CESifo WORKING PAPERS

7894
2019
October 2019

Rising to the Challenge: Bayesian Estimation and Forecasting Techniques for Macroeconomic Agent-Based Models

Domenico Delli Gatti, Jakob Grazzini

CESifo

Impressum:

CESifo Working Papers
ISSN 2364-1428 (electronic version)
Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH
The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute
Poschingerstr. 5, 81679 Munich, Germany
Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de
Editor: Clemens Fuest
www.cesifo-group.org/wp

An electronic version of the paper may be downloaded
· from the SSRN website: www.SSRN.com
· from the RePEc website: www.RePEc.org
· from the CESifo website: www.CESifo-group.org/wp

CESifo *Working Paper No. 7894*
Category 6: Fiscal Policy, Macroeconomics and Growth

Rising to the Challenge: Bayesian Estimation and Forecasting Techniques for Macroeconomic Agent-Based Models

Abstract

We propose two novel methods to “bring ABMs to the data”. First, we put forward a new Bayesian procedure to estimate the numerical values of ABM parameters that takes into account the time structure of simulated and observed time series. Second, we propose a method to forecast aggregate time series using data obtained from the simulation of an ABM. We apply our methodological contributions to a medium-scale macro agent-based model. We show that the estimated model is capable of reproducing features of observed data and of forecasting one-period ahead output-gap and investment with a remarkable degree of accuracy.

JEL-Codes: C110, C130, C530, C630.

Keywords: agent-based models, estimation, forecasting.

Domenico Delli Gatti
Department of Economics and Finance
& Complexity Lab in Economics
Catholic University of Milan / Italy
domenico.delligatti@unicatt.it

Jakob Grazzini
Department of Economics and Management
University of Pavia / Italy
jakob.grazzini@unipv.it

1 Introduction

Agent-based models (ABMs) have been around for decades. At present they are routinely used in hard sciences and in a number of soft sciences. Applications to economics, on the contrary, are relatively recent. In fact, most macroeconomic ABMs have seen the light after the 2007 financial crisis. During its short life, the agent-based economic literature has grown at a dramatic pace, gaining the attention of both academics and policy makers. While the former have raised a few objections (on which we will elaborate momentarily), the latter – burdened with the daunting task of taming the beast of financial instability – have been eager to add ABMs to their modeling toolbox (Haldane and Turrell, 2018).

The main advantage of the AB approach is "flexibility", which allows macroeconomists to easily investigate the consequences of "complexity" – due essentially to the interaction of agents in a bounded rationality setting – on macroeconomic outcomes.

This same flexibility, however, is the main source of criticisms. According to the critics, macro ABMs "lack discipline" both theoretically and empirically. Theoretically, AB microfoundations are not (generally) derived from "first principles" so that any behavioral rule (*heuristic*) seems to be admissible. In other words, the AB modeller runs the risk of getting lost in the "wilderness of bounded rationality".[1]

Empirically, it is admittedly very difficult to "bring these models to the data" due to their built-in complexity. So far, the common strategy to validate an ABM empirically has consisted in launching a large number of simulations based on calibrated parameter values to generate a mass of artificial data. The ABM is empirically validated if the statistical properties of the artificial data replicate the "stylized facts" extracted from the empirical data.[2]

There are essentially two lines of attack to this empirical strategy which motivate our contributions in this paper. The critics argue, first of all, that ABMs are characterized by too many parameters, which are set by means of dubious "calibration" techniques. We vividly remember an interesting discussion over dinner with friends (who happen to be also DSGE critics

[1] There are (quite obvious) counterobjections. A discussion of this issue, however, is beyond the scope of the present paper.

[2] It has been verified *ad abundantiam* that ABMs can replicate not only the aggregate evidence (e.g., the dynamics of GDP) but also the cross-sectional evidence (e.g. the distribution of firm's size). Standard DSGE models relying on the representative agent hypothesis can replicate only aggregate stylized facts. A new wave of Heterogeneous Agents New Keynesian (HANK) DSGE models is rapidly overcoming this limit. A discussion of the pros and cons of this approach is beyond the scope of the present paper

of ABMs) a few years ago. When we objected that also medium and big-sized DSGE models have plenty of parameters, one of them shrugged and said: “No problem: we can estimate them!”. Since that dinner, we have started a long (and not yet completed) journey to take up this challenge. The estimation of the parameters of macro ABMs is indeed a daunting task but it is well worth tackling.

Second, so far the empirical strategy to validate an ABM has not developed methods to forecast future macroeconomic developments. Forecasting the future is key in understanding the impact of policy measures. Moreover, by looking at the forecasting errors of different models, it is possible to compare their performances.

In this paper, therefore, we make two contributions. First, we present and discuss at length a new Bayesian method to estimate the parameters of a macro ABM. This new method is obtained by augmenting the technique proposed in Grazzini et al. (2017) with an appropriate exploitation of the time structure of observed and simulated time series in the estimation process. We have applied this new method to a set of simple estimation problems and shown that it improves the identification of estimated parameters with respect to the original technique.

We use this new method to estimate a medium-sized *macroeconomic agent-based model.*To the best of our knowledge, this is the first time such an attempt is carried out. We compare the properties of the output of simulations of the estimated model to those of the observed empirical counterparts and show that they follow similar patterns. Moreover, we perform impulse-response analysis to investigate the behavior of the model.

Second, we provide a simulation-based forecasting method for macroeconomic ABMs (which exploits the Bayesian estimation techniques mentioned above). To the best of our knowledge, this is the first attempt to use an agent-based model to forecast macroeconomic aggregate time series. Essentially, we use the simulated time series of the estimated model to compute the conditional probability distribution of a vector of aggregate future (and therefore unknown) variables given the variables observed in the present period. This allows to use the model to compute one-period ahead forecast. Moreover we show that the estimated model can be used to provide a measure of macro risk, i.e., a quantitative assessment of the probability of a recession. This last exercise opens to the possibility of using ABMs as early-warning signaling tools. In fact, one of the strengths of ABMs is that they are often able to reproduce and explain endogenous business-cycle fluctuations, and, possibly, the emergence of crises. Therefore, our methodology could potentially provide a way to measure the probability of a future crisis, or more in general, of a future change of the business-cycle. We leave this area to future research.

The paper is organized as follows. In section 2 we put the contribution of this paper in context presenting in a succinct way the related literature. In section 3 we overview the building blocks and the main equations of the macroeconomic ABM to which we will apply our estimation and forecasting techniques. Section 4 is devoted to a discussion of the estimation procedure. In section 5 we present the evidence on moment matching by comparing the statistical properties of the time series originated by the estimated model's simulations to the properties of the real world macroeconomic time series. In Section 6 we investigate the behavior of the model using impulse response functions and in Section 7 we describe the forecasting method and apply it to the agent-based model. Section 8 concludes. Finally, Appendix A and Appendix B apply, respectively, the estimation and forecasting methods to simple models to evaluate their performances in known environments.

2 Related literature

In the last decade or so, a few competing macro AB frameworks have emerged. They are extensively analyzed in Dawid and DelliGatti (2018). Delli Gatti, Gallegati and co-authors in Ancona and Milan have proposed one of these frameworks, based on the notion of *Complex Adaptive Trivial Systems* (CATS). Complex aggregate behaviour stems from the *interaction* of simple (almost Trivial) behavioural heuristics. The single most important CATS framework, which is at the core of a wave of subsequent models, is described in chapter 3 of the book "Macroeconomics from the Bottom Up" (DelliGatti et al., 2011). This model, identified by the acronym MBU, features households, firms and banks. In MBU there is no capital: firms use only labor to produce consumption goods. In order to apply our estimation and forecasting techniques, we build a model belonging to the CATS class with capital goods, similar in spirit to previous models described in Assenza et al. (2015) and Assenza et al. (2018). For this reason, we will refer to this model with the acronym CC-MABM (Macroeconomic ABM with Capital and Credit). The model is described in section 3.

The empirical strategy to validate an ABM, as we said in the introduction, has coincided so far essentially with the replication of "stylized facts", both at the macro level (e.g., irregular fluctuations of GDP), and at the cross-sectional level (e.g., power law distribution of firm's size) by means of artificial data generated by simulations based on calibrated parameters (see for example Dawid et al., 2012; Dosi et al., 2013; Assenza et al., 2015).

To the best of our knowledge, the first attempt at proposing a technique to estimate the parameters of an ABM is Gilli and Winker (2003). The issue has been underesearched for a decade or so after this seminal paper. The rapidly growing AB literature has revived interest in the issue so that in the last half-decade a fair number of new papers has been published that proposed estimation techniques. In particular, Grazzini and Richiardi (2015), Grazzini et al. (2017), Kukacka and Barunik (2017) and Lux (2018) apply a number of econometric procedures to estimate simple agent-based models while Lamperti et al. (2018) and Barde and van der Hoog (2017) develop calibration techniques for agent-based models. In a recent paper, Platt (2019) compares and contrasts ABM estimation techniques to find that the Bayesian estimation method generally out-performs alternative procedures. In this paper, we propose an improved Bayesian estimation method.

Despite its importance, forecasting is extremely rare in agent-based modeling literature. The only example we are aware of is Recchioni et al. (2015), who uses a calibrated financial model a lá Brock and Hommes (1998) to forecast stock market prices. In this paper, we aim at filling this gap by proposing a simulation-based forecasting technique for macroeconomic agent-based models.

3 The Model

The estimation and forecasting methods we propose can be applied to any ABM. However, to make the case in favour of these methods more specific, we apply them to an actual medium-scale macro ABM. In this section we present the main features of the model.

The economy consists of four sectors: firms, households, the banking system and the public sector. There are N_F firms, of which N_F^k produce capital goods (K-firms) and N_F^c produce consumption goods (C-firms). K-firms use labor to produce and sell capital goods to C-firms. C-firms employ labour and capital to produce and sell consumption goods to households.

There are N_H households, of which N_W are workers and N_F are"capitalists", i.e., firm owners.[3] Workers supply labor to firms. If employed, each worker earns a wage; if unemployed, she will receive an unemployment subsidy equal to a fraction (the replacement rate) of the wage of the employed. The market for labour is characterized by *search and matching*: unemployed workers search for a job at firms and stop searching when a match occurs. Firm owners earns dividends proportional to the firm's profit if the latter is positive.

For simplicity, only wages are taxed. Hence the disposable income of the employed worker is a fraction of the wage. Unemployed workers and capitalists do not pay taxes so that their disposable income coincides with the unemployment subsidies and dividends respectively.

Both workers and capitalists are consumers, i.e., buyers on the market for consumption goods (C- goods). The market for C-goods is characterized by search and matching: households search for trading opportunities at C-firms and stop searching when a match occurs. Also the markets for capital goods (K-goods) is characterized by search and matching: C-firms search for trading opportunities at K-firms and stop searching when a match occurs.

For simplicity, the banking sector consists of only one bank. Households and firms hold deposits at the bank, which, for simplicity, are not remunerated. The bank also extends loans to firms which need to fill the financing gap (production costs net of internally generated funds). Since there is only one bank, by construction there cannot be search and matching on the market for credit. The bank sets the price (interest rate on loans) and the quantity of credit supplied to firms. The price/quantity decision of the bank is based on the assessment of the borrowing firm' s financial fragility, which is a proxy of the credit risk run by the bank. In particular, the interest rate on loans is set adding a mark up (external finance premium) to the risk free interest rate (i.e., the interest rate on Government bonds). The external finance premium, in turn, is increasing with the borrower's leverage. In this setting, a firm may well face a limit on the amount of credit it can get (credit rationing).

The public sector collects taxes on wage income and provides unemployment subsidies. In case of a public sector deficit, bonds are issued and sold to the bank. The interest rate on Government bonds is equal to the risk free interest rate. Figure 1 depicts agents' interactions on the five markets: deposits, credit, labor, K-goods and C-goods. The way in which markets work will be described in the following.

[3]By assumption, there is one owner per firm.

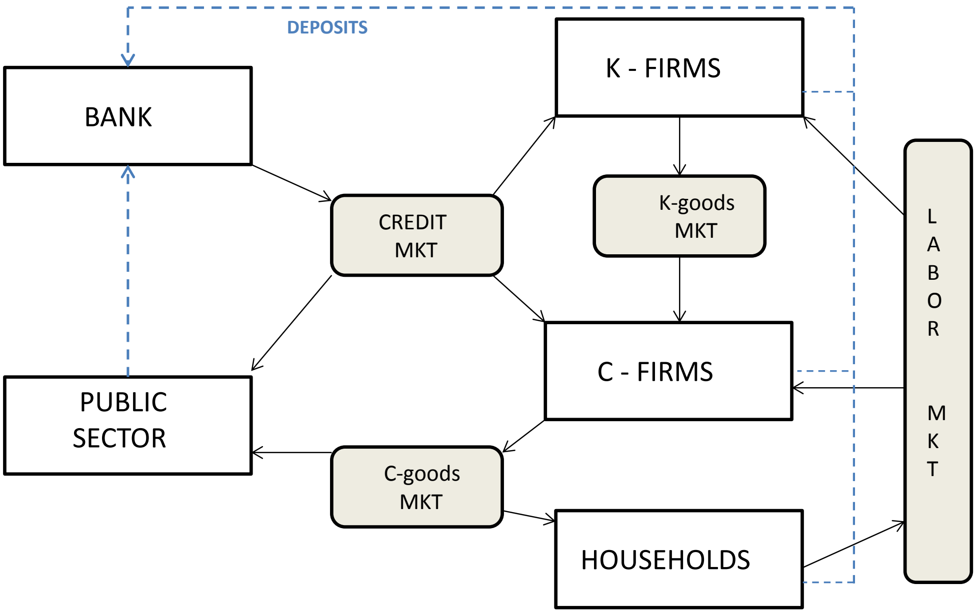

Figure 1: Agents and markets

3.1 Households

As we said, households can be either workers or firm owners.

Each worker supplies 1 unit of labour inelastically. If employed, she receives the nominal wage w_t and pays out a fraction t_w (the tax rate) of this wage to the Government. The tax rate is one of the parameters we will estimate (see table 2). If unemployed, the worker searches for a job visiting a subset Z_e of firms (chosen at random among the population of firms) and applies to the first one who has posted vacancies. Since the wage is uniform across firms and labour is homogeneous, once an unemployed worker finds a firm with an unfilled vacancy she stops searching and the match occurs. Unemployed workers who have not succeeded in finding a job (because firms in their subset did not post vacancies or because they have already filled all the vacancies), receive an unemployment subsidy from the Government equal to a fraction of the nominal wage.

The owner of the f-th firm receives a fraction τ (the pay-out ratio) of the current profit π_{ft}, if the latter is positive. If a firm faces a loss, it will not distribute dividends and will reduce its net worth correspondingly. Whenever a firm goes bankrupt, another one will replace it. We assume that the initial equity of the entrant firm is provided by the capitalist that owned the bankrupt firm. The capitalist's wealth, therefore, will be reduced correspondingly.

Households (workers and capitalists) are consumers/savers. Each consumer sets her *consumption budget* equal to the sum of her *permanent income* and a fraction of her wealth. Permanent income is a weighted average of current and past incomes (wages and unemployment subsidies in the case of workers, dividends in the case of firm owners). In computing the weighted average of past incomes, the consumer uses a memory parameter Ξ. The sensitivity of consumption to wealth is captured by the parameter χ. These two parameters will be estimated (see table 2).

Once a consumer has defined her consumption budget, she visits a subset Z_c of C-firms chosen at random and ranks them in ascending order of price: the consumer starts purchasing goods from the firm which posts the lowest price; if she still has resources to be spent on C-goods, she will buy from the second one (the firm which posts the second lowest price) and so on until the consumption budget is exhausted.

If the consumer's demand has not been completely satisfied after Z_c visits, she is forced to save the unspent portion of the consumption budget. Hence, savings are equal to the difference between actual disposable income and the budget allocated to consumption plus the involuntary savings possibly deriving from unsatisfied demand.

Savings are deposited at the bank. By assumption households do not hold Government bonds. Therefore households' wealth takes the form only of deposits at the bank.

3.2 C-firms

Each C-firm has some *market power* on its own local market (i.e. there are as many local C-markets as there are C-firms).

The firm has to set individual price and quantity under uncertainty. It knows from experience that if it charges higher prices it will get smaller demand but it does not know the actual demand schedule (i.e., how much the consumers would buy at any given price). Hence the firm is unable to maximize profits since the marginal revenue is unknown. The best the firm can do in this setting consists in charging a price as close as possible to the average price (approximately equal to the average price set by competitors) and producing a quantity as close as possible to (expected) demand. In this way the firm minimizes involuntary inventories (in case of excess supply) or the queue of unsatisfied customers (in case of excess demand).

In t, the i-th C-firm, $i = 1, 2..., N_F^c$, must choose the price and desired output for $t+1$ $\left(P_{it+1}, Y_{it+1}^*\right)$. Desired output is determined by expected demand $Y_{it+1}^* = Y_{it+1}^e$. The firm's information set in t consists of (i) the average price level P_t and (ii) excess demand

$$\Delta_{it} := Y_{it}^d - Y_{it} \tag{1}$$

where Y_{it}^d is actual demand and Y_{it} is actual output in t. Δ_{it} shows up as a queue of unsatisfied customers if positive; as an inventory of unsold goods if negative. By assumption C-goods are not storable. Therefore involuntary inventories cannot be employed to satisfy future demand. Notice that Δ_{it} is a proxy of the *forecasting error* $\epsilon_{it} := Y_{it}^d - Y_{it}^e$ where Y_{it}^e is expected demand formed in t-1 for t.[4]

A firm can decide either to update the current price or to vary the quantity to be produced. The decision process is based on two *rules of thumb* which govern price changes and quantity changes respectively.

The *price adjustment* rule is:

$$P_{it+1} = \begin{cases} P_{it}(1+\eta_{it}) & \text{if } \Delta_{it} > 0;\ P_{it} < P_t \\ \\ P_{it}(1-\eta_{it}) & \text{if } \Delta_{it} \leq 0;\ P_{it} > P_t \end{cases} \tag{2}$$

where η_i is a random positive parameter drawn from a distribution with support $(0, \bar{\eta})$.

The signs of Δ_{it} and of the difference $P_{it} - P_t$ dictate the direction of price adjustment but the *magnitude* of the adjustment is stochastic and bounded by the length of the support of the distribution. We also assume that the firm will never set a price lower than the average cost.

Since the quantity to be produced is equal to expected demand, the *quantity adjustment* rule takes the form of an updating rule for expected demand:

$$Y_{it+1}^* = Y_{it+1}^e = \begin{cases} Y_{it} + \rho \mathbf{1}_{[P_{it} > P_t]} \Delta_{it} & \text{if } \Delta_{it} > 0 \\ \\ Y_{it} + \rho \mathbf{1}_{[P_{it} < P_t]} \Delta_{it} & \text{if } \Delta_{it} \leq 0 \end{cases} \tag{3}$$

[4]In fact $\Delta_{it} = \epsilon_{it} + (Y_{it}^e - Y_{it})$ where the expression in parentheses is a non-negative discrepancy between expected demand and actual production.

where ρ is a positive parameter, smaller than one.

$\mathbf{1}_{[P_{it}>P_t]}$ is an *indicator function* equal to 1 if $P_{it} > P_t$, 0 otherwise. Analogously, $\mathbf{1}_{[P_{it}<P_t]}$ is an *indicator function* equal to 1 if $P_{it} < P_t$, 0 otherwise.

The *magnitude* of the quantity adjustment is not stochastic (as in the case of price adjustment) but determined by excess demand. If we assume that excess demand is close to the forecasting error, we can interpret Eq. (3) as a standard *adaptive mechanism* to update demand expectations. By iteration, as it is well known, desired production in $t+1$ will be determined by the weighted average of past quantities with exponentially decaying weights. The price and quantity adjustment parameters $\bar{\eta}$ and ρ will be estimated (see table 2).

Technology is represented by a Leontief production function: $Y_{it} = \min(\alpha N_{it}, \kappa\omega_{it}K_{it})$ where α and κ represent labor and capital productivity respectively and $\omega_{it} \in (0,1]$ is the rate of capacity utilization at firm i. When capital is employed at full capacity – i.e. when $\omega_{it} = 1$ – output will be $\hat{Y}_{it} = \kappa K_{it}$. This is "full capacity" output. Given undepreciated capital, actual capital in $t+1$, K_{it+1} is given – being determined by investment carried out in t, I_{it} (to be discussed momentarily) – and cannot be modified in $t+1$. Hence in period $t+1$ the maximum attainable output is $\hat{Y}_{it+1}$.

Once a decision has been taken on desired output in $t+1$, the firm retrieves from the production function how much capital it needs in t+1 to reach that level of activity (capital requirement): $K^*_{it+1} = Y^*_{it+1}/\kappa$. If actual capital is greater than the capital requirement, the desired *rate of capacity utilization* will be smaller than one. If actual capital is smaller than the capital requirement, the former will be utilized at full capacity (the rate of capacity utilization will be one) but desired output will not be reached.

Whatever the scenario, if actual employment in t N_{it} is smaller than labor required to reach the feasible level of activity in t+1, the firm will post vacancies. If the opposite holds true the firm will fire workers.

Firms set the nominal wage on the basis of labour market conditions captured by the distance between the current unemployment rate u_t and a threshold $\hat{u}$. Whenever the unemployment rate is above (below) the threshold firms will cut (increase) the wage. The wage updating mechanism therefore is:

$$w_{t+1} = \begin{cases} w_t \left[1 + u_{up} \left(\hat{u} - u_t\right)\right]; & \hat{u} - u_t > 0 \\ w_t \left[1 + u_{down} \left(\hat{u} - u_t\right)\right] & \hat{u} - u_t < 0 \end{cases}$$

where u_{up} and u_{down} are positive parameters. We will assume that $u_{up} > u_{down}$ to capture the downward stickiness of nominal wages.

As mentioned above, the firm determines in t the capital stock which will be available for use in production in $t+1$ by means of investment I_{it}. By assumption, in planning investment, the firm sets a *benchmark* equal to the capital stock used in production "on average" since the beginning of activity $\bar{K}_{it}$. This, in turn, is computed by means of an adaptive algorithm, i.e., the weighted average of past utilized capital from the beginning of activity until $t-1$ with exponentially decreasing weights. In computing this weighted average, the firm employs a memory parameter $\upsilon \in (0,1)$. Capital depreciates at the rate δ. Moreover we assume that C-firms may invest in each period with a probability γ. Hence investment necessary "on average" to replace worn out capital is $\frac{\delta}{\gamma}\bar{K}_{it}$. The parameters υ and γ will be estimated (see Table 2).

We assume, moreover, that the firm plans to maintain, in the long run, a capital stock buffer. Therefore the *target* capital stock is equal to $K^T_{it+1} = \frac{1}{\bar{\omega}}\bar{K}_{it}$ where $\bar{\omega} \in (0,1)$ is the desired long run capacity utilization rate (this parmaeter also will be estimated). Net investment is $K^T_{it+1} - K_{it-1}$. Therefore gross investment in t is:

$$I_{it} = \left(\frac{1}{\bar{\omega}} + \frac{\delta}{\gamma}\right)\bar{K}_{it} - K_{it-1}$$

Once investment has been determined, the i-th C-firm visits a subset Z_k of K-firms chosen at random to purchase capital goods. Visited K-firms are ranked in ascending order of price and the C-firm starts buying capital goods from the K-firm which has posted the lowest price. If this purchase does not exhaust planned investment, the C-firm will purchase capital goods also at the second firm in the ranking and so on. If the C-firm's demand for K-goods has not been completely satisfied after Z_k visits, it is forced to "save" the unspent portion of the investment budget. Therefore actual investment may turn out to be lower than planned investment.

3.3 K-firms

In setting the price of capital goods, K-firms follow essentially the same heuristic adopted by C-firms (see equation (2)). The quantity adjustment rule departs from the one adopted by C-firms (see equation (3)) to take into account the fact that K-goods are durable and therefore storable: inventories of capital goods can be carried on from one period to another and sold in the future. The quantity adjustment rule of the j-th K-firm, $j = 1, 2, ..., N_F^k$ therefore is:

$$Y_{jt+1}^* = Y_{jt+1}^e - Y_{jt}^k = \begin{cases} Y_{jt} + \rho \mathbf{1}_{[P_{jt} > P_t^k]} \Delta_{jt} - Y_{jt}^k & \text{if } \Delta_{jt} > 0 \\ Y_{jt} + \rho \mathbf{1}_{[P_{jt} < P_t^k]} \Delta_{jt} - Y_{jt}^k & \text{if } \Delta_{jt} < 0 \end{cases} \tag{4}$$

where Y_{jt+1}^* is the desired scale of activity, Y_{jt+1}^e is expected demand, Y_{jt}^k is the fraction of the inventory of capital goods held by firm j at time t which can be used to face demand in $t+1$, Δ_{jt} is excess demand, P_{jt} is the individual price and P_t^k is the average price of capital goods. Y_{jt}^k is computed applying a rate of depreciation δ^k to the stock of unsold machine tools accumulated until t. This parameter will be estimated. Since K-firms are endowed with a linear production function whose only input is labour, once the price-quantity configuration has been set, a K-firm may post vacancies or fire workers in order to fulfill labor requirements.

3.4 Credit

Once the quantity to be produced has been set and the cost of inputs determined, the firm has to deal with financing. Consider a generic firm, indexed by $f = 1, 2..., N_F$. If the firm's internal liquidity (i.e., the current deposits held at the bank) M_{ft-1} is "abundant", i.e., greater than the costs to be incurred, the firm can self-finance production. If, on the other hand, liquidity is not sufficient to carry out production up to the desired level, the firm applies for a loan to fill its financing gap:

$$F_{ft} = \max(wN_{ft} + \mathbf{1_c} P_{t-1}^k I_{ft} - M_{ft-1}, 0)$$

where $\mathbf{1_c}$ is an indicator function which assigns value 1 to C-firms and 0 to K-firms. In fact only C-firms purchase capital goods. By definition, the financing gap is the demand for loans of the f-th firm.

For simplicity we assume there is only one bank. The bank collects deposits from all the firms and households, supplies credit to firms and purchases Government bonds. The bank decides (i) the interest rate to be charged to each borrower and (ii) the size of the loan that will be actually extended (which may be different from the borrower's financing gap). As we will see momentarily, both decisions will be affected by the borrower's leverage λ_{ft}:

$$\lambda_{ft} = \frac{L_{ft}}{E_{ft} + L_{ft}}$$

where L_{ft} is the firm's debt and E_{ft} is equity or net worth.

The interest rate charged by the bank to each firm is determined as a mark up μ on the risk free interest rate r. Adopting the expression pioneered by Bernanke and Gertler, the firm is charged an *external finance premium* increasing with the probability of default which in turn is (non-linearly) increasing with leverage. The probability of default for the i-th C-firm is:

$$p(\lambda_{it}) = \frac{e^{b_{0c}+b_{1c}\lambda_{it}}}{1 + e^{b_{0c}+b_{1c}\lambda_{it}}}$$

Analogously, the probability of deafult for the j-th K-firm is:

$$p(\lambda_{jt}) = \frac{e^{b_{0k}+b_{1k}\lambda_{jt}}}{1 + e^{b_{0k}+b_{1k}\lambda_{jt}}}$$

In the end, therefore the interest rate charged to the generic f-th firm is a function of the risk-free interest rate and of the firm's leverage:

$$r_{ft} = \mu f(r, \lambda_{ft}) \tag{5}$$

where the function $f(.)$ is increasing with all the arguments.[5]

In order to determine the size of the loan, the bank sets first a tolerance level for the potential loss Γ_b on credit extended (to any borrower) as a fraction ϕ of its net worth: $\Gamma_b = \phi E_{bt}$. The borrower's total debt in t will be $\Phi_{ft} + L_{ft-1}$ where Φ_{ft} is the new credit line to be supplied in t. We assume the bank sets the new credit line in order to equate the expected loss on loans extended to the f-th firm to the tolerance level: $(\Phi_{ft} + L_{ft-1})p(\lambda_{ft}) = \phi E_{bt}$. Therefore the new

[5]For the specification of $f(.)$ see Assenza et al. (2015).

credit line is:

$$\Phi_{ft} = \frac{\phi}{p(\lambda_{ft})} E_{bt} - L_{ft-1}$$

Given the current exposure of the bank to the firm, the new credit line is increasing with the bank's net worth and decreasing with the firm's leverage. The size of the loan actually granted to firm f at time t will be

$$\hat{L}_{ft} = \min(\Phi_{ft}; F_{ft})$$

i.e., the minimum between new credit line and the financing gap. If the latter is greater than the former the firm will face a *borrowing constraint* and therefore will be forced to scale down production. Finally, firms in each period repay a fraction ϑ of the total debt to the bank.

The parameters ϕ, μ and ϑ will be estimated (see Table 2).

3.5 Net Worth Updating

In every period, the firm's net worth E_f is updated by means of retained profits:

$$E_{ft+1} = E_{ft} + (1-\tau)\pi_{ft}$$

where τ is the dividend payout ratio and π_{ft} is the firm's profit:

$$\pi_{ft} = P_{ft} \min(Y_{ft}, Y_{ft}^d) - (wN_{ft} + \mathbf{1_c}\omega_{ft}\delta K_{ft} + r_{ft}L_{ft})$$

Whenever the firm's equity turns negative, the firm goes bankrupt and is replaced by a new one. The owner of the bankrupt firm confers the initial net worth of the entrant firm (out of her own private wealth). Hence, the population of firms is kept constant.

Also the bank's net worth is updated by means of retained profits:

$$E_{bt+1} = E_{bt} + (1-\tau)\pi_{bt} - BD_t$$

where π_{bt} is the bank's profit and BD_t is *bad debt*, i.e., the book value of non-performing loans. We assume that the bank does not remunerate deposits while it earns interests on loans (if borrowers are solvent) and the risk free interest rate on Government bonds. Hence, in every period the bank's profit is

$$\pi_{bt} = \sum_{s=1}^{N_F^s} r_{st} L_{st} + rB_{t-1}$$

where N_F^s is the number of solvent firms. Hence

$$BD_t = \sum_{n=1}^{N_F^n} L_{nt}$$

where N_F^n is the number of insolvent firms.

3.6 The Public Sector

The public sector raises tax revenues on wage income $TA_t = t_w w_t N_t$ where N_t is total employment and extend unemployment subsidies to households and interest payments on outstanding Government bonds to the bank. Total unemployment subsidies are $US_t = zw_t(N_W - N_t)$, where z is the "replacement rate". Interest payments are $INT_t = rB_{t-1}$ where r is the cost of public debt, equal, by assumption, to the risk free interest rate.

A public sector deficit occurs when taxes turn out to be lower than transfers (i.e., the sum of unemployment subsidies and interest payments). In this case, the Government will issue new bonds. For simplicity, we assume that the Government sells its bonds only to the bank. We assume moreover that regulation (a portfolio constraint) forces the bank to purchase Government bonds.

4 Can we estimate the ABM? A Bayesian proposal

The estimation of a generic model aims at finding the set of numerical values of the parameters such that the "output" of the model is "as close as possible" – according to some metric – to the observed empirical data. Contrary to standard DSGE models, ABMs cannot be solved analytically but must be simulated. The output of an ABM consists of the artificial time series generated by a simulation with given numerical values of the parameters and given random seeds. The properties of the output may be affected by the random seeds and the configuration of parameters: small changes in parameter values might lead to big changes in model behavior. This clearly poses huge challenges when the ABM has to be estimated.

4.1 Bayesian estimation of ABM parameters

In this section we elaborate on the Bayesian techniques for ABM estimation proposed in Grazzini et al. (2017). Bayesian estimation follows three steps: i) simulation of the ABM; ii) estimation of the likelihood function and the posterior distribution and iii) sampling from the posterior distribution of the parameters. We improve on the Bayesian method proposed in Grazzini et al. (2017) by taking into account also the time structure of observed data. In appendix A, we show how the improved Bayesian estimation method proposed in this paper provides a more precise posterior distribution with respect to Grazzini et al. (2017) in a few simple estimation models.

Step 1: Simulating the model. The first step consists in simulating the ABM. We write the period-t values of the macroeconomic variables generated by each simulation in compact form as follows:

$$y_t = F(x_0, \Psi_{t_0:t}; \theta)$$

where y_t is the $M \times 1$ vector of period-t values[6] of M aggregate variables, $F(.)$ represents the model, x_0 denotes a set of initial conditions, $\Psi_{t_0:t}$ are all the random shocks from period $t_0 > 0$ to the current period[7] and θ is the $k \times 1$ vector of numerical values of the parameters used in the simulation. Since the (relevant) time span of the simulation goes from period t_0 to period T, the output of the model is represented by the $M \times T$ matrix of simulated data $Y = [y_{t_0}, y_{t_1}, ..., y_T]$. In general the behavior of the model depends on the parameters in θ, on the random seed s governing the set of random shocks $\Psi_{t_0:t}$ and on the initial conditions x_0. With a slight abuse of notation, therefore, we can characterize the matrix of simulated data as $Y = Y(\theta, s, x_0)$.

If the time series $y^m = [y^m_{t_0}, y^m_{t_1}, ..., y^m_T]$ represented by the m-th row of this matrix ($m = 1, 2, ...M$) were *stationary* and *ergodic*, changes in x_0 and s would not affect the output of the model. In fact if y^m is stationary and ergodic, *long run* simulations with different initial conditions and different random seeds will generate statistically equivalent time series. In this case the $M \times T$ matrix Y of simulated data (with T very large) can be characterized as $Y = Y(\theta)$ because the properties of the matrix depend only on the set of parameters θ.

Since time series can be easily made stationary, the assumption of stationarity is usually not problematic. On the contrary, ergodicity is a strong assumption for ABMs: simulated time

[6] We interpret one period as a quarter.

[7] We consider the output of the model from t_0 to get rid of the transient period (i.e., the interval ranging from 0 to t_0).

series may well be non-ergodic. When y^m is non-ergodic, its properties depend on the specific draw from the distribution of shocks and on the particular initial conditions. For the sake of simplicity and tractability, in the following we assume that artificial time series are characterized by a "weak" form of non-ergodicity. More precisely, first of all we assume that the output of the ABM is unaffected by changes in initial conditions so that Y depends only on the vector of parameters θ and the random seed s: $Y = Y(\theta, s)$. Second, we simulate the model $\mathcal{S}$ times for T periods with the same parameter values but different random seeds obtaining a matrix $Y(\theta, \mathcal{S}) = [Y(\theta, s_1), Y(\theta, s_2), ..., Y(\theta, s_{\mathcal{S}})]$ of M stacked time series each with a length of $T \times S$ periods. We assume that the stacked artificial time series are approximately ergodic. We will refer to $Y(\theta, \mathcal{S})$ as the *ensemble* distribution and we will use it to characterize the behavior of the model as a function of θ in the estimation procedure. In principle non-ergodicity has important implications for the interpretation of estimation results. For the specific ABM discussed in this paper, however, we will show that the results of estimation are not biased by a potential non-ergodicity issue.

We will now describe the steps of the estimation procedure.

Step 2: Estimating the likelihood and posterior distribution The second step consists in determining the likelihood of θ. The empirical evidence (observed macroeconomic time series) can be organized in a $M \times n$ matrix $\bar{Y} = [\bar{y}_1, \bar{y}_2, ..., \bar{y}_n]$ where $\bar{y}_t$ is the $M \times 1$ vector of period-t observed values of M aggregate variables, n is the length of the time span over which data have been collected. We want to compute the likelihood function:

$$\mathcal{L}(\theta, \bar{Y}) = \prod_{t=2}^{n} P(\bar{y}_t, \bar{y}_{t-1}; \theta)$$

where $P(\bar{y}_t, \bar{y}_{t-1}; \theta)$ is the joint distribution of $\bar{y}_t$ and $\bar{y}_{t-1}$. By using this distribution, we take into account the time structure of the simulated and observed time series in the estimation procedure.[8]

The most difficult task in estimating an agent-based model – whose data generating process is unknown to the modeller – is to determine the density of the observed time series for each possible value of θ, i.e., to determine $P(.; \theta)$. Our strategy to overcome this difficulty consists in using the simulated data to estimate a non-parametric pdf by means of a kernel density

[8]The number of lags in the joint distribution is arbitrary. In this application we use only one lag.

estimator with a Gaussian kernel. Of course, the distribution of simulated data depends on the vector of parameters θ. We compute the density of each observed data point $\{\bar{y}_t, \bar{y}_{t-1}\}$ under the kernel density estimated distribution. The likelihood $\mathcal{L}(\theta, \bar{Y})$ is the product of the density of each observed data point. Using the prior $p(\theta)$, we can compute the posterior using Bayes' theorem:

$$p(\theta|\bar{Y}) \propto p(\theta)\mathcal{L}(\theta, \bar{Y}) \tag{6}$$

Eq.(6) gives the posterior probability of a parameter vector θ, given the prior $p(\theta)$ and the matrix of observed data $\bar{Y}$. Of course, it is possible to take the log and compute the log-posterior.

Step 3: Sampling from the posterior distribution. The third task we have to carry out consists in investigating the k-th dimensional posterior $p(\theta|\bar{Y})$. To compute the posterior, for each value of θ we need to simulate the model and estimate the likelihood. Exploring the k dimensional space, therefore, can be computationally very heavy. Grazzini et al. (2017) use the MCMC method with Metropolis-Hastings algorithm, Lux (2018) uses the particle filters estimation method.

A search algorithm is generally efficient in exploring highly dimensional spaces to find the areas of high posterior when the the latter is not too irregular. In ABMs, due to the complexity of interactions and the pervasiveness on non-linearities, small changes in the values of the parameters may imply huge changes of model behavior – i.e., of the properties of simulated time series. The posterior, therefore, is highly irregular and possibly non-continuous. Search algorithms such as MCMC can be very inefficient in this case. They may end up sampling the posterior from flat areas of local maximum. Moreover, even with relatively simple models, search algorithms usually need a great number of simulations to eventually converge toward the correct posterior. To avoid these problems in this paper we have adopted a different strategy. We draw $\mathcal{N}$ samples using latin hypercube sampling (a similar strategy is adopted in Barde and van der Hoog, 2017) and compute the posterior for each sample θ. If $\mathcal{N}$ is big enough, this simple sampling procedure gives an exhaustive overview of the parameters space. More efficient sampling procedures could be applied. For example, Salle and Yıldızoğlu (2014) propose to use kriging and Lamperti et al. (2018) use *Machine Learning Surrogates*. In this paper, to keep the estimation procedure intuitive we use the simple sampling method described above. However, extending the estimation algorithm to improve the sampling of the parameter space should be straight forward.

4.2 Application of the Bayesian estimation method to our ABM

We apply now the proposed technique to the actual estimation of our MABM. Our economy is populated by 1120 households (1000 workers and 120 capitalists) and 120 firms, of which 1/6 in the K-sector. Overall there are 28 parameters. To reduce the computational burden we set the numerical values of 17 parameters (see Table 1) and estimate the remaining 11 parameters (listed in Table 2).

We sample $\mathcal{N} = 5000$ combinations of the above mentioned 11 parameters using latin hypercube sampling. For each parameter combination we simulate the model $\mathcal{S} = 30$ times. Therefore we run 150 000 simulations. It is important to stress that we select a set of $\mathcal{S}$ random seeds and use the same random seeds for each parameter combination. Each simulation runs for $T = 1500$ periods so that the total number of simulated periods is 225 000 000. Notice however that we discard a transient of $t_0 = 500$ periods per simulation.

We estimate only a subset of parameters to alleviate the computational strain of exploring a high dimensional posterior space. To estimate more parameters we should increase $\mathcal{N}$ adequately which would affect computing time significantly also in Step 2. In fact, more parameter combinations imply more densities to estimate and more posteriors to compute.[9] We record the artificial time series of output, gross investment (including changes in inventories), consumption (all measured at constant prices), the aggregate price level and the unemployment rate.

We estimate the model using US data. We have downloaded (from FRED) GDP, personal consumption, gross private investment (all in real terms), the implicit price deflator and the civilian unemployment rate from Q1-1948 to Q1-2018.[10]

The model is simulated in levels. To make simulated and observed data comparable, we consider the cycle component estimated by an HP-filter of both simulated and observed real

[9]To get an idea of the computational time involved, consider that 1 simulation on our computer takes on average 69 seconds. The total computational time spent in the simulation and parameter sampling phase was 10,350,000 seconds, i.e., 2875 hours or 119 days. Using a computer with 24 cores, we were able to parallelize simulations and reduce total simulation time to approximately 5 days. As mentioned above, some recent contributions – e.g. Lamperti et al. (2018) – put forward proposals to improve sampling efficiency by means of machine learning and surrogate models. These methods can improve efficiency in exploring the parameter space to a large extent. We plan to use these methods in future developments of our estimation project.

[10]FRED codes: GDPC1, PCECC96, GPDIC1, GDPDEF, UNRATE. We transformed monthly data for the unemployment rate into quarterly data by taking the average of monthly data over a quarter.

Table 1: Calibrated parameters.

Parameter	Description	Value
Z_c	C-firms visited by each consumer	2
Z_k	K-firms visited by each C-firm	2
Z_e	Firms visited by each unemployed worker	5
α	Productivity of labor	0.5
k	Productivity of capital	1/3
τ	Dividend payout ratio	0.2
δ	Capital depreciation rate	0.03
$\bar{\omega}$	Long run capacity utilization rate	0.85
b_{0c}	Parameters for bank's risk evaluation	-15
b_{1c}	"	13
b_{0k}	"	-5
b_{1k}	"	5
r	Risk free interest rate	0.01
z	Replacement rate (unemployment subsidy)	0.5
u_{up}	Upward wage adjustment	0.1
u_{down}	Downward wage adjustment	0.01
$\hat{u}$	Unemployment threshold	0.1

Table 2: Priors and posteriors' mode of estimated parameters.

Parameter	Description	Prior	Post. mode
Ξ	Memory parameter in cons.	$\mathcal{U}(0.5, 1)$	0.7382
χ	Wealth parameter in cons.	$\mathcal{U}(0, 0.1)$	0.0172
ρ	Quantity adjustment	$\mathcal{U}(0.5, 1)$	0.7301
$\bar{\eta}$	Price adjustment	$\mathcal{U}(0, 0.5)$	0.1649
μ	Bank's gross mark-up	$\mathcal{U}(1, 1.5)$	1.007
ϕ	Bank's leverage	$\mathcal{U}(0, 0.1)$	0.0024
δ^k	Inventories depreciation rate	$\mathcal{U}(0, 1)$	0.0781
γ	Fraction of investing C-firms	$\mathcal{U}(0, 0.5)$	0.3260
ϑ	Rate of debt reimbursement	$\mathcal{U}(0, 0.15)$	0.0328
υ	Memory parameter in inv.	$\mathcal{U}(0, 1)$	0.1591
t_w	Tax rate	$\mathcal{U}(0, 0.4)$	0.0594

GDP, consumption and investment. Moreover, we use the simulated and observed price deflator to compute de-meaned inflation rates. Finally, we consider the unemployment rate in levels.

Table 1 lists the numerical values of 17 parameters, calibrated on the basis of empirical evidence and/or to generate a "plausible" output of the model. Table 2 lists 11 estimated parameters, the prior and the mode of marginal posteriors.

In this estimation exercise, we have chosen uniform priors to impose a "reasonable" restriction on the values of estimates, thereby defining a closed parameter space. If available, different priors could be easily implemented and possibly improve estimation results. In the last column of Table 2 we show the combination of parameters with maximum posterior. Figure 2 shows the marginal posterior distributions of the estimated parameters. Posteriors are very narrow.

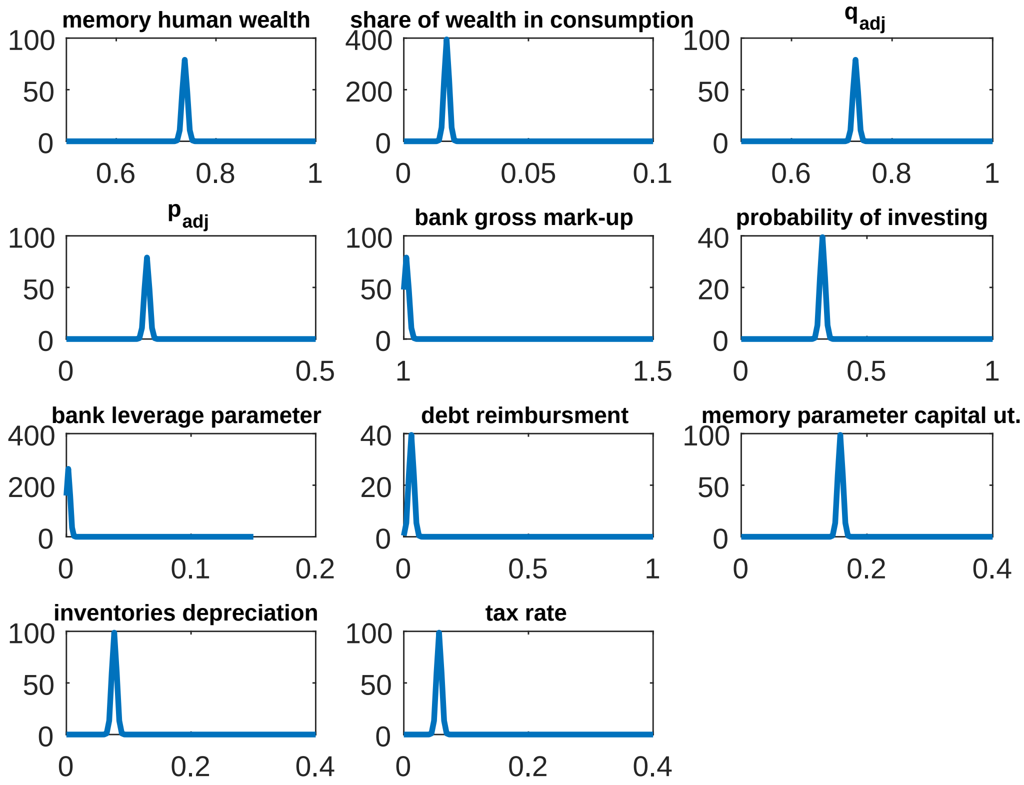

Figure 2: Marginal distributions of the posterior

Intuitively, this means that, among the $\mathcal{N}$ parameter combinations, the estimation procedure finds that one parameter combination has a much higher posterior density than all the others. This parameter combination coincides with the mode of the posterior distributions reported in Table 2. However, in order to correctly interpret the posterior marginal distributions, it is important to recall that by using the latin hypercube sampling, we have discretized the parameter space.

To evaluate the reliability of the procedure we have implemented a pseudo-estimation exercise. We choose the modes of the posteriors as the point estimates of the parameters, and use these point estimates to simulate the model 400 times with different random seeds for 800 time periods. Discarding the first 500 periods as transient, we are left with 400 pseudo-observed time series, i.e. artificial time series, of a length comparable to that of observed time series. We then use these pseudo-observed time series to estimate the model following the procedure outlined above. We find that 72.5% of the estimates deliver the correct vector of parameters. This result suggests that, assuming the model is well-specified and the selection of the parameter space is accurate, the probability that the estimation procedure yields the correct estimates of the

	Y	C	I	π	u
Simulated					
standard deviation	0.0163	0.0133	0.0709	0.0067	0.0329
first lag autocorrelation	0.5722	0.3408	0.4191	0.3254	0.7958
Empirical					
standard deviation	0.0161	0.0125	0.0665	0.0062	0.0163
first lag autocorrelation	0.8469	0.8162	0.8249	0.7795	0.9665

Table 3: Artificial and empirical moments. Moments are computed on the cyclical component using HP-filter of real output (Y), real consumption (C), real gross investments (I), on demeaned inflation (π) and on unemployment (U).

parameters is high. This result suggests also that non-ergodicity does not impair the efficacy of the procedure to estimate the model. In fact, if the model were strongly non-ergodic, simulations based on the same parameter set but different seeds would deliver very different parameter estimates. We cannot rule out non-ergodicity but we can safely claim that non-ergodicity, if present, is not strong enough to influence the ability of the estimation procedure to find the correct parameter vector.

5 Does the ABM fit the data? Matching moments

Does the estimated model "behave" approximately as the observed economic system? To answer this question, in this section we investigate the statistical properties of a number of simulated aggregate time series and compare them with their empirical counterparts. We simulate the model 96 times using the combination of parameters listed in Table 1 and Table 2 as point estimates. The model is able to approximately reproduce the behavior of output, the price index, consumption, investment and the unemployment rate observed in the US economy. Figure 3 shows the densities of the empirical and simulated time series.[11] In general the statistical properties of observed time series are matched by those of the artificial time series generated by simulations. Some estimated densities, however, are slightly shifted to the left and all the estimated time series are more volatile than the observed counterparts. To quantify the distance between simulated and observed time series, we compute the standard deviation and the first lag autocorrelation for output, consumption, investment, (demeaned) inflation rate

[11] In interpreting the comparison it is important to note that the estimation procedure employs the joint density of period t and period $t-1$ variables to compute posteriors, not the unconditional densities plotted in the figure.

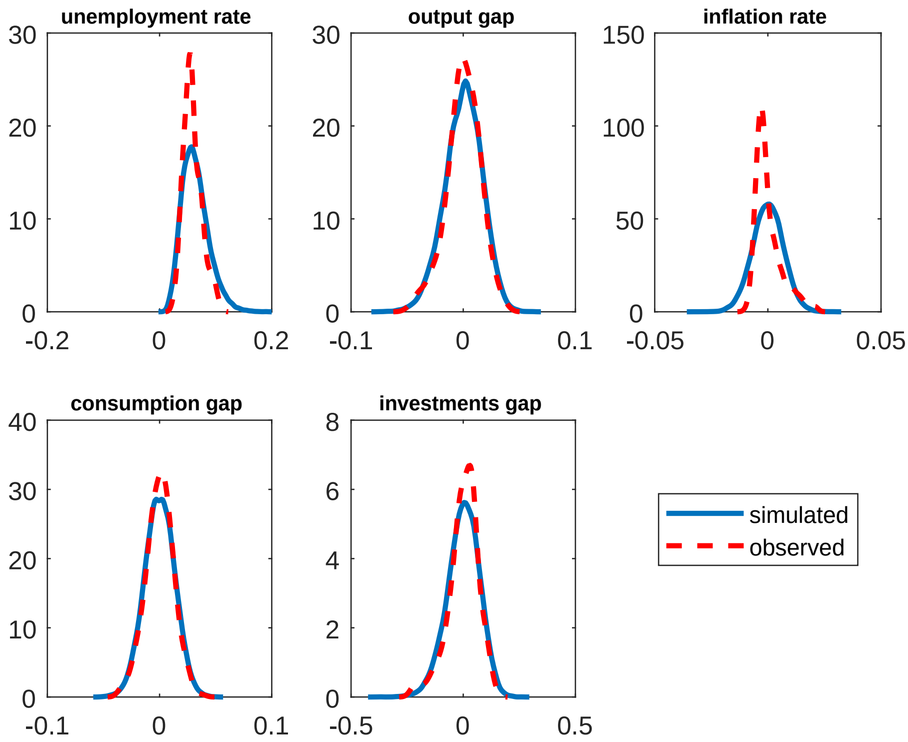

Figure 3: Estimated densities of empirical (dashed red line) and simulated (blue line) cyclical component of output, consumption, investment and price index and of the unemployment rate.

and unemployment rate (see Table 3). The standard deviation of simulated time series is very close to that of the empirical counterparts but the observed persistence is only partially matched by the model: the autocorrelation of simulated time series is non-negligible but is always significantly lower than that of observed time series. It is worth noting, however, that the persistence of simulated series is entirely built-in (due to assumptions on agents' behaviour and their interactions). In fact, the ABM under scrutiny does not feature any exogenous auto-regressive aggregate shocks.

6 How does the model behave after a shock? Impulse Response functions

How does the behaviour of the model respond to aggregate shocks? For DSGE models, the obvious answer to this question is provided by Impulse-Response analysis. Following this well-established methodology the DSGE modeller can explore in a neat and generally unambiguous way the deterministic trajectory of a macroeconomic variable after a shock to an aggregate variable.

Of course, one can investigate the effects of an aggregate shock also in a macro ABM. In this case, however, it is impossible to completely shut off after-shock randomness. Moreover, it is impossible to isolate the effects of the shock: the model as a whole will be affected in highly unpredictable ways. To get the flavour of the argument, consider the search and matching process in the goods market. As we have shown above, there is an unavoidable random ingredient of this process, which drives the adaptation of prices and quantities to changing market circumstances. The same aggregate shock may trigger different dynamic paths of prices and quantities in the presence of different random seeds. Of course the modeller can at least fix the random seed so that the simulated search and matching sequence does not change after the shock. This does not solve the problem, however, as shown by the following mental experiment.

Suppose we simulate the model for T periods with a given set of parameter values, exogenous variables and random seed. This is the benchmark simulation. Then we simulate again the model over the same horizon (T periods) with the same random seed, but in period T/2 we generate an exogenous uniform increase of permanent income. This is an aggregate shock as its impact (first round effect) is the same for all the household.[12] This is the shocked simulation. How does the shock affect the search and matching mechanism?

Let's assume that in the benchmark simulation the first (randomly selected) consumer visits Z_c (randomly selected) firms and spends the entire consumption budget at firm A, that sets the lowest price. The consumer therefore does not buy at the firm setting the second lowest price, which we denote with B. Let's assume, moreover, that firm B goes bankrupt due to lack of demand.

Since we have fixed the random seed, in the shocked simulation the first randomly selected consumer will be matched with exactly the same Z_c firms as in the benchmark simulation. However, since the consumer's after shock permanent income is now higher, the goods available at firm A are not sufficient to exhaust her consumption budget: she will turn to firm B to buy goods. Other things being equal, firm B will increase production and/or price in the next period. Most likely its profits will increase. It may even avoid bankruptcy, changing the course of events in the model.

The presence of non-linearities and discontinuities (such as bankruptcies) in the ABM implies that a small change in an exogenous variable (or parameter) can have a huge impact on aggregate

[12] We will analyze the effect of such a shock below.

behavior. In our mental experiment, a small change in demand can make the difference between life (survival) and death (bankruptcy) for firms. Therefore, after the shock the overall behavior of the model changes and the differences between the benchmark simulation and the shocked simulation will not disappear.

6.1 Seed-specific Impulse-Response (SIR) functions

To illustrate this behavior, we analyze the response of GDP to an aggregate shock in our ABM. Let y_t represent period-t GDP (at constant prices) generated by the benchmark simulation ($t = 1, 2, ..., 1000$) and $\tilde{y}_t$ period-t GDP generated by the shocked simulation, characterized by the same random seed but an exogenous increase of permanent income in period 500. To analyze the macroeconomic effect of the shock we compute the percent deviation of the shocked GDP from the benchmark GDP:

$$\hat{y}_t = \log \tilde{y}_t - \log y_t$$

The time series of $\hat{y}_t$ is shown in figure 4. Of course, the benchmark and shocked simulations generate the same GDP in the first half of the simulation horizon so that $\hat{y}_t$ is exactly equal to zero for any $t \leq 500$. After the shock $\hat{y}_t$ is generally different from zero. We can conceive of $\hat{y}_t$ over the second half of the simulation horizon as the equivalent of an Impulse-Response (IR) function for our ABM. Notice that $\hat{y}_t$ oscillates irregularly around zero, i.e., *on average* the shocked GDP is equal to the benchmark GDP over the interval $500 < t < 1000$. Since this IR is associated with a specific random seed, we will refer to $\hat{y}_t$ as a *Seed-specific IR (SIR) function.*

The dynamics of $\hat{y}_t$ however is difficult to interpret. It does not behave nicely as in DSGE models, with a smooth monotonic departure from and return to the steady state. This is due at least in part to the specific random seed. Different random seeds generate different SIRs.

6.2 Robust Impulse-Response (RIR) functions

To form a "robust opinion" of the response of a macroeconomic variable to an aggregate shock in an ABM, we must carry out a number of simulations with different random seeds and compute the average response of the variable in question to the aggregate shock.[13]

[13] Gobbi and Grazzini (2019) and Guerini et al. (2018) propose similar methods to produce impulse response functions in agent-based models.

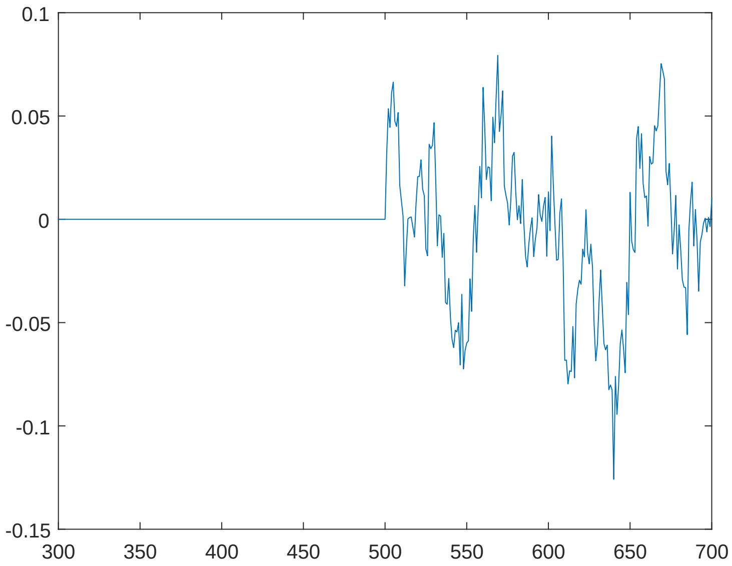

Figure 4: Log difference of real GDP between the benchmark simulation and a simulation with a shock to permanent consumption in period 500.

To build *Robust Impulse-Response (RIR) functions* we propose the following method. Consider an ABM which can generate the artificial time series of M macroeconomic variables (e.g. GDP, consumption etc.) by means of a simulation over T periods for a given set of parameter values and exogenous variables and a given random seed. We want to quantify and interpret the *response* of each macroeconomic variable to an *impulse* generated by an aggregate shock.

1. We simulate the ABM (over T periods) with a given set of parameters and exogenous variables with S different random seeds. These are the S benchmark simulations. We get S benchmark time series for each variable of interest. We denote the benchmark time series of variable y^m (where $m = 1, 2, ..., M$) in simulation s ($s = 1, 2, ..., S$) with $y^m_{s,t}$ where $t = 1, 2, ..., T$.

2. Then, we simulate the ABM with the same seeds as above but with an aggregate shock in period $\tau < T$. These are the S shocked simulations. We get S shocked time series for each variable of interest. We denote the shocked time series of variable y^m in simulation s with $\tilde{y}^m_{s,t}$.

3. We compute S Single-seed Impulse-Response functions, i.e., log differences of the shocked and benchmark time series, for each variable. We denote the SIR of variable y^m in simulation s with $SIR^m_{s,t} = \log \tilde{y}^m_{s,t} - \log y^m_{s,t}$.

4. Finally, we compute the Robust Impulse-Response function for each variable as the mean

of the single seed impulse response functions across seeds:

$$RIR_t^m = \frac{\sum_{s=1}^{S} SIR_{s,t}^m}{S} \quad t = 1, 2, ..., T$$

6.3 Interpreting the transmission of shocks by means of RIRs

In this section we apply the methodology proposed above to our ABM in order to build RIRs for $M = 9$ variables: GDP, inflation, gross capital formation, consumption, the unemployment rate, total debt of the corporate sector, the bankruptcy rate (fraction of firms that go bankrupt), total capital stock and bank net worth. We will consider $\Sigma = 9$ aggregate shocks: 3 real shocks (exogenous change of consumption, investment and capital stock), 2 nominal shocks (exogenous change of the prices of capital goods and of consumption goods) and 4 distributional shocks.

To build RIRs we simulate the ABM over $T = 1000$ periods S=192 times with different seeds. We generate S=192 benchmark time series for each of the 9 variables.[14]

Then, we simulate the ABM S=192 times with the same seeds but hitting the economy with an aggregate shock in period $\tau = 500$. We repeat the procedure for each of the Σ shocks. We get $S = 192$ shocked time series for each variable and each shock. Therefore we compute $S = 192$ Seed specific IR functions (SIRs) for each variable and each shock.

Finally, we compute the Robust IR functions (RIRs). We take the average of SIRs across random seeds for each variable and each shock. Overall therefore we have $M \times \Sigma = 81$ RIRs. They are shown in figures 5, 6 and 7, together with the 20th and 80th percentiles to account for the variability of the response to the shock.

The dynamic pattern of RIRs is generally similar to those of IRs in DSGE models. The shock generates a departure of the variable from the benchmark statistical equilibrium and a return to it with the passing of time. In some cases the return to equilibrium may occur by means of long dampening swings.

6.3.1 Real shocks

The first set of RIRs in the upper-left panel (block of diagrams) in Figure 5 captures the response of the 9 variables of interest to a (positive) consumption shock generated by a uniform

[14]The number of simulations $S = 192$ has been chosen to use efficiently parallel computing over our 24 cores computer. In fact, 192 is a multiple of 24.

exogenous increase of permanent income in period 500.[15] Consumption increases on impact, since households feel richer and allocate a bigger budget to consumption, and then goes back to the benchmark statistical equilibrium.[16] The increase in consumption has three effects on firms. First higher demand for consumption goods implies that firms charging higher prices will sell more. This leads to an increase of inflation on impact. Second, higher demand translates into higher revenues, leading to higher profits and a slight decrease of the bankruptcy rate. Third, firms react to higher consumption by increasing production and employment, so that the unemployment rate decreases (slightly). The effects of the consumption shock on investment and the capital stock are negligibile. The upper-right panel of Figure 5 shows the RIRs

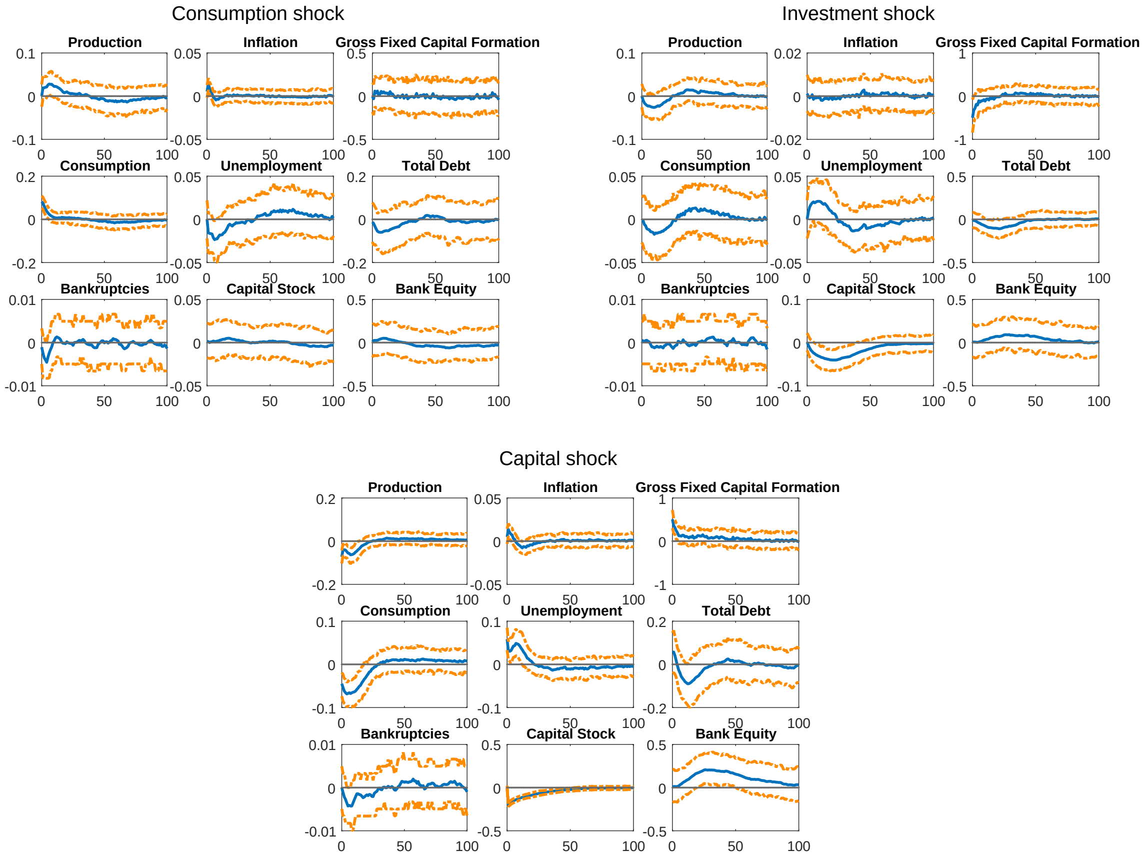

Figure 5: **Upper-left**: RIRs to a positive persistent shock to permanent income. **Upper-right**: RIRs to a negative persistent shock to the probability to invest (γ). **Bottom**: RIRs to a shock to the capital stock. Blue lines are average impulses, orange dashed lines are the 20th and 80th percentiles. Period 0 corresponds to the period in which the shock hits the economy.

following a (negative) investment shock, i.e., a reduction of the probability to invest γ in period

[15]We assume that the permanent income of all workers goes up by 20% in period 500. This could be interpreted as an increase of expected future income.

[16]In the comments that follow we will focus on the short run effects of the shock. To avoid repetitions we will not reiterate that RIRs go back to zero over the long run.

500.[17] Thanks to the law of large numbers, γ can be conceived as the fraction of firms able to adjust their capital stock in each period. It is worth noting that, given the investment decision described in Eq. (3.2), when γ goes down firms that can adjust their capital stock will invest more, anticipating that it will be less likely for them to be able to invest in the future. Therefore, a reduction of γ makes the number of investing firms shrink, but boosts the investment of those firms that are able to invest. The first effect prevails so that the shock will lead to a reduction of gross fixed capital formation and of the capital stock. Production goes down and the unemployment rate increases. Moreover, the reduction of investment leads to a decrease of corporate debt.

The bottom panel of Figure 5 shows the RIRs in case of a shock to the capital stock, namely an exogenous "disruption" that destroys 20% of the capital stock in period 500. Other things being equal, after the shock capacity utilization goes up and becomes higher than desired. C-firms therefore have an incentive to increase investment to restore their capital stock. Higher demand for capital goods leads K-firms to increase their prices causing an increase in inflation. Due to perfect complementarity in production, a lower capital stock implies also lower employment and therefore the unemployment rate increases. Labour income goes down and so does consumption. Following a capital disruption, therefore, investment and consumption move in opposite direction. The latter effect prevails so that aggregate demand and GDP decrease. Following the shock, moreover, corporate debt and the bankruptcy rate go down and bank's net goes up.

6.3.2 Nominal shocks

The left panel of Figure 6 shows the effect of a shock to the prices of K-goods, namely a 10% increase of the price set by each and every K-firm.[18] Since K-goods become more expensive, C-firms transfer the sudden increase of their costs to their prices. This implies a rise in the prices of C-goods. Inflation goes up due to the joint effect of the increase in the prices of C-goods and K-goods. Higher prices of K-goods, moreover, depress investment, employment, production and consumption. Since C-firms need more funds to buy K-goods, their debt increases.

[17] We assume that the probability to invest goes down by 80% in period 500. To make the effect of the shock more interesting, we assume that after the shock γ converges back gradually to its benchmark value. In other words, we assuem that there is a persistent, rather than temporary, shock to the probability to invest.

[18] To magnify the effects of the shock, we assume that the shock is persistent, i.e., it vanishes gradually.

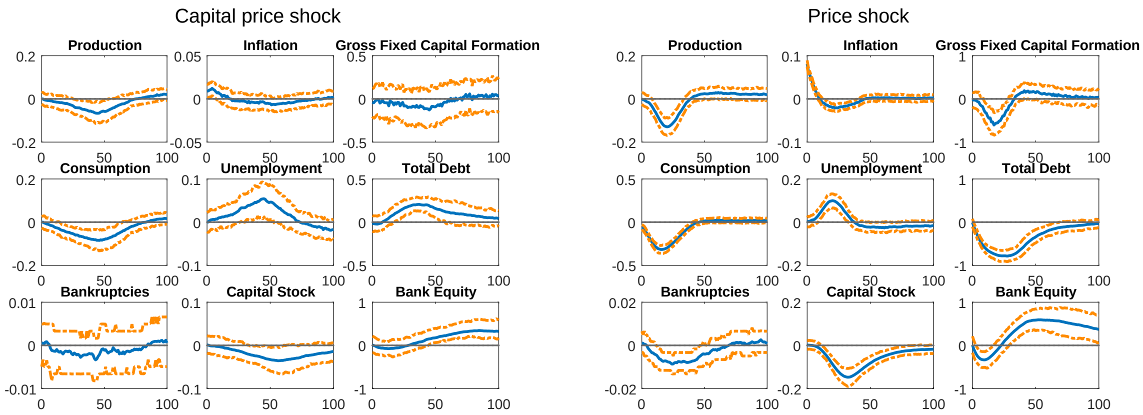

Figure 6: Price shocks. Left (right) panel: RIRs to a persistent shock to the price of capital goods (consumption goods). Period 0 corresponds to the period in which the shock hits the economy.

The right panel of Figure 6 shows the consequences of a shock to the prices of C-goods, namely a 10% increase of the price charged by each and every C-firm. The magnitude of the shock on C-prices is the same as that on K-prices but the difference in terms of macroeconomic outcomes is spectacular. The obvious reason is the relative size of the two industries, the C-sector being much bigger than the K-sector. The increase in C-prices has a sizable direct effect on inflation, whch shoots up. Higher C-prices imply a smaller demand for C-goods and a contraction of production and employment. The capital requirement goes down and so does investment. Moreover higher inflation implies lower real wages and lower consumption. Lower demand for consumption goods leads firms to reduce their production, reducing employment, investment and the capital stock. This will further reduce consumption. The recession induced by the shock leads to a reduction of firms' exposure to banks and of the bankruptcy rate.

6.3.3 Distributional shocks

The distribution of wealth across sectors affects aggregate economic performance.[19] If the wealth of the corporate sector (firms' equity) increases, firms become more financially robust, their financing gap goes down and the bankruptcy rate becomes smaller. If the bank's equity increases, the supply of loans will increases and firms will be able to produce more.

In this section we explore the macroeconomic consequences of distributional shocks, i.e., exogenous changes of the allocation of wealth to sectors. Each of these shocks can be interpreted as a policy shock, i.e., a lump-sum tax on one sector whose revenue finances a lump-sum subsidy

[19] Of course, the distribution of wealth can change also *within* each sector, e.g. between rich and poor households. We will not deal with inequality in this paper.

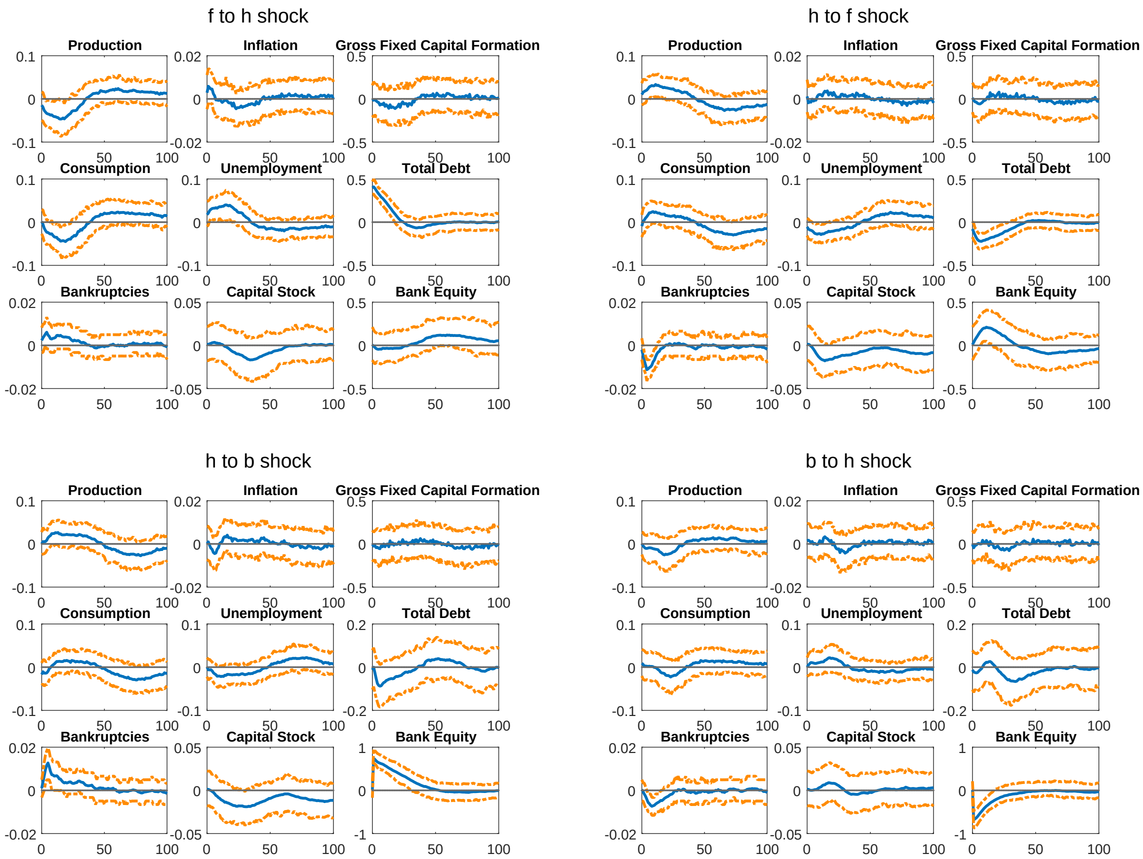

Figure 7: **Upper left**:RIRs of redistribution from firms to households. **Upper right**:RIRs of redistribution from households to firms. **Lower left**: RIRs of redistribution from household to the bank. **Lower right**: RIRs of redistribution from the bank to households. Period 0 corresponds to the period in which the shock hits the economy.

to another sector. The distributional shock, therefore, modifies the amount of resources (wealth) available to each sector.

The upper-left panel of Figure 7 shows the RIRs in case of a redistribution of wealth from firms to households. The distributional shock takes place in period 500 and consists in a 50% reduction of the equity of all C-firms (by means of a levy on deposits) and uniform redistribution of this wealth to all workers in the economy. In other words, the shock consists in transferring liquidity from firms' deposits to households' deposits. The shock affects households' liquidity but does not change their permanent income. Therefore, consumption does not increase visibly on impact. On the other hand, the shock leaves firms in need of liquid resources. Firms need additional loans to finance production and investment. Corporate debt, therefore, shoots up by almost 50% on impact. Basically, firms substitute the internally generated liquidity with borrowed liquidity. Leverage increases and the borrowing constraint becomes tighter. Since the liquidity lost due to the shock cannot be substituted completely with bank loans, firms do not have the financial resources to keep desired employment. Hence unemployent increases, production, income and consumption decrease.

The upper-right panel of Figure 7 shows the the RIRs in case of a redistribution of wealth from households to firms. The distributional shock takes place in period 500 and consists in a 50% reduction of the wealth of each and every worker (by means of a levy on deposits) and uniform redistribution of this wealth to C-firms. Since this distributional shock mirrors the previous one, the response of the economy is symmetric. After the shock, firms have liquid resources and therefore borrow less from the bank. Corporate debt and the bankruptcy rate go down. Despite fewer loans, the bank's equity increases due to the lower bankruptcy rate. Financially robust firms are able to hire more, reducing unemployment, and increasing production and income. The latter effect reverberates on consumption.

The lower-left panel shows the RIRs in case of a redistribution of 50% wealth from households to the bank. When the bank's equity goes up, the supply of loans increases relaxing the borrowing constraint on risky firms. These firms, which were already burdened with relatively high debt, will borrow more. Some of them will become overindebted and go bankrupt: the bankruptcy rate goes up. After the shock total corporate debt shrinks.[20]The effects are symmetric in case of a redistribution of wealth from the bank to households, as shown in the lower-right panel. A smaller loan supply will hinder the ability of more fragile firms to access credit leading to a smaller bankruptcy rate. The effect of the redistribution of wealth between households and the bank on the other variables are not sizable.

7 Can we use the ABM to forecast?

We propose a new method to generate macroeconomic forecasts using simulation models. We apply this method to our estimated ABM and evaluate its ability to forecast a set of aggregate variables. To the best of our knowledge, this is the first attempt to use an ABM to forecast macroeconomic times series.[21] To evaluate the performance of the ABM to forecast US data, we estimate a simple VAR model and use it as a benchmark. It is well known that it is very difficult for any standard macroeconomic model to outperform VARs in terms of forecasting

[20]Once a firm goes bankrupt, its debt represents a loss for the bank and is written off

[21]As we anticipated in the introduction, Recchioni et al. (2015) use a heterogeneous agents model a lá Brock and Hommes (1998) to forecast financial time series.

capability. This is true also of CC-MABM: forecasts produced by the VAR model are more accurate than the forecasts generated by CC-MABM.[22] The issue, therefore, is to assess by how much CC-MABM underperforms with respect to the VAR benchmark.

7.1 Forecasting method

In period $t-1$, the value y_t^m that variable m will take on in period t is uncertain. To capture this uncertainty we assume that y_t^m is a discrete random variable which takes on a finite number of values $y_t^{i,m}$ where $i = 1, 2, ..., \mathcal{I}$. The index i denotes a certain "state" of the economy. We arrange the levels of the M variables in period t and in state i in the vector y_t^i. Hence $y_t^{i,m}$ is the $m-th$ element of this vector, $m = 1, 2, ...M$. The probability of each state $P(y_t^i|\bar{y}_{t-1})$ is conditional on the vector $\bar{y}_{t-1}$ of real data in $t-1$. We define the forecast for period t of variable m as follows:

$$E[y_t^m|\bar{y}_{t-1}] = \sum_{i=1}^{\mathcal{I}} P(y_t^i|\bar{y}_{t-1}) y_t^{i,m} \ , \quad m = 1, ..., M \tag{7}$$

In words, the forecast of variable m for period t is the expectation of the random variable y_t^m, conditional on real data in period $t-1$. This conditional expectation, in turn, is the weighted average of the period t values that the variable m can assume in $\mathcal{I}$ states where the weights are the probabilities $P(y_t^i|\bar{y}_{t-1})$ of each state conditional on the vector of real data in period $t-1$. In order to compute this forecast therefore, we need to define the conditional probability $P(y_t^i|\bar{y}_{t-1})$ for a generic vector y_t^i and then choose (i.e. quantify) the vectors which characterize the states. In this subsection we describe the two-step procedure to produce forecasts.

Step 1: Estimation of the conditional probability. The first step consists in estimating the conditional probability $P(y_t^i|\bar{y}_{t-1})$. In order to do so, we simulate the ABM for T periods with $\mathcal{S}$ different random seeds using the parameter values in Table 1 and the point estimates in Table 2. We discard the first t_0 periods to get rid of the transient phase. For each variable of interest we get $\mathcal{S}$ time series of length $T - t_0$. We use these $M \times \mathcal{S} \times (T - t_0)$ artificial data to estimate (by means of a non-parametric kernel density estimator) the joint probability density $P(y_t, y_{t-1})$ where y_t represents a generic M-dimensional vector of period t values of the

[22] Our ABM (as NK-DSGE models), however, have the advantage – at least in the eyes of the macroeconomist – of being a "structural models", while VARs are atheoretical (purely statistical) frameworks.

variables of interest.

Given the vector of observed (real) period $t-1$ data, which we denote with $\bar{y}_{t-1}$, we use the estimated distribution $P(.,.)$ to compute the joint density $P(y_t, \bar{y}_{t-1})$, i.e. the density of any arbitrary M dimensional vector y_t and vector of true data $\bar{y}_{t-1}$. We use the estimated distribution also to evaluate the probability of the observed data $P(\bar{y}_{t-1})$. By definition, the probability of y_t *conditional* on $\bar{y}_{t-1}$ is

$$P(y_t|\bar{y}_{t-1}) = \frac{P(y_t, \bar{y}_{t-1})}{P(\bar{y}_{t-1})} \tag{8}$$

To wrap up, by simulating the model we generate artificial data which we use to estimate the joint density $P(y_t, y_{t-1})$. Using real data $\bar{y}_{t-1}$ and the definition of conditional probability, we compute the probability of observing an arbitrary vector y_t given that we have observed $\bar{y}_{t-1}$. Notice that Eq. (8) defines the *generic* conditional probability, since the vector y_t has not been specified.

Step 2: Specification of the states. The next step consists in specifying quantitatively the states that characterize the economy in period t. As we said above, state i is represented by the i-th vector y_t^i. Each element of this vector $y_t^{i,m}$ is the value that variable m can take on in period t and state i. To specify state i, we have to quantify $y_t^{i,m}$, $i = 1, 2, ..., \mathcal{I}$, $m = 1, 2, ..., M$. To do so, we exploit the fact that most macroeconomic variables are very persistent, i.e. they change "very little" from one quarter to the next. Therefore we specify $y_t^{i,m}$ adding a set of $\mathcal{I}$ possible variations $\{\Delta y^{i,m}\}$ of our choice to the value observed for variable m in period t-1 $\bar{y}_{t-1}^m$:

$$y_t^{i,m} = \bar{y}_{t-1}^m + \Delta y^{i,m}$$

The possible states that variable m can take on in period t are therefore characterized by the values $\Delta y^{i,m}$ that we must calibrate in order to explore, as much as possible, the space of possible changes between $t-1$ and t in the proximity of (i.e., at a reasonably small distance from) $\bar{y}_{t-1}^m$.

State i in period t therefore can be written in vectorial form as

$$y_t^i = \bar{y}_{t-1} + \Delta y^i \tag{9}$$

where $\bar{y}_{t-1} = \left[\bar{y}^1_{t-1} \ ... \ \bar{y}^M_{t-1}\right]'$; $\Delta y^i = \left[\Delta y^{i,1} \ ... \ \Delta y^{i,M}\right]'$.

7.2 Application of the forecasting method to our ABM

This forecasting method is computationally very intensive. To reduce the computational burden we limit the number of variables to $M = 2$, namely output and investment. We simulate the model $\mathcal{S} = 96$ times for $T = 3000$ periods and discard the first $t_0 = 1000$ periods. For each variable of interest we get 96 time series that last 2000 periods. We apply the HP filter to these artificial data to obtain $2 \times 96 \times 2000 = 384000$ artificial cyclical components of output and investment. We use these artificial data to estimate the joint probability density $P(y_t, y_{t-1})$. The generic y_t vector consists of two elements. The first element y^1_t is the cyclical component of output, i.e. the output gap. Analogously, the second element y^2_t is the cyclical component of investment.

Observed (real) data are the cyclical components of US real GDP ($\bar{y}^1_{t-1}$) and gross investment ($\bar{y}^2_{t-1}$) retrieved from the FRED database over the period Q1-1948 to Q1-2018.

We now characterize the possible states of the economy. We choose 100 possible values of output variation in the range $-0.07 < \Delta y^{i,1} < 0.07$ and 100 values of investment variation in the range $-0.3 < \Delta y^{i,2} < 0.3$. This calibration is grounded in the well know business cycle fact according to which investment is much more volatile than output. We then take all possible combinations of these changes ending up with $\mathcal{I} = 10000$ 2-dimensional Δy^i vectors. We use these vectors and the observed lagged vector $\bar{y}_{t-1}$ to generate $\mathcal{I} = 10000$ 2-dimensional y_t vectors following the procedure described by Eq. (9). Overall therefore there are 10000 states the economy can take on in each period t given the observed data in period $t-1$. Each of these states is captured by a vector y^i_t, $i = 1, 2, ..., 10000$. The probability of each state is given by $P(y_t, \bar{y}_{t-1})$ defined by Eq. (8). Our one period ahead forecasts for output and investment in each period are the elements of the vector $E[y_t|\bar{y}_{t-1}]$, namely[23]

$$\begin{bmatrix} y^{1,abm}_t \\ y^{2,abm}_t \end{bmatrix} = \begin{bmatrix} \sum_{i=1}^{10000} P(y^i_t|\bar{y}_{t-1}) y^{i,1}_t \\ \sum_{i=1}^{10000} P(y^i_t|\bar{y}_{t-1}) y^{i,2}_t \end{bmatrix}$$

To provide a benchmark in evaluating the forecasting performance of the ABM, we estimate

[23] See equation (7).

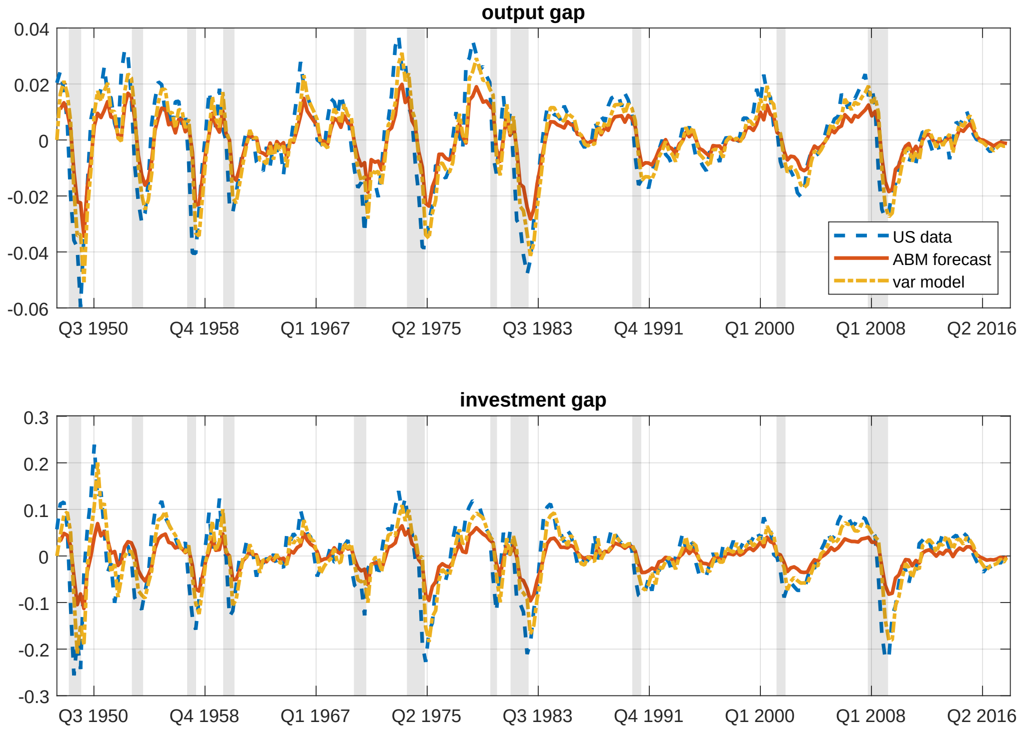

Figure 8: Observed data and one period ahead forecasts for output (upper panel) and investment (lower panel) obtained from the ABM (solid line) and from the VAR model. Shaded areas are recession bands as determined by the National Bureau of Economic Research.

the following simple VAR(1) model on US data:

$$\begin{bmatrix} \bar{y}_t^1 \\ \bar{y}_t^2 \end{bmatrix} = A \begin{bmatrix} \bar{y}_{t-1}^1 \\ \bar{y}_{t-1}^2 \end{bmatrix} + \begin{bmatrix} \varepsilon_t^1 \\ \varepsilon_t^2 \end{bmatrix}$$

where $\bar{y}_t^1$ is the observed (cyclical component of) output and $\bar{y}_t^2$ is observed investment. The one period ahead forecasts produced by the VAR model therefore are:

$$\begin{bmatrix} y_t^{1,var} \\ y_t^{2,var} \end{bmatrix} = A \begin{bmatrix} \bar{y}_{t-1}^1 \\ \bar{y}_{t-1}^2 \end{bmatrix}$$

In the upper (lower) panel of figure 8, for each quarter of the time window 1948-2018 we show: a) the actual (observed) output $\bar{y}_t^1$ (investment $\bar{y}_t^2$); b) the one period ahead forecast for output $y_t^{1,abm}$ (investment $y_t^{2,abm}$) generated by our ABM and c) the one period ahead forecast for output $y_t^{1,var}$ (investment $y_t^{2,var}$) generated by the VAR model.

We evaluate the forecasts of the VAR and of the ABM model (separately for output and investment) by means of the Diebold-Mariano test which compares the forecast errors. The test rejects the null-hypothesis for both set of forecasts with test statistics of -4.18 and -4.28 for output and investment respectively. Moreover, the strong negative value of the test favors the VAR model. One possible explanation is that the artificial time series generated by the ABM display less persistence than the observed times series and the time series generated by the VAR model. Of course, the persistence generated by the model is the result of the mechanisms embedded in the model. Therefore, we view this results as promising and as a starting point for future research.[24]

Out-of-sample forecasts and T-periods ahead forecasts. In figure 8 we have shown one-period ahead forecasts over the horizon Q1-1948 to Q1-2018 using the CC-MABM and VAR models estimated on same time window. In this section we evaluate the forecasting performance of the ABM out-of-sample and T-periods ahead.

In the upper (lower) left panel of figure 9 we show one-period ahead forecasts of output gap (investment) out-of-sample, i.e. from Q2-2018 to Q2-2019.[25] The forecasts generated by the ABM (estimated on 1948-2018 data) are close to the actual data in the case of investment: they capture the increase in investment from Q2-2018 to Q3-2018 and the substantial stability of investment in Q4-2018 and in Q1-2019. On the contrary, the forecasts systematically underestimate the output gap. To interpret these results it is useful to recall that in the third quarter of 2017 the US Congress passed the Tax Cuts and Jobs Act of 2017, that has introduced the "most drastic changes to US tax code in 30 years"[26], worth over a trillion dollars. The tax shock may have affected the US economy from Q1 2018 onward. The underestimation of the output gap therefore may be due to the effect of the tax shock on the real economy which is only gradually incorporated in the one-period ahead forecasts. The right panel of figure 9 shows the 4-period ahead forecasts of output from Q2-2018 to Q2-2019, conditional on the observed output gap in Q1-2018. The computing burden necessary to generate these forecasts is very heavy. Therefore, we save on computing time resorting to a short cut: we compute the forecast in $t+k$ as the expected value conditional on the forecast in $t+k-1$. In other words, we use the expected value in $t+k-1$ as a proxy for the observed value in the same period. This shortcut underestimates the uncertainty around the forecasts (which we don't show). The sequence of forecasts captures

[24] The algorithm used to compute the forecasts (and the computation of the probability of specific events presented in the next section) is explained in appendix B and applied to two simple models.

[25] At the time of writing, the output gap for Q2-2019 is not available yet, therefore we show only the forecast.

[26] The Guardian, 19 December 2017, https://www.theguardian.com/us-news/2017/dec/19/donald-trump-tax-bill-plan-house-approves-senate (retrieved on 14 June 2019)

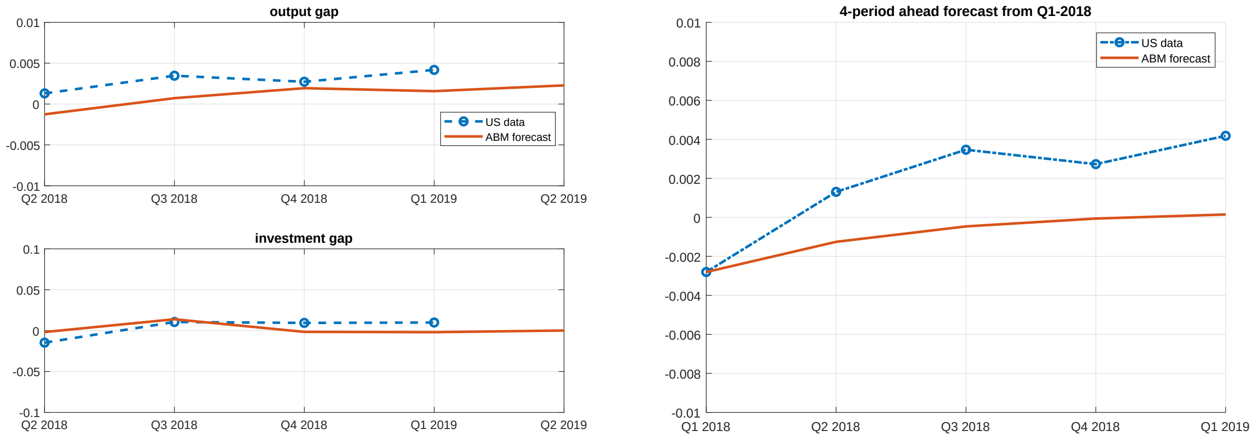

Figure 9: **Left panel**: one-period ahead out-of-sample expectations for output gap and investmant gap from Q2-2018 to Q2-2019. **Right panel**: four-periods ahead out-of-sample expectation for output gap.

the tendency of output to increase over the horizon Q1-2018 to Q1-2019 but in each quarter the forecast systematically underestimates actual output. This may be due, as explained above, to the reform of the tax code introduced in 2017. Moreover, the 4-period ahead forecast seems to converge back towards a statistical equilibrium. This is due to the fact that, as shown by the impulse response functions in section 6, the model is characterized by the tendency of converging back to its statistical equilibrium. In order to forecast richer dynamics, it is necessary increase the number of time series considered. The explanatory power of ABMs usually lies in the rich interactions between the financial and real sectors of the economy. Taking simultaneously into account the level of economic activity and financial variables such as banks' or firms' leverage would likely improve the explanatory power of the model. We are missing these characteristics of our ABM in this simple exercise. However, a part from concerns regarding computational time, the forecasting method presented in this section is perfectly capable of handling higher dimensional forecasts.

Is the recession coming? An ABM-based early warning indicator. Another interesting exercise consists in measuring the conditional probability of a *specific event*. For example, we can compute the conditional probability of a simultaneous decrease of output and investment in t conditional on the observed data in $t-1$. Given the density $P(y_t|\bar{y}_{t-1})$ we can compute the cumulated density of negative variations of output and investment, i.e., the probability of an *output plus investment* contraction in t given the state of the economy in $t-1$. This is shown in Figure 10. Results show that the probability of a joint contraction of output

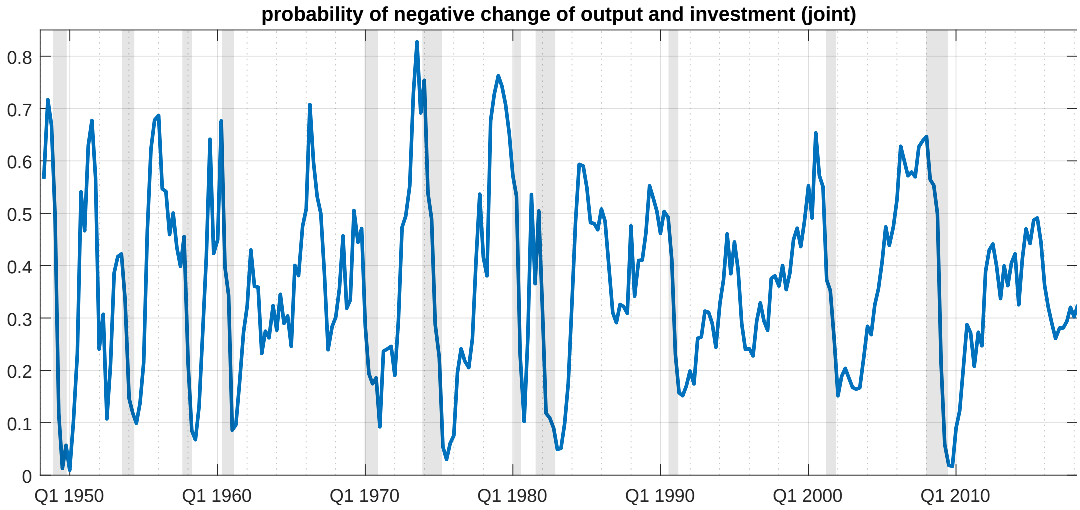

Figure 10: Probability of a joint negative change of output and investment in t conditional on observed output and investment in $t-1$. The shaded areas are recession bands identified by the National Bureau of Economic Research.

and investment estimated by means of the ABM peaks often before or at the beginning of a recession in the US economy. This measure therefore could be used as an early-warning signal of a recession. The pattern is particularly evident before the last two recessions. In both cases, the probability of a joint decrease of output and investment displays a long and steep increase, followed by a sharp reduction once the recession actually takes place. As argued above, to have better forecasts, and better measures of the likelihood of future events, we would probably need to take into account also the behavior of financial variables (and possibly a more sophisticated ABM).

8 Conclusions

In this paper we have proposed two novel methods to bring ABMs to the data.

First, we have put forward a new Bayesian estimation procedure that builds upon a former method proposed in Grazzini et al. (2017) and augment it by taking into account the time structure of simulated and observed time series. This improvement has the clear advantage of increasing the amount of information used to estimate the parameters of the model.

Second, we have proposed a fully simulation-based methodology to forecast time series. This method allows to forecast aggregate time series using data obtained from the simulation of an ABM.

We strongly believe that finding a convincing estimation method and a reliable forecasting procedure to the ABM toolbox is key in making these models usable as empirically corroborated frameworks for policy analysis. Moreover, comparing the forecasting performances of different models is a possible way to compare models.

We apply our methodological contributions to CC-MABM, a medium-scale macro agent-based model. The focus of the paper is not the model itself. However, to make the case in favor of the proposed estimation and the forecasting methods, it is important to show that they work when applied to a fully-fledged macro agent-based model. We show that the estimated model is able to reproduce some features of observed data and to forecast one-period ahead output-gap and investment.

These results should be interpreted in two ways. First, the methods we propose can be applied to macroeconomic agent-based models. Second, future research should focus on building "more useful" models. In this context, this means to design models that can better reproduce observed time series and possibly out-perform VAR (and other models) in the forecasting exercise. In other words, the estimation and the forecasting methods may offer a measure of performance/validity of models and possibly guide the design of new models.

A Comparison of Bayesian estimation techniques

Grazzini et al. (2017) proposed a Bayesian technique to estimate (fairly simple) heterogeneous agents models. The procedure proposed in section 4 of the present paper improves on the former one as it exploits the time structure of data (and of the model). In this appendix we compare the posterior obtained by means of the Grazzini et al. (2017) technique to the one generated by the improved technique proposed in this paper. We show that the procedure proposed here increases the precision of the estimate to a significant extent.

Let's start from the simplest univariate AR(1) model:

$$y_t = a_{11} y_{t-1} + \varepsilon_t \tag{10}$$

where $\varepsilon_t \sim N(0, 1)$. The model is characterized by the parameter a_{11}.

Let's assume that the real economy is completely characterized by (11) with $\bar{a}_{11} = 0.4$. We run the model with this parameter value for 1000 periods. The resulting simulated time series

– which we label the *pseudo-observed time series* – will play the role of the real benchmark. We can represent the pseudo-observed time series as follows:

$$\bar{y}_t = 0.4\bar{y}_{t-1} + \varepsilon_t \tag{11}$$

In the following we will refer to the pseudo-observed data with $\bar{Y}$. Then, we select 50,000 parameter values for a_{11} using the latin-hyper-cube function in Matlab. We restrict the admissible parameter values[27] to the interval $a_{11} \in [0,1]$. We simulate the model for 5000 periods for each of the selected 50,000 parameter values. We use these data to estimate the joint density $P(y_t, y_{t-1})$.

For each value of the parameter we compute the likelihood

$$\mathcal{L}(a_{11}|\bar{Y}) = \prod_{t=2}^{n} P(\bar{y}_t, \bar{y}_{t-1}; a_{11}) \tag{12}$$

where $P(\bar{y}_t, \bar{y}_{t-1}; a_{11})$ is the probability of observing the pair $(\bar{y}_t, \bar{y}_{t-1})$ under the estimated joint density and the specific parameter value a_{11}. We get 50,000 likelihood functions.[28] We assume uniform prior, so that the prior probability $p(a_{11})$ is given and uniform across the set of feasible values of a_{11}.

Finally, we compute the posterior probability as:

$$p(a_{11}|\bar{Y}) \propto p(a_{11})\mathcal{L}(a_{11}, \bar{Y}) \tag{13}$$

The solid line in figure 11 shows the posterior distribution computed using the likelihood function (12). For the sake of comparison, we show also the posterior distribution generated by the estimation method proposed in Grazzini et al. (2017) (dashed line). We remind that Grazzini et al. (2017) do not take into account the time structure in computing the likelihood function. The accuracy of the improved Bayesian method is much higher. The posterior peaks very close to the true value of the parameter while the posterior generated by the former method peaks around 0.35 and is characterized by a much higher dispersion. In models with strong time dependency as the AR(1) the time structure is clearly an important information

[27]In a real estimation setting any parameter restriction should be justified by economic arguments or prior information.

[28]It is important to notice that the number of lags in the joint distribution is arbitrary.

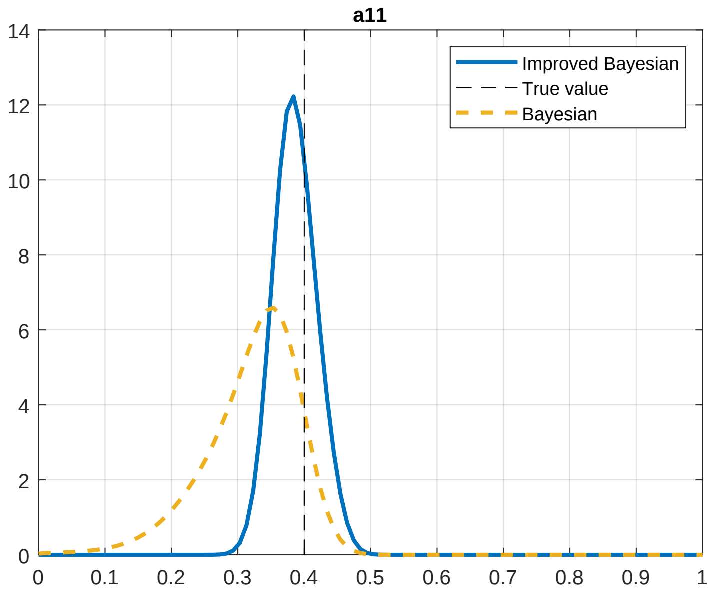

Figure 11: Posterior distribution of a_{11} using the Bayesian estimation including one lag in the likelihood (continuous blue line) and using the Bayesian estimation without lags. The vertical dashed line represents the true value of the parameter.

for the estimation process.

We perform the same exercise in a slightly more challenging environment. Consider the following bivariate VAR model:

$$\begin{pmatrix} y_{1t} \\ y_{2t} \end{pmatrix} = \begin{pmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{pmatrix} \cdot \begin{pmatrix} y_{1t-1} \\ y_{2t-1} \end{pmatrix} + \begin{pmatrix} \varepsilon_{1t} \\ \varepsilon_{2t} \end{pmatrix} \tag{14}$$

In this second exercise we assume to know a_{12} and a_{21}, and we estimate a_{11} and a_{22} using the improved Bayesian method and the method put forward by Grazzini et al. (2017). Figure 12 displays the marginal posteriors of a_{11} and a_{22}. Both methods perform well: the posterior distributions peak more or less at the same value of the paramater and fairly close to the true value. However, taking into account the time structure makes the variance of the posterior distribution shrink.

Let's now assume that all the 4 parameters are unknown and must be estimated. Figure 13 shows the 4 marginal posteriors obtained by means of the two methods. Comparing the solid line and the dashed yellow line, it is clear that the method proposed in the present paper outperforms Grazzini et al. (2017). Finally, we increase the length of the pseudo-observed time series to 5000 and perform again the posterior using the Bayesian method described in Grazzini et al. (2017). The resulting marginal posteriors are represented by the green dotted line in

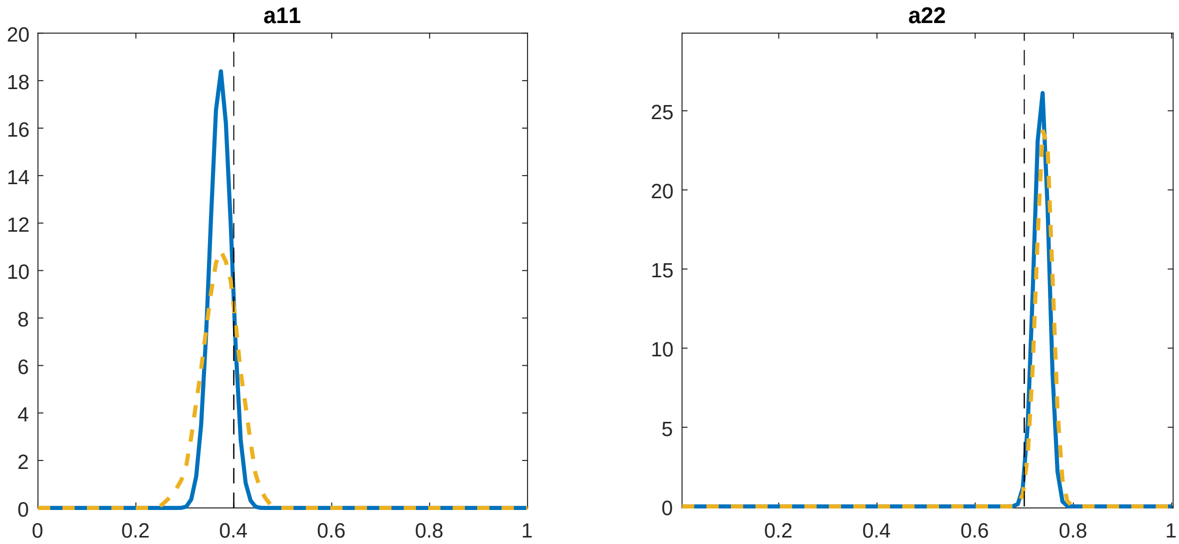

Figure 12: Posterior distribution of a_{11} and a_{22} using the Bayesian estimation including one lag in the likelihood (continuous blue line) and using the Bayesian estimation without lags. The vertical dashed lines represent the true value of the parameters.

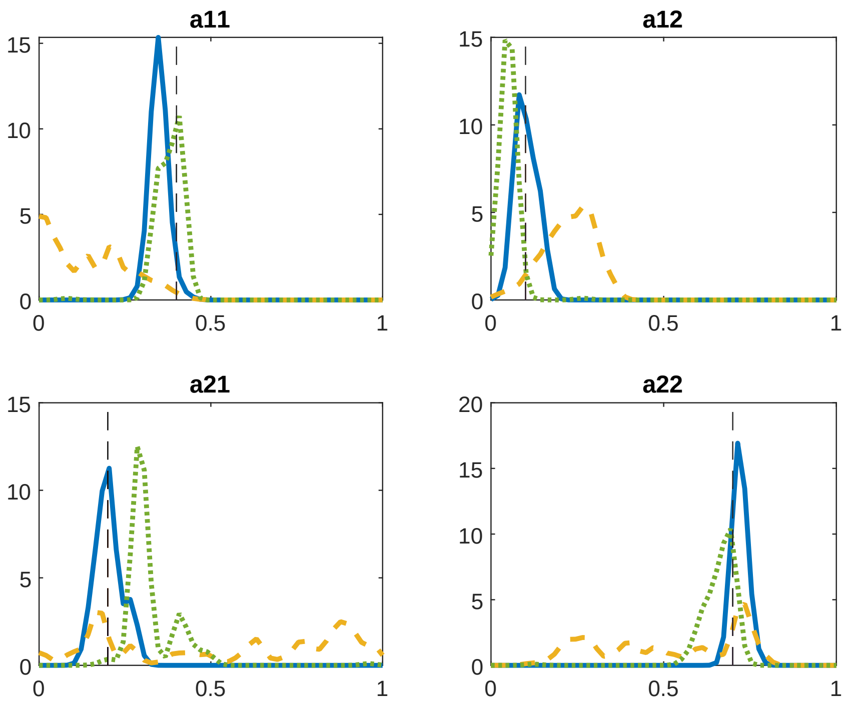

Figure 13: Posterior distribution of a_{11}, a_{12},a_{21} and a_{22} using the Bayesian estimation including one lag in the likelihood (continuous blue line), using the Bayesian estimation without lags (yellow dashed line) and using the Bayesian estimation without lags and 5 times more observations (green dotted line).

Figure 13. Essentially, this exercise shows that the improved version of the Bayesian method, by taking into account the time structure, is more efficient when the data generating process has a strong auto-correlation structure.

B Forecasting procedure

In this appendix we apply the forecasting procedure described in section 7 to two VAR models to assess in the simplest and clearest setting the performance of the method.

Consider first the following univariate AR(1) model: $y_t = a_{11}y_{t-1} + \varepsilon_t$ with $\varepsilon_t \sim N(0, \sigma)$ and $\sigma = 0.01$. Suppose the dynamics of the macroeconomic variable of interest (e.g. GDP) is correctly represented by the law of motion above with $a_{11} = 0.4$. The true law of motion of GDP therefore is

$$y_t = 0.4y_{t-1} + \varepsilon_t \tag{15}$$

We simulate the model (15) for 200 periods to generate the pseudo-observed time series.

Suppose that we have correctly estimated the model and retrieved an estimated parameter identical to the true value. We simulate the model with the "estimated" parameter (identical to the true value) for $T = 10,000$ periods and use the resulting artificial time series to determine the probability distribution of the variable.

Then, we generate a set of $\mathcal{I} = 100$ "variations" – denoted with Δy^i, $i = 1, ..., 100$ – evenly spaced in the interval $-0.05 < \Delta y^i < 0.05$. Therefore, in period t GDP can take on the following values: $y_t^i = y_{t-1} + \Delta y^i$ where y_{t-1} is the pseudo observed value in period t-1. These values characterize the possible "states" of the economy in t.

Finally, we estimate the density $P(y_t^i|y_{t-1})$ using the kernel density estimator mentioned above. By means of an appropriate normalization, we can interpret the density as a probability. Thanks to this procedure we can associate a probability to each state. To generate a forecast for GDP in t we simply take the expected value of y_t^i, conditional on y_{t-1}:

$$E(y_t^i) = \sum_{i=1}^{\mathcal{I}=100} y_t^i P(y_t^i|y_{t-1}) \tag{16}$$

By assumption, the economy is correctly represented by equation (15). Therefore, the

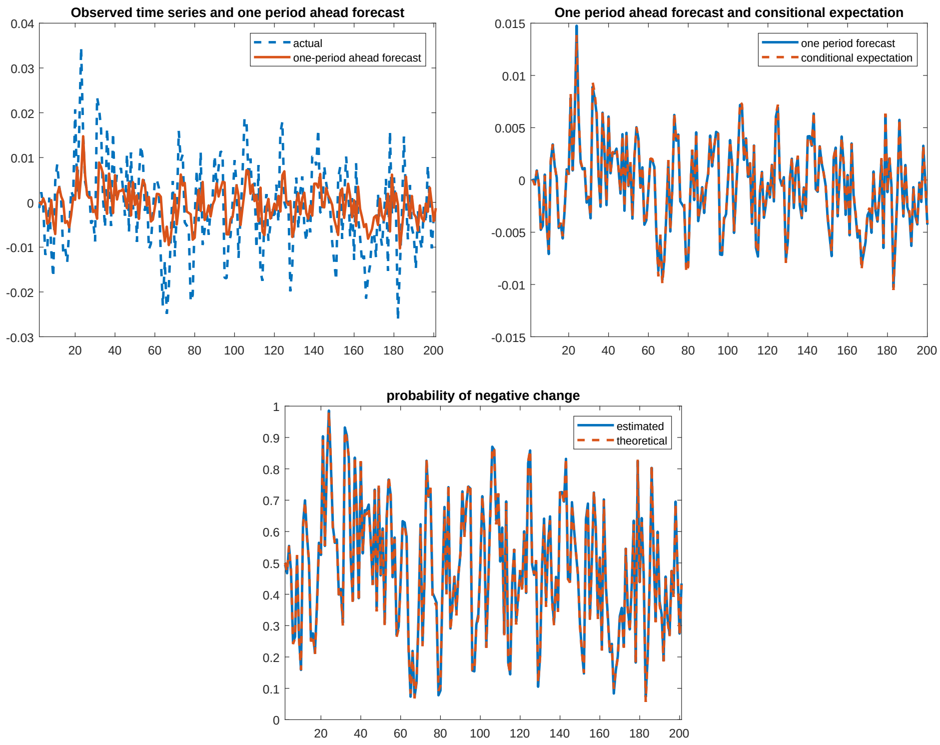

Figure 14: **Upper left**: pseudo-observed time series and one period ahead forecast. **Upper right**: one-period ahead forecast and theoretical conditional expectation. **Bottom**: probability of negative change in t given the value of the variable in t estimated and theoretical.

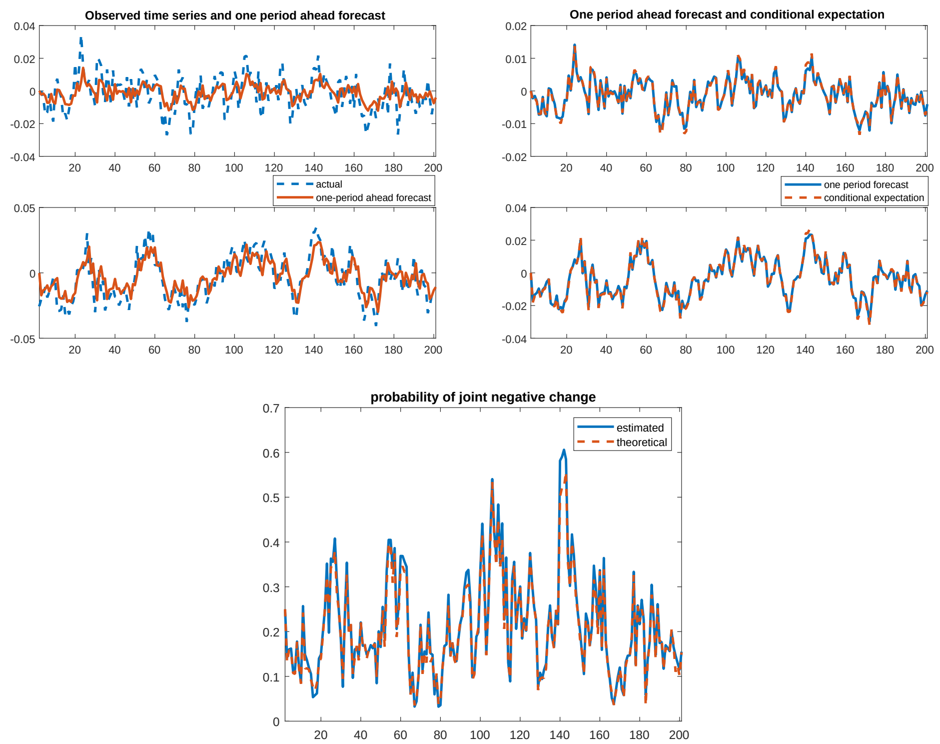

Figure 15: **Upper left**: pseudo-observed time series and one period ahead forecast. **Upper right**: one-period ahead forecast and theoretical conditional expectation. **Bottom**: probability of negative change in t given the value of the variable in t estimated and theoretical.

natural candidate to play the role of forecast in this setting is the expectation

$$E(y_t) = 0.4y_{t-1} \tag{17}$$

which is, by construction, proportional to y_{t-1}.

In the upper-right panel of figure 14 we show the one period ahead forecast $E(y_t^i)$ and the pseudo-observed GDP y_t over the life time of our economy (200 periods). The forecast is highly synchronized with actual GDP but is less volatile. In other words, a change in GDP from t-1 to t is associated with a change of the forecast in the same direction in the same interval but the amplitude of each fluctuation of actual GDP is larger than the corresponding change of the forecast.

In the upper-left panel we show the one period ahead forecast $E(y_t^i)$ and the expectation $E(y_t)$. The two time series essentially overlap which means that our forecasting method is capable of delivering the correct conditional expectation of the variable in t.

Finally, we can use the estimated conditional distribution to compute the probability of a specific event. In particular, we are interested in assessing the probability of a recession. We denote with $\Psi = \{i|\Delta y^i < 0\}$ the set of recessions, i.e., states such that $\Delta y^i < 0$. Following the procedure described above, the probability of a recession can be computed as follows

$$P(y_t^i < y_{t-1}|y_{t-1}) = \sum_{i \in \Psi} P(y_t^i|y_{t-1}) \tag{18}$$

Notice that from the model (15), we can retrieve the correct probability of a recession:

$$P(y_t < y_{t-1}|y_{t-1}) = P\big(\varepsilon_t < (1 - 0.4)y_{t-1}\big) \tag{19}$$

The estimated probability $P(y_t^i < y_{t-1}|y_{t-1})$ and the correct probability $P(y_t < y_{t-1}|y_{t-1})$ are shown in the bottom panel of figure 14. The two measures essentially overlap indicating that the procedure we propose to compute the probability of a recession is essentially correct.

We perform exactly the same exercise using a VAR model of the following form:

$$\begin{pmatrix} y_{1t} \\ y_{2t} \end{pmatrix} = \begin{pmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{pmatrix} \cdot \begin{pmatrix} y_{1t-1} \\ y_{2t-1} \end{pmatrix} + \begin{pmatrix} \varepsilon_{1t} \\ \varepsilon_{2t} \end{pmatrix} \tag{20}$$

where we set $a_{11} = 0.4$, $a_{12} = 0.1$, $a_{21} = 0.2$, $a_{22} = 0.7$. In this case the forecast for each of the two time series is computed by using the marginal distributions. The upper-left panel of figure 15 shows the two pseudo-observed time series and the one-period ahead forecasts. The upper-right panel shows the one-period ahead forecast and the correct conditional expectation for each of the two series. The bottom panel shows the probability of joint decrease of both variables as estimated using our procedure and as computed using the correct model. The forecasts and the probability measures have very good performances.

References

Assenza, T., Colzani, P., Delli Gatti, D., and Grazzini, J. (2018). Does Fiscal Policy Matter? Tax, Transfer and Spend in a Macro ABM with Capital and Credit. *Industrial and Corporate Change*, 27(6):1069–1090.

Assenza, T., Delli Gatti, D., and Grazzini, J. (2015). Emergent Dynamics of a Macroeconomic Agent Based Model with Capital and Credit. *Journal of Economic Dynamics and Control*, 50(1):5–28.

Barde, S. and van der Hoog, S. (2017). An empirical validation protocol for large-scale agent-based models. *Bielefeld Working Papers in Economics and Management*, 04-2017.

Brock, W. A. and Hommes, C. H. (1998). Heterogeneous beliefs and routes to chaos in a simple asset pricing model. *Journal of Economic Dynamics and Control*, 22(8):1235 – 1274.

Dawid, H. and DelliGatti, D. (2018). Chapter 2 - agent-based macroeconomics. In Hommes, C. and LeBaron, B., editors, *Handbook of Computational Economics*, volume 4 of *Handbook of Computational Economics*, pages 63 – 156. Elsevier.

Dawid, H., Gemkow, S., Harting, P., Van der Hoog, S., and Neugart, M. (2012). The eurace@ unibi model: An agent-based macroeconomic model for economic policy analysis.

DelliGatti, D., Desiderio, S., Gaffeo, E., Cirillo, P., and Gallegati, M. (2011). *Macroeconomics from the Bottom-up*, volume 1. Springer Science & Business Media.

Dosi, G., Fagiolo, G., Napoletano, M., and Roventini, A. (2013). Income distribution, credit and fiscal policies in an agent-based keynesian model. *Journal of Economic Dynamics and Control*, 37(8):1598 – 1625. Rethinking Economic Policies in a Landscape of Heterogeneous Agents.

Gilli, M. and Winker, P. (2003). A global optimization heuristic for estimating agent based models. *Computational Statistics & Data Analysis*, 42(3):299 – 312. Computational Ecomometrics.

Gobbi, A. and Grazzini, J. (2019). A basic new keynesian dsge model with dispersed information: An agent-based approach. *Journal of Economic Behavior & Organization*, 157:101 – 116.

Grazzini, J. and Richiardi, M. (2015). Estimation of ergodic agent-based models by simulated minimum distance. *Journal of Economic Dynamics and Control*, 51:148 – 165.

Grazzini, J., Richiardi, M. G., and Tsionas, M. (2017). Bayesian estimation of agent-based models. *Journal of Economic Dynamics and Control*, 77:26 – 47.

Guerini, M., Napoletano, M., and Roventini, A. (2018). No man is an island: The impact of heterogeneity and local interactions on macroeconomic dynamics. *Economic Modelling*, 68:82–95.

Haldane, A. G. and Turrell, A. E. (2018). An interdisciplinary model for macroeconomics. *Oxford Review of Economic Policy*, 34(1-2):219–251.

Kukacka, J. and Barunik, J. (2017). Estimation of financial agent-based models with simulated maximum likelihood. *Journal of Economic Dynamics and Control*, 85:21 – 45.

Lamperti, F., Roventini, A., and Sani, A. (2018). Agent-based model calibration using machine learning surrogates. *Journal of Economic Dynamics and Control*, pages 366–389.

Lux, T. (2018). Estimation of agent-based models using sequential monte carlo methods. *Journal of Economic Dynamics and Control*, 91:391 – 408.

Platt, D. (2019). A comparison of economic agent-based model calibration methods. *arXiv preprint arXiv:1902.05938*.

Recchioni, M. C., Tedeschi, G., and Gallegati, M. (2015). A calibration procedure for analyzing stock price dynamics in an agent-based framework. *Journal of Economic Dynamics and Control*, 60:1 – 25.

Salle, I. and Yıldızoğlu, M. (2014). Efficient sampling and meta-modeling for computational economic models. *Computational Economics*, 44(4):507–536.