Expectations and Social Influence
Dynamics in Economics

ExSIDE Working Paper Series

No. 25-2020
May 2020

From CATS to CAOS: Fiscal multipliers and agents' expectations in a macroeconomic agend-based model

Severin Reissl

www.exside-itn.eu

Marie Sklodowska-Curie Innovative Training Network "Expectations and Social Influence Dynamics in Economics (ExSIDE)"
This project has received funding from the European Union Horizon 2020 Research and Innovation action under grant agreement No 721846.

From CATS to CAOS:
Fiscal multipliers and agents' expectations in a macroeconomic agent-based model

Severin Reissl*
Università Cattolica del Sacro Cuore & Universität Bielefeld

May 18, 2020

Abstract

This paper uses a macroeconomic agent-based model building on Delli Gatti et al. (2011) to investigate the influence of agents' expectations and consumption choices on government expenditure multipliers. Following a thorough investigation of the size of the multiplier in the pre-existing baseline model, a modification is introduced, allowing agents to engage in intertemporal optimisation of consumption subject to a budget constraint which is based on estimates of future income. Compared to the baseline, the fiscal multiplier is strongly affected by this alternative consumption behaviour, becoming significantly smaller. In a further step, agents' beliefs about the effects of government expenditure shocks are explicitly introduced. In the case of exogenously imposed beliefs coupled either with adaptation of individual beliefs or switching behaviour between different types of beliefs, it is shown that both optimistic and pessimistic expectations can be temporarily self-fulfilling and either increase or decrease the value of the multiplier. Both forms of belief dynamics also allow for the incorporation of announcement effects of fiscal policy. In a final experiment, agents are allowed to engage in least-squares learning in order to gain an estimate of the effect of government expenditure shocks on future income. It is shown that under least squares learning, beliefs are 'rational' insofar as they lead to broadly correct predictions on average. The paper hence contributes to addressing aspects of the Lucas critique as applied to macro-ABMs, since agents react systematically (and reasonably) to announcements of changes in fiscal policy.

Keywords: Agent-based models, Expectations, Learning, Fiscal policy
JEL-Classification: C63, E62, D84, E71

*Email: severindavid.reissl@unicatt.it; This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 721846 (ExSIDE ITN). The author would like to thank Herbert Dawid, Domenico Delli Gatti, Alex Grimaud, Jakob Grazzini, Domenico Massaro, as well as conference participants in London, Berlin and Darmstadt for useful hints and comments. The usual disclaimer applies.

1 Introduction

Over the past decade, various different macroeconomic agent-based modelling (ABM) frameworks have been developed and refined by different research groups, and there is a growing literature applying these frameworks to a range of policy questions. Despite these advances, the role of agents' beliefs and expectations formation on simulation results, and especially their impact on the outcomes of policy experiments, is to date not well-investigated in the ABM literature (cf. Dosi et al., 2017a). While agents in conventional macroeconomic models are arguably too rational and forward-looking, agents in ABMs frequently form expectations in a way which implies behaviour that seems excessively naive. While forward-looking behaviour in the way it is inserted into conventional models typically cannot be implemented in an agent-based context, it nevertheless appears desirable to work towards the incorporation of somewhat more sophisticated and in some sense forward-looking behavioural rules. This is particularly so since, despite their various advantages over more conventional and aggregative modelling paradigms, ABMs are potentially susceptible to the Lucas critique (Lucas, 1976). Although, as argued by Haldane and Turrell (2018, p. 2), "no model is fully Lucas critique-proof; it is a matter of degree",[1] it nevertheless appears reasonable to argue that agents in a macroeconomic model should be capable of reacting to policies and policy announcements in a systematic and sensible manner. Moreover, the construction of agent-based frameworks which are relatively similar to their mainstream counterparts, or even attempts at merging the paradigms (cf. Assenza and Delli Gatti, 2013; Gobbi and Grazzini, 2019) can be helpful in aiding aiding comparability of results and in gaining a better understanding of the fundamental differences between modelling approaches.

The focus of the present paper is on the consumption behaviour of households and its effect on government expenditure multipliers. Beginning from a variant of the 'complex adaptive trivial systems' (CATS) family of models, namely the empirically estimated macroeconomic ABM framework presented by Delli Gatti and Grazzini (2019) (based on Assenza et al., 2018b), the paper provides a thorough investigation of the government expenditure multiplier. It is shown that in the unmodified framework this multiplier is significantly greater than one and that the model gives rise to regime-dependent multipliers. In a next step, the consumption behaviour of households is modified through the implementation of a framework in which households intertemporally optimise consumption subject to an to an estimated budget constraint, giving rise to what may, somewhat provocatively, be termed a 'complex adaptive optimising system' (CAOS). This modification is found to lead to a significant improvement in the welfare households derive from consumption. It is also demonstrated that under the alternative consumption behaviour, government ex-

[1] Indeed, even the models which according to Lucas and his followers should be robust to the critique are in certain ways vulnerable to it (see e.g. Kirman, 1992; Altissimo et al., 2002).

penditure multipliers are significantly smaller than in the baseline model, although they are still positive and still exhibit regime-dependence, particularly due to the presence of frictions on the goods market and liquidity-constrained households. Extending the modified consumption behaviour, heterogeneous and ex-ante exogenous beliefs of households about the effects of government expenditure shocks on their budget constraints are introduced. These beliefs may change endogenously as agents learn either through adaptation or switching behaviour following an actual shock. In both cases it is shown that beliefs may be temporarily self-fulfilling in that their presence may increase or reduce the magnitude of the fiscal multiplier. In a final extension, households use a least squares learning algorithm in order to form an estimate of the effect of fiscal policy shocks. The addition of expectation and learning dynamics allows for an incorporation of announcement effects of fiscal policy whereby agents react systematically to announcements of future changes in government expenditure. Moreover, it is shown that the expectations resulting from least-squares learning appear to be broadly correct on average.

The paper contributes to the literature on agent-based macroeconomic models in three main ways. Firstly, it provides a thorough assessment of the government expenditure multiplier and the factors affecting it in a canonical framework. Such an analysis is novel to the macro-ABM literature. Secondly, the paper shows how more forward-looking consumption behaviour closer to common mainstream formulations can be introduced into a relatively simple macro-ABM for comparative purposes. Finally, the paper contributes to a partial addressal of the Lucas Critique in an agent-based framework through the incorporation of learning and announcement effects of fiscal policy, enabling agents to react systematically to policy announcements.

The paper is structured as follows: section 2 provides a motivation of the present paper, placing it within the wider literature. Section 3 provides a brief overview of the pre-existing model used. Section 4 explores the effects of government expenditure shocks in the baseline model. Section 5 introduces the modification in households' consumption behaviour and investigates its effect on the government expenditure multiplier both in the absence and presence of beliefs. Section 6 presents the least squares learning experiment. Section 7 summarises and discusses the results of the previously shown exercises. Section 8 concludes. Parameter values can be found in appendix A. Appendix B tests some assumptions underlying the alternative consumption model. The code necessary to replicate all simulations shown below is available at `https://github.com/SReissl/CATS`.

2 Policy, expectations and learning in agent-based models

Dawid and Delli Gatti (2018) provide an excellent survey of the major existing macroeconomic ABM frameworks including the Keynes + Schumpeter model (Dosi et al., 2010), the Eurace models (Cincotti et al., 2010; Dawid et al., 2012) as well as the CATS family of models (Delli Gatti et al., 2011; Assenza et al., 2015), a version of which is utilised in the present study. They compare the institutional features and behavioural assumptions of the respective baseline models while identifying a range of common elements across the existing model families. While the agent-based approach allows for a realistic depiction of macroeconomic dynamics and their micro- or mesoeconomic origins, it also makes macroeconomic ABMs very complex and their output difficult to analyse. One important area of work has therefore been to address issues of transparency and reproducibility (Dawid et al., 2019), along with attempts to establish techniques and standards for the empirical validation of models, a process which is still ongoing (Windrum et al., 2007; Guerini and Moneta, 2017). Related to this, various authors have worked on possible ways to empirically calibrate or estimate macroeconomic agent-based models (Gilli and Winker, 2003; Grazzini and Richiardi, 2015; Grazzini et al., 2017) and on techniques to deal with the typically very high-dimensional parameter spaces and the high computational cost involved (Salle and Yildizoglu, 2014; Barde and van der Hoog, 2017; Lamperti et al., 2018).

At the same time, the major macroeconomic ABMs surveyed by Dawid and Delli Gatti (2018) are also being applied to an increasing range of policy questions (see Fagiolo and Roventini, 2017 for a survey). For instance, versions of the framework being used in the present paper have been applied to the study of monetary policy as well as the effects of prudential regulation. Delli Gatti and Desiderio (2015) show that a central bank following a Taylor rule can significantly abate the effects of a negative productivity shock compared to a scenario in which the central bank rate is fixed. Assenza et al. (2018a) show that a Taylor rule can decrease the incidence of crises in simulated time-series and that a maximum leverage ratio for banks can also have a positive effect on macroeconomic stability. Various other aspects of monetary policy, including the appropriate choice of inflation targets (Ashraf et al., 2016), variations in the central bank's mandate (Dosi et al., 2015; Chiarella and Di Guilmi, 2017) and the interactions between or joint effects of monetary policy and prudential regulation (Krug, 2018) have been addressed using other models. Considerations of appropriate financial regulation and prudential policy to decrease the risk of serious financial crises have also been explored by Popoyan et al. (2017) and van der Hoog and Dawid (2019). Other areas in which macroeconomic ABMs have been applied to policy questions include labour markets and income distribution, with various studies comparing the results emerging from different degrees of labour market flexibility (e.g. Seppecher, 2012; Dosi et al., 2017b), analysing regional disparities in labour market regimes (Dawid et al., 2018a), or exploring possible links between inequality and economic growth (Caiani et al., 2019).

As one would expect, fiscal policy, which is the focus of the present paper, has also featured prominently in the ABM literature. Assenza et al. (2018b) use the same model which is utilised here to compare different fiscal policy regimes (including one in which a Stability and Growth Pact-type rule is implemented) showing that in many cases, the model will converge to an unemployment rate consistent with a balanced government budget, allowing policy-makers to target quasi-equilibrium unemployment rates through varying fiscal policy parameters. Indeed, fiscal rules such as the Stability and Growth Pact or the Fiscal Compact have been examined in a range of studies (e.g. Dosi et al., 2015; Teglio et al., 2019) which generally conclude that such rules which artificially constrain fiscal manoeuvring space have negative effects on economic performance. Caiani et al. (2018) study the effects of changes in fiscal rules in the setting of a monetary union, concluding that more restrictive rules will have a negative effect on GDP, employment and productivity. Dawid et al. (2018b) similarly consider a context akin to that of the European Monetary Union, showing that fiscal transfers from the core can alleviate crises in the periphery region. Harting (2015) focuses on the effects of different forms of fiscal policy during economic downturns, showing that a fiscal policy which promotes technological innovation not only augments economic growth but also stabilises cyclical fluctuations. Napoletano et al. (2017) present what is to my knowledge the only existing paper calculating fiscal multipliers in a macroeconomic ABM. Building a strongly simplified ABM of an endowment economy with time-varying financial constraints, they examine the effects of a bankruptcy shock under different fiscal policy regimes. They find that fiscal multipliers in their model are state-dependent and can comfortably exceed one depending on the exact specification of the fiscal policy intervention.

In general it can hence be said that fiscal policy tends to be highly effective in existing macroeconomic ABMs. One reason for these results is that agents in ABMs, in contrast to most of their counterparts in conventional macroeconomic models, tend to act according to heuristics and rules of thumb based on bounded rationality. Expectations are typically formed using naive, adaptive or extrapolative rules and agents' planning horizons are rather limited. The heuristics used by agents may be quite sophisticated (as e.g. some which can be found in Dawid et al., 2012) or the model may purposely rely on fairly simplistic behavioural rules (as tends to be the case e.g. with models building on Delli Gatti et al., 2011) but in all cases the implied degree of rationality is below that which can be found in dynamic general equilibrium models. One reason for this is certainly methodological; even the simplest macroeconomic ABMs are typically too complex to allow for a rational expectations solution to be found. On the other hand, the use of bounded rationality and heuristics is also frequently justified on theoretical grounds, building on the arguments of authors such as Simon (1982) or Gigerenzer (2008), the idea being that simple heuristics may be quite successful and 'rational' in complex environments - a result which does appear to bear out e.g. in the analysis of Dosi et al. (2017a) and the simulations conducted in chapter **??**. Nevertheless, the use of behavioural rules which do not allow agents to react systematically and in a forward-looking manner to policy interventions does make macroeconomic ABMs vulnerable to the Lucas critique.[2] While an implementation of forward-looking and rational behaviour in the same sense as it is used in conventional models will not be possible in most macroeconomic ABMs, it nevertheless appears sensible to allow agents to plan and pro-actively adapt their behaviour to what they believe the effects of policy may be. One way to do so is by allowing agents to learn about their environment and the effects of policy in some fashion and to let them endogenously adapt their behaviour as necessary.[3] In a series of papers, Salle and co-authors (Salle et al., 2013; Salle, 2015; Salle et al., 2019) emphasise the themes of expectations, learning and adaptation in the context of policies and policy changes, using a macroeconomic ABM framework designed to be comparable to a conventional New Keynesian macroeconomic model. The model includes feedback channels for inflation expectations on households' reservation wages and consumption propensities, and agents are engaged in learning behaviour to update their behavioural rules. This allows the authors to examine the effects of different central bank communication strategies for managing inflation expectations and the associated impacts on the effectiveness of monetary policy. Salle (2015) places a particular focus on the issue of robustness, letting households use artificial neural networks to form inflation expectations based on data generated by the model and information communicated by the central bank, allowing them to adapt endogenously to changes in central bank policy. This series of papers is close in spirit to the present work, the latter part of which will focus on allowing agents to form expectations about the impacts of fiscal policy and to adapt their behaviour accordingly. In contrast to the aforementioned works, the present paper makes use of an existing, fairly well-known framework rather than a purpose-built model and shows a possible way to bring the behaviour of agents as close as possible to that of their counterparts in conventional models. This in turn also provides for an improved understanding of the role of heuristics in driving policy effects.

3 Model overview

The model I choose to carry out the analysis is the one presented by Delli Gatti and Grazzini (2019), who estimate the model of Assenza et al. (2018b) using the Bayesian estimation

[2]As already pointed out in the introduction, this statement does not imply that conventional models are immune to the Lucas critique in a broad sense. One particular weakness of such models is that the representative agent may not be robust to policy interventions, potentially invalidating the results of welfare analyses (Jerison, 1984; Geweke, 1985).

[3]Indeed, the use of learning algorithms is closely linked to the complexity and agent-based economics literature, and various techniques such as genetic algorithms (Dawid, 1999), evolutionary learning (Arifovic, 2000) and classifier systems (Holland, 1975) have been proposed as possible ways to model learning in an ABM. In the more conventional literature, least squares learning algorithms as outlined by Evans and Honkapohja (2001) have been widely applied.

algorithm for macroeconomic ABMs outlined in Grazzini et al. (2017).[4] Parameter values for the baseline are set equal to those calibrated and estimated by Delli Gatti and Grazzini (2019). The model is one of an economy with no long-term growth and features five sectors, namely households (capitalists/firm owners and workers), consumption goods firms, capital goods firms, a public sector and a single, representative bank. Figure 1 provides an overview of the model in terms of sectoral balance sheets and transactions, while figure 2 depicts the different markets on which agents interact in the model.

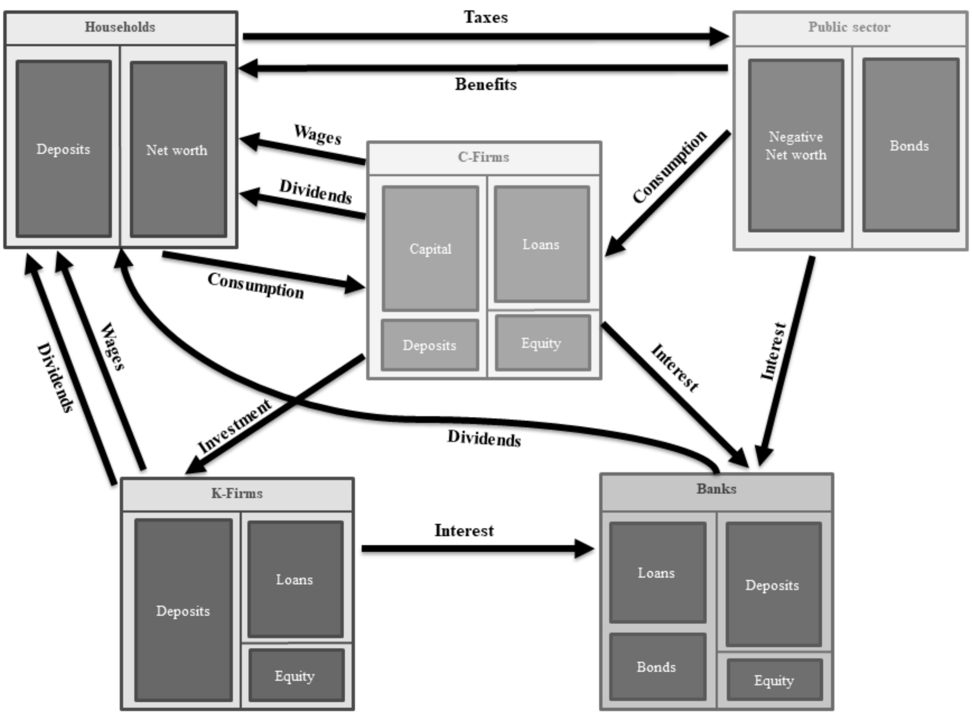

Figure 1: Balance sheet and transactions overview

Households, as indicated above, can be divided into workers and capitalists, the latter of which own the consumption and capital goods firms (there is one capitalist for each firm). Capitalist households also jointly own the single bank in the model. Both types of households hold unremunerated deposits with the banking sector as their only assets and by assumption do not borrow for any purpose.[5] Workers inelastically offer one unit of labour to firms at a uniform wage determined by an aggregate Phillips-curve type equation with downward stickiness. If they are unemployed, they receive a certain fraction of the current wage in benefits and visit a fixed number of randomly drawn firms in each period looking for a job. Wage income is taxed at a fixed rate. Capitalists receive dividend

[4]For a more detailed summary of the model than can be provided below, the reader is referred to Assenza et al. (2018b) and Assenza et al. (2015).

[5]This assumption will become important below since it implies that households may be liquidity-constrained in their consumption decisions.

income from the firms they own as well as from the bank, in the case that profits are positive. Dividend income is not taxed. Each household determines a level of desired consumption using a weighted average of past incomes along with a fixed fraction of its current wealth. It then visits a fixed number of randomly drawn firms and demands consumption goods, starting from the cheapest firm.

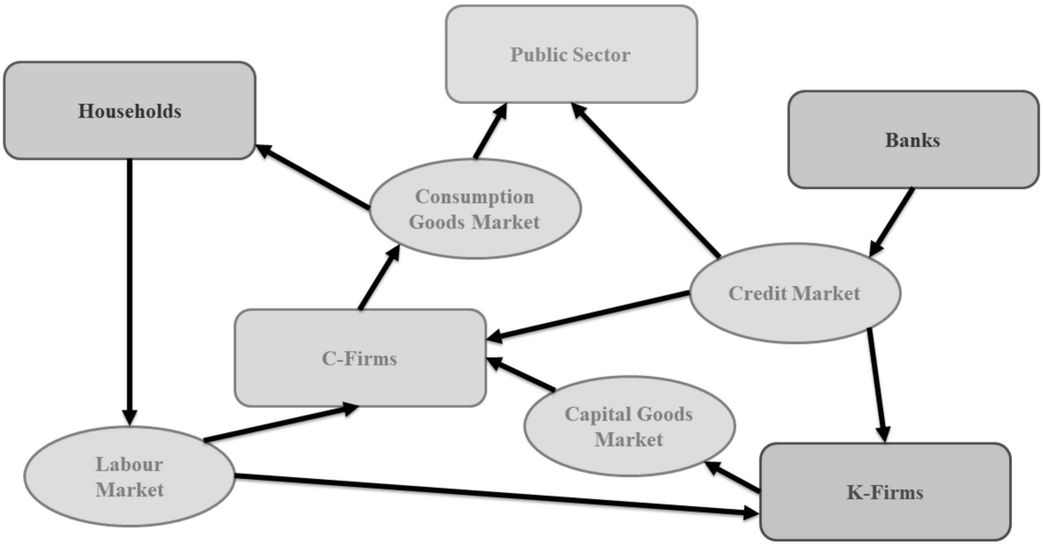

Figure 2: Market interaction overview

C-firms produce a perishable consumption good using a Leontieff production function. They form expectations about demand and set their prices using an adaptive process based on their current relative price as well as past excess demand or supply as shown in figure 3. For instance, a firm in quadrant a, which has experienced excess demand and the previous price of which is below the average price, will increase its price. A firm in quadrant c with a higher than average price and excess demand will increase its production. Firms hire or fire workers based on their current production plans and existing labour force. They demand capital goods based on essentially Kaleckian investment behaviour, reacting to trends in a weighted average of past values of capacity utilisation in order to keep utilisation close to an exogenous target. C-firms are allowed to undertake investment at stochastic intervals and, as on all other markets, each C-firm visits a fixed number of randomly drawn K-firms to shop for capital goods. C-firms demand loans from the banking sector if their liquidity is insufficient to finance production and investment. Liquidity reserves are held in the form of deposits.

K-firms are highly similar to C-firms, the major differences being that they produce using a labour-only technology (and hence do not invest) and that their output is not instantly perishable but stored as a gradually depreciating inventory. If a C-firm's or K-firm's equity turns negative, it is declared bankrupt. A new firm then enters the market (new C-firms receive the capital stock of the bankrupt firm) and receives an injection of equity and liquidity from the capitalist household which owns it.

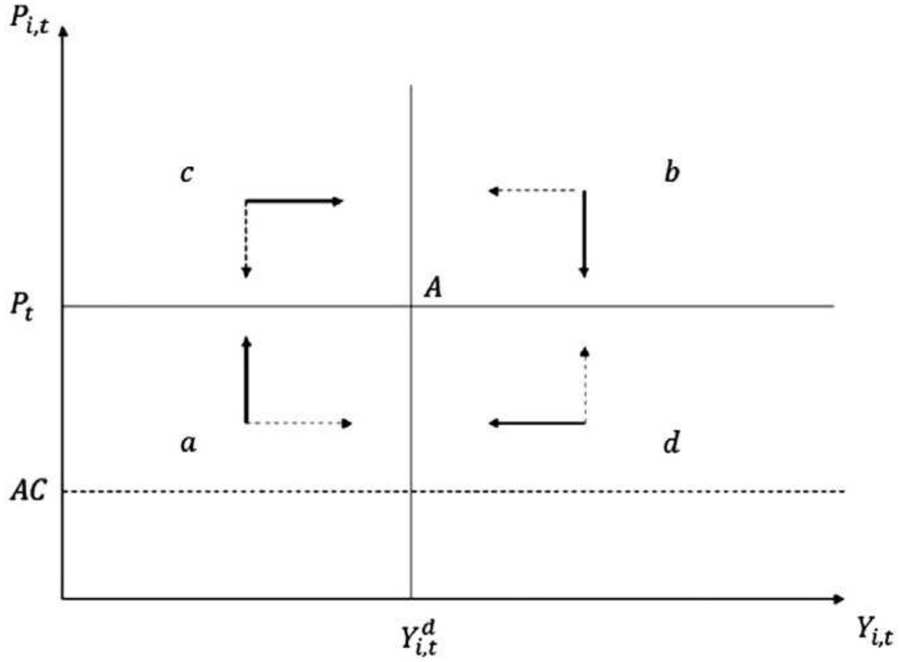

Figure 3: Firms' price and quantity adjustment (reproduced from Assenza et al., 2015, p. 12)

The **Government** levies a constant rate tax on labour income and pays unemployed workers a benefit equal to a fixed fraction of the current wage rate. Deficits are covered by the issuance of bonds at a fixed risk-free interest rate. It is assumed that all bonds are purchased by the bank. In the baseline, the government does not undertake any spending on firms' output. Such spending will be introduced as a fiscal shock below.

The **Bank** in the model takes deposits from firms and households and makes loans to the firm sector. In addition it purchases all bonds issued by the government. The bank provides all loans demanded by firms provided that the total exposure of the bank to any individual firm remains below a given fraction of its equity. This exposure is calculated based on a perceived bankruptcy probability of the firm, which is an increasing function of its leverage. The bankruptcy probability also feeds into the interest rate offered by the bank to each borrowing firm, with more risky firms paying a higher rate, i.e. a higher mark-up over the risk-free rate (see Assenza et al., 2015, 2018b for details). Importantly for the present paper, it should be noted that the model *does not* feature monetary policy as there is an exogenous and fixed risk-free interest rate, such that there is no interaction between monetary and fiscal policy.

Overall, model dynamics are strongly driven by frictions in goods, labour and credit markets stemming from both incomplete and imperfect information of agents. Firms will frequently experience periods of excess demand or supply, they may be unable to undertake planned production or investment due to capacity constraints, matching failures on the labour market or credit rationing, and consumers may be unable to fully exhaust their consumption budget due to the limited number of firms they visit. Nevertheless, as is the case with all other existing major macro-ABM frameworks, the model is essentially demand-driven, although there exist a range of potential supply-side constraints including limited productive capacity and a finite labour force. The fundamentally demand-driven nature of the model unsurprisingly gives rise to a positive macroeconomic effect of expansionary government expenditure shocks, as is outlined in the next section. Macroeconomic dynamics are characterised by persistent and irregular fluctuations of GDP and its components around a long-term stochastic quasi-steady state. In order to reserve space for the thorough analysis of the effects of government expenditure shocks carried out below, the reader is referred to Assenza et al. (2015), Assenza et al. (2018b) and Delli Gatti and Grazzini (2019) for an extensive presentation of simulations as well as validation exercises of the model.

4 Fiscal policy shocks in the baseline model

While the analysis presented by Assenza et al. (2018b) also focuses on changes in taxes and transfer payments in addition to government consumption expenditure, and does not investigate the size of fiscal multipliers, my analysis solely considers the introduction of government expenditure on consumption goods with a focus on the size of the resulting multiplier. This focus on a single policy tool enables me to carry out a thorough analysis, running a range of different experiments. Rather than looking at different policy regimes (e.g. the combination of a certain tax rate, together with a certain replacement ratio and some *permanent*, exogenous level of government spending on consumption goods for the entire duration of a simulation), as do Assenza et al. (2018b), I consider only *increases* in government expenditure from the baseline level of 0, for a *single period* (one period being one quarter), at a certain point during the simulation. Unless otherwise stated, the additional government demand for consumption goods is always equal to 10% of real GDP in the period prior to the shock.

Government expenditure on consumption goods enters the model in a simple fashion. As with all AB macro models, each simulation period proceeds in sequential fashion, with a series of distinct events taking place one after the other rather than simultaneously. At the beginning of the period, C-firms plan their production and it is assumed that in the period in which the shock is introduced, C-firms which wish to adjust their production relative to the previous period (in accordance with figure 3) take into account the government demand they will receive, as calculated by equation (1) below. Government expenditure on consumption goods is introduced after the closing of the consumption goods market on which C-firms and households interact. A desired level of real government consumption g_d is distributed to each C-firm according to a share calculated as

$$s_g^f = \frac{R^f}{\sum_{f=1}^{F} R^f} \tag{1}$$

where R^f is the (previous) revenue of firm f. The government then purchases from each firm f the minimum between its desired expenditure $s_g^f g_d$ and the remaining output produced earlier in the period which has not already been sold to households. The effect of government expenditure shocks is analysed on the one hand through the calculation of cumulative multipliers and through the construction of 'robust impulse-response functions' similar to those proposed by Delli Gatti and Grazzini (2019). The cumulative multiplier at time t is defined as the cumulative change in output relative to the baseline divided by the *actual* government consumption expenditure induced by the shock (which may differ from the desired government expenditure). In Delli Gatti and Grazzini (2019), the robust IRFs are constructed by calculating and plotting the average response (in percentage terms relative to the baseline) of a macro time-series generated by the model (e.g. GDP) to a shock across all Monte Carlo repetitions. Since the model displays a strong tendency to return to a stochastic stationary state practically identical to the previous one following transitory shocks, the construction of such functions gives useful information about the magnitude and duration of the effects of temporary fiscal policy shocks.

Since the shocks considered in the present paper are not persistent and, at least compared to those shown by Delli Gatti and Grazzini (2019), relatively small, their effects are at times difficult to discern from looking at Monte Carlo averages only. I therefore draw on Dawid et al. (2018a) in using spline regressions (see e.g. Ruppert et al., 2003) in order to smooth out and depict the effects of the fiscal shocks. In particular, for each outcome variable X (such as GDP or consumption), I estimate the model

$$X_{t,p,i} = s(t) + I_{[shock=1]} s_{shock}(t) + \varepsilon_{t,p,i}, \tag{2}$$

where t is the simulation period, i is the MC repetition and p indicates the presence or absence of the policy shock. $s(t)$ is the baseline spline, to which the policy spline $s_{shock}(t)$ is added through the dummy I if the run in question includes a government expenditure shock. The policy spline then depicts the smoothed reaction of the model to the policy shock (in the figures below reactions are expressed as % of baseline levels), hence serving the same purpose as the robust IRFs suggested by Delli Gatti and Grazzini (2019). Both the baseline model and the model including the expenditure shock are simulated for 1500 periods, each 200 times with different reproducible seeds, and the shock is introduced in the same period in each individual run. In this way, possible cyclical effects (which are investigated below) are averaged out as booms and downturns do not coincide temporally across different runs.

Figures 4 and 5 summarise the effect of a government expenditure shock in the baseline model in terms of the cumulative multiplier and responses of key aggregate model variables. In the case of cumulative multipliers, the black and grey lines represent the median and 90% confidence intervals from a Wilcoxon signed rank test across Monte Carlo repetitions respectively. In the case of the policy splines the black line represents the estimated spline while the grey lines are standard error bands.

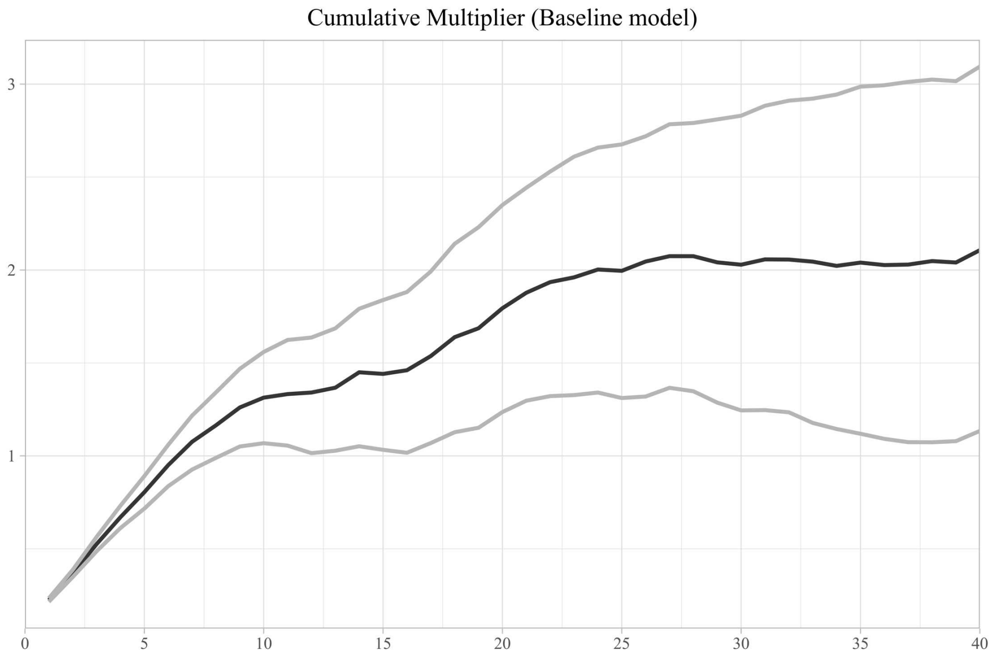

Figure 4: Cumulative government expenditure multiplier in the baseline model

It can be seen that the model generates a small but positive multiplier on impact,[6] which increases to eventually become significantly larger than one. Figure 5 plots the policy splines for real GDP, real consumption, real consumption demand, as well as gross fixed capital formation. It shows that the model reaction to the fiscal shock is mainly driven by consumption whilst capital investment does not react strongly to the non-persistent shock. Real GDP increases on impact as firms increase their output in response to the additional government demand. Actual consumption increases in the first instance due to a reduction in frictions on the consumption goods market as a consequence of the additional output made available by firms which reduces the incidence of coordination

[6]The impact multiplier is smaller than one since it is defined as the *additional* output produced relative to the baseline, divided by the amount of actual government expenditure. Due to the algorithm by which firms adjust their production as well as the various constraints to which they may be subject at any given time, firms will usually not increase their production by the full amount of government demand which they receive, and some of the output purchased by the government will be goods which would have been produced even in the absence of the shock.

failures on the market, as well as through a loosening of liquidity constraints on households. Consumption demand (i.e. desired consumption) reacts somewhat more sluggishly and persistently, causing a hump-shaped response of actual consumption in later periods.

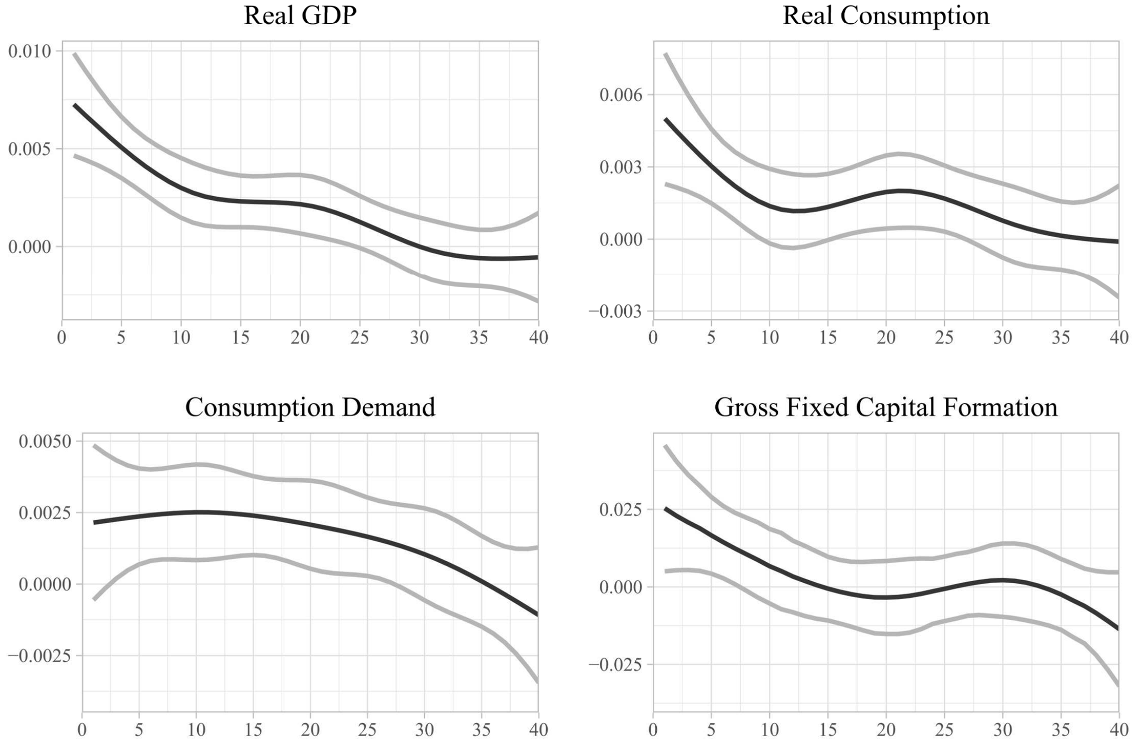

Figure 5: Policy splines for the baseline model

Another interesting experiment can be conducted by investigating whether government expenditure multipliers are affected by cyclical conditions, i.e. whether they are larger during downturns and smaller during booms. Following the global financial crisis and the associated policy responses, a great deal of evidence in favour of such regime-dependent multipliers has been presented (see e.g. Auerbach and Gorodnichenko, 2012a; Fazzari et al., 2014; Qazizada and Stockhammer, 2014; Gechert and Rannenberg, 2018). There are good reasons to suppose that such cyclical effects might be present in the model used here. Firstly, during downturns the unemployment rate will be high and C-firms' capacity utilisation may be low, hence removing two of the potential supply-side constraints which can limit the effectiveness of fiscal policy. Secondly, recall that the model by construction includes a liquidity constraint on households. During recessions, this constraint may be binding for a greater number of households and may be eased by the introduction of government expenditure as an additional source of demand, an effect which has frequently been invoked to explain the empirical regime-dependence of multipliers (Mittnik and Semmler, 2012; Warmedinger et al., 2015).[7] At the same time however, credit

[7]The way in which the liquidity constraint is modelled here differs somewhat from that typically

constraints on firms may be more stringent during downturns, possibly preventing them from expanding production. To test for a cyclical effect on the government expenditure multiplier, I identify the post-transient periods of minimum and maximum GDP in each of the 200 baseline runs without the policy shock, and then run two experiments, one in which the shock is introduced in the period of minimum GDP in each run and one in which it is introduced at the maximum.

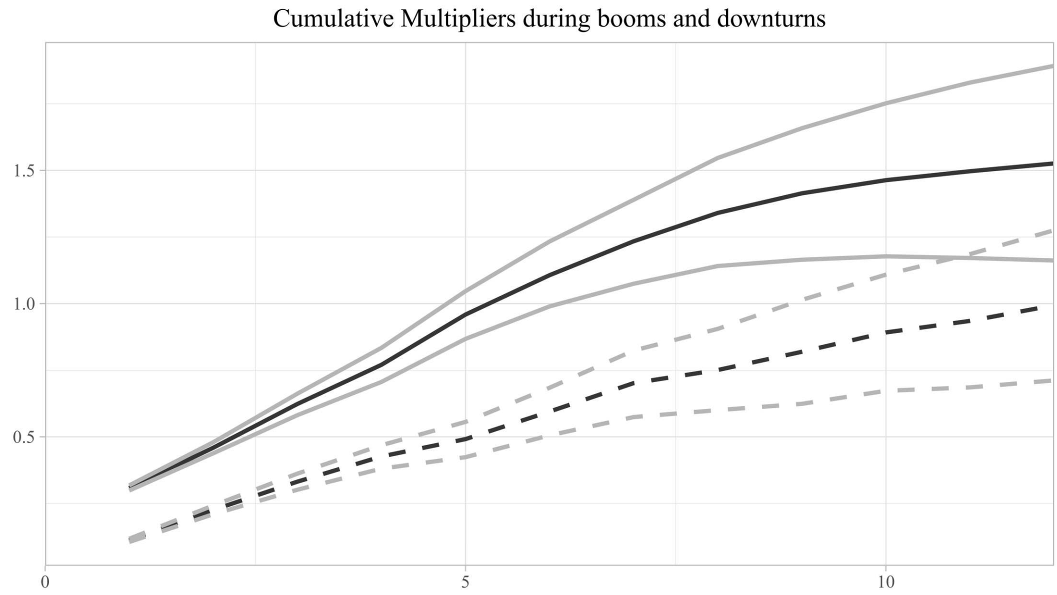

Figure 6: Cumulative government expenditure multipliers during downturns (solid) and booms (dashed)

Figure 6 shows that there is indeed a difference between multipliers during booms and downturns, particularly so during the initial periods following the shock. Impact multipliers differ as firms are more capable of quickly increasing output during downturns. As firms produce additional output to meet the demand from the government, they hire additional workers, leading to the liquidity constraint being binding for fewer households, which can consequently increase their consumption. Just like the model analysed by Napoletano et al. (2017), the framework used here hence gives rise to state-dependent fiscal multipliers. The analysis of government expenditure shocks in the baseline model hence already reveals some interesting and arguably realistic effects. The following sections investigate to what extent the effects of government expenditure shocks depend on households' consumption behaviour, including their expectations and beliefs about the effects of fiscal policy.

used in the New Keynesian DSGE literature (e.g. Galí et al., 2007) in that in the present model, while households are not allowed to borrow, they may save if willing and able to do so, i.e. they are not hand-to-mouth consumers.

5 Time for CAOS - agent-based intertemporal optimisation

In order to more closely investigate the influence of consumption behaviour on fiscal multipliers, the baseline model is modified with respect to Delli Gatti and Grazzini (2019). In the baseline version of the model, households formulate a plan for desired consumption based on a weighted average of current and past incomes (dividends for the owners of firms and wages/unemployment benefits for worker households), i.e. each household i calculates

$$\overline{Y_t^i} = \xi \overline{Y_{t-1}^i} + (1 - \xi) Y_t^i, \tag{3}$$

(where Y_t^i is their income in the current period) and then determines their desired consumption according to

$$C_d^i = \overline{Y^i} + \chi D^i \tag{4}$$

where D^i are their deposits with the bank. This rule serves to impart a degree of persistence to consumption demand, as is empirically warranted, inducing agents to react less than one-for-one to transitory shocks to disposable income. $\overline{Y^i}$ is understood as "a proxy for future expected income" (Assenza et al., 2015, p. 9). The modification of the baseline model proposed here alters this modelling of future expected income to introduce a behavioural rule which is somewhat more forward-looking and brings households' consumption behaviour closer to conventional permanent income/life-cycle models of consumption (Hall, 1978).

In various respects which are important for the purpose of this paper, the structure of the baseline model is very simple relative to some other well-established macro-ABM frameworks. Workers are all identical with regard to their productivity and there is a uniform wage determined via a Phillips-curve equation at the aggregate level. Unemployed workers accept the first open position they find during the search and matching process on the labour market, and if a firm is firing workers, it does so in a random order. All unemployed workers receive the same level of benefits (which are a fixed fraction of the current wage). Similarly, any consumption goods firm is structurally identical to all others in the sense that they all produce the same good using the same technology (and the same is true for capital goods firms), although any particular firm may at any given time of course be in a particularly good or bad financial situation. All firms (as well as the representative bank, which is owned in equal parts by all capitalists) have very simple dividend heuristics whereby they pay out a fixed fraction of profits if the latter are positive. In addition, workers' and capitalists' *real* incomes do not exhibit long-term trends but minor fluctuations around a stochastic stationary state, as do all other model variables in the real dimension (cf. Assenza et al., 2015).

Making use of the above observations, households' behaviour is modified by assuming that they base the estimation of their future income on transition probabilities between different states and a long-term weighted average of the associated real incomes. For instance, given the parsimonious modelling of the labour market, it appears to be at least a reasonable approximation to represent workers' employment status over time by a Markov chain with a transition matrix **T** of the form[8]

	E	**U**
E	π_{EE}	$1 - \pi_{EE}$
U	$1 - \pi_{UU}$	π_{UU}

Transition data is collected centrally during the simulation and used to estimate the transition probabilities, using data of the past $\mathcal{H}$ periods. It is assumed that all workers have access to this aggregated information so that for the moment, they all have a common estimate of the transition probabilities. Moreover, it is assumed that workers calculate an *unweighted* average of the $\mathcal{H}$ past values of the uniform real wage and the real unemployment benefit, resulting in a payoff vector $\mathbf{P} = \begin{bmatrix} \overline{w} \\ s \cdot \overline{w} \end{bmatrix}$. Given that the real wage shows no long-term upward or downward trends in model simulations, this appears to be a reasonable way to forecast the long-run average future real wage and benefit. **T** and **P** can in turn be used to provide an estimate of future income, as they describe the likelihood of future state transitions along with an estimate of the associated incomes. In addition to being broadly familiar from various economic applications, this way of formulating the values of agents' states draws closely on the machine learning literature, specifically that on reinforcement learning (Sutton and Barto, 1998). In the case of capitalists, the matrix **T** represents the transitions between the states of receiving and not receiving a dividend (again calculated centrally and uniform across firms), and **P** contains a long-run average of an individual firms' dividends as its first element and zero as the second (i.e. there is a **P** for each individual firm owner). The transition probabilities and payoff vectors form the basis of an alternative consumption behaviour involving intertemporal optimisation.

Consider the simplest case of a consumer holding a stock of assets and earning an income who attempts to maximise utility from consumption over a horizon $\mathcal{H}$. Their intertemporal budget constraint can be written as

$$\sum_{t=0}^{\mathcal{H}} R^{-t} c_t \leq d_0 + \sum_{t=0}^{\mathcal{H}} R^{-t} y_t, \tag{5}$$

[8] The validity of this claim is assessed through simulation experiments presented in appendix B.

where d_0 is the real value of their assets in period 0 and R is a gross interest rate. Based on the transition matrices and payoff vectors calculated as described above, an estimate of income in some future period t from the perspective of period 0 can be calculated as

$$y_t = \prod_{i=1}^{t} \mathbf{T}_i \mathbf{P}_t, \tag{6}$$

where y_t is a vector with a number of rows equal to the number of states. The estimated sum of future incomes, conditional on a series of expected transition matrices $\mathbf{T}_i$ and payoff vectors $\mathbf{P}_t$,[9] is then given by

$$\sum_{t=0}^{\mathcal{H}} y_t = \mathbf{P}_0 + \sum_{t=1}^{\mathcal{H}} \left(\prod_{i=1}^{t} \mathbf{T}_i \mathbf{P}_t \right), \tag{7}$$

so that

$$\sum_{t=0}^{\mathcal{H}} R^{-t} y_t = \mathbf{P}_0 + \sum_{t=1}^{\mathcal{H}} \left(R^{-t} \prod_{i=1}^{t} \mathbf{T}_i \mathbf{P}_t \right) = YP. \tag{8}$$

By iterating forward from $t = 0$ to $t = \mathcal{H}$ (in the simplest case assuming that both $\mathbf{T}$ and $\mathbf{P}$ are constant), the transition matrices and payoff vectors hence allow for the derivation of an estimated budget constraint (or, more precisely, a set of budget constraints with cardinality equal to the number of possible states contained in $\mathbf{T}$), giving an estimate of the net present value of the sum of resources available for consumption over the future $\mathcal{H}$ periods. In a slight abuse of terminology, I call this sum 'permanent income', YP. The consumer can then optimise, maximising a utility function subject to this *estimated* budget constraint. The consumer will hence face the problem

$$\begin{aligned} &\max_{c_t} \sum_{t=0}^{\mathcal{H}} \beta^t u(c_t) \\ s.t.\ &\sum_{t=0}^{\mathcal{H}} R^{-t} c_t \leq d_0 + \mathbf{P}_0 + \sum_{t=1}^{\mathcal{H}} R^{-t} \left(\prod_{i=1}^{t} \mathbf{T}_i \mathbf{P}_t \right), \end{aligned} \tag{9}$$

where β is a discount factor. The solution to this problem will be a set of consumption paths with cardinality equal to the number of states contained in $\mathbf{T}$[10] or, to put it differently, there will be one maximisation problem for each possible current state. Writing

[9] Neither $\mathbf{T}$ nor $\mathbf{P}$ need necessarily be constant. Households may well expect transition probabilities or payoffs to change in some future period and incorporate this expectation into their calculations.

[10] For instance, in the context of the model presented above, one path for a worker who is employed in $t = 0$ and one for a worker who is unemployed in $t = 0$.

the Lagrangian and taking derivatives gives a set of first order conditions of the form

$$u'(c_t) = (\beta R)^{-t}\lambda. \tag{10}$$

The case in which $\beta R = 1$ gives rise to the familiar result that the household plans to consume a fixed fraction of their estimated total budget in each period, as in the simplest versions of the life-cycle or permanent income models of consumption (Friedman, 1957; Modigliani and Brumberg, 2005). If $\beta R \neq 1$, $u(c_t)$ must be specified in order to solve for c_t. In the case of the present model, $R = 1$ by construction since, while there is a base interest rate in the model, households cannot earn it as their only asset are *unremunerated* bank deposits. β on the other hand is calibrated below to a value producing sensible results, meaning that in the present framework, $\beta R \neq 1$. Utility from consumption is assumed to be given by

$$u(c_t) = \frac{c^{(1-\theta)} - 1}{1-\theta}, \tag{11}$$

such that the first order condition, given that $R = 1$, produces

$$c_t = \left(\frac{\lambda}{\beta^t}\right)^{-\frac{1}{\theta}}. \tag{12}$$

Substituting this into the budget constraint to solve for λ and letting $B = \sum_{t=0}^{\mathcal{H}} \beta^{\frac{t}{\theta}}$, the consumption demand of a households j becomes

$$c_t^j = \left(\beta^{-t}\left(\frac{d_0^j + \sum_{t=0}^{\mathcal{H}} y_t^j}{B}\right)^{-\theta}\right)^{(-\frac{1}{\theta})}, \tag{13}$$

which in period $t = 0$ simplifies to

$$c_0^j = \frac{d_0^j + \sum_{t=0}^{\mathcal{H}} y_t^j}{B}. \tag{14}$$

This consumption behaviour is followed by both worker and capitalist households, with the **T**'s and **P**'s being re-estimated in each period in order to gain an updated estimate of the budget constraint and re-optimise. Note that since the **T**'s and **P**'s are identical for all worker households, the only source of heterogeneity in the consumption demand of worker households stems from heterogeneity in asset holdings d, which in turn reflects worker households' idiosyncratic employment history. Firm owner households, on the other hand, also each have an individual **P** vector as firm dividends are less uniform than worker earnings (recall that there is a uniform wage rate).

The purpose of this modification is to introduce what might be viewed as an agent-based version of the permanent income or life-cycle model of consumption within a relatively simple macroeconomic framework. The values of the variables entering into equation (7) are of course still calculated based on past data and in this sense the formulation is backward-looking. Nevertheless, the consumption behaviour thus induced can be regarded as more forward-looking than the baseline one in that agents take into account their current state as well as the estimated likelihood of future ones and formulate a consumption *path* conditional on this expectation. Moreover, while in the simplest case, $\mathbf{T}$ and $\mathbf{P}$ are assumed constant and equal to the most recent estimate for each agent over the horizon $\mathcal{H}$, the framework also allows for the incorporation of *expected future variations* in transition probabilities and payoffs which would have an immediate effect on consumption demand.

Before moving on to an analysis of the effects of this modification on the government expenditure multiplier, two more points should be emphasised. Firstly, the alternative consumption behaviour outlined above does not eliminate the presence of the liquidity constraint on households, i.e. I assume that they are *not* allowed to borrow against their estimated future income. Secondly, fiscal policy will have an impact on agents' consumption demand to the extent that it impacts their current state, the estimated transition probabilities, and their estimate of the income vectors. There is no notion of Ricardian equivalence in this model, i.e. agents do not, for instance, conjecture a future rise in income tax following a temporary shock to government expenditure (and indeed such rises do not take place as the tax rate is fixed in all simulations).[11]

In all experiments shown below, the discount factor β is set to a value of 0.9999.[12] $\mathcal{H}$, which gives both the length of the time-window used to estimate income vectors and transition probabilities, and the horizon over which households optimise consumption, is set to 400 (implying a horizon of 100 years). θ is set to a value of $\frac{2}{3}$, implying an intertemporal elasticity of substitution equal to 1.5, which is in line with values commonly used in the macroeconomic literature (e.g. Fernández-Villaverde and Rubio-Ramírez, 2005; Smets and Wouters, 2007; Del Negro et al., 2015). The results presented below are qualitatively robust to changes in these parameters, though changes to one of them necessitate concurrent changes in the others to obtain sensible results. Relative to the baseline model, the implementation of the 'permanent income' and optimisation mechanism significantly re-

[11] An analysis of simulation data shows that the brief positive government expenditure shocks considered in this paper in fact do not give rise to a significantly higher level of government debt in the long run.

[12] This discount factor is admittedly very high, even compared to the DSGE literature in which high discount factors are relatively common (see e.g. Smets and Wouters, 2007; Gerali et al., 2010), although the mainstream literature does at times consider the case of $\beta = 1$ (Ascari and Rossi, 2012, p.1123). This issue could be overcome by the introduction of an interest rate on the households' assets. However, this would represent an additional modification of the original baseline model I wish to avoid for ease of comparability between the baseline and the model featuring the modified consumption behaviour. The focus of the present paper lies exclusively on exploring fiscal multipliers and their relationship with households' consumption behaviour rather than on extending the baseline model itself.

duces the volatility of consumption and (to a somewhat lesser extent) GDP in simulations. Table 1 shows the standard deviations of consumption demand, actual consumption and output in the modified model as a ratio of those obtained from the baseline, demonstrating the reduction in volatility.

Table 1: Comparison of standard deviations in the modified model and the baseline

	Consumption Demand	**Consumption**	**Output**
Ratio	0.3820	0.6586	0.7133

The introduction of a utility function for households also opens up the possibility of conducting a simple form of welfare analysis. Recall that instantaneous utility from consumption for some household j is assumed to be given by

$$u(c_t^j) = \frac{(c_t^j)^{(1-\theta)} - 1}{1-\theta}. \tag{15}$$

A measure of the welfare of an individual household j over a simulation can be constructed as

$$U^j = \frac{\sum_{t=0}^{T} u(c_t^j)}{T}, \tag{16}$$

where T is the length of the simulation and $t = 0$ is the first post-transient simulation period. An aggregate measure of welfare from consumption can then be calculated as

$$W = \sum_{j=1}^{J} U^j, \tag{17}$$

where J is the totality of all households in the model. Table 2 shows a comparison of welfare from consumption between the baseline model (where consumption following the baseline heuristic is inserted into the same utility function used in the modified framework) and the model featuring the alternative consumption behaviour. Differences in utility are calculated agent-by-agent and then summed up. Line one shows the Monte-Carlo average difference in welfare, both for consumption demand (which is the object being optimised by agents in the modified model) and actual consumption (which may differ from consumption demand as outlined above), along with 95% confidence intervals. Line two expresses the difference in welfare in terms of consumption units[13] and the third lines gives this difference as a percentage of average per-period consumption/consumption demand in the baseline model.

[13]This is calculated for each agent as $u^{-1}(U_m^j) - u^{-1}(U_b^j)$ where m signifies the modified model and b the baseline. The differences are then summed up across agents.

Table 2: Difference of welfare between the modified model and the baseline

	Consumption Demand	Actual Consumption
Welfare	6.51384	8.98573
	(5.84885; 7.17046)	(8.63143; 9.32912)
Consumption units	4.24096	4.90486
	(3.81849; 4.66100)	(4.72348; 5.08150)
% of mean cons.	0.8314	1.0511

As can be seen, the implementation of the optimising behaviour leads to a small but significant increase in welfare from actual consumption which is equivalent to roughly 1% of aggregate consumption in every period. Interestingly, the increase in welfare derived from consumption is even greater than that implied by the increase in welfare calculated on desired consumption. This suggests that the optimising behaviour, which leads to a significant reduction in macroeconomic volatility, is also able to somewhat reduce welfare losses from frictions on the consumption goods market. Since the optimisation behaviour is purposely calibrated such that the level of average consumption under optimisation is almost equal to that under the baseline heuristic, the gain in welfare is indeed due to the reduction in the volatility of consumption demand and actual consumption. Despite the effects outlined above, the model is qualitatively unaffected by the modification to consumption behaviour insofar as model variables still exhibit fluctuations around a real stationary state, the difference being that these fluctuations tend to be smaller.

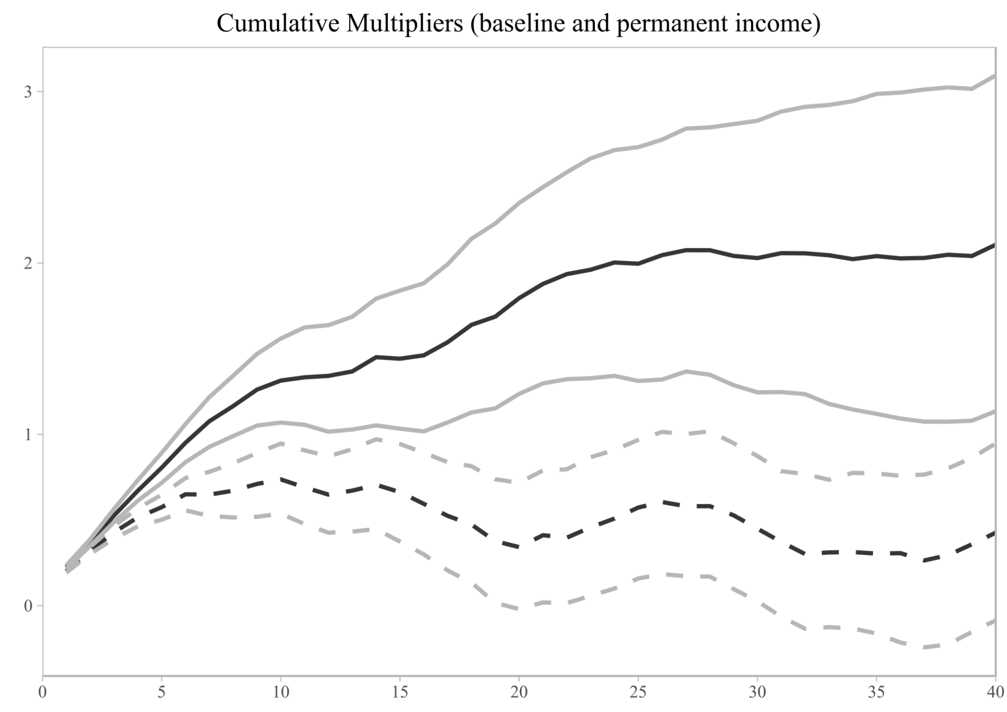

Figure 7: Cumulative multipliers in the baseline (solid) and modified (dashed) model

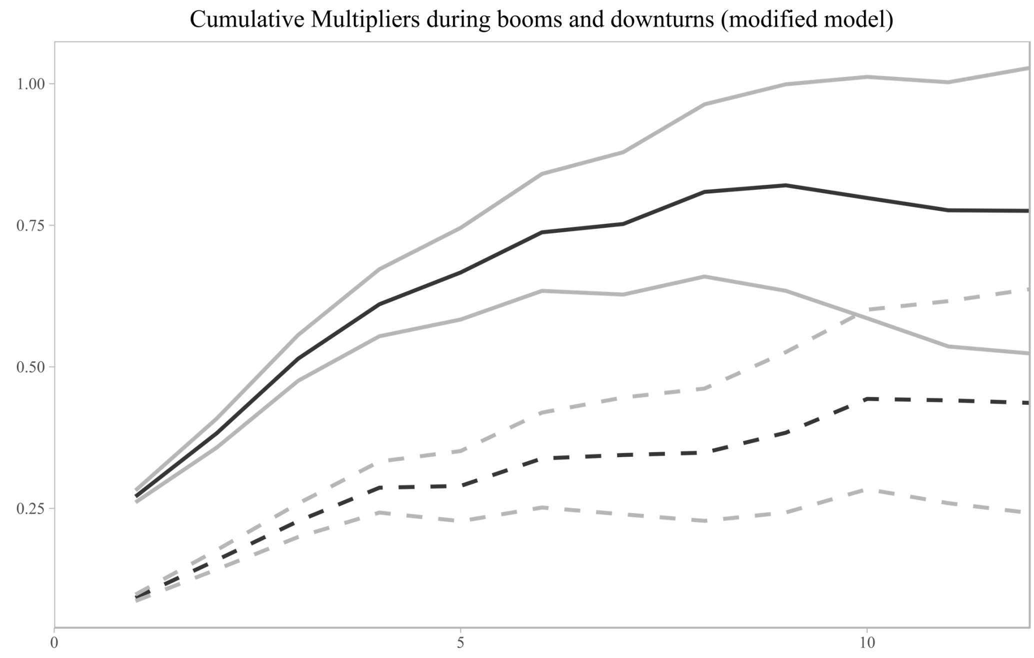

Figure 8: Cumulative government expenditure multipliers during downturns (solid) and booms (dashed)

Nevertheless, the reaction of the altered model to a government expenditure shock, shown in figure 7, is substantially different from that of the baseline model. It can be seen that in the modified model, the cumulative multiplier - particularly in later post-shock periods - is much lower than in the baseline version and that the median always remains well below one. Figure 8 demonstrates that under the modified consumption behaviour, the effects of cyclical conditions on the government expenditure multiplier remain intact.

Figure 9 compares the policy splines produced by both versions of the model, showing that the more muted response of the modified model is driven by a smaller response of consumption, and that the duration of the impact is shorter than in the baseline model. Desired consumption reacts to a fiscal shock in a very muted fashion. The reaction is chiefly due to the induced increase in employment, which has a direct impact on workers' estimated budget constraint as the latter is state-dependent (i.e. there is one budget constraint for each possible current state), as well as to state-transitions of firm owners. The estimated transition probabilities and payoff vectors are based on a large amount of past data and are hence hardly affected by the one-period shock to government expenditure. Despite the implementation of the alternative consumption behaviour, households may still be liquidity-constrained and this constraint may be eased by the fiscal shock, leading to an increase in actual consumption greater than that observed in desired consumption.

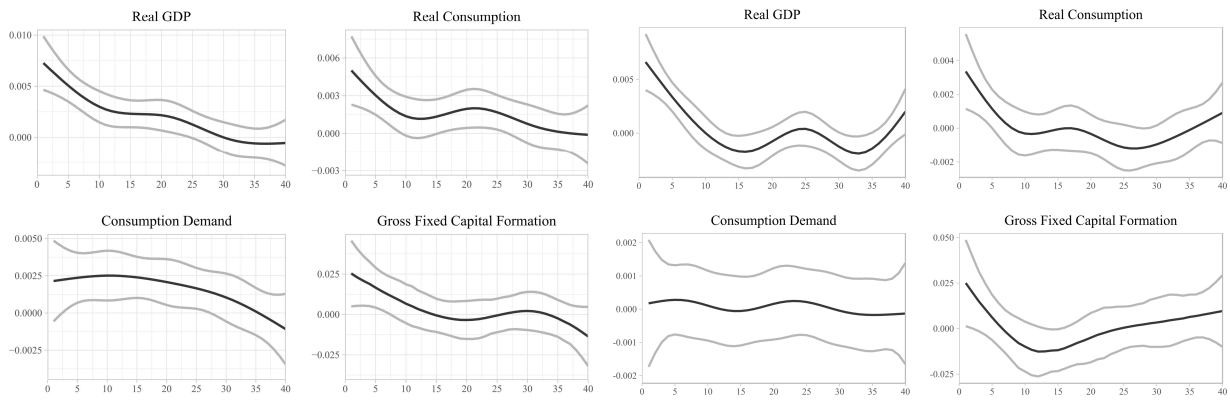

(a) Baseline model (b) Modified model

Figure 9: Comparing quasi-IRFs

The above experiment demonstrates the importance of households' income expectations in driving the model's response to government expenditure shocks and in particular highlights that the longer-term reaction of the baseline model to such shocks appears to be substantially driven by the relatively simple heuristic households draw on for their consumption decision in the baseline model. A major effect of the introduction of the alternative consumption behaviour is to substantially strengthen the degree to which households aim to smooth consumption which in turn contributes to a smaller cumulative multiplier. The present experiment, however, does not yet fully exhaust the possibilities afforded by the alternative modelling of consumption behaviour as expectations of future income do not explicitly account for the effects of policy. The following experiments, therefore, retain the modified consumption behaviour introduced above but aim to show how it can be extended to investigate the role of households' beliefs and expectations about the effects of fiscal policy.

As a first experiment, I suppose that households hold initially exogenous and heterogeneous beliefs about the effects of a government expenditure shock. In particular, I assume that households may take either an optimistic or a pessimistic view of the prospective effects of a positive government expenditure shock on their estimated budget constraint. One class of households which I term 'Keynesian' or optimist believe that the effect of the shock will be expansionary, whilst another class which I term 'Classical' or pessimist may be thought of as holding a strong belief in favour of 'expansionary fiscal consolidations' (Giavazzi and Pagano, 1990; Alesina et al., 2019) such that they expect the effect of the shock to be negative. Importantly, it is assumed for the moment that agents cannot switch between the optimistic and pessimistic types; rather they are assigned a type at the beginning of the simulation which remains constant, but the strength of their belief adapts according to post-shock developments. Recall that households' estimates of their budget constraint or 'permanent income' are based on two components, namely their assessments of the transition probabilities (from employed to unemployed, dividend paid to no dividend paid, etc.) and their pay-off vectors (containing long-run averages of the real wage, firm dividends etc.), which in turn they project forward using equation (7). It is now assumed that from the period in which the shock is introduced, households believe that for the following 40 periods, the transition probabilities and payoff vectors are going to differ from their previous estimate by some factor, the value of which depends on their type. For instance, the expectations of a worker household j are modified as follows:

$$
\begin{aligned}
\hat{\pi}^j_{EU,t} &= \pi_{EU,t} \cdot \iota^j_{EU,t} \\
\hat{\pi}^j_{EE,t} &= 1 - \hat{\pi}^j_{EU,t} \\
\hat{\pi}^j_{UU,t} &= \pi_{UU,t} \cdot \iota^j_{UU,t} \\
\hat{\pi}^j_{UE,t} &= 1 - \hat{\pi}^j_{UU,t} \\
\hat{w}^j_t &= \overline{w}_t \cdot \iota^j_{w,t},
\end{aligned}
\tag{18}
$$

where $\pi_{EU,t}$ is the estimate (based on past data) of the transition probability (from employed to unemployed in this example) at t, $\hat{\pi}^j_{EU,t}$ is worker j's belief about the post-shock probability, and so on. The initial values of the ι's are all identical across households of the same type, with the value depending on whether the household is Keynesian (in which case $\iota^j_{EU,t} < 1$, $\iota^j_{UU,t} < 1$ and $\iota^j_{w,t} > 1$) or Classical ($\iota^j_{EU,t} > 1$, $\iota^j_{UU,t} > 1$ and $\iota^j_{w,t} < 1$). An equivalent set of beliefs about changes in transition probabilities and payoffs is introduced for firm owners. During the period in which the shock is introduced, I hence impose on each household an exogenous belief which either increases or decreases the estimate of its future income and hence modifies its budget constraint and consumption behaviour. As indicated above, each household initially believes that this change will have a duration ℓ of 40 post-shock periods[14] and adapts the projection of its expected future income accordingly. Over the periods following the shock, the household observes whether and to what extent the actual transition probabilities and payoffs, calculated using only data from post-shock periods, have in fact changed as implied by its initial beliefs. The ι's are then adjusted adaptively. For instance, for each worker household j they change according to

$$
\begin{aligned}
\iota^j_{EU,t} &= \psi \cdot \iota^j_{EU,t-1} + (1-\psi)\frac{\pi^p_{EU,t}}{\hat{\pi}^j_{EU,t-1}} \\
\iota^j_{UU,t} &= \psi \cdot \iota^j_{UU,t-1} + (1-\psi)\frac{\pi^p_{UU,t}}{\hat{\pi}^j_{UU,t-1}} \\
\iota^j_{w,t} &= \psi \cdot \iota^j_{w,t-1} + (1-\psi)\frac{\overline{w}^p_t}{\hat{w}^j_{t-1}},
\end{aligned}
\tag{19}
$$

where the p-superscripts denote values calculated using only post-shock data as inputs. The adaptive process ceases when the ι's reach a value of 1, at which point the household

[14]The choice of 40 periods is made for presentational reasons and affects the strength of the belief-effect discussed below quantitatively but not qualitatively.

no longer holds any belief and may be regarded as neutral. To gain an understanding of the effects of these exogenously imposed beliefs on the fiscal multiplier, I re-run the fiscal policy experiment for different shares of Keynesian households (with the rest being Classical). In all experiments, the initial values of the ι's for worker households are set as follows: $\iota^j_{EU} = 0.9$ (Keynesian), $\iota^j_{EU} = 1.1$ (Classical); $\iota^j_{UU} = 0.9$ (Keynesian), $\iota^j_{UU} = 1.1$ (Classical) and $\iota^j_w = 1.1$ (Keynesian), $\iota^j_w = 0.9$ (Classical). The corresponding ι's for firm owners are set to the same values. The value of ψ is set to 0.95 implying some persistence in beliefs.[15]

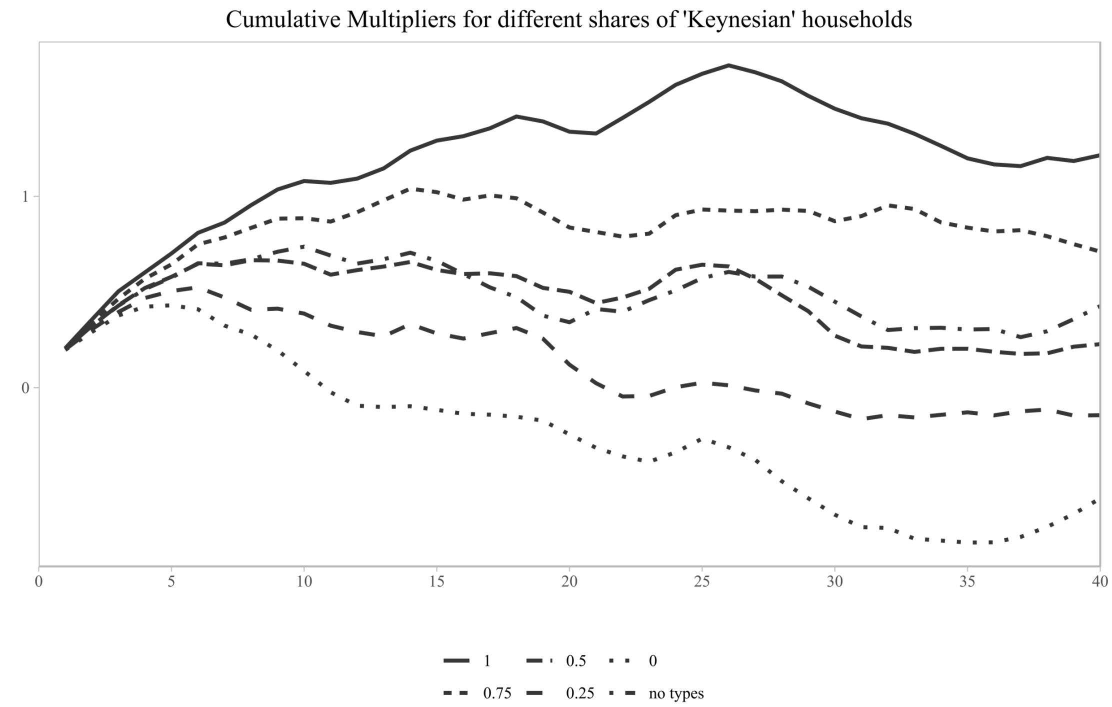

Figure 10: Cumulative multipliers for the case of adaptive beliefs

Figure 10 demonstrates how the impact of the government expenditure shock is affected by the presence of beliefs, showing the cumulative multiplier for different shares of 'Keynesian' households in the economy. It can be seen that for high shares of 'Keynesian' or optimist households, the cumulative multiplier is much higher than in the absence of beliefs. In the case of only 50% optimist households, with the rest being pessimist, the cumulative multiplier is almost identical to that obtained in the baseline without types, meaning that the two beliefs appear to cancel each other out. High shares of pessimist or 'Classical' households, on the other hand, can in fact give rise to *negative* cumulative

[15]Note that these values, together with the assumed duration ℓ of 40 periods, imply fairly extreme and persistent beliefs considering that the shock itself is neither exceptionally large nor persistent. Qualitatively similar results of shorter duration are obtained for smaller values of the respective parameters but for expositional reasons the specification shown above is chosen here to emphasise the differences produced by the model under different beliefs.

multipliers after some time.[16] Beliefs about the impact of government expenditure shocks can hence to some degree be self-reinforcing in the present model, even being potentially able to change the sign of the effect of shocks on aggregate output. The fact that cumulative multipliers stabilise after some time (rather than tend to $\pm\infty$) shows that beliefs do return to neutrality eventually however, meaning that the self-reinforcing dynamic is only temporary.

The introduction of belief dynamics as described above also allows for the incorporation of a simple announcement effect of government expenditure, as expectations about *future* changes in transition probabilities and payoff vectors can enter into households' *present* estimates of future income. Suppose for instance that at time t the government announces a one-period increase in government consumption from 0 to some positive level in period $t + x$. As long as $x < \mathcal{H}$ (the projection horizon used by households), and assuming that households believe the announcement to be credible,[17] households' beliefs about how the shock in $t+x$ will affect their income from $t+x$ onwards will be incorporated immediately into their estimated budget constraint. Figure 11 illustrates this announcement effect, for simplicity using the case in which all households are Keynesian. It can be seen that desired and actual consumption increase as soon as the shock is announced (period 0 in the plots), with output rising in response. There is then a further increase in output and consumption as the actual shock occurs, since additional output is being produced by firms in response to anticipated government demand.

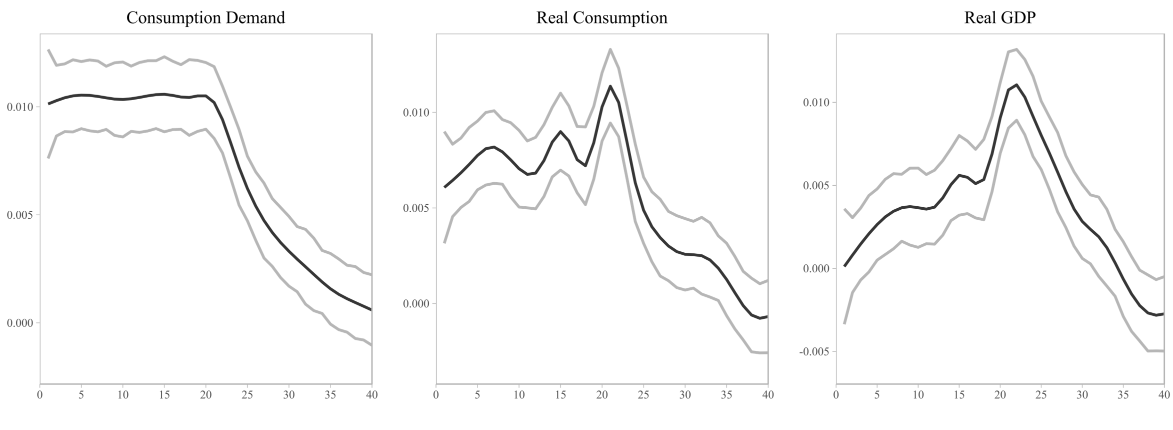

Figure 11: Announcement effect of government expenditure

A further variation on the belief dynamics introduced above is the incorporation of switching behaviour between types (see e.g. Franke and Westerhoff, 2017). In a modified version of the model, households are still initially classified as 'Keynesian' or 'Classical', but instead of adapting the strength of their belief according to equation (19), they instead

[16]The presence of this effect however depends on the strength of the initial belief and the speed of adaptation.

[17]The present paper completely abstracts from the issue of credibility of policy announcements, but this would be an interesting dimension to explore in further research.

switch between types with fixed beliefs based on the value of a switching index. The value of the latter is a function of the post-shock development of estimated 'permanent income' as well as the relative concentration of types within the population of households. For instance, for some household j which is currently an optimist, the switching index is calculated as

$$
\begin{aligned}
deviation_t^j &= \frac{YP_t^{j,p} - \widehat{YP_t^j}}{\widehat{YP_t^j}} \\
switch_t^j &= \epsilon \cdot deviation_t^j + (1-\epsilon)(opt_{t-1} - (neut_{t-1} + pes_{t-1})) \\
index_t^j &= \frac{1}{1 + \exp(\sigma_1\, switch_t^j + \sigma_2)},
\end{aligned}
\tag{20}
$$

where $\widehat{YP_t^j}$ is j's expected 'permanent income' (incorporating both their belief and pre-shock data), $YP_t^{j,p}$ is their 'permanent income' estimated on *post-shock data only* and *opt*, *pes* and *neut* are the shares of optimists, pessimists and neutral households in the system. The values of all parameters are given in appendix A. $index_t^j$, which is bounded between 0 and 1, gives the probability that household j will switch from being an optimist to being neutral. In the following period, j reassesses their opinion and may switch from being neutral back to being an optimist or to becoming a pessimist, using appropriately adapted versions of equation (20) and depending on the developments of $YP_t^{j,p}$ relative to estimates incorporating pre-shock data and the new relative numbers of the different types within the population. The model incorporating switching behaviour is run for different initial shares of optimist and pessimist households (it is always assumed that there are initially no neutral households), in each case recording the model reaction to a government expenditure shock.

Figure 12 shows that the effect on cumulative multipliers is even stronger than that shown in figure 10 for the case of adaptive beliefs without switching. The absence of adaptation in the switching case appears to make beliefs more persistent, producing more pronounced differences in multipliers and more strongly negative multipliers for low initial shares of 'Keynesian' households. As with the case of adaptive beliefs, the case of switching of course also allows for the introduction of an announcement effect of policy. Both cases, however, suffer from the deficiency that households' initial beliefs about the effects of the government expenditure shock are completely exogenous rather than being based on a quantitatively reasonable assessment of what the effect of the shock might be. For the moment, there has also not been any attempt to assess to what extent households' beliefs are in some sense 'correct', though results suggest that beliefs which are strong enough can become at least partly and temporarily self-fulfilling. For these reasons, the next section explores a variant of the model in which households attempt to learn about the effects of government expenditure on their budget constraint from repeated shocks over time using least squares learning.

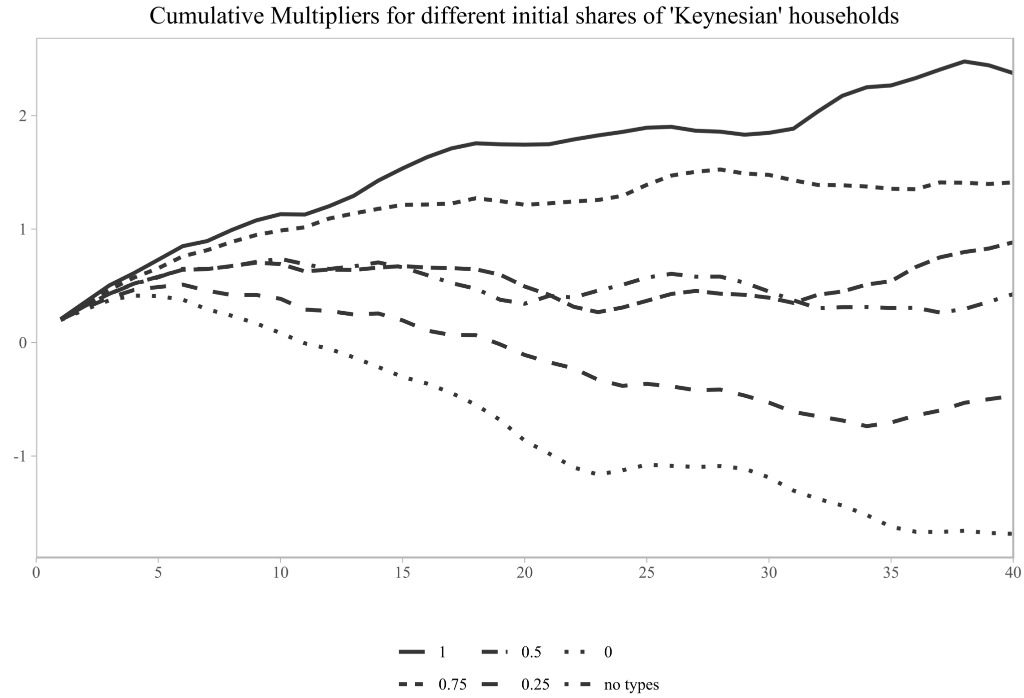

Figure 12: Cumulative multipliers for the case of switching

6 Learning about policy

The application of least squares learning algorithms as presented by Evans and Honkapohja (2001) is a relatively long-standing component of the macroeconomic literature (e.g. Marcet and Sargent, 1989; Bullard and Mitra, 2002; Branch, 2004) but has featured only sparsely in the agent-based literature (see Chapter **??** and Dosi et al., 2017a). The goal of the present exercise is similar to applications of least squares learning in conventional models, although (due to the complexity of the underlying framework) its scope is much more limited. The question I wish to answer is simply whether households are able, on average, to correctly learn the *impact effect*[18] of a government expenditure shock on their budget constraint.

Recall once more that households base their estimates of future income on state transition matrices and income vectors calculated using $\mathcal{H}$ periods of past simulation data (where $\mathcal{H} = 400$ in the simulations shown in this paper). As argued above, the estimates of these transition matrices and income vectors hardly change in response to one-period government expenditure shocks due to the large amount of past data used to calculate them. Instead, the impact of government expenditure shocks on consumption demand

[18] That is, the effect only in the period in which the shock is applied.

stems chiefly from households temporarily moving to a different state as a consequence of the shock, which in turn alters their budget constraint as, for instance, $YP^E > YP^U$, i.e. the 'permanent income' of a currently employed worker household is greater than that of a currently unemployed household.[19] Ignoring the negligible impacts of government expenditure shocks on households' estimates of transition probabilities and income vectors, an approximation of the expected impact-effect of a shock on the aggregate value of households' 'permanent income' (let this be denoted by YP^{agg}) can hence be expressed as

$$\Delta YP^{agg} = \Delta E(\overline{YP^E} - \overline{YP^U}) + \Delta D(\overline{YP^D} - \overline{YP^{ND}}) + \Delta D_k(\overline{YP^{D_k}} - \overline{YP^{ND_k}}), \tag{21}$$

where E, D and D_k denote the numbers of employed workers, C-firms paying dividends and K-firms paying dividends respectively and the bars over the YP's represent averages across agents. In order to obtain a numerical estimate of this impact, values for ΔE, ΔD and ΔD_k must be estimated, which is done using the least squares learning algorithm. Estimation is performed on aggregate data, i.e. learning does not take place at the level of individual agents (for computational reasons). The three estimated models are

$$E = \gamma_1^E E_{-1} + \gamma_2^E c + \gamma_3^E g_d + \varepsilon^E, \tag{22}$$

where c is aggregate real private consumption and g_d is aggregate real demand for consumption goods coming from the government,

$$D = \gamma_1^D + \gamma_2^D D_{-1} + \gamma_3^D g_d + \gamma_4^D w + \varepsilon^D, \tag{23}$$

where w is the real wage rate, and

$$D_k = \gamma_1^{D_k} + \gamma_2^{D_k} D_{k,-1} + \gamma_3^{D_k} g_d + \varepsilon^{D_k}. \tag{24}$$

The goal is to gain a 'correct' estimate of the γ_3's in the three models, which can in turn be used to predict the changes in employment and the number of additional firms paying dividends as a consequence of a government expenditure shock. In particular, if the estimates are broadly correct, it should be the case that

$$\Delta E = \lfloor \gamma_3^E g_d \rceil, \tag{25}$$

and equivalently for dividends.[20] When a government expenditure shock is announced, and recalling the structural similarity between different individual workers (as well as

[19] This is simply due to the fact that the wage is higher than the unemployment benefit, such that a difference in the sums produced by equation (7) is preserved as long as $\mathcal{H} < \infty$.

[20] Recall that in non-shock periods, $g_d = 0$.

firms) in the model, an individual worker household j, for instance, would hence expect its budget to change by

$$\frac{\lfloor \gamma_3^E g_d \rceil}{W}(YP^{E,j} - YP^{U,j}) \tag{26}$$

in the period in which the shock takes place, where W is the number of worker households in the economy. This expected change, in discounted form, is added to households' estimated budget constraint in the periods after the shock has been announced and before it actually takes place, and subtracted in the period in which the shock takes place (in order to avoind a double-counting of the effect). This in turn affects households' consumption behaviour as the perceived budget constraint entering their optimisation problem is altered. Agents hence react systematically to the announcement of a government expenditure shock and, provided that the estimates of the γ_3's are broadly correct, do so in a model-consistent fashion. If this is the case, households' behaviour in reaction to a government expenditure shock becomes at least partly robust to the Lucas critique, though due to the complexity of the model I limit myself on to the *impact effect* of the shock, disregarding any lagged effects which occur in following periods.

In order to implement least-squares learning about the effect of government expenditure shocks, the simulation set-up must be changed slightly. In particular, agents must be exposed to repeated expenditure shocks in order to collect data that can be used for estimation. For the present experiment, the model is hence simulated 200 times for a duration of 5000 instead of 1500 periods. After the transient and up to period 2500, the model is exposed to a series of random one-period increases in government expenditure from the baseline level of 0, which occur on average every 10 periods. The resulting data is collected as input for the learning algorithm. Prior to period 2500, however, the results of the estimations are not incorporated into households' behaviour in the way described above, in order to allow the estimated coefficients to converge to reasonable values before allowing for feedback effects. From period 2500 onwards, shocks are applied at *regular* intervals of 10 periods and each shock is announced 5 periods before it takes place (i.e. both the timing and the size of the shock are disclosed). Households then immediately incorporate this information into their consumption behaviour as described above, using the results of the learning algorithm which in turn continue to be updated, now incorporating any potential feedback effects from households' altered behaviour.

Parameters are estimated using ordinary least squares. The dataset consists of a rolling window of the past $\mathcal{H}$ periods of simulation data out of which only shock periods and periods immediately preceding shocks are used for estimation, such that estimates are only updated following shocks.[21] Due to the limited amount of data used for each estimation,

[21] Using the full set of $\mathcal{H}$ observations appears to lead to a downward bias in the estimated parameters as in non-shock periods g_d always equals zero.

the algorithm is hence similar to a constant-gain one (cf. Evans and Honkapohja, 2001, Ch. 3). In order to assess to what extent parameter estimates take 'correct' values, given that the actual law of motion of employment and dividend payouts is unknown, I use a relatively simple protocol. For the 500 final periods of each individual run, I calculate the actual changes in employment and dividend payouts following shocks, i.e.

$$\begin{aligned} \Delta E_t &= E_t - E_{t-1} \\ \Delta D_t &= D_t - D_{t-1} \\ \Delta D_{k,t} &= D_{k,t} - D_{k,t-1}, \end{aligned} \tag{27}$$

where t is the shock period, as well as the predicted changes given by

$$\begin{aligned} \widehat{\Delta E_t} &= \lfloor \gamma^{E}_{3,t-x} g_{d,t} \rceil \\ \widehat{\Delta D_t} &= \lfloor \gamma^{D}_{3,t-x} g_{d,t} \rceil \\ \widehat{\Delta D_{k,t}} &= \lfloor \gamma^{D_k}_{3,t-x} g_{d,t} \rceil, \end{aligned} \tag{28}$$

where $t-x$ is the period in which the shock is announced. I then compare both the within-run means and confidence intervals of the respective actual and predicted changes, as well as those calculated across Monte Carlo repetitions. The results of this are summarised in table 3. Lines one and two show the MC average actual and predicted effects along with the associated 95% confidence intervals, whilst line three gives the percentage of runs in which the within-run confidence intervals of the predicted and actual effects overlap. It can be seen that the estimation of the effects of government expenditure shocks on both employment and C-firm dividend payouts appears to be quite accurate. Results are less accurate for K-firm dividend payouts, though both actual and predicted effects are centred around zero and regressions on the entire dataset suggest that the impact effect of government expenditure on K-firm dividend payouts is both small and statistically insignificant, meaning that an estimated effect close to zero appears broadly correct.

Table 3: Actual and estimated effects of government expenditure shocks

	Employment	**Dividends (C-firms)**	**Dividends (K-firms)**
Actual (MC average)	10.98235 (9.45099; 12.51371)	5.88915 (5.06797; 6.71032)	−0.03323 (−0.03787; −0.02860)
Predicted (MC average)	11.49983 (9.89632; 13.10335)	6.18203 (5.32002; 7.04404)	−0.01058 (−0.01205; −0.00910)
% overlap	97.5	100	37

Overall agents are hence able to successfully learn the impact effect of government expenditure shocks on their budget constraint and incorporate this information into their optimisation problem. It should be noted, however, that relative to the exogenously imposed beliefs discussed in section 5, the estimated effects calculated according to equation (28) are quite small. Figure 13 compares the cumulative multiplier in the presence of learning effects to that produced by the basic 'permanent income' model without learning. It can be seen that under learning the multiplier tends to be slightly lower in early periods but these differences are neither statistically significant nor persistent. Consequently, there is also no significant change in households' welfare under least squares learning.

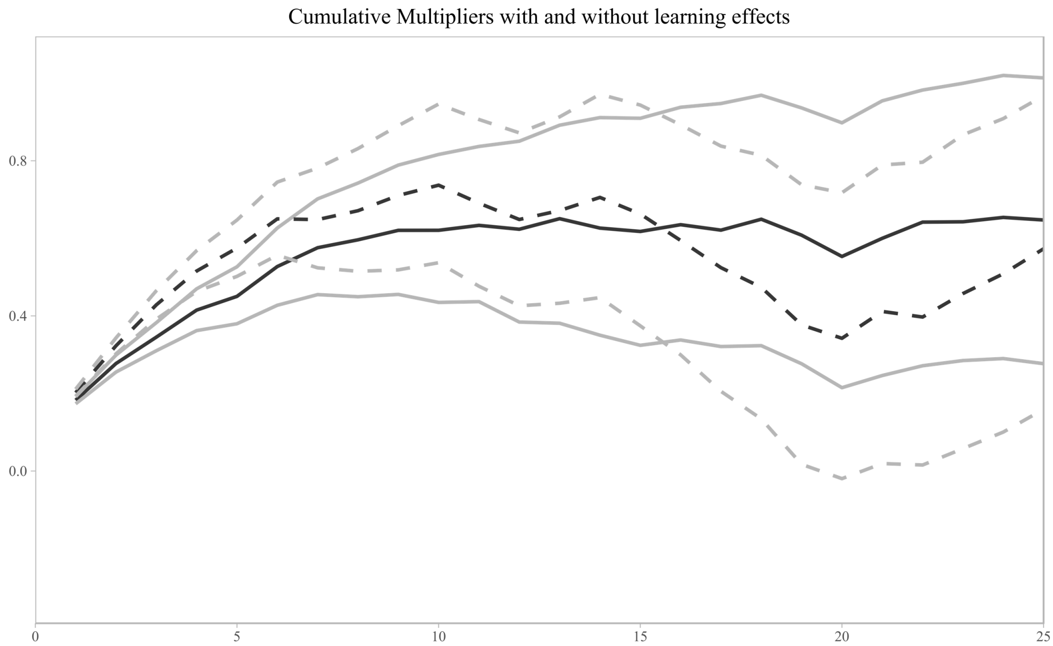

Figure 13: Cumulative multipliers under learning (solid) compared to the absence of learning effects (dashed)

The reason for this result is that, as already outlined in section 5, the majority of the impact of government expenditure shocks on GDP in the 'permanent income' version of the model (beyond the additional output consumed by the government itself) stems from the easing of liquidity constraints and the reduction of frictions on the goods market. Desired consumption as derived from households' optimisation problem only reacts very weakly to non-persistent government expenditure shocks such that the difference produced by the learning effect is very small.

7 Discussion

Overall, the results presented in the preceding sections emphasise the important role of households' consumption behaviour in determining the effects of government expenditure shocks. Positive government expenditure shocks were shown to be strongly expansionary in the baseline featuring a relatively simple, backward-looking consumption heuristic.

Whilst still expansionary, their effect is much weaker if households use a greater amount of past data to formulate a forward-looking optimal consumption plan subject to an estimated budget constraint. A similar dichotomy can be drawn with respect to explicit expectations about the effects of government expenditure shocks on future income. While exogenously imposed beliefs can strongly increase the magnitude of the multiplier, its size hardly changes if agents can learn from repeated shocks and beliefs become broadly 'rational'. In general, the size of the multiplier hence appears to strongly depend on the degree of rationality incorporated into agents' beliefs and decision-making processes. Nevertheless it was argued that even exogenously imposed (i.e. ex-ante non-rational) optimistic or pessimistic beliefs can be partly self-fulfilling, meaning that governments can potentially increase the impact of their fiscal policy measures if they are able to credibly manage agents' expectations and generate confidence about the effectiveness of fiscal policy.

As discussed in section 5, a simple welfare analysis shows that households are able to significantly increase their welfare from consumption by engaging in optimising behaviour, producing a welfare gain equivalent to around one percent of average per-period consumption. Due to the smaller overall effect on consumption, however, the welfare gain households experience from expansionary government expenditure shocks is somewhat lower in the 'permanent income' version of the model. In the baseline, the average welfare gain of each agent from an expansionary government expenditure shock equal to 10% of real GDP in the 20 periods following the shock is equivalent to around 0.9 consumption units per period, whereas under the modified consumption behaviour the welfare gain is equivalent to only around 0.39 consumption units. While an expansionary government expenditure shock is hence welfare-improving under both specifications, as it leads to a reduction in frictions and an easing of liquidity constraints, the effect is stronger if households are not optimising their consumption plans.

Beyond emphasising the dependence of the multiplier on consumption behaviour, beliefs and expectations, the paper also showed that the model gives rise to state-dependent multipliers, i.e. multipliers that are larger during recessions and smaller during booms. This empirically well-documented finding is robust under both versions of households' consumption behaviour. The model used in this study highlights a large range of potential channels for this effect, including liquidity constraints of households, credit constraints of firms, capacity constraints, and labour market tightness. Some of these may merit additional attention in empirical studies as an improved understanding of what precisely drives the state-dependence of multipliers could be important in designing fiscal policy interventions to be as effective as possible.

The results presented in this paper may be viewed as being broadly in line with the existing literature. As indicated in section 2, fiscal policy tends to be highly effective in macroeconomic ABMs, and the fairly large cumulative multiplier produced by the baseline model is certainly in line with this general tendency. It should be noted, however, that expenditure multipliers of similar or even greater magnitude can also be obtained in suitably specified DSGE models,[22] particularly when monetary policy does not react to changes in fiscal policy as is indeed also the case in the present model (Auerbach and Gorodnichenko, 2012b). By contrast, the cumulative multiplier produced by the modified model, which remains well below one and in later periods fluctuates around 0.5, is closer to those produced in RBC-type frameworks (Gechert and Will, 2012; Mitra et al., 2019) and consistent with the lower end of the range of available empirical estimates (cf. Batini et al., 2014). As documented by Gechert and Will (2012), the range of empirical and model-based estimates for all types of fiscal multipliers is relatively large, such that there is no consensus in the existing literature. Some of the most widely cited empirical estimates for government expenditure multipliers, such as those presented in Barro (1981), Hall (2009) and Ramey (2011a,b), fall roughly within the interval bounded by the multiplier in the baseline model and that in the modified one, although both frameworks could of course be calibrated to produce smaller or larger multipliers than those shown here. The general conclusion in this regard should hence be that both the baseline model featuring strongly simplified and heuristics-based consumption behaviour and the modified framework featuring optimising and forward-looking households are capable of producing realistic reactions to government expenditure shocks.

The pure focus on positive government expenditure shocks and households' consumption behaviour in the present paper, while enabling a very detailed analysis, clearly implies some limitations. For instance, it is likely that the effects of negative government expenditure shocks would not be precisely symmetrical to positive ones due to the various constraints and non-linearities incorporated in the model. Moreover, it would be interesting to examine how government expenditure can be optimally distributed among firms to maximise its impact, and to what extent fiscal multipliers are dependent on the type of instrument (as the present model would also allow for tax changes and transfers) or the type of expenditure (the government could conceivably also purchase capital goods for instance). It should also be noted that, despite the implementation of the alternative consumption behaviour, the majority of the modified model is still driven by relatively simple behavioural heuristics such that the gap to conventional frameworks remains wide. In addition, the absence of active monetary policy and the complete abstraction from considerations of Ricardian equivalence limits the comparability between the model used here and the DSGE literature. Finally, the least squares learning exercise considered in this paper is very limited in that households only attempt to learn the *impact effect* of government expenditure shocks on their budget constraints rather than dynamic effects, not to mention the law of motion of all other state variables relevant to them. Despite the limited nature of both the optimising behaviour and the learning dynamics, their in-

[22]In fact, for particular specifications of utility, multipliers may be infinite (Auclert and Rognlie, 2017).

corporation into a pre-existing macroeconomic ABM was fairly roundabout and implied a large modification. Nevertheless I believe that the paper represents an interesting addition to the literature which may serve to inspire similar approaches in other frameworks, contribute to an improvement in comparability across modelling paradigms, and provide a first step in addressing some common criticisms raised against macroeconomic ABMs.

8 Conclusion

This paper has provided an investigation of the effects of government expenditure shocks and the dependence of such effects on household consumption behaviour in a canonical macroeconomic ABM framework. Cumulative government expenditure multipliers were demonstrated to strongly depend on the way households form their demand for consumption goods; in particular it was shown that the inclusion of intertemporal utility maximisation behaviour subject to an estimated budget constraint leads to a multiplier that is significantly smaller than its baseline counterpart. Furthermore, the magnitude of cumulative multipliers can be increased or decreased through the introduction of exogenously imposed beliefs about the effect of government expenditure shocks on future income, highlighting the potentially important role of expectation management in determining the effects of government policy. In a final experiment, it was shown that under least squares learning, households are able to form broadly correct expectations about the impact effect of government expenditure shocks on their budget constraint and incorporate this expectation into their behaviour.

Beyond the thorough analysis of government expenditure multipliers, which is novel in the field of macroeconomic ABMs, the paper contributes to the literature by demonstrating how more forward-looking behavioural rules can be introduced into a well-known existing model. This both contributes to a better understanding of the role of heuristics in determining simulation results and may facilitate comparisons between macroeconomic ABMs and their mainstream counterparts. Finally, the combination of forward-looking behaviour with beliefs or learning makes it possible to incorporate announcement effects of fiscal policy into the model. Since agents are made to react systematically (and, under least squares learning, broadly correctly), to government expenditure shocks, the paper demonstrates a potential way in which the Lucas critique as applied to macroeconomic ABMs can be addressed at least partially.

Bibliography

Alesina, A., C. Favero, and F. Giavazzi (2019): *Austerity - When It Works and When It Doesn't*, Princeton, NJ: Princeton University Press.

Altissimo, F., S. Siviero, and D. Terlizzese (2002): "How Deep Are the Deep Parameters?" *Annales d'Économie et de Statistique*, 67/68, 207–226, https://doi.org/10.2307/20076348.

Arifovic, J. (2000): "Evolutionary algorithms in macroeconomic models," *Macroeconomic Dynamics*, 4, 373–414, https://doi.org/10.1017/s1365100500016059.

Ascari, G. and L. Rossi (2012): "Trend Inflation and Firms Price-Setting: Rotemberg Versus Calvo," *The Economic Journal*, 122, 1115–1141, https://doi.org/10.1111/j.1468-0297.2012.02517.x.

Ashraf, Q., B. Gershman, and P. Howitt (2016): "How inflation affects macroeconomic performance: An agent-based computational investigation," *Macroeconomic Dynamics*, 20, 558–581, https://doi.org/10.1017/S1365100514000303.

Assenza, T., A. Cardaci, D. Delli Gatti, and J. Grazzini (2018a): "Policy experiments in an agent-based model with credit networks," *Economics: The Open-Access, Open-Assessment E-Journal*, 12, 1–17, https://doi.org/10.5018/economics-ejournal.ja.2018-47.

Assenza, T., P. Colzani, D. Delli Gatti, and J. Grazzini (2018b): "Does fiscal policy matter? Tax, transfer, and spend in a macro ABM with capital and credit," *Industrial and Corporate Change*, 27, 1069–1090, https://doi.org/10.1093/icc/dty017.

Assenza, T. and D. Delli Gatti (2013): "E Pluribus Unum: Macroeconomic modelling for multi-agent economies," *Journal of Economic Dynamics & Control*, 37, 1659–1682, https://doi.org/10.1016/j.jedc.2013.04.010.

Assenza, T., D. Delli Gatti, and J. Grazzini (2015): "Emergent dynamics of a macroeconomic agent based model with capital and credit," *Journal of Economic Dynamics & Control*, 50, 5–28, https://doi.org/10.1016/j.jedc.2014.07.001.

Auclert, A. and M. Rognlie (2017): "A note on multipliers in NK models with GHH preferences," *Draft Paper*, web.stanford.edu/~aauclert/ghh_nk.pdf, accessed 24th April 2020.

Auerbach, A. and Y. Gorodnichenko (2012a): "Fiscal multipliers in recession and expansion," in *Fiscal policy after the financial crisis*, ed. by A. Alesina and F. Giavazzi, Chicago, IL: University of Chicago Press, 63–98.

——— (2012b): "Measuring the Output Responses to Fiscal Policy," *American Economic Journal: Economic Policy*, 4, 1–27, https://doi.org/10.1257/pol.4.2.1.

Barde, S. and S. van der Hoog (2017): "An empirical validation protocol for large-scale agent-based models," *ISIGrowth Working Paper*, 29/2017.

Barro, R. (1981): "Output effects of government purchases," *Journal of Political Economy*, 89, 1086–1121, https://doi.org/10.1086/261024.

Batini, N., L. Eyraud, L. Forni, and A. Weber (2014): "Fiscal Multipliers: Size, Determinants, and Use in Macroeconomic Projections," *IMF Technical Notes and Manuals*, September.

Branch, W. (2004): "The theory of rationally heterogeneous expectations: evidence from survey data on inflation expectations," *The Economic Journal*, 114, 592–621, https://doi.org/10.1111/j.1468-0297.2004.00233.x.

Bullard, J. and K. Mitra (2002): "Learning about monetary policy rules," *Journal of Monetary Economics*, 49, 1105–1129, https://doi.org/10.1016/S0304-3932(02)00144-7.

Caiani, A., E. Cattulo, and M. Gallegati (2018): "The effects of fiscal targets in a monetary union: a multi-country agent-based stock flow consistent model," *Industrial and Corporate Change*, 27, 1123–1154, https://doi.org/10.1093/icc/dty016.

Caiani, A., A. Russo, and M. Gallegati (2019): "Does inequality hamper innovation and growth? An AB-SFC analysis," *Journal of Evolutionary Economics*, 29, 177–228, https://doi.org/10.1007/s00191-018-0554-8.

Chiarella, C. and C. Di Guilmi (2017): "Monetary Policy and Debt Deflation: Some Computational Experiments," *Macroeconomic Dynamics*, 21, 214–242, https://doi.org/10.1017/S1365100515000450.

Cincotti, S., M. Raberto, and A. Teglio (2010): "Credit Money and Macroeconomic Instability in the Agent-based Model and Simulator Eurace," *Economics: The Open-Access, Open-Assessment E-Journal*, 4, 1–32, https://doi.org/10.5018/economics-ejournal.ja.2010-26.

Dawid, H. (1999): *Adaptive Learning by Genetic Algorithms - Analytical Results and Applications to Economic Models*, Berlin: Springer.

Dawid, H. and D. Delli Gatti (2018): "Agent-Based Macroeconomics," in *Handbook of Computational Economics, Vol. 4*, ed. by C. Hommes and B. LeBaron, London: Elsevier/North-Holland, 63–156.

Dawid, H., S. Gemkow, P. Harting, S. van der Hoog, and M. Neugart (2012): "The Eurace@Unibi Model - An Agent-Based Macroeconomic Model for Economic Policy Analysis," *University of Bielefeld Working Papers in Economics and Management*, 05-2012.

Dawid, H., P. Harting, and M. Neugart (2018a): "Cohesion policy and inequality dynamics: Insights from a heterogeneous agents macroeconomic model," *Journal of Economic Behavior & Organization*, 150, 220–255, https://doi.org/10.1016/j.jebo.2018.03.015.

——— (2018b): "Fiscal transfers and regional economic growth," *Review of International Economics*, 26, 651–671, https://doi.org/10.1111/roie.12317.

Dawid, H., P. Harting, S. van der Hoog, and M. Neugart (2019): "Macroeconomics with heterogeneous agent models: fostering transparency, reproducibility and replication," *Journal of Evolutionary Economics*, 29, 467–538, https://doi.org/10.1007/s00191-018-0594-0.

Del Negro, M., M. Giannoni, and F. Schorfheide (2015): "Inflation in the Great Recession and New Keynesian Models," *American Economic Journal: Macroeconomics*, 7, 168–196, https://doi.org/10.1257/mac.20140097.

Delli Gatti, D. and S. Desiderio (2015): "Monetary policy experiments in an agent-based model with financial frictions," *Journal of Economic Interaction and Coordination*, 10, 265–286, https://doi.org/10.1007/s11403-014-0123-7.

Delli Gatti, D., S. Desiderio, E. Gaffeo, P. Cirillo, and M. Gallegati (2011): *Macroeconomics from the bottom up*, Milano: Springer.

Delli Gatti, D. and J. Grazzini (2019): "Rising to the Challenge: Bayesian Estimation and Forecasting Techniques for Macroeconomic Agent-Based Models," *CESifo Working Paper*, 7894.

Dosi, G., G. Fagiolo, M. Napoletano, A. Roventini, and T. Treibich (2015): "Fiscal and monetary policies in complex evolving economies," *Journal of Economic Dynamics & Control*, 52, 166–189, https://doi.org/10.1016/j.jedc.2014.11.014.

Dosi, G., G. Fagiolo, and A. Roventini (2010): "Schumpeter meeting Keynes: A policy-friendly model of endogenous growth and business cycles," *Journal of Economic Dynamics & Control*, 34, 1748–1767, https://doi.org/10.1016/j.jedc.2010.06.018.

Dosi, G., M. Napoletano, A. Roventini, J. Stiglitz, and T. Treibich (2017a): "Rational Heuristics? Expectations and Behaviors in Evolving Economies with Heterogeneous Interacting Agents," *Sciences Po OFCE Working Paper*, 32.

Dosi, G., M. Pereira, A. Roventini, and M. Virgillito (2017b): "When more flexibility yields more fragility: the microfoundations of Keynesian aggregate unemployment," *Journal of Economic Dynamics & Control*, 81, 162–186, https://doi.org/10.1016/j.jedc.2017.02.005.

Evans, G. and S. Honkapohja (2001): *Learning and Expectations in Macroeconomics*, Princeton NJ: Princeton University Press.

Fagiolo, G. and A. Roventini (2017): "Macroeconomic policy in DSGE and agent-based models redux: New developments and challenges ahead," *Journal of Artificial Societies and Social Simulation*, 20, https://doi.org/10.18564/jasss.3280.

Fazzari, S. M., J. Morley, and I. Panovska (2014): "State-dependent effects of fiscal policy," *Studies in Nonlinear Dynamics & Econometrics*, 19, 285–315, https://doi.org/10.1515/snde-2014-0022.

Fernández-Villaverde, J. and J. Rubio-Ramírez (2005): "Estimating dynamic equilibrium economies: linear versus nonlinear likelihood," *Journal of Applied Econometrics*, 20, 891–910, https://doi.org/10.1002/jae.814.

Franke, R. and F. Westerhoff (2017): "Taking stock: A rigorous modelling of animal spirits in macroeconomics," *Journal of Economic Surveys*, 31, 1152–1182, https://doi.org/10.1111/joes.12219.

Friedman, M. (1957): *A Theory of the Consumption Function*, Princeton: Princeton University Press.

Galí, J., J. D. López-Salido, and J. Vallés (2007): "Understanding the effects of government spending on consumption," *Journal of the European Economic Association*, 5, 227–270, https://doi.org/10.1162/JEEA.2007.5.1.227.

Gechert, S. and A. Rannenberg (2018): "Which fiscal multipliers are regime-dependent? A meta-regression analysis," *Journal of Economic Surveys*, 32, 1160–1182, https://doi.org/10.1111/joes.12241.

Gechert, S. and H. Will (2012): "Fiscal Multipliers: A Meta Regression Analysis," *IMK Working Paper*, 97.

Gerali, A., S. Neri, L. Sessa, and F. Signoretti (2010): "Credit and Banking in a DSGE Model of the Euro Area," *Journal of Money, Credit & Banking*, 42, 107–141, https://doi.org/10.1111/j.1538-4616.2010.00331.x.

Geweke, J. (1985): "Macroeconometric Modeling and the Theory of the Representative Agents," *American Economic Review*, 75, 206–210, https://www.jstor.org/stable/1805597.

Giavazzi, F. and M. Pagano (1990): "Can Severe Fiscal Contractions be Expansionary? Tales of Two Small European Countries," *NBER Working Paper*, 3372.

Gigerenzer, G. (2008): *Rationality for Mortals - How People Cope with Uncertainty*, Oxford: Oxford University Press.

Gilli, M. and P. Winker (2003): "A global optimization heuristic for estimating agent based models," *Computational Statistics & Data Analysis*, 42, 299–312, https://doi.org/10.1016/S0167-9473(02)00214-1.

Gobbi, A. and J. Grazzini (2019): "A basic New Keynesian DSGE model with dispersed information: An agent-based approach," *Journal of Economic Behavior & Organization*, 157, 101–116, https://doi.org/10.1016/j.jebo.2017.12.015.

Grazzini, J. and M. Richiardi (2015): "Estimation of ergodic agent-based models by simulated minimum distance," *Journal of Economic Dynamics & Control*, 51, 148–165, https://doi.org/10.1016/j.jedc.2014.10.006.

Grazzini, J., M. Richiardi, and M. Tsionas (2017): "Bayesian estimation of agent-based models," *Journal of Economic Dynamics & Control*, 77, 26–47, https://doi.org/10.1016/j.jedc.2017.01.014.

Guerini, M. and A. Moneta (2017): "A method for agent-based models validation," *Journal of Economic Dynamics & Control*, 82, 125–141, https://doi.org/10.1016/j.jedc.2017.06.001.

Haldane, A. and A. Turrell (2018): "An interdisciplinary model for macroeconomics," *Oxford Review of Economic Policy*, 34, 219–251, `https://doi.org/10.1093/oxrep/grx051`.

Hall, R. (1978): "Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence," *Journal of Political Economy*, 86, 971–987, `https://doi.org/10.1086/260724`.

——— (2009): "By how much does GDP rise If the government buys more output?" *Brookings Papers on Economic Activity*, 2, 183–231.

Harting, P. (2015): "Stabilization policies and long term growth: Policy implications from an agent-based macroeconomic model," *Bielefeld Working Papers in Economics and Management*, 06-2015.

Holland, J. (1975): *Adaptation in Natural and Artificial Systems*, Ann Arbor, MI: University of Michigan Press.

Jerison, M. (1984): "Social Welfare and the Unrepresentative Representative Consumer," *State University of New York at Albany Discussion Paper*.

Kirman, A. (1992): "Whom or What Does the Representative Individual Represent?" *Journal of Economic Perspectives*, 6, 117–136, `https://doi.org/10.1257/jep.6.2.117`.

Krug, S. (2018): "The interaction between monetary and macroprudential policy: should central banks 'lean against the wind' to foster macro-financial stability?" *Economics: The Open-Access, Open-Assessment E-Journal*, 12, 1–69, `https://doi.org/10.5018/economics-ejournal.ja.2018-7`.

Lamperti, F., A. Roventini, and A. Sani (2018): "Agent-based model calibration using machine learning surrogates," *Journal of Economic Dynamics & Control*, 90, 366–389, `https://doi.org/10.1016/j.jedc.2018.03.011`.

Lucas, R. E. (1976): "Econometric policy evaluation: A critique," in *The Phillips Curve and Labor Markets; Carnegie-Rochester Conference Series on Public Policy*, ed. by K. Brunner and A. Meltzer, New York: Elsevier, 19–46.

Marcet, A. and T. Sargent (1989): "Convergence of least squares learning mechanisms in self-referential linear stochastic models," *Journal of Economic Theory*, 48, 337–368, `https://doi.org/10.1016/0022-0531(89)90032-X`.

Mitra, K., G. Evans, and S. Honkapohja (2019): "Fiscal policy multipliers in an RBC model with learning," *Macroeconomic Dynamics*, 23, 240–283, `https://doi.org/10.1017/S1365100516001176`.

Mittnik, S. and W. Semmler (2012): "Regime Dependence of the Fiscal Multiplier," *Center for Quantitative Risk Analysis Working Paper*, 5.

Modigliani, F. and R. Brumberg (2005): "Utility Analysis and the Consumption Function: An Interpretation of Cross-Section Data," in *The Collected Papers of Franco Modigliani, Vol. 6*, Cambridge, Mass.: The MIT Press, 3–46.

Napoletano, M., A. Roventini, and J. Gaffard (2017): "Time-varying fiscal multipliers in an agent-based model with credit rationing," *Kiel Institute for the World Economy Economics Discussion Paper*, 2017-112.

Popoyan, L., M. Napoletano, and A. Roventini (2017): "Taming macroeconomic instability: Monetary and macro-prudential policy interactions in an agent-based model," *Journal of Economic Behavior & Organization*, 134, 117–140, https://doi.org/10.1016/j.jebo.2016.12.017.

Qazizada, W. and E. Stockhammer (2014): "Government spending multipliers in contraction and expansion," *International Review of Applied Economics*, 29, 238–258, https://doi.org/10.1080/02692171.2014.983053.

Ramey, V. (2011a): "Can government purchases stimulate the economy?" *Journal of Economic Literature*, 49, 673–685, https://doi.org/10.1257/jel.49.3.673.

——— (2011b): "Identifying government spending shocks: It's all in the timing," *The Quarterly Journal of Economics*, 126, 1–50, https://doi.org/10.1093/qje/qjq008.

Ruppert, D., M. Wand, and R. Carroll (2003): *Semiparametric Regression*, Cambridge: Cambridge University Press.

Salle, I. (2015): "Modeling expectations in agent-based models - An application to central bank's communication and monetary policy," *Economic Modelling*, 46, 130–141, https://doi.org/10.1016/j.econmod.2014.12.040.

Salle, I., M. Sénégas, and M. Yildizoglu (2019): "How transparent about its inflation target should a central bank be? An agent-based model assessment," *Journal of Evolutionary Economics*, 29, 391–427, https://doi.org/10.1007/s00191-018-0558-4.

Salle, I. and M. Yildizoglu (2014): "Efficient Sampling and Metamodeling in Computational Economic Models," *Computational Economics*, 44, 507–536, https://doi.org/10.1007/s10614-013-9406-7.

Salle, I., M. Yildizoglu, and M. Sénégas (2013): "Inflation targeting in a learning economy: An ABM perspective," *Economic Modelling*, 34, 114–128, https://doi.org/10.1016/j.econmod.2013.01.031.

Seppecher, P. (2012): "Flexibility of wages and macroeconomic instability in an agent-based computational model with endogenous money," *Macroeconomic Dynamics*, 16, 284–297, https://doi.org/10.1017/s1365100511000447.

Simon, H. (1982): *Models of bounded rationality*, Cambridge MA: MIT Press.

Smets, F. and R. Wouters (2007): "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach," *American Economic Review*, 97, 586–606, https://doi.org/10.1257/aer.97.3.586.

Sutton, R. and A. Barto (1998): *Reinforcement Learning: An Introduction*, Cambridge, MA: MIT Press.

TEGLIO, A., A. MAZZOCCHETTI, L. PONTA, M. RABERTO, AND S. CINCOTTI (2019): "Budgetary rigour with stimulus in lean times: Policy advices from an agent-based model," *Journal of Economic Behavior & Organization*, 157, 59–83, https://doi.org/10.1016/j.jebo.2017.09.016.

VAN DER HOOG, S. AND H. DAWID (2019): "Bubbles, crashes, and the financial cycle: The impact of banking regulation on deep recessions," *Macroeconomic Dynamics*, 23, 1205–1246, https://doi.org/10.1017/S1365100517000219.

WARMEDINGER, T., C. CHECHERITA-WESTPHAL, AND P. HERNÃ¡NDEZ DE COS (2015): "Fiscal multipliers and beyond," *ECB Occasional Paper Series*, 162.

WINDRUM, P., G. FAGIOLO, AND A. MONETA (2007): "Empirical Validation of Agent-Based Models: Alternatives and Prospects," *Journal of Artificial Societies and Social Simulation*, 10, http://jasss.soc.surrey.ac.uk/10/2/8.html.

Appendix A: Parameter values

The parameter values of the baseline model used in this paper are identical to those provided by Delli Gatti and Grazzini (2019). Table 4 below lists all parameters together with an explanation and their value.

Table 4: Model parameters

Symbol	Description	Value
W	Number of workers	1000
F	Number of C-firms	100
N	Number of K-firms	20
z_c	Number of C-firms visited by consumers	2
z_e	Number of Firms visited by unemployed	5
z_k	Number of K-firms visited by C-firms	2
ξ	Memory parameter for baseline 'permanent income'	0.7382
χ	Propensity to consume out of wealth	0.0172
ρ_q	Quantity adjustment parameter	0.7301
ρ_p	Price adjustment random parameter	0.1649
μ	Bank's gross mark-up	1.007
η_k	Capital depreciation	0.03
ν	Probability to invest	0.3260
ϕ	Bank's leverage parameter	0.0024
ζ	Debt repayment rate	0.0328
δ	Memory parameter for capacity utilisation	0.1591
α	Labour productivity	0.5
κ	Capital productivity	1/3
ω	Dividend payout ratio (firms)	0.2
ω_b	Dividend payout ratio (bank)	0.3
$\overline{x}$	Normal/target capacity utilisation	0.85
η_i	Inventory depreciation	0.0781
λ_1	Bank's risk evaluation parameter (C-firms)	-15
λ_2	Bank's risk evaluation parameter (C-firms)	13
λ_1^k	Bank's risk evaluation parameter (K-firms)	-5
λ_2^k	Bank's risk evaluation parameter (K-firms)	5
r	Risk-free interest rate	0.01
s	Unemployment benefit ratio	0.5
τ	Income tax rate	0.0594
u^{up}	Upward wage adjustment parameter	0.1
u^{down}	Downward wage adjustment parameter	0.01
u^T	Unemployment threshold	0.1

Table 5 contains the parameters which are introduced into the original model by the modifications and subsequent experiments described in the main body of the paper.

Table 5: Parameters used in experiments

Symbol	Description	Value
$\mathcal{H}$	Length of estimation window and projection horizon	400
β	Discount factor	0.9999
$\frac{1}{\theta}$	Intertemporal elasticity of substitution	1.5

Table 5 – continued from previous page

Symbol	Description	Value
ι_0	Initial values of the exogenous belief shock	1.1; 0.9
ψ	Belief adaptation parameter	0.95
ℓ	Ex-ante duration of belief shock	40
ϵ	Weighting parameter for switching index	0.95
σ_1	Switching parameter (workers)	20
σ_1	Switching parameter (owners)	15
σ_2	Switching parameter	4

Appendix B: Testing the alternative consumption model

Recall that in section 5, households' consumption behaviour is modified by assuming that they maximise utility from consumption subject to an estimated budget constraint. This budget constraint is calculated based on two components, namely unweighted long-run averages of income flows (wages and dividends) and state transition probabilities (e.g. from employed to unemployed or from receiving to not receiving a dividend). Sequences of state transitions are represented by 2×2 transition matrices. For instance, the matrix for workers is given by

	E	**U**
E	π_{EE}	$1 - \pi_{EE}$
U	$1 - \pi_{UU}$	π_{UU}

This formulation implicitly assumes either that the stochastic process described by the matrix above satisfies the Markov property (i.e. that, for instance, the probability of being employed in the next period depends only on the state in the present period), or else that agents use this representation as an approximation of a process which is in fact not Markov. In either case it appears appropriate to test whether transition probabilities produced by the model can be sensibly represented in this form. Model simulation data can be used to test to what extent the stochastic state transition processes of workers and firm owners satisfy the Markov property. If the property holds, it should be true, for instance, that

$$p(E|EE) = p(E|EU) = p(E|E), \tag{29}$$

i.e. that the probability of an agent to be employed, given that they were employed in the previous two periods (EE) is equal to the probability given that they were employed in the previous period and unemployed two periods ago (EU). Simulated transition data across 200 Monte Carlo runs of the model is used to calculate the respective probabilities which are then compared. The results of this exercise are summarised by figures 14 to 16 in which boxplots are used to compare the probabilities.

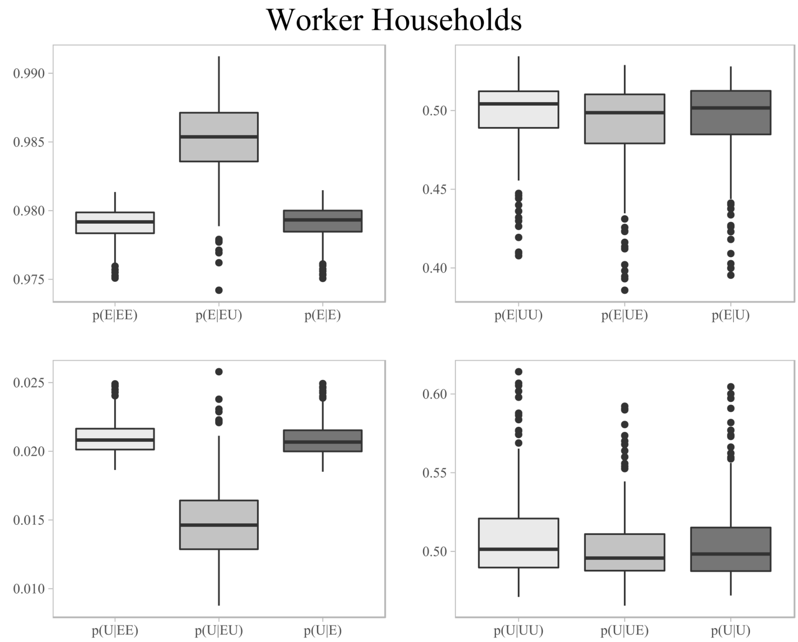

Figure 14: Transition probabilities of worker households

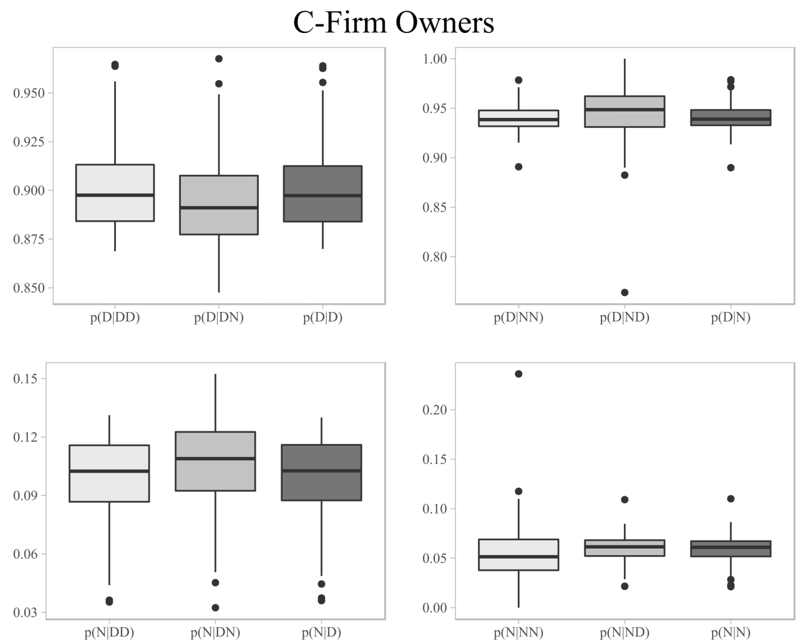

Figure 15: Transition probabilities of C-firm owner households

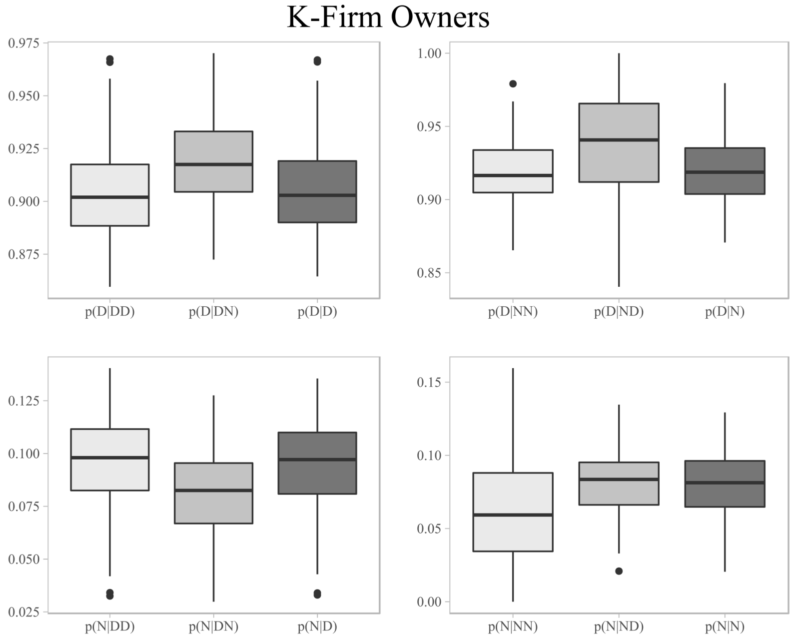

Figure 16: Transition probabilities of K-firm owner households

It can be seen that in all cases the differences between the calculated probabilities are very small and in almost all cases there is a strong overlap between the distributions. This suggests that while the Markov property may not hold exactly in all cases, the representation of probabilities using matrices of the form given above is a reasonably good approximation.

Recall that in the model, households use aggregate transition data to estimate the transition probabilities rather than looking only at their own history of transitions (meaning that in the absence of exogenously imposed beliefs, they all hold a common estimate). Figures 17 to 19 show the distributions of individual transition probabilities together with the population mean and the estimates resulting from the backward-looking calculation of aggregate transition probabilities. While the mean and the aggregate estimate are very close to each other in all cases, the distributions do show some amount of dispersion among agents' individual transition probabilities. This might be taken as an argument against the use of aggregate transition probabilities as an input into the behaviour of individual agents. As outlined in the model description, however, agents of the same class (workers, C-firm owners and K-firm owners) are identical in all factors which determine their transition probabilities. If agents are assumed to be informed enough to be aware of this fact, the rational implication would be for them to ignore their own idiosyncratic transition history as they should expect their own transition probabilities to converge to the population mean in the limit.

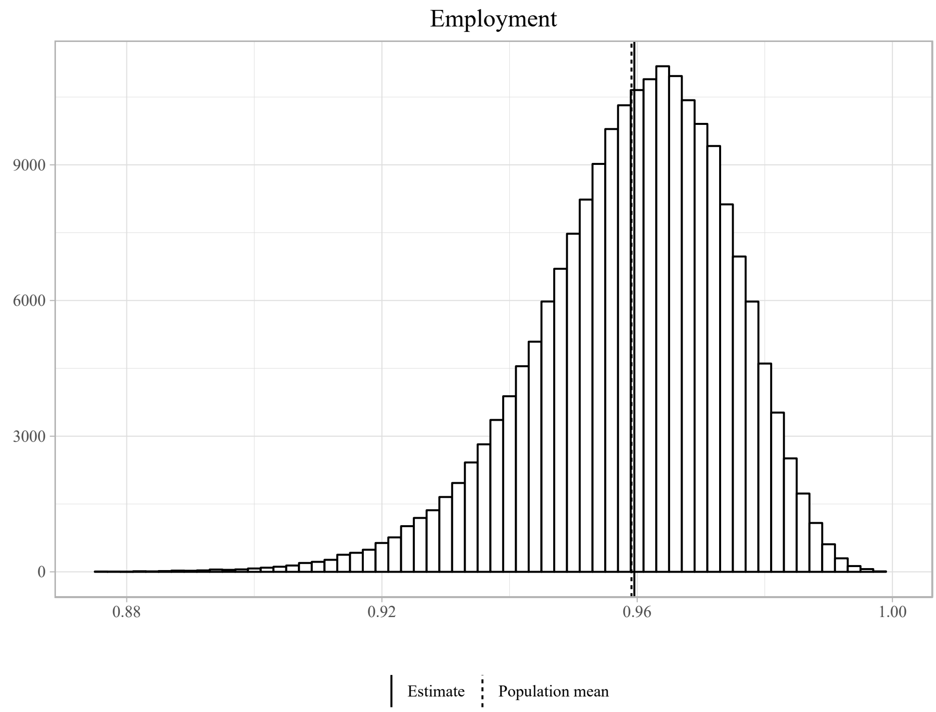

Figure 17: Distribution of individual worker transition probabilities

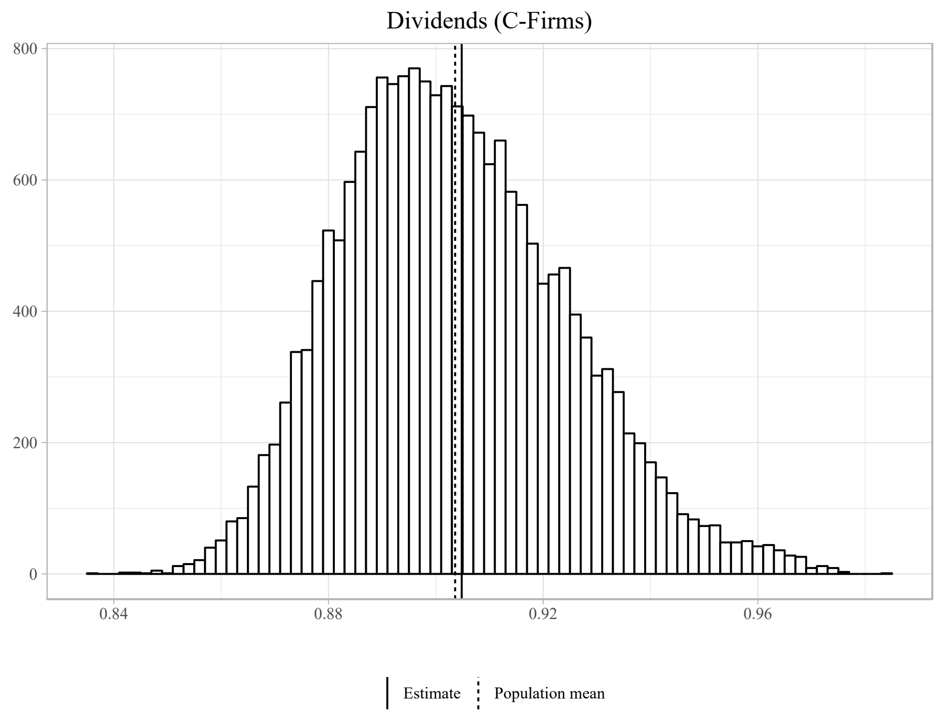

Figure 18: Distribution of individual C-firm owner transition probabilities

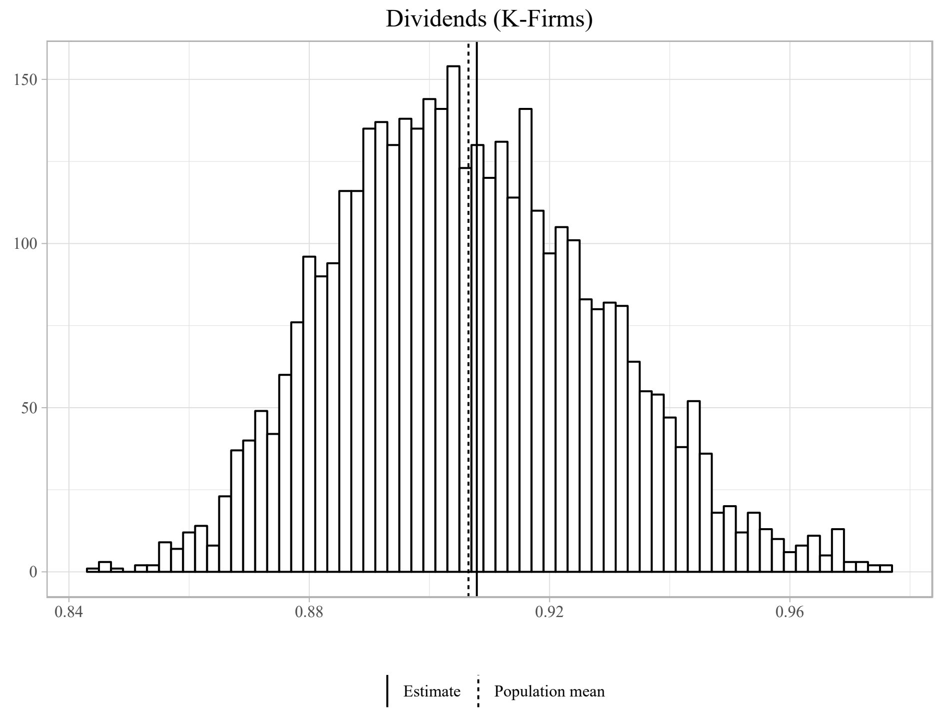

Figure 19: Distribution of individual K-firm owner transition probabilities