alpha^g G_t)^{-\sigma-1}\alpha^g. \quad (6)$$

Since σ is positive in the model calibration and government consumption is less than private consumption,[5] this means that when $\alpha^g < 0$ the derivative of $U_{c_{jt}^p}$ with respect to G_t is positive. Private and public consumption are therefore complementary. Conversely, when $\alpha^g > 0$ this derivative is negative, and private and public consumption are therefore substitutes. If $\alpha^g = 0$ then government consumption drops out of the utility function, and the model delivers the standard real business cycle result of government consumption completely crowding out private consumption.

Karras (1994) notes that simply introducing government spending into the economy by decomposing consumption as in equation 2, means that the marginal utility of government

---

[5] In the Federal Reserve Bank of St. Louis and Bureau of Economic Analysis data used in section 7, government consumption is always much lower than private consumption.

consumption is negative. He outlines a change to the utility function that can be used to ensure that the marginal utility of government spending is positive. One can augment the utility function of the household with a function of government consumption that ensures the marginal utility of government consumption becomes positive.

Currently in this model households maximise

$$U\left(c_{jt}^{p}, G_t\right) = E\sum_{t=0}^{\infty}\beta^t\, u(c_{jt}^{p} + \alpha^g G_t). \quad (7)$$

With the augmentation this would change to

$$U\left(c_{jt}^{p}, G_t\right) = E\sum_{t=0}^{\infty}\beta^t\left\{u(c_{jt}^{p} + \alpha^g G_t) + \phi(G_t)\right\}, \quad (8)$$

where $\partial\phi \;/\; \partial G_t \; > 0$. This means that government consumption now has a positive marginal utility. Since individuals are assumed to have no control over government spending, the household problem can be solved whilst 'ignoring' the effect of government consumption on utility (Karras, 1994, p. 10). I can ignore the issue of negative marginal utility of government spending because the household problem, and not the utility function, enters into the equation block in Dynare.

3.3.2. Household Labour Supply and Assets

Households inelastically supply ε_{jt} efficiency units of labour, where ε_{jt} is independently distributed across j households. For each household, ε_{jt} follows a two-state Markov process $\varepsilon_{jt} \in \{\varepsilon_0 = 0, \varepsilon_1 = 1\}$ with a probability of transitioning from one state to another of $\pi(\varepsilon', \varepsilon)$. If $\varepsilon_{jt} = 0$ then the household receives an unemployment benefit (b) which is financed by a tax (τ) on labour income, therefore unemployment benefits are bw_t. If $\varepsilon_{jt} = 1$ then the household receives wage w_t, resulting in an after-tax labour income of $(1-\tau)w_t$. The government operates with a balanced budget each period. Because the probability of transitioning from $\varepsilon_0 \; to \; \varepsilon_1$ is constant over time, the mass of households that is employed (L) will always be the same.

Asset markets are incomplete in that the households' next period capital must always be non-negative, that is, $a_{jt+1} \geq \underline{a}$ where $\underline{a} = 0$. The real return on capital is denoted r_t. Households have different employment histories based on their ε_{jt} process, which means that they each accumulate a different level of capital depending on their realisation of the Markov process. This process therefore generates an endogenous distribution of assets across households. As denoted in Winberry (2016b) households are distributed over productivity (ε) and asset (a) levels. This distribution is represented by $\mu_t(\varepsilon_{jt}, a_{jt})$.

The household optimisation problem is then

$$\max_{\{c_{jt}, a_{j+1}\}} E \sum_{t=0}^{\infty} \beta^t \frac{(c_{jt}^p + \alpha^g G_t)^{1-\sigma} - 1}{1-\sigma}$$

$$s.t.\ c_{jt} + a_{jt+1} \leq w_t(1-\tau)\varepsilon_{jt} + b\big(1 - \varepsilon_{jt}\big) + (1 + r_t)a_t \qquad (9)$$

$$and\ a_{jt+1} \geq \underline{a}.$$

That is, households maximise utility subject to the constraint that labour income or the unemployment benefit (depending on the realisation of the Markov process) plus income from asset holdings are not less than spending on consumption and asset purchases for next period. Household assets cannot be negative.

3.2. Firms

Assume a representative firm that produces output (Y_t) using the following Cobb-Douglas technology

$$Y_t = K_t^{\alpha} L^{1-\alpha}, \qquad (10)$$

K_t is the aggregate capital stock and L is the aggregate labour supply, which due to the constant transition probabilities and inelastic labour supply, is invariant over time.

The firm's profit maximisation problem is then

$$\max_{\{K_t, L\}} \pi = K_t^{\alpha} L^{1-\alpha} - (r_t + \delta)K_t - w_t L, \qquad (11)$$

which gives the following factor prices

$$r_t = \alpha K_t^{\alpha-1} L^{1-\alpha} - \delta, \qquad (12)$$

$$w_t = (1-\alpha)K_t^{\alpha}L^{-\alpha}, \quad (13)$$

where δ is the capital depreciation rate.

3.3. Equilibrium

Similar to Winberry (2016b) a recursive competitive equilibrium is a list of functions

1. Household optimisation. Households take factor prices r_t and w_t, as given and $a_{j,t+1}$ satisfies the consumption Euler

$$\left(c_{jt}^{p} + \alpha^{g}G_t\right)^{-\sigma} \geq \beta E\left[(1+r_{t+1})\left(c_{j,t+1}^{p} + \alpha^{g}G_{t+1}\right)^{-\sigma}\right]. \quad (14)$$

This holds with equality if $a_{t+1} > \underline{a}$, that is; $a_{t+1} > 0$.

Optimal consumption is then given by

$$c_{jt}^{p} = w_t\left((1-\tau)\varepsilon_{jt} + b\left(1-\varepsilon_{jt}\right)\right) + (1+r_t)a_t - a_{t+1}. \quad (15)$$

2. Firm optimisation and market clearing. Factor prices satisfy (12) and (13)

$$r_t = \alpha K_t^{\alpha-1}L^{1-\alpha} - \delta,$$

$$w_t = (1-\alpha)K_t^{\alpha}L^{-\alpha}.$$

Aggregate capital is given by $K_t = \sum_{\varepsilon}\int a_{jt}d\mu_t(\varepsilon_{jt}, a_{jt})$. (16)

3. Evolution of the household distribution. For all measurable sets Δ_{a_t},

$$\mu_{t+1}(G_t,\mu_t)\left(\varepsilon_{jt},\Delta_{a_t}\right) = \sum_{\tilde{\varepsilon}} \pi(\varepsilon|\tilde{\varepsilon}) \int 1\left\{a_{t+1}(\tilde{\varepsilon}, a_t;\ G_t,\mu_t) \in \Delta_{a_t}\right\} \mu_t\left(\tilde{\varepsilon}, da_{jt}\right). \quad (17)$$

3.4. Solving the Model

Winberry (2016b)'s computational method for solving heterogeneous agent models is used to solve the model. This section therefore very closely follows his paper, with minor adjustments to take into account the inclusion of government spending, and the removal of the total factor product shock. Winberry notes that the most challenging part of solving the model is approximating the equilibrium with 'finite-dimensional objects' (p.3). What he means by this is both the distribution of households across productivity and assets, and the households' decision rules (which incorporate this distribution), are infinite dimensional objects that must be approximated in order to calculate the steady-state and dynamics of the model.

3.4.1. Why Winberry?

Previous algorithms designed to approximate these objects, such as in Krusell and Smith (1998), use approximate aggregation where they normally only use the mean of the distribution to approximate it (Winberry, 2016a). Ma (2017) for example solves his model following Krusell and Smith, using the mean of the asset distribution. Winberry (2016a) develops a new method which parsimoniously approximates the entire distribution. He uses a parametric family to approximate the infinite dimensional distributions with a finite object.

There are several reasons for using the Winberry algorithm over others. The first is that rather than relying on approximate aggregation as in Krusell and Smith and Ma, Winberry's method approximates the entire distribution which enables me to easily compute impulse responses for the higher moments of the distribution (for example variance and skewness). This more richly captures the dynamics of how wealth inequality evolves in response to the government consumption shock. The second reason is that it is (relatively) easy to adjust the Winberry algorithm to the model one is using, and it does not take long to solve the model. In all my simulations it never took more than five minutes to calculate the steady-state and dynamics of the model.

A useful paper written recently is Terry (2017) which looks at five different methods for solving heterogeneous agent firm models, including the Krusell and Smith approximate aggregation approach and Winberry's algorithm. Terry finds that all five methods calculate 'broadly similar' results for business cycle fluctuations (p. 1083). The Krusell and Smith method is found to be most accurate, however it is also very computationally intensive. The Winberry algorithm is favoured for its speed. In a similar approach to Winberry, also evaluated by Terry, Reiter (2009)'s algorithm uses a nonparametric histogram to model how the distribution changes over time. In contrast, Winberry's method approximates this histogram with a 'flexible parametric family' (Winberry, 2016a, p. 1). Terry notes that by doing this, Winberry ensures that the distribution over time is fully characterised by its moments (and the idiosyncratic productivity level in the firm model) instead of an entire histogram. The Winberry method is therefore much faster and less computationally intensive, whilst still delivering similar results to the Krusell and Smith method (Terry, 2017).

3.4.2. Approximation

As a result of the borrowing constraint ($a_{jt+1} \geq \underline{a}$) Winberry (2016b) approximates the mass of households at the borrowing constraint and the distribution of households away from the borrowing constraint separately.[6] Once the distribution of household assets is approximated, aggregate capital can be computed.

The household consumption decision rule has a kink in it due to the borrowing constraint. Winberry notes that there are two options in this case: approximation of the savings rule can be done with linear splines, or approximation of expected future consumption using polynomials. He finds polynomials to be more efficient, therefore polynomials are used here to approximate the expectation of future consumption in this model. Chebyshev polynomials are used to approximate the expectation of consumption. The conditional expectation function is

$$\psi_t(\varepsilon_{jt}, a_{jt}) = E\left[\beta(1 + r_{t+1})\left(c_{j,t+1}^p(\varepsilon_{jt+1}, a_{jt+1}(\varepsilon_{jt}, a_{jt})) + \alpha^g G_{t+1}\right)^{-\sigma}\right]. \quad (18)$$

Winberry then notes that if the borrowing constraint is not binding, optimal consumption next period will follow the Euler equation such that

$$\psi_t(\varepsilon_{jt}, a_{jt}) = \left(w_t\left((1-\tau)\varepsilon_{jt} + b(1-\varepsilon_{jt})\right) + (1 + r_t)a_t - a_{t+1} - \alpha^g G_t\right)^{-\sigma}, \quad (19)$$

and the saving and consumption policies are derived through the conditional expectation to be

$$a_{jt+1}(\varepsilon_{jt}, a_{jt}) =$$

$$max\left\{\underline{a}, w_t\left((1-\tau)\varepsilon_{jt} + b(1-\varepsilon_{jt})\right) + (1 + r_t)a_t - \alpha^g G_t - \psi_t(\varepsilon_{jt}, a_{jt})^{-\frac{1}{\sigma}}\right\}, \quad (20)$$

[6] See Winberry, 2016b, p. 4-5 for details on the laws of motion for households at and away from the borrowing constraint.

$$c_{jt}^{p}(\varepsilon_{jt}, a_{jt}) = w_t\left((1-\tau)\varepsilon_{jt} + b(1-\varepsilon_{jt})\right) + (1+r_t)a_t - a_{t+1}(\varepsilon_{jt}, a_{jt}). \qquad (21)$$

The conditional expectation function is then approximated using Chebyshev polynomials as in Winberry (2016b).[7] At this point it is now possible to calculate the equilibrium: the distribution of households and conditional expectation of function have been approximated.

4. Calibration

The baseline specification of the model is calibrated for the US as in Winberry (2016b). Following Winberry I set the subjective discount factor of households to $\beta = 0.96$. Capital depreciation is $\delta = 0.1$, that is, ten percent per year. Household productivity if employed (ε_1) is set to 1 and productivity if unemployed (ε_0) is 0. The capital share of income is set to $\alpha = 0.36$ This is consistent with the estimates income share estimates of Gollin (2002). The unemployment benefit is fifteen percent of the employed wage, that is, $b = 0.15$. The tax rate (τ) is derived by setting taxes on labour equal to the amount paid in benefits

$$bw_t(1-L) = \tau w_t L, \quad (22)$$

$$\tau = \frac{b(1-L)}{L}. \quad (23)$$

The probability of an unemployed household becoming employed is 0.5, whilst the probability of an employed household becoming unemployed is 0.038. Markets are incomplete in this model because households are not allowed to borrow. This means that the asset constraint will be $\underline{a} = 0$. σ is set to 1 which means that the elasticity of intertemporal substitution is one. This paper considers three values for the complementarity of private and public consumption spending: α^g: -0.5, -1, and -1.5 with -1 being the main scenario. Steady-state government consumption expenditure is set to eleven percent of output. Following Gali et al. (2007) ρ^g (the persistence of the government consumption shock) is set to 0.9. This is consistent with the estimates of Finn (1998) which set $\rho^g = 0.86$.

[7] See Winberry, 2016b, p. 7 for details.

4.1. Government Spending Calibration

It is difficult to find estimates for α^g as entered into equation 2 because recent papers tend to use a CES aggregate of public and private consumption (see Ambler, Bouakez and Cardia, 2017; Brown and Wells, 2008; and Bouakez and Rebei, 2007 for examples).

Feve, Matheron and Sahue (2013) use US data and likelihood estimation to pin down a value of -0.634 for α^g when endogenous government spending policy is set to zero. McGrattan (1994) estimates α^g to be negative but statistically insignificant from zero for the US. Karras (1994) estimates the value of α^g for many countries (excluding the US) and finds that it is either negative, or not significantly different from zero. Karras concludes that private and public consumption are therefore complementary. He describes substitutability between private and public consumption (that is $\alpha^g > 0$) as 'the exception and not the rule' for the countries he examined (p. 10). Fiorito and Kollintzas (2004) also find that public and private consumption are complements.

Due to the lack of available estimates, and the inconsistent results using US data, I use the average of Karras (1994)'s results for developed countries, with insignificant α^gs set to zero. This comes out at -0.979, which I round up to -1. For robustness, since many countries have values much larger and smaller in absolute value than -1, I also consider -1.5 and -0.5. The case calculated by McGrattan (1994) where $\alpha^g = 0$ is not considered because this is just the standard real business cycle result. A positive value is not considered because none of the attempts to estimate α^g (that I have come across) find a positive and significant value for α^g.

Calibrating steady-state government spending share of output is more complicated than usual due to the type of government spending I am examining. Gali et al. (2007) and Feve, Matheron and Sahue, (2013) set government spending to 20 percent of output in steady-state. This is consistent with empirical estimates such as Ramey and Shapiro (1998) who calculate that government spending consists of 21 percent of output in steady-state for the US, and Christiano and Eichenbaum (1992) who estimate the government share of output to be 17.7 percent.

It is important to note that in this model it is government consumption expenditures (minus defence spending) which enters into the utility function, not overall government expenditure. Finn (1998) calibrates her model for the US and sets the government consumption share of output at seven percent. I calculate the output share of government consumption to try and mirror the merit good spending in Fiorito and Kollintzas (2004). I take nominal Federal and state government consumption expenditures, excluding defence spending, and divide their sum by nominal GDP to obtain the steady-state output share of government consumption. I use observations from quarter one 1975 to quarter three 2017 because during this period government consumption oscillates around eleven percent of GDP. Prior to this period US government consumption spending increases as a share of GDP from seven to eleven percent of GDP. Including these observations might result in fifty year old government spending behaviour downwardly biasing my calculation for steady-state government consumption. This results in an output share for government spending of eleven percent. The data is taken from the Bureau of Economic Analysis and FRED databases.

This value is consistent with Leeper, Walker and Yang (2010) who calibrate steady-state government consumption in a real business cycle model. They find steady-state government consumption is 14.4 percent of GDP. Their model is calibrated for the US using data on government assets from 1960-2007 in the Fixed Assets Accounts tables (from the Bureau of Economic Analysis).

Table 1: Values assigned to each parameter of the model.

Parameter	**Value**	**Interpretation**
β	0.96	Discount Factor
δ	0.1	Depreciation Rate
ε_0	0	Unemployed Efficiency Units
ε_1	1	Employed Efficiency Units
α	0.36	Capital Share of Output
b	0.15	Unemployment Benefit
τ	$\frac{b(1-L)}{L}$	Income Tax
σ	1	Utility Curvature (Elasticity of intertemporal substitution $= 1/\sigma$)
Emp. to Unemp. probability	0.038	Probability of becoming employed each period
Unemp. to Emp. probability	0.5	Probability of unemployed becoming employed each period
$\underline{\alpha}$	0	Borrowing Constraint
$\frac{G}{Y}$	0.11	Steady-state government consumption share
α^g	-0.5, -1, -1.5	Complementarity of government and private consumption
ρ^g	0.9	Persistence of government consumption shock

Notes: All parameters except government consumption parameters are calibrated as in Winberry (2016b). The calibration of the government consumption parameters is discussed in sections 4 and 4.1.

5. Model Simulations.

The main case that is considered in this paper is $\alpha^g = -1$. Since the estimates for α^g discussed in the calibration section are so inconsistent, I use the average of Karras (1994)'s estimates for developed countries. As a robustness check, in the next section I compare simulations with $\alpha^g = -1.5, -1$ and -0.5.

5.1. Steady-State

It is instructive to compare the steady-state distributions of the model with and without government consumption. The model without government consumption, that is $\alpha^g = 0$, has the same steady-state distribution as Winberry (2016b). Figure 1 (page 38) plots the steady-state asset distributions for the unemployed households with $\alpha^g = 0$ and $\alpha^g = -1$. Figure 2 (page 38) does the same but for the employed. For ease of interpretation the initial guess at the distribution is excluded and only the finite parametric approximation that is used to solve the model is shown. The dotted red line shows the steady-state from the $\alpha^g = 0$ (and Winberry) model. The solid black line shows the model when government consumption is calibrated to impact private consumption.

It is interesting to note that when government consumption is complementary to private consumption, in steady-state the mean of the asset distribution is higher. This is true of both the employed and unemployed households. The difference is a result of the way government consumption is entered into the model. When $\alpha^g \neq 0$, adding government spending to the model changes the steady-state asset distribution by directly altering the savings rule used in the initial guess and approximation of the distribution. The savings rule is derived from the optimality condition of each household, and because government consumption enters into the utility function, it directly impacts the savings decision. When government consumption is a complement to private consumption, it directly increases household savings by entering positively into the savings function. Higher savings means more asset accumulation by households, resulting in higher average asset holdings when private and public consumption are complements.

Other features to note are that the variance of the distributions when $\alpha^g = -1$ is smaller, with households more tightly distributed around the mean. For the unemployed households, there is a much smaller mass at the borrowing constraint when government consumption is complementary compared to when it does not impact private consumption.

Figure 1: Steady-state asset distribution of unemployed households. This figure shows the wealth distribution for unemployed households when there is no complementarity between private and public consumption, and when there is ($\alpha^g = -1$).

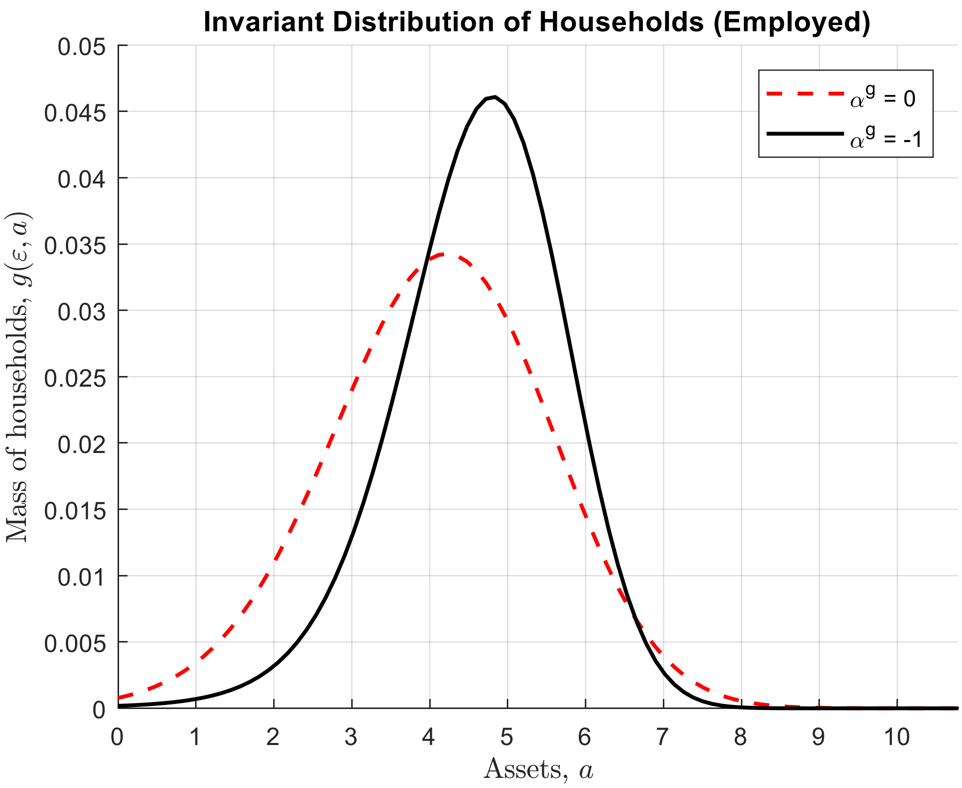

Figure 2: Steady-state asset distribution of employed households. This figure shows the wealth distribution for employed households when there is no complementarity between private and public consumption, and when there is ($\alpha^g = -1$).

5.2. Dynamics

I now calculate impulse response functions for a temporary one percent shock to government consumption expenditure. Figure 3 shows the IRFs for the main aggregate variables.

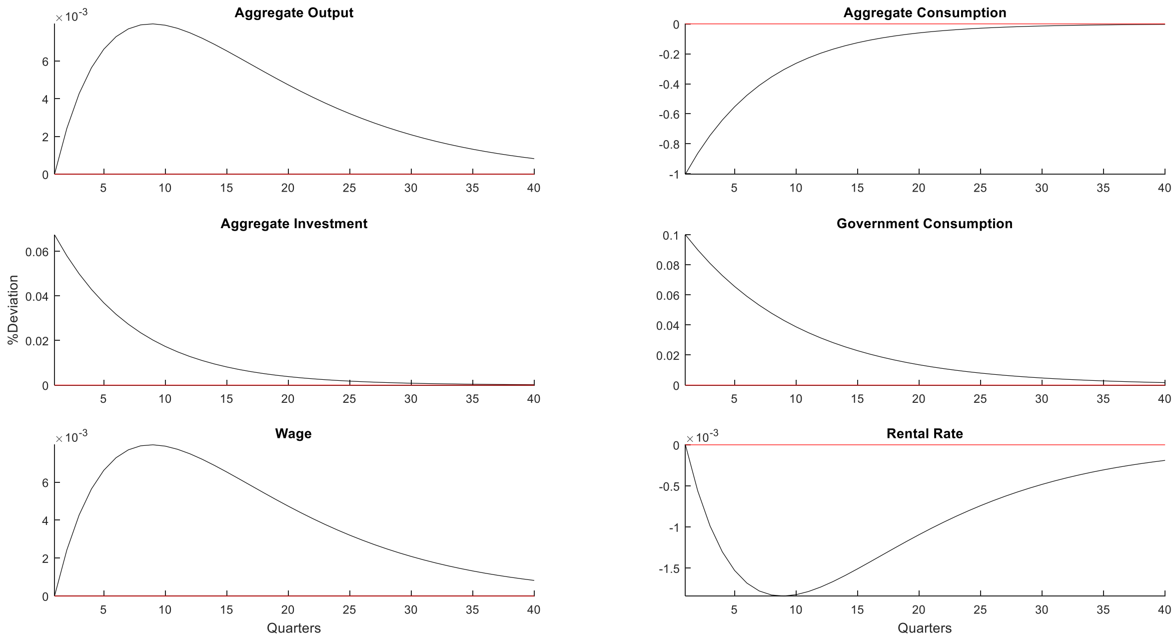

Figure 3: Impulse responses of aggregate variables to a temporary, positive 1% shock to government consumption spending. Responses are percentage deviations from steady-state at each quarter after the shock.

5.2.1. Aggregate Variables

In response to the shock in quarter zero, each household immediately decides to save more which increases quarter one asset holdings of all households, regardless of employment status. In quarter zero, the wage and real rental rate are unchanged, whilst households are accumulating more assets and cannot borrow (see Figure 6 on page 47 for the IRFs of the moments of the asset distributions). They must therefore reduce their consumption. This is seen in Figure 3 where aggregate consumption immediately falls by one percent.

Higher asset purchases by all households in the initial quarter mean that aggregate capital in the economy will be higher in the first quarter after the shock. This results in higher aggregate output by the representative firm. This can be seen in the IRFs where aggregate output does not respond immediately to the government spending shock, only increasing in quarter one as the higher aggregate capital enters the production function of the firm. Because there is more capital in the economy, the real return on capital falls in the second quarter, whilst wages increase as each worker now has more capital to work with, increasing their productivity. The model is constructed such that labour is supplied inelastically if the household is employed, so the higher wage does not lead to any changes in aggregate employment.

The temporary nature of the shock to government consumption means that government consumption is lower in the second quarter which *ceteris paribus* would lead to asset holdings of employed and unemployed households to fall compared to quarter one. However, there are several other components of the household savings function that have been altered. Wages are higher, and the change in household income from assets is positive because the percentage increase in asset holdings by households is far larger than the reduction in the real return to capital. Overall the changes in income from assets, and labour/the unemployment benefit, more than offset changes in government consumption and expected private consumption. Asset accumulation by employed and unemployed households thus continues to increase. This then continues to drive up aggregate capital, wages and output, and drive down the interest rate.

This process continues until the eighth quarter. At this point, the net effect of changes in household income, government consumption and private consumption expectations is negative, and both household types decide to accumulate less assets than quarter seven. The mean asset holdings of employed and unemployed households peak at in quarter seven and then return to steady-state.

Output and wages peak in quarter eight (as they are determined by aggregate capital in period seven). With asset holdings declining after quarter eight, aggregate capital falls, which results in lower aggregate output, lower wages and higher real return on capital.

Table 2: Peaks/Troughs of the IRFs of each variable and what quarter after the government consumption shock they occur.

Variable	**Quarter of Peak/Trough**	**Size of Peak/Trough** (Percentage deviation from steady-state).
Output	8	0.008%
Consumption	0 (Immediate)	-1.0025%
Investment	0 (Immediate)	0.0674%
Rental Rate	8	-0.0018%
Wage	8	0.008%
Mean Assets (Unemployed)	7	0.0935%
Mean Assets (Employed)	7	0.1018%
Variance Assets (Unemployed)	5	0.0284%
Variance Assets (Employed)	4	0.0078%
Skewness Assets (Unemployed)	4	-0.0576%
Skewness Assets (Employed)	2	-0.0172%

5.2.2. Private Consumption Response

The response of private consumption to the government consumption shock warrants more discussion. As shown in the model section of this paper, government and private consumption are complements in this calibration of the model due to the way increases in government consumption increase the marginal utility of private consumption. The question then, is why the large negative response of private consumption to government consumption?

One argument is that the use of a real business cycle model is the problem. As Gali, Lopez-Salido and Valles (2007) note, the standard real business cycle framework suggests a negative response of private consumption to government spending shocks. In real business cycle models, the consumption decision of households is based on an intertemporal budget constraint, therefore if government spending increases, this lowers their present-value of after tax income, thus reducing consumption. Gali, Lopez-Salido and Valles instead use a New Keynesian model with Calvo pricing in which a fraction of households are Ricardian, and a fraction are non-Ricardian.[8] They find that it is possible in the New Keynesian framework to observe a positive response of private consumption to government spending shocks. The non-Ricardian households reduce the impact of the wealth effect on aggregate demand because they do not anticipate future tax increases, whilst having sticky-prices allows real wages to increase. Combined, these two impacts raise the consumption of the non-Ricardian households.

The issue with this argument is that the negative response of consumption in this model does not come from anticipation of future taxation to pay for government spending. Government spending is financed with lump sum taxation, which is non-distortionary, and the negative response of private consumption actually comes from the fact that government consumption enters into the asset accumulation decision of households. An increase in government consumption triggers higher asset accumulation by households, lowering private consumption.

Bouakez and Rebei (2007) develop an explanation of the positive response of private consumption to government consumption within the real business cycle framework. This explanation helps untangle the response of private consumption in my model. In their model, as in this one, government spending and private consumption are complementary, with

[8] By Ricardian they refer to households that optimise, whereas non-Ricardian households do not, for example due to ‘myopia, lack of access to capital markets, fear of saving, ignorance of intertemporal trading opportunities etc’ (Gali, Lopez-Salido and Valles, 2007, p.236).

increases in government spending raising the marginal utility of private consumption. Government spending in their model is also financed by lump sum taxation. The key point in their paper is that by raising the marginal utility of private consumption, government consumption increases household incentives to work so that they can consume more. When the complementarity between government spending and private consumption is high enough, a positive response of private consumption to government consumption is observed.

The model in my paper is very stylised. The fact that households supply one efficiency unit of labour if employed, zero otherwise, means that it cannot capture the labour supply effect of complementarity between private and public consumption. I observe a negative response of private consumption to government consumption because households do not adjust their labour supply to finance more private consumption spending. Although they might consume more due to the income effect of higher wages, this is not enough to counteract the consumption they sacrifice to accumulate more assets.

Another reason why my model might not be generating a crowding-in effect of government consumption is that household consumption is a simple linear combination of private and public consumption. In a more recent strand of the literature, government and private consumption have been combined using a CES aggregate. Brown and Wells (2008) offer an excellent overview of the two different approaches in their literature review. They note that in this later work, authors attempt to determine the relationship between private and public consumption by looking at their elasticity of substitution. In Ambler, Bouakez and Cardia (2017), Brown and Wells, (2008), and Bouakez and Rebei (2007) for example, total household consumption, $\tilde{c}_t$, is

$$\tilde{c}_t = \left(\theta c_t^{-\sigma} + (1-\theta) C_{gt}^{-\sigma}\right)^{1/\sigma},$$

where c_t is private consumption and C_{gt} is government consumption spending. With this formulation, the elasticity of substitution between private and public consumption spending is $\upsilon \equiv 1/(1+\sigma)$. Bouakez and Rebei (2007) also use a CES aggregate to introduce government consumption to the utility function. They show that when υ is low enough (they set $\upsilon = 0.25$) the complementarity between private and public consumption is large enough to result in the crowding in of private consumption in response to government spending shocks. An interesting extension of my simple model would be to incorporate an endogenous labour supply and a CES aggregate of private and public consumption in order to activate the crowding-in mechanism described by Bouakez and Rebei (2007).

A final point to note regarding the response of private consumption is that it depends on the persistence of the government spending shock. Gali, Lopez-Salido and Valles (2007) note that the more persistent is government consumption, the more negatively it will affect the present value of discounted household income. This means for example that when the coefficient on the lagged value of government consumption is 0.9, the positive response of aggregate consumption is 'relatively small' (p. 253). In the context of my paper, this means that when government consumption shocks are quite persistent, their impact on asset accumulation lingers meaning that private consumption takes much longer to recover to steady-state, and is more negative in later quarters than it would be if shocks were less persistent. I still use 0.9 as the value for the persistence of government consumption because this is what they estimate using a VAR.

5.3. Evolution of the Asset Distribution

The key contribution of this paper is to examine the impact of government consumption on wealth inequality over the business cycle. To do this, the response of the asset distribution to the government consumption shock will now be analysed. To help visualise how the wealth distributions change over time, Figures 4 and 5 (on page 46) report three-dimensional plots of each one.

The variance of the asset distribution is used as the measure of inequality. The variance is one of the most common measures of how widely dispersed a distribution is (Atkinson, 1970; Heshmati, 2004). One problem with measuring inequality is that different measures generate different rankings, and each measure entails its own judgement about societal preferences for given income profiles (Chakravarty, 1990; Dalton, 1920). Atkinson (1970) notes that using the mean and variance to rank income (in our case wealth) distributions assumes that preferences exhibit 'increasing relative and absolute inequality-aversion' (p. 253). This assumption may not be desirable. To test the robustness of this model (and therefore the measure of inequality used), in Section Seven of this paper a VAR analysis of government consumption and wealth inequality is carried out. The estimated impulse response for inequality is consistent with the story told by the simulated impulse response for the variance. I am therefore confident that the variance of the wealth distribution acceptably characterises the response of inequality to government consumption shocks in this model.

Figure 4 shows the asset distribution of households that are currently unemployed at each quarter after the initial government consumption shock. Figure 5 does the same for the employed households at each period after the shock. Figure 6 (page 47) shows the IRFs for the first three moments of the asset distribution. Both asset distributions respond in similar ways to the government consumption shock, however as Figure 6 illustrates, the magnitude of the response is different.

Once the shock hits the economy, the asset distributions immediately shift to the right, whilst their variances both increase. This means that on average households are becoming wealthier, but inequality is increasing in tandem. As asset holdings by both types of household continue to rise, the percentage increase in the variance also rises, further exacerbating wealth inequality. Once average asset holdings peak, the distribution then gradually shifts left again as asset holdings return to steady-state. The variance of the asset distribution returns to steady-state as well, reducing inequality. As Figure 6 shows, the average response of employed household's asset holdings is larger than unemployed. The response of the variance is much larger for the unemployed households. This merits further discussion.

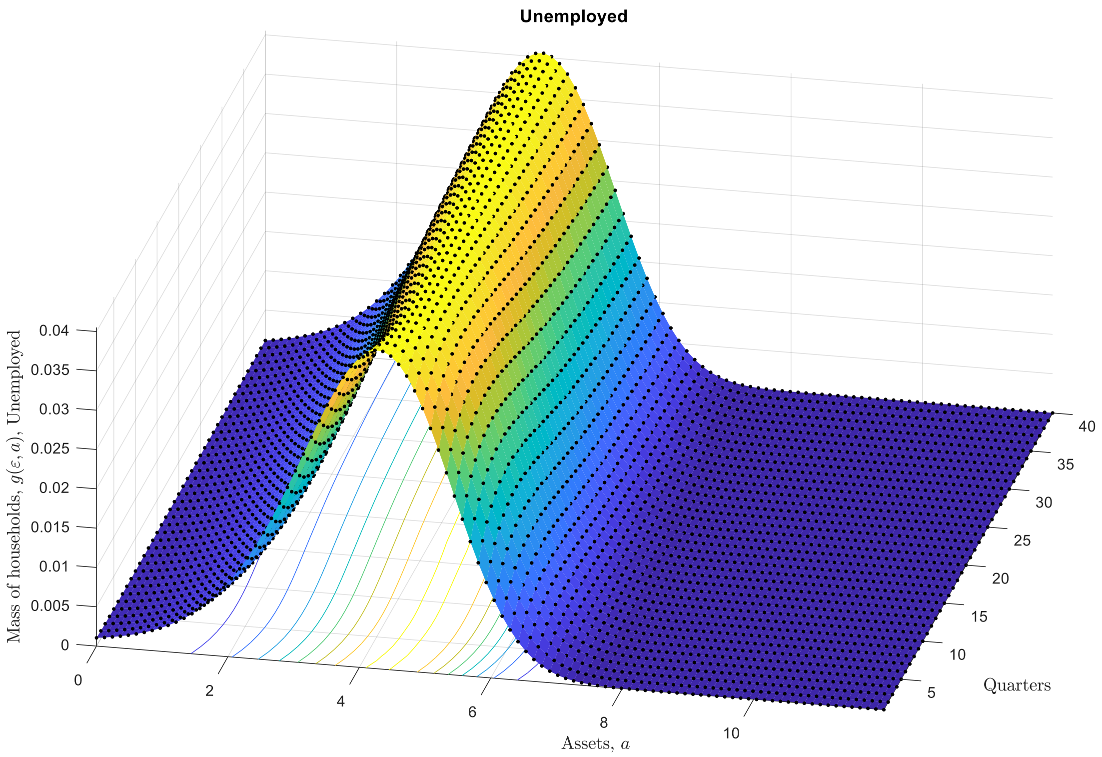

Figure 4: Three-dimensional plot of the response of unemployed asset distribution after the government consumption shock. The distribution immediately shifts right and widens, indicating higher wealth inequality, before returning to steady-state.

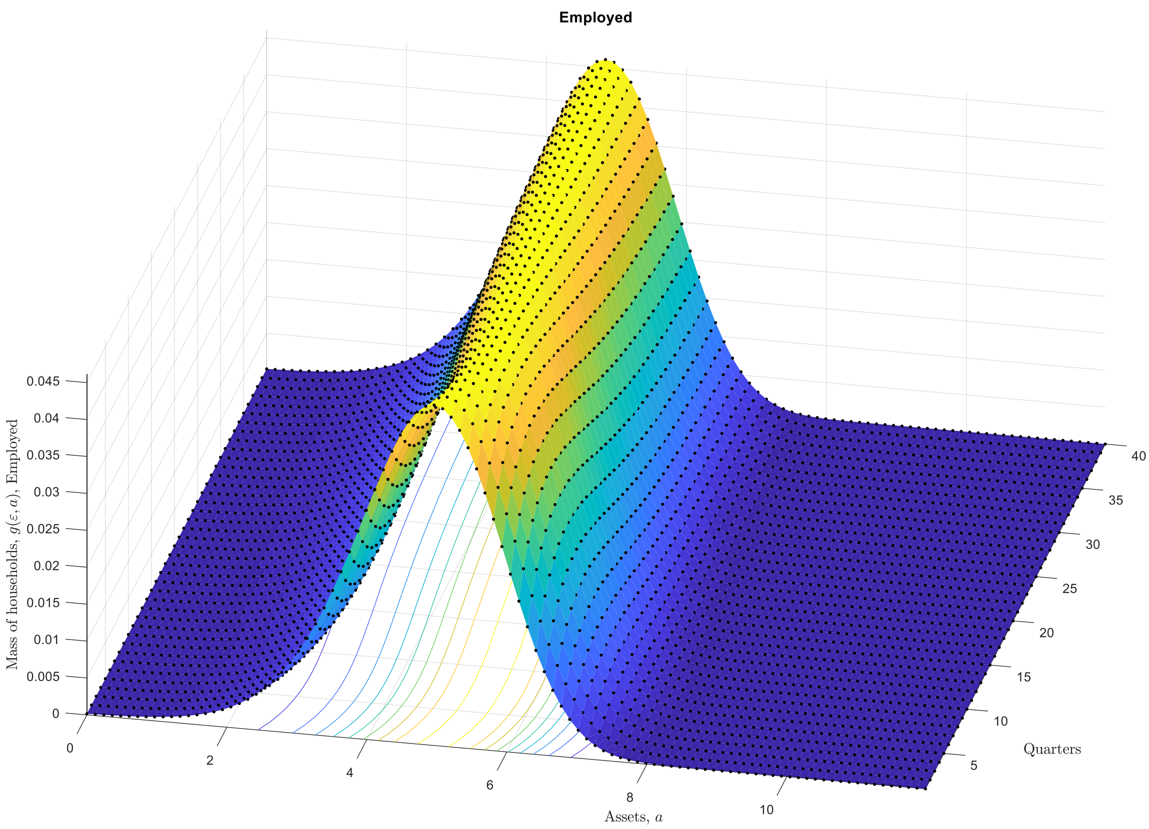

Figure 5: Three-dimensional plot of the response of the employed asset distribution after the government consumption shock. As with the unemployed distribution, it immediately shifts right and becomes more unequal, before returning to steady-state.

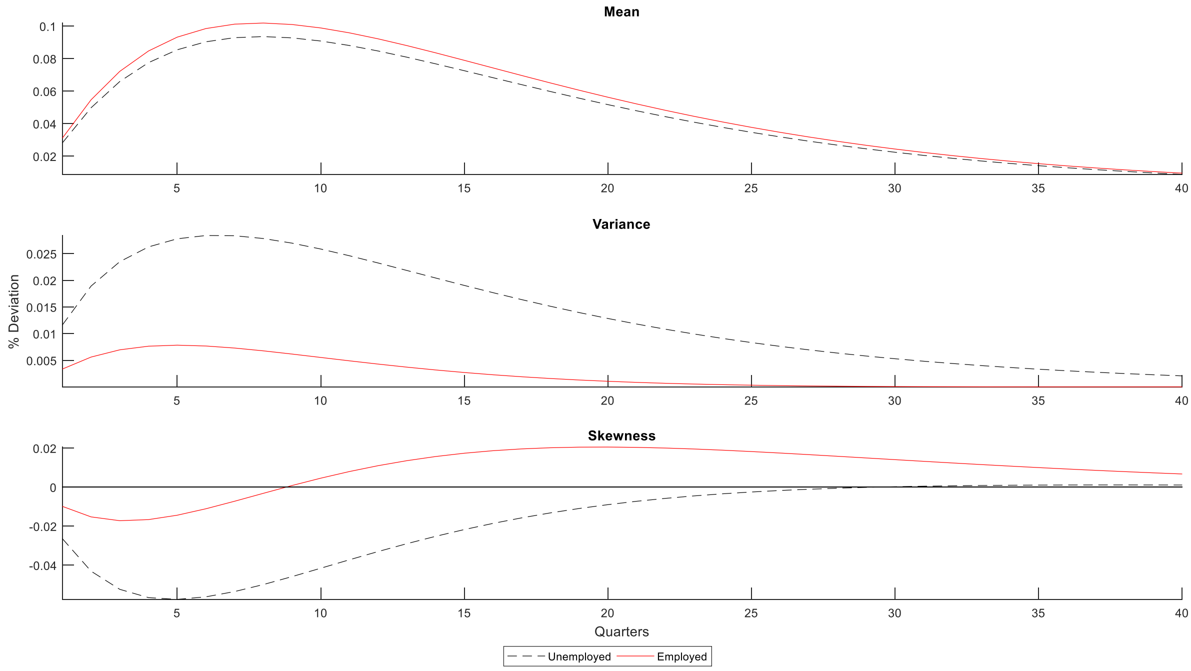

Figure 6: Impulse responses of the moments of the asset (wealth) distribution to a temporary, positive 1% shock to government consumption spending. Responses are percentage deviations from steady-state at each quarter after the shock.

5.3.1. Variance

The variance of the unemployed asset distribution increases until the fifth quarter after the government consumption shock. This increase happens because in steady-state, the unemployed households have fewer assets on average, as a result of living off the unemployment benefit. In the first quarter, half of these unemployed households become employed, whilst 3.8 percent of employed households become unemployed, and these freshly unemployed households have more assets on average than the rest of the unemployed since they were living off the wage in steady-state. Thus the variance of the unemployed asset distribution will increase as asset-rich households join the unemployed each period.

The percentage increase in the mean of the employed asset distribution is larger than the unemployed distribution, which means that each period, households are becoming unemployed who have accumulated more and more assets than those who stayed unemployed, which explains why the variance keeps increasing until quarter five. At this point the increase in the mean of the asset distribution is starting to reach its maximum and the households entering the unemployed state are similarly wealthy to the currently unemployed, thus the percentage increase in the variance of the unemployed asset distribution reaches its maximum.

After this point the process reverses, and the variance of the asset distribution of unemployed households returns to steady-state.

The IRF of the variance of the employed distribution is driven by a similar process, with asset poor households entering the employed state where households have accumulated far more assets. The variance increases until quarter four, before returning to steady-state. 93 percent of households are employed in all periods due to the constant transition probabilities (Winberry, 2016b). This explains the much smaller response of the variance employed asset distribution: the unemployed households entering the employed state are a much smaller proportion of the employed than the employed entering the unemployed state are of the unemployed, therefore their impact on the variance of the asset distribution is smaller.

To summarise, the government consumption shock temporarily increases inequality amongst both types of household, with the variances of the asset distributions both increasing.

5.3.2. Skewness

In steady-state, both the employed and unemployed household asset distributions are left-skewed with values of -0.85 and -1.07 respectively. The responses of the skewness of both distributions are non-monotonic. In response to the government consumption shock the left-skew of both distributions increases immediately. The skewness value of the employed asset distribution has a trough of 0.0172 percent below steady-state in quarter two after the shock. Between quarters seven and eight the skewness returns to steady-state, before increasing to 0.0205 percent above steady-state in quarter nineteen after the shock. Finally it returns to steady-state. The skewness value of the unemployed distribution follows a similar pattern, except the trough is more negative, the return to steady-state occurs later, and the peak above steady-state is smaller.[9]

These IRFs are a result of the mean of the asset distributions shifting to the right whilst the variance increases. As the asset distribution shifts to the right, the left tail gets longer. For the employed distribution, this results from unemployed households (who have lower assets on

[9] To be precise, the skewness of the unemployed asset distribution reaches a trough of 0.0576 percent below steady-state in quarter four after the shock, returns to steady-state during quarters twenty-eight and twenty-nine, reaches a peak of 0.011 percent above steady-state in quarter thirty-seven and ten returns to steady-state.

average than the other employed households) becoming employed and having fewer assets on average than the rest of the employed households, thus dragging the left tail of the distribution out. In the unemployed distribution by contrast, newly unemployed households (who have higher assets on average than the other unemployed households) shift the unemployed asset distribution to the right, leaving behind the households in the left tail.

5.4. Persistence and Size of the Impulse Responses

The results of the simulations have thus far shown that government consumption shocks do impact the mean and variance of the distribution and cause wealth inequality to increase above steady-state. I next examine the persistence of the response of the asset distribution compared to the other macroeconomic variables. This is an important question because it might be the case that even after the policy variables have returned to steady-state, the asset distribution might still be moving in response to the government consumption shock.

As an illustration, Figure 7 (page 50) plots the IRFs for the moments of the asset distribution along with the IRF for government consumption. It is clear in Figure 7 that the means of the employed and unemployed asset distributions are returning to steady-state many quarters after the government consumption shock dissipates. What is not as clear from the IRFs is that all the moments of the employed and unemployed are twice as persistent as government consumption.

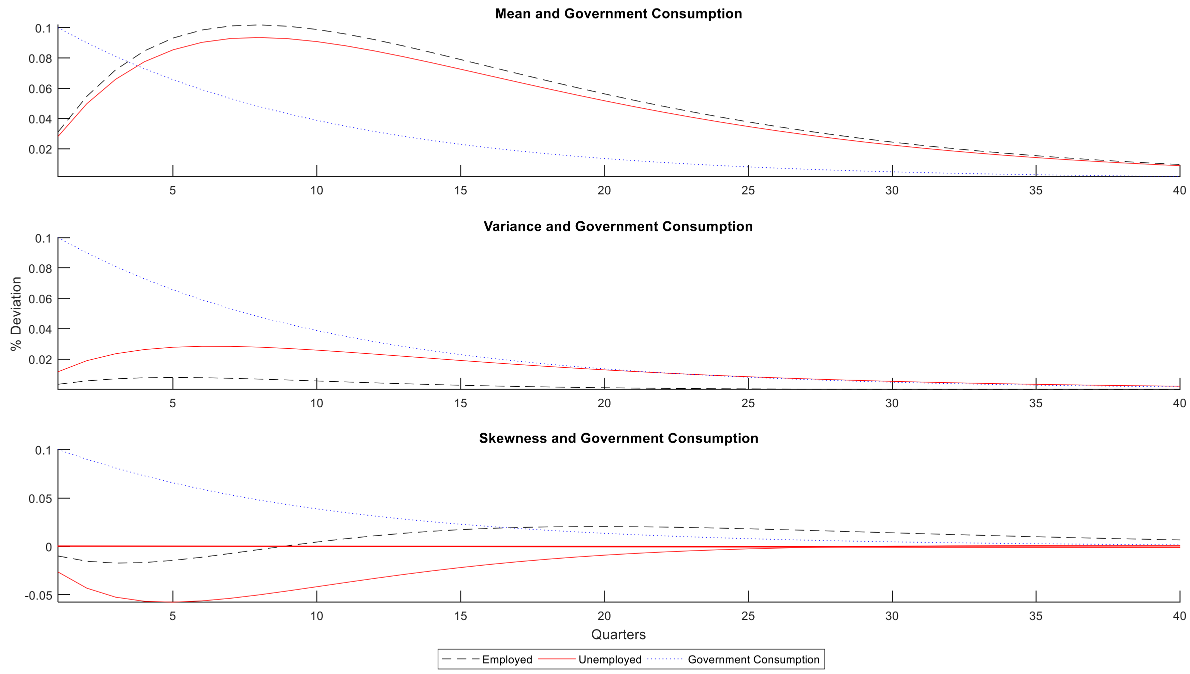

Figure 7: Impulse responses of the moments of the asset distributions for the employed and unemployed households compared with the response of the government consumption shock. Responses are percentage deviations from steady-state.

The simulated first order autocorrelation coefficient of government spending is 0.4635, whilst for all the moments of the asset distributions, the autocorrelation is 0.801 or higher. This result indicates that even after the government consumption shock has dissipated, it is still impacting inequality.

Figure 7 also illustrates the size of the response of the moments of the asset distribution to the government consumption shock. The peak of the IRF for the mean of the employed asset distribution is approximately the same size as the government consumption shock, whilst the unemployed peak is slightly lower. Government consumption shocks in this model therefore result in a roughly one for one response in the mean of the asset distributions after a lag of seven quarters. The higher moments of the distributions, despite being more persistent than government consumption, have smaller responses.

6. Robustness

6.1. Alternative Calibrations

As discussed in the calibration section, a number of developed countries in the estimates of Karras (1994) have values of α^g that are significantly different from minus one. This section will therefore analyse the impact of government consumption shocks when there is stronger complementarity ($\alpha^g = -1.5$) and weaker complementarity ($\alpha^g = -0.5$).

6.1.1. Steady-State

Government spending directly impacts the savings decision of households. The initial asset distribution should therefore be different when the complementarity between private and public consumption changes. This is indeed what is observed. In Figure 8 (next page) the steady-state asset distributions from all three calibrations of α^g are plotted for the employed and unemployed households.

The plot shows that for both types of households, as the degree of complementarity between private and public consumption increases, the asset distribution shifts to the right and becomes tighter around the mean, with a larger mass of households at the mean of the distribution. This again is because of the way that government consumption impacts the savings decision of households. When α^g increases, this stimulates more asset accumulation by both types of household. As a result, the steady-state asset distribution shifts to the right as average asset holdings increase.

Another interesting point to note in steady-state is that the mass of unemployed households who have no assets (that is, are at the borrowing constraint) becomes insignificant. This is because as government consumption becomes an increasingly strong complement for to private consumption, households accumulate more assets rather than consume privately, which means that few households will have zero assets.

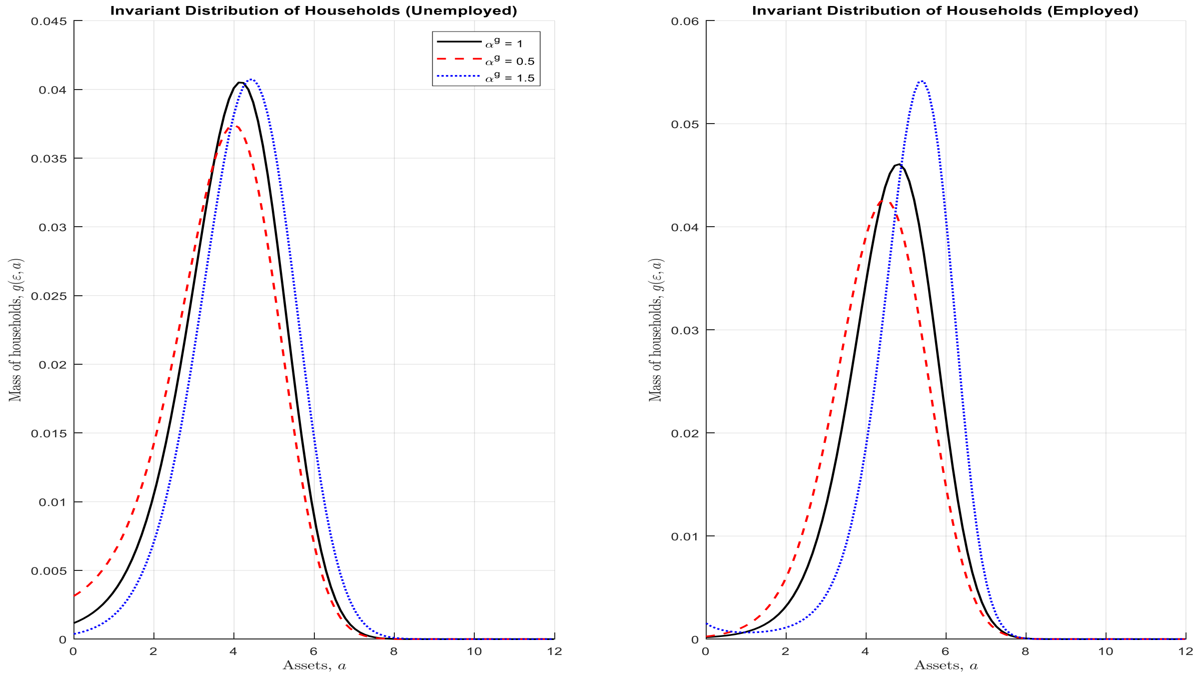

Figure 8: Steady-state asset distributions for employed and unemployed households with three different levels of complementarity between private and public consumption spending.

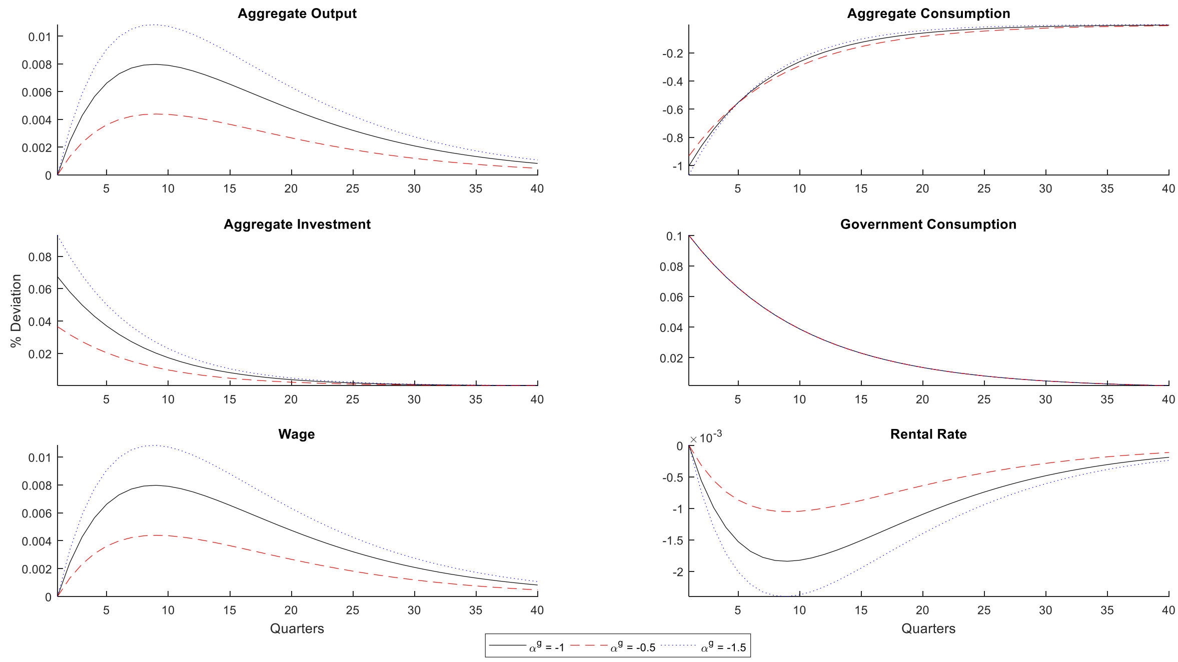

Figure 9: Impulse responses of the aggregate variables to the positive, temporary, 1% government consumption shock with three different levels of complementarity between private and public consumption spending.

6.1.2. Dynamics

Next the model is simulated again, but with different values for the complementarity between private and public consumption. Figure 9 (page 52) reports to IRFs for the main variables in the economy. It is immediately clear from the IRFs that the degree of complementarity has a significant impact on how the economy reacts to the government spending shock. This is all driven by the response of household asset accumulation to government spending. When there is weak complementarity ($\alpha^g = -0.5$) government consumption shocks result in less asset accumulation by households than when $\alpha^g = -1.5$ or -1. Next period there is less capital in the economy, resulting in changes to output, investment, the rental rate and wages that are all of a smaller magnitude than with more negative values of α^g. Aggregate consumption also falls by a smaller amount when private and public consumption are less complementary.

The evolution of the asset distribution also varies depending on how complementary private and public consumption are (see Figure 10 on the next page). For employed and unemployed households when $\alpha^g = -0.5$, the peak response of the mean of the asset distribution is half the size compared to when $\alpha^g = -1$, and a third of the size compared to $\alpha^g = -1.5$. For the unemployed, the impulse response of the variance is much smaller when $\alpha^g = -0.5$ than the other two values. When government and private consumption are less complementary, wealth inequality amongst the unemployed does not increase as much in response to government consumption shocks. This suggests that the larger α^g is in absolute value, the larger the increase in inequality amongst the unemployed in response to government spending shocks.

This impact on inequality amongst the unemployed is a result of the mechanism discussed in the main model simulations. The spike in inequaltiy amongst the unemployed is driven by richer employed households becoming unemployed and therefore driving up wealth inequality. When private and public consumption are less complementary mean asset accumulation is lower, therefore the newly unemployed have fewer assets, and their impact on wealth inequality is smaller.

The response of the variance of the employed asset distribution is quite different depending on the value of α^g and it is not clear why. When $\alpha^g = -1.5$ the increase in wealth inequality is the smallest and peaks the earliest. When $\alpha^g = -0.5$ wealth inequaltiy peaks higher than when it is -1.5 but lower than the baseline case. It then returns to steady-state and overshoots so that in later quarters, inequality is lower than in steady-state.

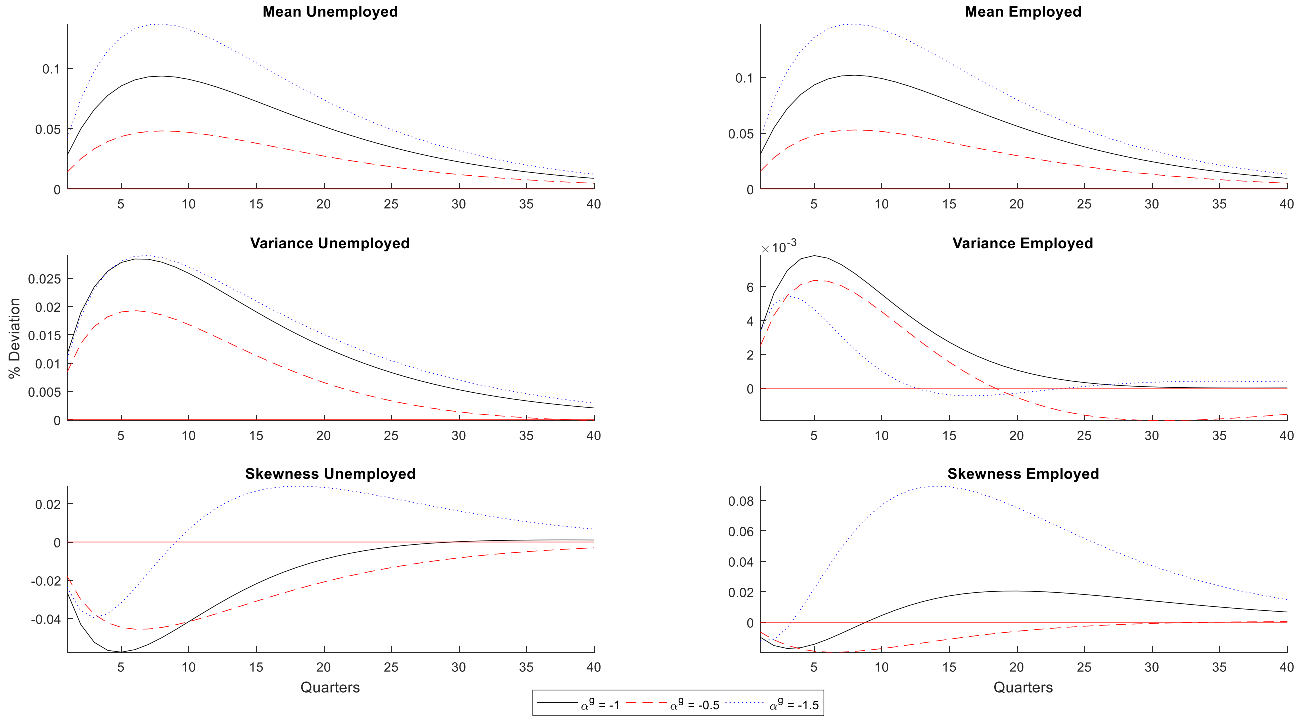

Figure 10: Impulse responses of the moments of the unemployed and employed asset distributions in response to a positive, temporary, 1% shock to government consumption spending. Three different levels of complementarity between private and public consumption spending.

The picture that emerges from this is that for the unemployed, higher complementarity between private and public consumption unambiguously increases wealth inequatliy. For the employed households, the impact of increasing complementarity is ambiguous. This is not an expected result: given that the only thing altered is the degree of complementarity and therefore how much households adjust savings in response to government consumption shocks, I would expect the variance impulse responses to have the same direction for all three cases, just with larger magnitudes. The sensitivity of the wealth inequality amongst the employed to different levels of α^g is something that could be investigated in future work.

6.2. Unemployment Duration

A final interesting question to consider is how changing the duration of unemployment changes the response of the economy to government spending shocks. This is a relevant question in the aftermath of the Global Financial Crisis, as the average duration of unemployment increased after 2008. The mean duration of unemployment between January 1948 and December 2007 was 13.5 weeks, between January 2008 and December 2017 it was 31 weeks (data from Federal Reserve Bank of St. Louis).

The duration of unemployment in the model thus far has been one quarter. In this section I change the duration to half a quarter and simulate the model again. The results from this illustrate how simply looking at the average impact of a policy (as in representative agent models) can hide other important effects. The IRFs for the variables in Figure 11 (page 56) suggest that changing the duration of unemployment makes very little difference. The responses of the means of the asset distributions are very similar for the different unemployment durations as seen in Figure 12 (page 56).

The IRFs for the variances of the asset distributions in Figure 12 highlight the main impact of changing the unemployment duration. When unemployment is half a quarter, the peak of the response of the variance is much smaller and occurs earlier than when unemployment is a full quarter. Interestingly for the employed households, when unemployment is only half a quarter, the variance actually gets smaller after quarter ten.

These IRFs suggest than when the duration of unemployment is shorter, wealth inequality increases by less. This is a sensible result. Households spend less time unemployed, therefore they do not fall as far behind the employed households in their asset accumulation. This simulation shows that governments should be aware that their consumption spending can have different impacts on inequality depending on the state of the macroeconomy. When the duration of unemployment spells for households is small, inequality might actually be reduced by increased government consumption. However, if the duration of unemployment is larger, then it can be exacerbated. Investigating the relationship between the duration of unemployment, and the impact of government consumption on wealth inequality could be another interesting extension for this model.

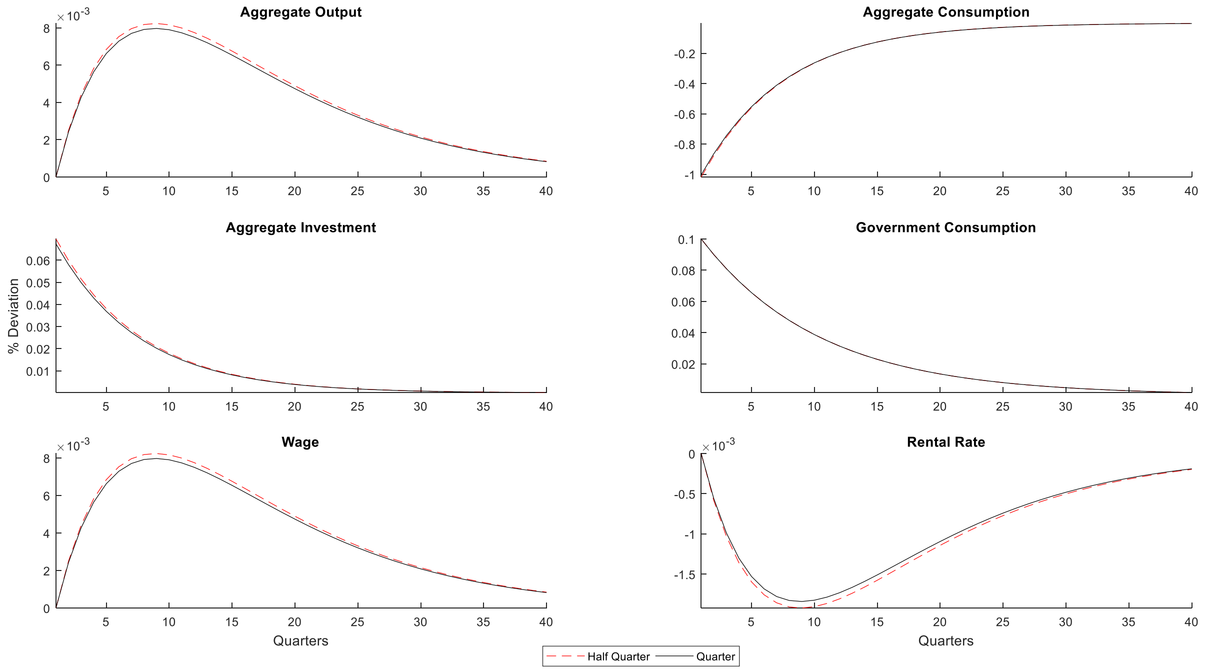

Figure 11: Impulse responses of the aggregate variables to a temporary, positive, 1% shock to government consumption spending when the unemployment duration is either a quarter, or half a quarter.

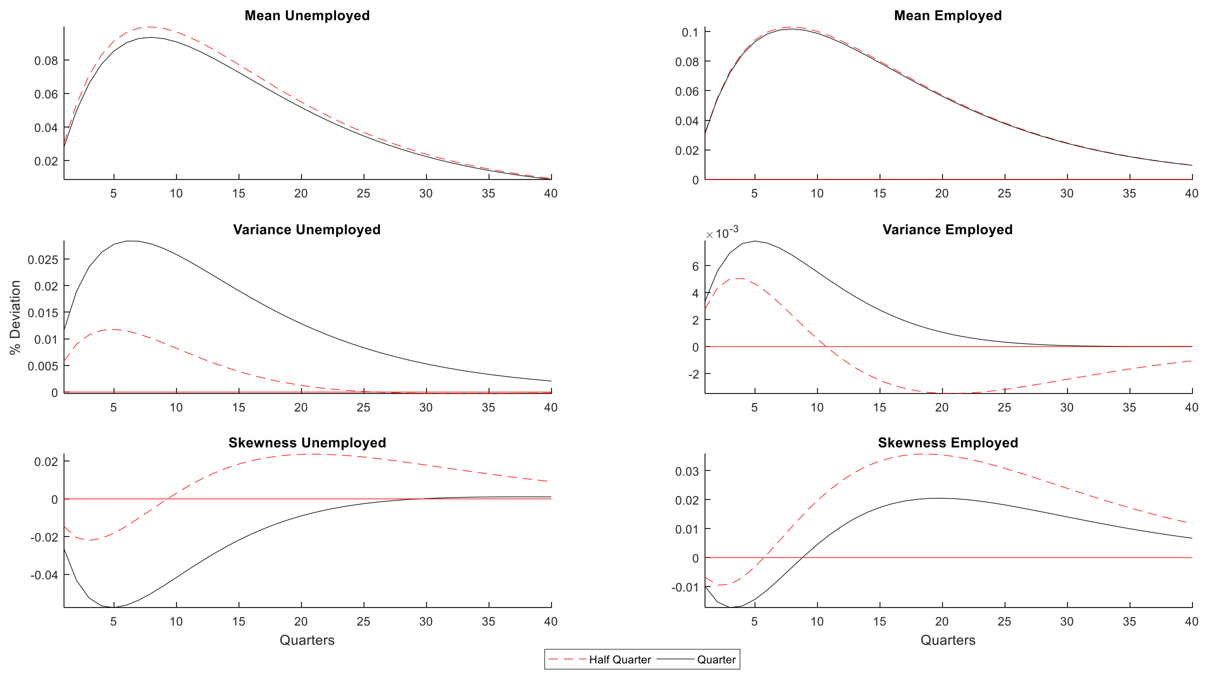

Figure 12: Impulse responses of the aggregate variables to a temporary, positive, 1% shock to government consumption spending when the unemployment duration is either a quarter, or half a quarter.

7. Mixed Frequency VAR

A potential concern with the model that has been discussed in the preceding sections is that it is very stylised. As such, it might not generate accurate predictions for the response of the wealth distribution to government consumption shocks. The following section therefore compares the theoretical model's predictions with empirical evidence. I estimate the response of wealth inequality to government consumption shocks in the US using a Vector Autoregression (VAR) model. The main issue in using a VAR to estimate the relationship between wealth inequality and government spending is that the VAR model assumes uniform sampling frequencies (Qian, 2010). Unfortunately, as a result of limited data availability, the macroeconomic variables are all at a quarterly frequency, whilst the data on wealth inequality is annual.

7.1. Data

The data are for the US economy for the period 1961:Q1-2008Q3. The limits to this range are dictated by two factors. The first is data availability. The second is that the years either side of this range are characterised by abnormal government spending behaviour. Blanchard and Perotti (2002) for example exclude the 1950s from their VAR analysis of US fiscal policy because events such as the Korean War meant that government spending shocks were far larger than in the latter half of the century. At the other end, in quarter four of 2008 the first round of Quantitative Easing by the US Federal Reserve took place. The US economy was also at the Zero-Lower Bound after 2008 which would introduce nonlinearity into the VAR. By keeping the sample between 1961 and 2008 I ensure that there are no missing observations for any of the series, and government consumption shocks are not abnormally large.

The data for GDP, private consumption, Federal Funds Rate, private investment and the CPI come from the Federal Reserve Economic Data (FRED) database. Government spending data is taken from the US Bureau of Economic Analysis. GDP, private consumption, government spending, and private investment are real and seasonally adjusted with 2009 as the base year.[10]

Government consumption is constructed by taking the sum of Federal, State and Local consumption and subtracting Federal defence consumption. The reason for subtracting defence consumption is that this is not the type of spending that is of interest in this paper. Government investment is constructed the same way. Total government spending is then the sum of these parts.

Finding detailed data on wealth inequality is challenging. As Saez and Zucman (2016) note, issues such as 'offshore wealth management… and indirect wealth ownership' make it difficult to accurately measure wealth inequality (p. 524). There is also a lack of data collection by governments. To try and get around these problems Saez and Zucman (2016) use income tax returns and data on household balance sheets to try and estimate how wealth is distributed in the US since 1913. As part of their research they estimate annual wealth shares for the top ten percent, as well as the bottom ninety percent of households. My measure of wealth inequality is constructed from these two wealth shares, and is the ratio of the wealth share of the top ten percent and bottom ninety percent. This is a different measure of wealth inequality to that used in the theoretical model (which uses the variance). This is a limitation of the empirical analysis which means that the results of the VAR should not be taken as definitive proof that the theoretical model is correct. A future avenue for work on the theoretical model should be to try and extract wealth shares from the code and calculate a simulated wealth ratio to compare with the VAR. Despite this limitation, it will still be encouraging to observe how wealth inequality responds to fiscal policy shocks in the data compared to in the model, bearing in mind the different measures.

To remove unit roots I first-difference GDP, private consumption, private investment and government spending. I test for unit roots on Matlab which uses the test and lag selection method by Kwiatkowski et al. (1992). The test starts with a small amount of lags then adds more to look at the sensitivity of the results. I reject the null that output, private consumption, private investment and government consumption are trend stationary, with each variable having a p-value of 0.01. For first-differenced GDP, private consumption and investment I fail to reject the null of stationarity with p-values of 0.1. The test of differenced government

[10] The US GDP deflator with base year 2009 from the FRED database is used to make nominal variables real for those that are not available in real terms online.

consumption fails to reject the null of stationarity when more lags are added. With the data fully described, I now describe the empirical model.

7.2. The Empirical Model

This paper uses the MFVAR developed by Qian (2010), which he calls the varied data sampling (VARDAS) model. There is a vector of variables $\{\boldsymbol{Y}_t\}_{t=1}^{T}$ with a reduced form $VAR(p)$ process,

$$\boldsymbol{Y}_t = \boldsymbol{c} + \sum_{i=1}^{p} \boldsymbol{\phi}_i \boldsymbol{Y}_{t-i} + \varepsilon_t$$

where $\varepsilon_t \sim N(0, \Omega)$ and t is at quarterly frequency. $\boldsymbol{c}$ is a vector of intercepts and $\boldsymbol{\phi}_i$ is a vector of coefficients on the lags of each variable. As Qian explains, $\{\boldsymbol{Y}_t\}_{t=1}^{T}$ is not observable yet because one component is annual (wealth inequality). The VARDAS model deals with this by aggregating the higher frequency (quarterly) series into annual. At this point, $\{\boldsymbol{Y}_t\}_{t=1}^{T}$ becomes observable because all series are now at the same (annual) frequency. At the same time, no information about the quarterly series is lost, because each quarterly observation is used in the annual aggregate series. For the details of this aggregation process see Qian (2010).

7.2.1. Ordering

$\{\boldsymbol{Y}_t\}_{t=1}^{T}$ is a vector consisting of government spending, G_t, gross domestic product, GDP_t, private consumption, C_t, private investment, I_t (all in real terms), the Federal Funds Rate, r_t, and the measure of wealth inequality, $10/90Ratio_t$. The measure of inequality is a ratio of the wealth shares of the top ten percent to the bottom ninety percent. These variables are chosen as they most closely reflect the key variables in the RBC model used previously.

The fiscal policy shock is identified in Qian's VARDAS model and code using the Cholesky decomposition. When using the Cholesky decomposition the ordering of the variables is important. Lin (2006) explains that the decomposition imposes the restriction that causality runs from the 'top variables to the bottom variables' but not the reverse (p. 1). Many papers have attempted to empirically describe the impact of fiscal policy shocks on the economy. They provide a useful guide for how best to order the variables. Ramey (2011), Blanchard and Perotti (2002), Ravn, Schmitt-Grohe and Uribe (2012), and Fisher and Peters (2010) all examine the impact of fiscal policy and in their VARs set the fiscal policy variable first, followed by GDP and private consumption. Ramey (2011) then adds private investment after private consumption.

Following these authors, the baseline MFVAR is ordered $\{\boldsymbol{Y}_t\}_{t=1}^{T} \equiv [G_t, r_t, GDP_t, C_t, I_t,\ 10/90Ratio_t]'$. In the first estimation, government investment and consumption is used as the variable that is shocked. For comparison, government investment and government consumption are then estimated separately. The MFVAR does not contain inflation, which might bias the results. As a robustness check, the MFVAR with government investment and consumption combined is run with inflation included. The impulse responses were very similar to the model when inflation is not included.

7.3. Estimation

7.3.1. Government Investment and Consumption

Figure 13 reports the impulse responses of the variables after a positive, transitory, one standard deviation shock to government investment and consumption spending. Ninety percent confidence bands are displayed, and the frequency of the VAR is quarterly.

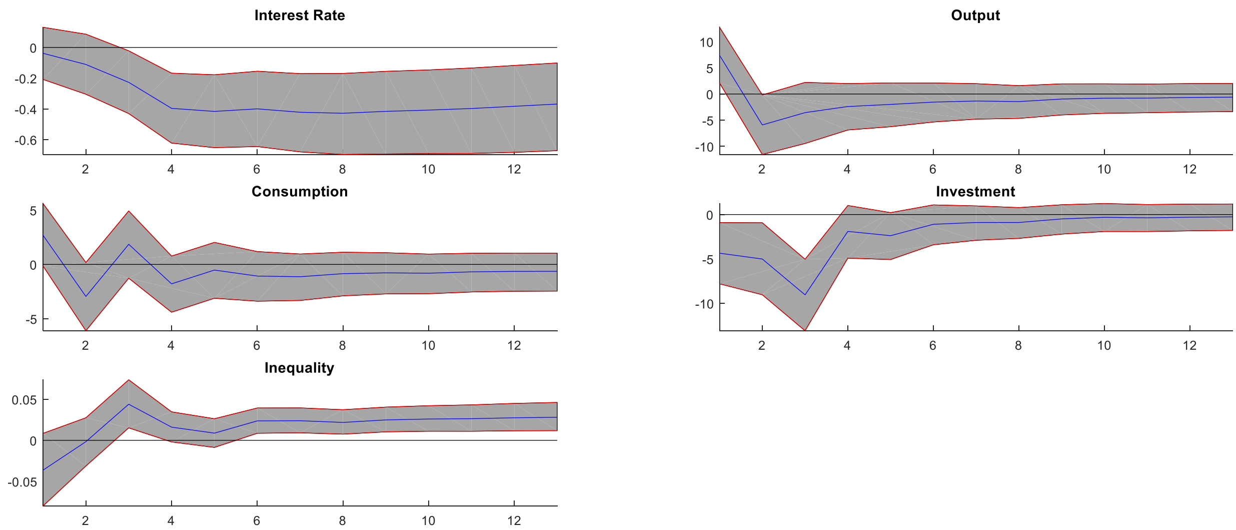

Figure 13: Baseline ordering IRFs. Estimated impulse responses for a positive, temporary, one standard deviation shock to government consumption and investment spending. IRFs reported over years.

The results are broadly consistent with the theoretical model. In response to the government consumption shock, output increases significantly in the first year, before returning to steady-state, as in the theoretical model. The interest rate falls significantly. Wealth inequality increases above steady-state, and consistent with the theoretical model, is very persistent. Wealth inequality does not return to steady-state even twelve quarters after the government spending shock.

Several differences between the model and MFVAR are notable. First is that private consumption has an insignificant response to the shock. In the model, private consumption falls in response to the shock. Second, private investment falls significantly in the MFVAR, whilst it increases in the model due to the higher asset accumulation by households. The response of private consumption in the model was discussed in detail in 5.2.2.: it would be expected to be insignificant or positive (depending on how it is calibrated) if government consumption increased. Therefore an insignificant estimate for the impulse response of private consumption is not inconsistent with the thoery. Ramey (2011) does note that in the empirical literature positive government spending shocks generally have a positive impact on output, a negative impact on private investment, and an uncertain effect on private consumption.

A negative impulse response of private investment does contradict the model. In the theoretical model, an increase in government consumption spending increases household asset holdings, which translates to higher aggregate capital, and therefore higher investment. Interestingly Ma (2017) does find that private investment is crowded out by government spending shocks, however he does interpret government spending as public capital. Blanchard and Perotti (2002) and Mountford and Uhlig (2009) also find that government spending shocks crowd out private investment.

The conflict between the MFVAR and theoretical models' predictions for investment likely stems from two sources. Firstly, the same issues discussed in 5.2.2. Since labour supply is inelastic, and household consumption is a linear combination of private and public consumption (rather than a CES aggregate), there is a large negative response of consumption to the government spending shock. If, however, we observed an increase in private consumption in response to a government spending shock, then household asset holdings would fall to accommodate this, leading to lower aggregate capital, and therefore lower investment.

Secondly, since government investment is included in the baseline regression, when government spending is shocked, government investment increases, which is likely to crowd out private investment.

The empirical results from the MFVAR are consistent with the theoretical model once the response of private consumption, and composition of the fiscal variable are accounted for. The MFVAR's estimated impulse responses are also consistent with the literature. Overall the MFVAR supports the model's prediction that in response to a positive, temporary government consumption shock, wealth inequality will increase.[11]

7.3.2. Separating Government Consumption and Investment

As a robustness check, I separate the components of government spending into consumption and investment, and IRFs for the same government spending shock (See Figures 14 and 15 on next page).

When government consumption is used as the spending variable, the results are broadly similar to the baseline regression, with the exception that output falls rather than increases. When government investment is the spending variable, the results are quite different. Output, consumption and the interest rate all increase significantly. Inequality does not respond significantly to the shock, and investment falls significantly. These results suggest that it is indeed changes in government consumption spending, and not investment, that are driving the changes in wealth inequality. This is encouraging given that the theoretical model interprets government spending as consumption, and not investment spending.

The negative response of private investment when only government consumption is used in the regression highlights that crowding out from government investment does not explain all of the negative response. However, when only using government investment as the fiscal variable, private investment also falls significantly. This is evidence in favour of my argument that in the data it is both crowding out from government investment, and increased consumption spending, that is driving private investment down.

[11] Keeping in mind that the use of different measures of wealth inequality undermines the robustness of this observation.

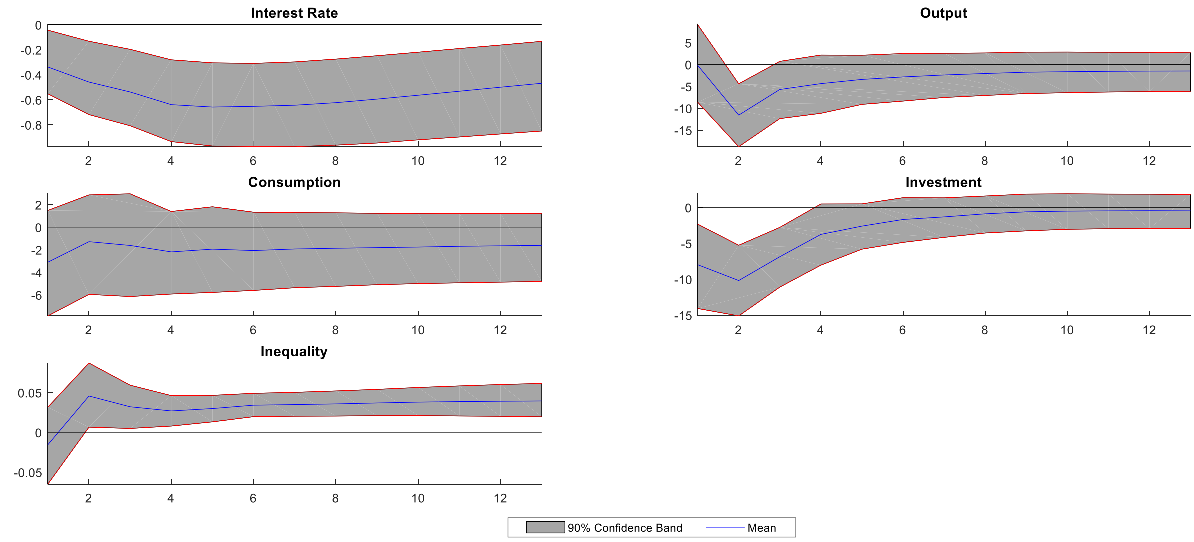

Figure 14: MFVAR with government consumption only. Estimated impulse responses with government consumption spending as the fiscal spending variable. Positive, temporary, one standard deviation shock to government consumption.

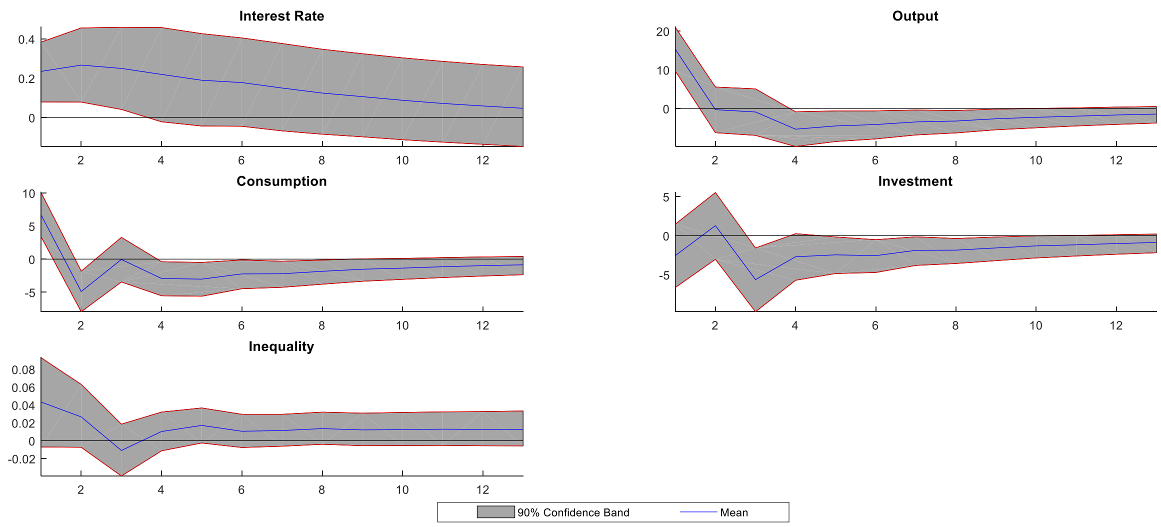

Figure 15: MFVAR with government investment only. Estimated impulse responses with government investment as the fiscal spending variable. Positive, temporary, one standard deviation shock to government consumption.

8. Conclusion

With wealth inequality becoming a staple of academic and political debate, it is increasingly important that we understand the dynamics of how government policy shocks impact the wealth distribution. To the best of my knowledge this is the first paper to use a heterogeneous agent DSGE model to fully characterise how the wealth distribution responds to government consumption shocks at business cycle frequencies.

The literature thus far has focussed on heterogeneity as a way of relaxing assumptions, and has not given much attention to the evolution of wealth/income distributions. The closest paper to mine is Ma (2017) which does look at how the mean of the asset distribution responds to government spending shocks as part of his analysis. However, as discussed earlier, by using the Krusell and Smith (1998) method, he is restricted to only look at the mean. My paper uses the Winberry algorithm to show how the mean, variance and skewness of the distribution evolve, and even create a three-dimensional plot of this evolution over time.

The model shows that in response to a positive, temporary government spending shock, wealth inequality increases amongst both employed and unemployed households. The impulse response of wealth inequality is far more persistent than the shock to government consumption, and can change in size and direction depending on how complementary private and public consumption are. The duration of unemployment spells also has a large impact on how inequality responds to government consumption shocks. These results all highlight that looking at the average impact of a policy is not enough to fully understand its impacts. To properly evaluate how a policy shock might feed through the economy, policy makers need to examine their distributional impacts too.

The paper used a MFVAR to test the robustness of the theoretical model. The data support the model's prediction that an increase in government consumption spending will result in higher wealth inequality.

There are many directions in which future research can go from here. The most obvious would be to increase the complexity of the model so that it is more consistent with the real world. First, one could include an endogenous labour supply and CES aggregate of private and public consumption in the model. As discussed in 5.2.2. this should result in the model generating crowding in of private consumption by public consumption. Second, incorporating distortionary taxation and a government budget constraint would lead government consumption shocks to impact the permanent income of households, generating further changes in their behaviour. The end goal would be to have a heterogeneous agent New Keynesian model such as Kaplan, Moll and Violante (2016), but based on the Winberry (2016a) algorithm. It would be quick to simulate and extend, could be used to fully characterise the distributional impact of policy shocks at business cycle frequencies, and therefore extend our understandings of the dynamics of wealth inequality.

On the empirical side, the simple MFVAR analysis carried out here shows that it is possible to use VARs to empirically test the predictions of heterogeneous agent models. This could be extended by using wealth inequality data from different sources, testing for and accommodating cointegrating relationships, and investigating the use of identification methods in the MFVAR that are more robust than the Cholesky decomposition.

The results from these exercises illustrate firstly that as Krusell and Smith (1998) noted twenty years ago, it is now possible to use general equilibrium models to analyse how inequality, business cycles and policy interact. Secondly, they show that it is important that we do. Without looking at the distributional impacts of government policy, we overlook many important dynamics.

9. References

Alesina, A. and Perotti, R. (1996) Income distribution, political instability, and investment. *European Economic Review*, *40*(6), 1203-1228.

Ambler, S. Bouakez, H. and Cardia, E. (2017) Does the crowding-in effect of public spending on private consumption undermine neoclassical models? *Research in Economics*, 71(3), 399-410.

Anderson, E. Inoue, A. and Rossi, B. (2016) Heterogeneous consumers and fiscal policy shocks. *Journal of Money, Credit and Banking*, *48*(8), 1877-1888.

Aschauer, D.A. and Greenwood, J. (1985) Macroeconomic effects of fiscal policy. In *Carnegie-Rochester Conference Series on Public Policy* 23, 91-138.

Atkinson, A. B. (1970) On the measurement of inequality. *Journal of Economic Theory*, 2(3), 244-263.

Barlevy, G. and Tsiddon, D. (2006) Earnings inequality and the business cycle. *European Economic Review*, *50*(1), 55-89.

Bassetto, M. (2005) Equilibrium and government commitment. *Journal of Economic Theory*, *124*(1), 79-105.

Biswas, S. Chakraborty, I. and Hai, R. (2017) Income inequality, tax policy, and economic growth. *The Economic Journal*, *127*(601), 688-727.

Blanchard, O. and Perotti, R. (2002) An empirical characterization of the dynamic effects of changes in government spending and taxes on output. *Quarterly Journal of Economics*, *117*(4), 1329-1368.

Bouakez, H. and Rebei, N. (2007) Why does private consumption rise after a government spending shock? *Canadian Journal of Economics/Revue canadienne d'économique*, *40*(3), 954-979.

Brown, A. and Wells, G. (2008) Substitution between public and private consumption in Australian states. *Australasian Meeting of the Econometric Society, Wellington, New Zealand*, 9-11.

Canzoneri, M.B. Cumby, R.E. and Diba, B.T. (2007) Euler equations and money market interest rates: A challenge for monetary policy models. *Journal of Monetary Economics*, *54*(7), 1863-1881.

Chakravarty, S.R. (1990) *Ethical social index numbers.* [Online] Berlin, Springer. Available from: https://link.springer.com/book/10.1007%2F978-3-642-75502-6#about [Accessed 7th February 2018].

Christiano, L.J. and Eichenbaum, M. (1992) Current real-business-cycle theories and aggregate labour-market fluctuations. *The American Economic Review*, 82(3), 430-450.

Collins, C. and Hoxie, J. (2015) Billionaire bonanza: The Forbes 400 and the rest of us. [Online] *Institute for Policy Studies*, Available from: https://inequality.org/wp-content/uploads/2017/11/BILLIONAIRE-BONANZA-2017-Embargoed.pdf/ [Accessed 7th February 2018].

Cozzi, M. (2014) Equilibrium heterogeneous-agent models as measurement tools: Some Monte Carlo evidence. *Journal of Economic Dynamics and Control*, *39*, 208-226.

Dalton, H. (1920) The measurement of the inequality of incomes. *The Economic Journal*, 30(119), 348-361.

de Nardi, M. (2015) Quantitative models of wealth inequality: A survey. *National Bureau of Economic Research,* Working Paper w21106, 1-43.

Fève, P. Matheron, J. and Sahue, J.G. (2013) A pitfall with estimated DSGE-based government spending multipliers. *American Economic Journal: Macroeconomics*, *5*(4), 141-178.

Finn, M.G. (1998) Cyclical effects of government's employment and goods purchases. *International Economic Review*, 39(3), 635-657.

Fiorito, R. and Kollintzas, T. (2004) Public goods, merit goods, and the relation between private and government consumption. *European Economic Review*, *48*(6), 1367-1398.

Fisher, J. D. and Peters, R. (2010) Using stock returns to identify government spending shocks. *The Economic Journal*, *120*(544), 414-436.

Galí, J. López-Salido, J.D. and Vallés, J. (2007) Understanding the effects of government spending on consumption. *Journal of the European Economic Association*, *5*(1), 227-270.

Ganelli, G. and Tervala, J. (2009) Can government spending increase private consumption? The role of complementarity. *Economics Letters*, *103*(1), 5-7.

Garcia-Penalosa, C. and Turnovsky, S.J. (2007) Growth, income inequality, and fiscal policy: What are the relevant trade-offs? *Journal of Money, Credit and Banking*, *39*(2-3), 369-394.

Gollin, D. (2002). Getting income shares right. *Journal of political Economy*, *110*(2), 458-474.

Gornemann, N. Kuester, K. and Nakajima, M. (2012) Monetary policy with heterogeneous agents. *FRB of Philadelphia, Working Paper No. 12-21*, 1-48.

Heathcote, J. (2005) Fiscal policy with heterogeneous agents and incomplete markets. *The Review of Economic Studies*, *72*(1), 161-188.

Heathcote, J. Storesletten, K. and Violante, G.L. (2009) Quantitative macroeconomics with heterogeneous households. *Annual Review of Economics*, *1*(1), 319-354.

Heshmati, A. (2004) Inequalities and their measurement. *MTT Economic Research and IZA Bonn*, Discussion Paper No. 1219, 1-17.

Hoover, G.A. Giedeman, D.C. and Dibooglu, S. (2009) Income inequality and the business cycle: A threshold cointegration approach. *Economic Systems*, *33*(3), 278-292.

Imrohoroğlu, A. (1989). Cost of business cycles with indivisibilities and liquidity constraints. *Journal of Political Economy*, *97*(6), 1364-1383.

Imrohoroğlu, A. (2008) Welfare costs of business cycles. In Durlauf, S.N. and Blume, L. (eds) *The New Palgrave Dictionary of Economics*. [Online] Basingstoke, UK, Palgrave Macmillan, 1-7. Available from: http://www-bcf.usc.edu/~aimrohor/My%20papers/Palgrave%202008.pdf [Accessed 7th February 2018].

Kaplan, G. Moll, B. and Violante, G.L. (2016) Monetary policy according to HANK. *National Bureau of Economic Research, Working Paper* w21897, 1-61.

Karras, G. (1994) Government spending and private consumption: some international evidence. *Journal of Money, Credit and Banking*, *26*(1), 9-22.

Keane, M.P. and Wolpin, K.I. (1997) The career decisions of young men. *Journal of political Economy*, *105*(3), 473-522.

Krusell, P. Mukoyama, T. Şahin, A. and Smith, A.A. (2009) Revisiting the welfare effects of eliminating business cycles. *Review of Economic Dynamics*, *12*(3), 393-404.

Krusell, P. and Smith, Jr, A.A. (1998) Income and wealth heterogeneity in the macroeconomy. *Journal of political Economy*, *106*(5), 867-896.

Kwiatkowski, D. Phillips, P. C. Schmidt, P. and Shin, Y. (1992) Testing the null hypothesis of stationarity against the alternative of a unit root: How sure are we that economic time series have a unit root?. *Journal of econometrics*, *54*(1-3), 159-178.

Leeper, E.M. Walker, T.B. and Yang, S.C.S. (2010) Government investment and fiscal stimulus. *Journal of Monetary Economics*, *57*(8), 1000-1012.

Le Grand, F. and Ragot, X. (2017) Optimal fiscal policy with heterogeneous agents and aggregate shocks. *Society for Economic Dynamics*, Meeting Paper 969, 1-50.

Lin, J.L. (2006) Teaching notes on impulse response function and structural VAR. [Online] *Institute of Economics, Academia Sinica, Department of Economics, National Chengchi University*. Available from: http://faculty.ndhu.edu.tw/~jlin/files/impulse.pdf. [Accessed 23rd January 2018].

Ma, Eunseong (2017) The heterogeneous response of consumption between poor and rich to government spending shocks. Working Paper. [Online] Available from: https://docs.google.com/a/tamu.edu/viewer?a=v&pid=sites&srcid=dGFtdS5lZHV8ZXVuc2Vvbmd8Z3g6NThjM2I5ZDhlMGZmY2Q1Zg [Accessed 20th January 2018].

Mayer, S.E. Lopoo, L.M. and Groves, L.H. (2016) Government spending and the distribution of economic growth. *Southern Economic Journal*, *83*(2), 399-415.

McGrattan, E.R. (1994) The macroeconomic effects of distortionary taxation. *Journal of Monetary Economics*, *33*(3), 573-601.

Misra, K. and Surico, P. (2014) Consumption, income changes, and heterogeneity: Evidence from two fiscal stimulus programs. *American Economic Journal: Macroeconomics*, *6*(4), 84-106.

Morin, A. (2013) Wage dispersion over the business cycle. *European University Institute*, *EUI Working Paper MWP 2013/28*, 1-38.

Perotti, R (2007) In search of the transmission mechanism of fiscal policy. In: Acemoglu, D. Rogoff, K and Woodford, M. (eds.) *NBER Macroeconomics Annual 2007*, University of Chicago Press, 169-226.

Piketty, T. and Goldhammer, A. (2014). *Capital in the Twenty-First Century*. [Online] Harvard University Press. Available from: http://www.jstor.org.ezproxy.otago.ac.nz/stable/j.ctt6wpqbc.

Qian, H. (2010) Vector autoregression with varied frequency data. *Munich Personal RePEc Archive*, 34682, 1-36.

Ramey, V.A. (2011) Identifying government spending shocks: it's all in the timing. *The Quarterly Journal of Economics*, *126*(1), 1-50.

Ravn, M.O. Schmitt-Grohe, S. and Uribe, M. (2012) Consumption, government spending and the real exchange rate. *Journal of Monetary Economics*, *59(3)*, 215-234.

Reiter, M. (2009) Solving heterogeneous-agent models by projection and perturbation. *Journal of Economic Dynamics and Control*, 33, 649-665.

Romer, C.D. and Romer, D.H. (2016) Transfer payments and the macroeconomy: The effects of social security benefit increases, 1952-1991. *American Economic Journal: Macroeconomics, 8(4),* 1-42.

Ruiz, S. V. L. Peralta-Alva, A. and Puy, D. (2017). Macroeconomic and distributional effects of personal income tax reforms; A heterogenous agent model approach for the US. *International Monetary Fund*. *Working Paper No. 17/192*. International Monetary Fund, 1-32.

Saez, E. and Zucman, G. (2016) Wealth inequality in the United States since 1913: Evidence from capitalized income tax data. *The Quarterly Journal of Economics*, *131*(2), 519-578.

Schulhofer-Wohl, S. (2008) Heterogeneous risk preferences and the welfare cost of business cycles. *Review of Economic Dynamics*, *11*(4), 761-780.

Song, J. Price, D.J. Guvenen, F. Bloom, N. and Von Wachter, T. (2015) Firming up inequality. *National Bureau of Economic Research.* Working Paper w21199,1-37.

Terry, S.J. (2017) Alternative methods for solving heterogeneous firm models. *Journal of Money, Credit and Banking*, *49*(6), 1081-1111.

Winberry, T. (2016a) A toolbox for solving and estimating heterogeneous agent macro models. *Quantitative Economics.* Work in progress, 1-36.

Winberry, T. (2016b) User guide for ‘A toolbox for solving and estimating heterogeneous agent macro models.’ [Online] Available:

http://faculty.chicagobooth.edu/thomas.winberry/research/index.html. [Accessed 15th Jan 2018].

Yogo, M. (2004) Estimating the elasticity of intertemporal substitution when instruments are weak. *The Review of Economics and Statistics*, *86*(3), 797-810.