

Forecasting in a complex environment: Machine learning sales expectations in a Stock Flow Consistent Agent-Based simulation model

Ermanno Catullo¹, Mauro Gallegati², and Alberto Russo^{*3,2}

¹Università degli Studi di Teramo, Italy

²Università Politecnica delle Marche, Italy

³Universitat Jaume I, Spain

Abstract

The aim of this paper is to investigate how different degrees of sophistication in agents' behavioral rules may affect individual and macroeconomic performances. In particular, we analyze the effects of introducing into an agent-based macro model firms that are able to formulate effective sales forecasts by using simple machine learning algorithms. These techniques are able to provide predictions that are unbiased and present a certain degree of accuracy, especially in the case of a genetic algorithm. We observe that machine learning allows firms to increase profits, though this result in a declining wage share and a smaller long-run growth rate. Moreover, the predictive methods are able to formulate expectations that remain unbiased when shocks are not massive, thus providing firms with forecasting capabilities that to a certain extent may be consistent with the Lucas Critique.

Keywords: agent-based model, machine learning, genetic algorithm, forecasting, policy shocks.

JEL classification: C63, D84, E32, E37.

*We would like to thank the participants to the 24th Annual Workshop on Economics with Heterogeneous Interacting Agents (WEHIA) on June 24-26 2019 (London, UK), the 31st Annual Conference of the European Association for Evolutionary Political Economy (EAEPE) on September 12-15 2019 (Warsaw, Poland), the 46th Annual Conference of the Eastern Economic Association (EEA) on February 27th-March 1st 2020 (Boston, USA), and three anonymous reviewers, for useful comments and suggestions. *Corresponding address:* Department of Economics, Universitat Jaume I, Campus del Riu Sec, Avda Vicent Sos Baynat s/n, Castellón de la Plana (Spain), E-mail: russo@uji.es

1 Introduction

The aim of the paper is studying the effects of introducing into an agent based simulation model firms that are able to formulate effective predictions on sales variations. In order to do that, we consider simple examples of machine learning, like a genetic algorithm and an auto-regressive model, which are sufficient for firms to make accurate predictions. In this framework, we test the effect of fiscal and macroprudential shocks (a variation of the capital requirement for banks) on the predictive capabilities of agents. Each firm makes expectations on the variations of its sales in order to orientate its production and price decisions. In this paper, we focus on sales expectations, thus on the features of the goods market in which production and innovation decisions are taken, whereas both the credit market and the labor market are quite simplified. In particular, in every period each firm formulates a forecast on the growth rate of its sales and uses it to determine its production choices.

The model is based on Caiani et al. (2018, 2019a). Given the focus of the paper on firms' forecasting, we slightly modified the parent model. We simulate a closed economy in which both the wage formation mechanism and the government fiscal policy rule are simplified. Finally, we modified the rule that determines selling prices and production quantities to link these decisions more directly to sales expectations. We test different methods to make sales forecasts: a genetic algorithm (*GA*), an autoregressive model (*AR*) and a naïve approach (*N*). The *GA* and the *AR* methods are able to provide expectations that are unbiased and that present a relative degree of accuracy. Thus, adopting *GA* and *AR* firms get correct insights on their sales dynamics. In a growing economy sales tend to rise, thus when firms are endowed with *GA* and *AR* they tend to set higher productive targets, thus increasing their size and, at the same time, their demand for credit. This may result in a less competitive market, with higher profits and a reduced wage share. The system modeled is a wage-led closed economy (Caiani et al., 2018, 2019a); therefore, a lower wage share tends to reduce the long-run growth trend.

We compare different scenarios featured by different learning rules, though within each scenario all firms behave according to the same rule. Basically, we aim at exploring how different learning rules shape agents' ability to forecast in a complex environment as well as which method performs better even in the case of large shocks. In order to keep the model as simple as possible regarding this aspect, our choice was to experiment one rule at a time for all firms in each scenario, so to make a comparison across scenarios and evaluate both firms' forecasting performance and the macroeconomic impact.

Simulation results show that micro is different from macro as for the impact of different

behavioral rules. Indeed, the basic result of the model points out that firms adopting more sophisticated forecasting rules are able to have a better understanding of the environment in which they operate, thus making higher profits. This results in a rising profit share that has as a counterpart a decrease of the wage share. This is what we can expect in a closed economy, which is the kind of economy we model in our paper. Indeed, closed economies tend to be “wage-led”, meaning that a rise in the wage share tends to result in a higher growth rate via an aggregate demand channel. Conversely, an increase in the profit share tends to hamper the growth performance. Accordingly, the main effect of a wage share change in the economy is related to the fact that in a closed economy the positive effect of an increasing wage share on the macroeconomic performance, that is more consumption and then an increase of the aggregate demand, tends to outperform the negative effect due to an increase of production costs for firms. The following papers stress this kind of result from both a theoretical and empirical perspective: Onaran and Galanis (2012), Onaran and Obst (2016), Stockhammer et al. (2009), Stockhammer and Sotiropoulos (2014).

Fiscal and macroprudential shocks have a similar impact on the economy either when firms use more sophisticated predictive methods or when they follow simple adaptive rule as in the baseline scenario. Nevertheless, macroprudential shocks tend to have a stronger effect in the *GA* and *AR* settings because firms’ leverage tends to be higher in such cases. Moreover, it is worth noting that when shocks are not massive, firms predictions are still effective, so predictive methods are able to provide firms with a forecasting capacity that leads them to consistent behaviors even in an environment that is affected by exogenous shocks. In a sense, then, firms endowed with more sophisticated forecasting rules exhibit a behavior which is, to a certain extent, Lucas critique proof. In other words, although agents are not characterized by rational expectations and they are not maximizing intertemporally an objective function, they can effectively adapt to a complex environment based on evolutionary principles as selection and experimentation, even in presence of unexpected shocks of non-negligible size.

In the next section, we provide a discussion of the literature on agent-based macro modeling, focusing on the role of expectations and the contributions coming from laboratory experiments. In the third section, we describe the model. The fourth section illustrates the micro and macro effects of adopting sales predictive methods. In the fifth section, we explore the impact of both macroprudential and fiscal shocks on the economy. The sixth section synthetically describes the effects of shocks in market with different degrees of competition. Then, we provide some concluding remarks.

2 Related literature

Mainstream macro models typically assume that agents are characterized by rational expectations (Muth, 1961). In a nutshell, agents' expectations on a given variable correspond to the mathematical conditional expectations implied by the model; consequently, only random errors may appear in individuals' forecasts, or in other words, agents do not make systematic mistakes and the mean error is zero. Agents are assumed both to know how “the model” of the economy works and to have access to all relevant information. Moreover, each agent knows that other agents know about that, according to a “common knowledge” assumption. This is equivalent to affirm that a universal model of the world exists and that events occur according to a well-known probability distribution. The rational expectations hypothesis is a central piece of the most important mainstream macro tool, namely DSGE models, and it is at the odd with the concept of “fundamental or strong uncertainty” (Knight, 1921; Keynes, 1937). In particular, rational expectations are not viable under strong uncertainty, when we cannot assign an objective probability to an event, and there are radical policy changes (Stiglitz, 2016), as well as when structural breaks invalidate agents' forecasts based on a given probability distribution (Hendry and Mizon, 2010).

Among the various direction of research triggered by the Great Recession, from financial frictions to heterogeneous agents and so on, mainstream macroeconomists have proposed some models for dealing with the actual limitations in agents' rationality, from statistical learning to restricted perceptions and rational inattention (see Woodford (2013), for a review). However, the impression is that this literature is providing some sort of quasi-rational expectations equilibrium models, which only partly deviate from the fully rational benchmark, perhaps underrating the macroeconomic implications of a different modeling of agents' behavior more in line with the original interpretation of bounded rationality (Simon, 1959). In our view, agent-based models are closer to Simon's bounded rationality in that they typically assume that agents follow (relatively) simple behavioral rules according to “procedural rationality” (Simon, 1976).

Rational expectations, as a model-consistent form of expectations, represent a too strong form of rationality to be realistic. The polar opposite is to model agents behavior based on naïve expectations (like static adaptive expectations) as it has been done in many macro agent-based models up to the present (similarly to the Keynesian macroeconometric models of the 1950s and 1960s). As opposed to rational expectations, this approach proposes a too weak form of rationality to be realistic. ¹ Experimental studies show that individuals tend to be characterized

¹In a sense, rational expectations can approximate the behavior of a few, well informed individuals with the

by adaptive expectations, rather than rational expectations, though they are for example able to understand that a trend is present in the evolution of a variable. This is the case of a Learning-To-Forecast (LtF) experiment, a kind of laboratory experiment firstly introduced by Marimon et al. (1993), proposed in Colasante et al. (2017) in which individuals have to provide one-period-ahead forecast of the price of a single asset, having the following information: dividend, interest rate, past realization of the market price, their own past predictions. Under the assumption that individuals' profit depends on minimizing the distance between their own prediction and the average forecast price, the result is that individuals hardly predict the fundamental price, while collectively they are not too far from it (though they systematically underestimate a positive trend). Thus, the paper rejects rational expectations (the zero mean condition is not satisfied) and points to adaptive expectations (with a certain degree of correction bias) as a more realistic scheme to describe the behavior of experimental subjects.²

Rejecting the rational expectation hypothesis, however, does not immediately provide an alternative model with which to interpret the actual behavior of agents. In other words, the absence of a precise or axiomatic basis for modeling the rational behavior of agents may result in the “wilderness of bounded rationality” (Sims, 1980). What we can say is that people follow some sort of adaptive expectations with a certain degree of forward-looking behavior, though this attitude is undoubtedly heterogeneous. For instance, three typical patterns in aggregate price behavior have been observed in the context of Learning-To-Forecast experiments (Hommes et al., 2005): slow monotonic convergence, permanent oscillations and dampened fluctuations. Anufriev and Hommes (2012) show that a simple model of individual learning can explain these kind of aggregate regularities emerging in experimental data (see also Hommes (2013), for a more general view): the model is based on evolutionary selection among heterogeneous expectation rules driven by their relative performance; in other words, agents can switch from a rule to another one based on the relative performance of each single rule. Therefore, agents learn from

attitude to reason on the future evolution of the economy, also trying to anticipate policy decisions, but they hardly represent the reasoning featuring the vast majority of people; by contrast, naïve expectations do not approximate the behavior of individuals with a forward-looking attitude, while they may provide a rough description of the adaptive behavior exhibited by a vast part of the population. While in the present paper we propose a simple framework in which all agents are endowed with the same behavioral rules (then comparing different scenarios with different rules), the idea of considering various degrees of rationality for different agents is something that we would address in an extension of the present computational framework. The experimental literature has already stressed the presence of heterogeneous expectations in various contexts; see, for instance, Hommes (2011, 2013). Indeed, expectations of participants in laboratory experiments can be highly heterogeneous, as found for example in Arifovic and Petersen (2017), where the degree of heterogeneity in subjects' expectations tends to increase when the game played is more complex.

²It is worth noting that individuals participating in a LtF experiment typically exhibit an adaptive behavior which results in following both a positive – as for instance in Colasante et al. (2017) – and a negative trend – as for instance in Anufriev and Hommes (2012) –, though not always converging exactly to it. The firms we consider in this paper, which are endowed with Machine Learning tools, are able to catch and follow the growth trend thus resulting in unbiased expectations. Also in the case of our model, agents (namely firms) are able to detect both a positive and a negative growth trend.

the past and, based on some “fitness” measure, they choose among different rules of behavior to better perform in a complex and uncertain environment. Hence, laboratory experiments may help us in detecting some specific forms of bounded rationality to be studied in analytical or computational models. We will see how in our agent-based model introducing learning according to a genetic algorithm affects both individual (positively, by increasing profits) and collective performance (negatively, by decreasing the wage share and then deteriorating the macroeconomic performance in a closed-economy setting).

As for macro agent-based models, endowing agents with very simple adaptive expectations has been a way to focus on other very relevant features of the economy, like heterogeneity and interaction, and provide a demonstration that complex properties at the macro level can result from the interaction of heterogeneous micro-entities, from the bottom up (Evans and Honkapohja, 1996). Consequently, an extended crisis like the Great Recession is not the consequence of a large exogenous shock, whereas it can be interpreted as the result of an endogenous process of debt accumulation, in a context of growing inequality, enlarged by the contagion among financially fragile (Delli Gatti et al., 2010; Russo et al., 2016) However, the task of showing that realistic macro properties can be the result of the interaction among (almost) zero-intelligence agents (Gode and Sunder, 1993) has been by now accomplished by a number of different models and research projects.

Russo et al. (2007) provide a macro setting with heterogeneous households and firms in which R&D expenditures boost long-run growth, with beneficial effects of fiscal incentives provided by the public sector; this paper has been then extended by introducing banks and credit relationships in Delli Gatti et al. (2011); Delli Gatti et al. (2010) explore the connection between the endogenous evolution of credit networks (firm-bank interlinkages) and business cycle fluctuations; Dosi et al. (2010) propose a Keynes plus Schumpeter model with endogenous growth able to reproduce a long list of micro, meso and macro facts, then employed for analysing fiscal and monetary policies (Dosi et al., 2015), or the destabilizing role of wage flexibility (Dosi et al., 2017); Riccetti et al. (2015) provide a fully decentralized matching protocol for simulating the working of different markets in an agent-based macro setting, then employed to analyze inequality and financial crises (Russo et al., 2016), financial regulation (Riccetti et al., 2018), monetary policy at the Zero Lower Bound (Giri et al., 2019); Caiani et al. (2016) propose a benchmark agent-based stock-flow consistent macro model, then extended to explore the interplay between inequality and long-run growth (Caiani et al., 2019b, 2020); Dawid et al. (2018) present the large-scale EURACE model, and discuss some general aspects of ABM, a framework already

used for analysing financial regulation (Cincotti et al., 2010), regional gaps and innovation policy (Dawid et al., 2014), and fiscal policy (Teglio et al., 2019); see Dawid and Delli Gatti (2018), for a comprehensive review.³ For the next future, there are two main challenges regarding macro agent-based models: enriching the modeling of agents' behavior in order to develop a new framework with heterogeneous boundedly rational expectations (here an integration between computational and experimental economics is in order); improve their empirical performance by developing new techniques for calibrating and estimating macro agent-based models (just to make a few examples: Bianchi et al. (2007); Bargigli et al. (2014, 2018); Grazzini et al. (2017)). The present paper contributes to the first challenge by comparing the performance of different types of behavioral rules, from naïve ones to machine learning, focusing on firms' sales expectations, in a closed economy macro agent-based model.

Some papers provided a study of different degrees of learning featuring agents' behavior to see how different rules can affect the macroeconomic performance. Sinitskaya and Tesfatsion (2015) propose a computational analysis based on a well-known New Keynesian three-equation models (Smets and Wouters, 2003). They remove the rational expectations hypothesis, thus allowing for out-of-equilibrium dynamics. Then, they endow agents with some specific rules of behaviors, from a simple reinforcement learning to more sophisticated schemes, up to approximate or adaptive dynamic programming (ADP). For each kind of behavioral rule, they analyze their macro impact, relative to a social planner benchmark solution, noting that it is not necessary that agents are featured with highly sophisticated, almost rational rules, to reach a better performance. They show that simpler rules outperform more sophisticated ones provided that the memory length is sufficiently long for an effective adaptive foresight. Dosi et al. (2017) analyze the individual and macroeconomic impact of heterogeneous expectations taken by Anufriev and Hommes (2012) within a macro agent-based setting. Accordingly, firms may choose among the following rules: naïve expectations (namely, the value expected in $t + 1$ corresponds to the value of the same variable at time t), adaptive expectations (that adjusts the previous period expectation by the difference between the realised and the expected value in the past period), weak and strong trend expectations (according to which firms behave like a “chartist”, and anchor and adjustment expectations (inspired by Tversky and Kahneman (1974), in which firms react to both their past demand and to an aggregate “anchor” like the GDP). Firms may switch from a rule to another one based on their relative performance, so producing an ecology of rules according to an evolutionary selection mechanism. The result of this computational investigation is that

³See also the review by Caverzasi and Russo (2018) that introduces the special issue entitled “Toward a new microfounded macroeconomics in the wake of the crisis”.

neither individual nor aggregate dynamics improve when firms replace myopic expectations with less naïve rules. Indeed, a sophisticated rule like recursive least squares give rise to less accurate individual predictions and worsen the macroeconomic performance, suggesting that rather than being rational, agents following “rational heuristics” may lead to more accurate individual forecasts and better aggregate results.

Both Sinitskaya and Tesfatsion (2015) and Dosi et al. (2017) point out that relatively low rational behavioral rules may perform better than more sophisticated schemes, resulting in better performance of the macroeconomy. Our results suggest something similar in that when firms abandon naïve rules and become to learn about how to forecast sales according to a genetic algorithm, the macroeconomic performance deteriorate (while the goods market becomes more concentrated). This is due, however, to the specific closed economy setting in which agents make their choices. The deterioration of macroeconomic conditions, indeed, is the consequence of the reduction of the wage share, which tends to decrease the aggregate demand. This is a necessary consequence of the shifting distribution of aggregate income from workers to firms in a Stock-Flow Consistent closed economy, with an unchanged fiscal stance. Indeed, the decline of the wage share is specular to an increase of the profit share. The latter is, in turn, due to an increase of firms’ profit as a consequence of the improved forecasting abilities of firms when more sophisticated rules are used. In the case of our paper, then, while individual forecasts ameliorate due to more rational behavioral rules (with zero mean errors), economic growth decelerates due to a lack of aggregate demand. In our view, the complex micro-macro relationship involving individual behavior, related to forecasting, and collective behavior, embedded in the economic and financial structure of the macroeconomy, will require more work, also focusing on the distributive impact of different degrees of sophistication regarding agents’ behavioral rules. Differently from Dosi et al. (2017), in which by using recursive least squares firms make worst forecast while the macro performance deteriorates, in our model we also observe a deterioration of the aggregate economic performance, but this is due to better forecasting provided by machine learning sales expectations that make firms’ profit to go up, thus lowering the wage share and then the aggregate demand. Moreover, the fact that relatively simple behavioral rules may enhance the economic performance within a macro model does not mean that this kind of expectations necessarily reflect the actual behavior of different types of agents like households, firms or banks. From this point of view, a great contribution can come from experimental economics and from an integration between the latter and computational methods like agent-based modeling.

Although our paper finds that less rational rules tend to be associated to better macroeco-

conomic performance, in line with the findings of Sinitskaya and Tesfatsion (2015) and Dosi et al. (2017), the focus is on the search for sufficiently rational behaviors, like auto-regressive rules and machine learning based on genetic algorithms, that can assure a zero mean error, as well as an ability to recover quite quickly after a (not too large) shock. In a sense, our aim is to provide a proof-of-concept: endowing firms with machine learning sales expectations allows us to demonstrate that agents are able to react to news, like shocks, returning to zero mean error in a reasonable time, thus passing to a certain extent the Lucas critique (Lucas, 1976). As we will see, the first principle to achieve the individual and collective adaptation to an ever-changing environment is not individual maximization but rather selection and experimentation along evolutionary lines. This is the main motivation for suggesting a firms' pricing strategy based on genetic algorithm machine learning.

Though agents are not fully rational, as in our model, they could be able to forecast in a quite accurate way, then avoiding systematic mistakes and effectively adapting to a changing and uncertain environment. The literature on learning in macroeconomics (Evans and Honkapohja, 2012) provides many examples that can be used for endowing agents with more sophisticated rules. In the present paper, we test some possibilities like auto-regressive forecasting and a genetic algorithm-based machine learning. Other examples involving the application of learning models in macroeconomics are Arifovic et al. (2010) and Salle (2015). Following Turrell (2015), we make a step in the direction of proposing an agent-based model which is as Lucas-critique proof as possible, with firms able to adapt quite effectively to changing environment and policy changes (unless the shock is too large). Nevertheless, behavioral rules remain tied to bounded rationality and then deliberately not resting on intertemporal constrained maximization of an objective function and model-consistent rational expectations. That people are not fully rational is a mere fact. Accordingly, we are not required to develop models that are not affected by the Lucas critique. It can certainly be that most people follow simple behavioral rules, but this does not mean that at least a minority cannot be characterized by more rational behaviors. According to what said above, expectations are heterogeneous as shown by experimental investigations. Perhaps only the behavior of a few individuals could be considered (almost) Lucas-critique proof, while many others fail by a little or completely in doing it. When we consider agents like firms and banks, especially large firms and banks, it is more likely that these kinds of agents exhibit an improved ability to forecast, according to more sophisticated rules. In the present paper we decided to focus on firms (and the goods market) through sales expectations, while leaving the other markets unchanged (with respect to the benchmark model). It is only a first step

toward the complex quest of understanding the individual and collective effects of heterogeneous boundedly rational expectations in complex economies.

3 The Model

3.1 Model description

Our model is based on the agent based-stock model proposed by Caiani et al. (2019a). The main difference with respect to Caiani et al. (2019a) is that the model simulates a closed economy. Moreover, we modify it to allow firms to formulate expectations on sales. Given that the present paper focuses on the effect of firms' expectations on sales, the wage rule has been simplified to avoid the necessity of formulating expectations on unemployment levels (Equations 4 and 13). Besides, the behavioral rule that governs the selling price and quantity produced by firms has been modified in order to take into account forecasting (Equations 3.2 and 4).

The model considers a 'pure labor' economy where firms' production is carried out by using labor only. The simulated economy is populated by a given number of households (H) and by an endogenously varying number of firms (I_t) and banks (Z_t), depending on defaults arising endogenously during the simulation as well as on households' investment in the creation of new firms and banks.

Government fiscal policy tries to meet a target deficit-to-GDP ratio. Indeed, government collects taxes on income and profits and provides public spending in the form of transfers to households. The Central Bank sets the discount rate and buys the possible residual bonds issued by the government which were not purchased by private banks.

Following the logic of Riccetti et al. (2015, 2016); Caiani et al. (2016, 2019b, 2018), model dynamics is driven by agents' adaptive reactions and decentralized interactions in different markets: goods markets, labor markets, deposit markets, credit and bond markets.

Next subsections sketch out the behaviors of agents and the structure of their interactions on the different markets (Fig. 1).

3.2 Firms

Firms' desired output level $q_{i,t}^P$ depends on both desired sales $q_{i,t}^S$ and the level of inventories inherited from the past $inv_{i,t}$. Firms try to maintain a certain amount of inventories, computed as a share θ of desired sales.

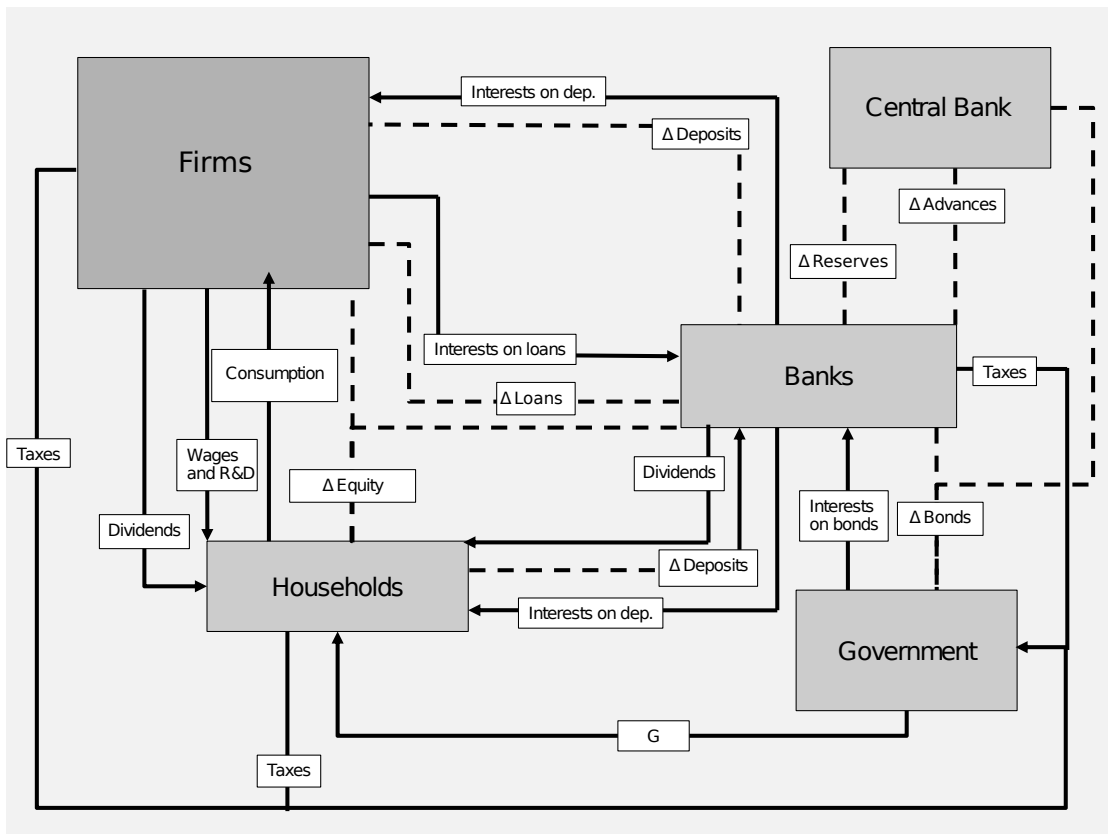


Figure 1: Flow chart of the simulated economy. Arrows indicate the direction of flows, dashed lines are balance sheet variations.

$$q_{i,t}^P = q_{i,t}^S(1 + \theta) - inv_{i,t} \quad (1)$$

In the baseline scenario, prices $p_{i,t}$ and desired sales $q_{i,t}^S$ are revised adaptively from period to period according to a simple scheme depending on $\hat{q}_{i,t-1}$, the quantities sold in $t - 1$.

$$\text{if } \hat{q}_{i,t-1} \geq q_{i,t-1}^S : \begin{cases} q_{i,t}^S = q_{i,t-1}^S(1 + U[0, \delta]) \\ p_{i,t} = p_{i,t-1}(1 + U[0, \delta]) \end{cases} \quad (2)$$

$$\text{if } \hat{q}_{i,t-1} < q_{i,t-1}^S : \begin{cases} q_{i,t}^S = q_{i,t-1}^S(1 - U[0, \delta]) \\ p_{i,t} = p_{i,t-1}(1 - U[0, \delta]) \end{cases} \quad (3)$$

Equation 2 states that if past sales exceeded desired sales, firms adaptively increase both the desired sales and their selling price. The opposite happens when past sales are lower than desired sales (equation 3). Moreover, prices have a lower bound represented by unit costs of production: $p_{i,t} \geq \frac{w_{i,t}}{\phi_{i,t}}$, where $\phi_{i,t}$ is firm's i labor productivity in period t . In the fourth section we modify the production and price decision to incorporate selling expectations.

Firm's labor demand is then computed as: $l_{i,t}^D = q_{i,t}^P / \phi_{i,t}$. Output can be lower than desired one when the firm is not able to employ all the worker needed because of financial constraints or scarcity of labor supply.

The salary $w_{i,t}$ offered by firm i is adaptively revised (Equation 4). If labor employed was below the demanded, firms will increase wages with a certain probability $Pr(w_{i,t}^+)$. The opposite happens when the labor demand was satisfied in the previous period. The wage offered cannot be lower than zero.

$$w_{i,t} = \begin{cases} w_{i,t-1}(1 + U[0, \delta]), & \text{if } l_{i,t-1}^D - l_{i,t-1} > 0 \text{ with } Pr(w_{i,t}^+) = v_F \\ w_{i,t-1}(1 - U[0, \delta]), & \text{if } l_{i,t-1}^D - l_{i,t-1} = 0 \text{ with } Pr(w_{i,t}^-) = 1 - v_F \end{cases} \quad (4)$$

With respect to Caiani et al. (2018), the probability of varying the offered wage does not depend on aggregate unemployment. We made this modification in order to simplify the model and avoid that the prediction methods we implement could influence the macro dynamics through the wage setting mechanism, while we are focusing on sales' choices. The same variation is made to equation (13). Moreover, firms try to introduce incremental innovation through $R\&D$ expenditure. Desired innovation expenditure is a given share γ of past sales, thus $R\&D_{i,t}^D$ is:

$$R\&D_{i,t}^D = \gamma \hat{q}_{i,t-1} p_{i,t-1} \quad (5)$$

The expenditure in R&D determines the probabilities of increasing productivity ($Pr_{success_{i,t}}$) through both genuine incremental innovation and imitation (Dosi et al., 2010):

$$Pr_{success_{i,t}} = 1 - e^{-\frac{\nu R\&D_{i,t}}{\Phi_t P_t}} \quad (6)$$

where P_t is the average level of prices and Φ_t is the average level of productivity; this correction through average productivity and prices is necessary because both prices and productivity tend to grow during the simulation. Thus, the higher the expenditure in R&D the higher the probability of innovating; when a firm succeeds in innovating, its productivity increases by:

$$\phi_{i,t+1} = \phi_{i,t}(1 + U[0, \delta]) \quad (7)$$

Moreover, if productivity is under the average level, a firm may try to imitate increasing productivity of others in order to reduce the gap with the average level (Eq. 8). For the sake of simplicity, also imitation depends on R&D expenditure with the same probability of genuine innovation (Eq. 6).

$$\phi_{i,t+1} = \phi_{i,t} + U[0, (\Phi_t - \phi_{i,t})] \text{ if } \phi_{i,t} < \Phi_t \quad (8)$$

Firms fund their productive and R&D expenditures first of all through their internal funds, $D_{i,t}$, and only after by asking credit to banks, given that external credit is more expensive than internal resources (Meyers, 1984); thus credit demand becomes $L_{i,t}^D$:

$$L_{i,t}^D = \begin{cases} w_{i,t} l_{i,t}^D + R\&D_{i,t}^D - D_{i,t}, & \text{if } w_{i,t} l_{i,t}^D + R\&D_{i,t}^D > D_{i,t} \\ 0, & \text{if } w_{i,t} l_{i,t}^D + R\&D_{i,t}^D \leq D_{i,t} \end{cases} \quad (9)$$

If internal funds, $D_{i,t}$, are not enough, the firm may ask for credit to several banks. Nevertheless, they may not receive all the funds they need ($L_{i,t} \leq L_{i,t}^D$): in this case, they firstly finance production and only after innovation. For simplicity, we assume that credit lasts for only one period (Delli Gatti et al., 2010; Dosi et al., 2010; Riccetti et al., 2015).

Firms' profits are the sum of revenues $p_{i,t} \hat{q}_{i,t}$, interests on deposits $r_{d,t} D_{i,t}$, where $r_{d,t}$ is the interest rate paid on the firm's deposit $D_{i,t}$, and the variation of inventories $\Delta INV_{i,t}$, minus

wages $w_{i,t}l_{i,t}$, R&D expenditure $R\&D_{i,t}$, and interests on loans $r_{i,t}L_{i,t}$, where $r_{i,t}$ is the interest rate charged by the bank on the firm's loan $L_{i,t}$:

$$\pi_{i,t} = p_{i,t}\hat{q}_{i,t} + r_{d,t}D_{i,t} + \Delta INV_{i,t} - w_{i,t}l_{i,t} - R\&D_{i,t} - r_{i,t}L_{i,t} \quad (10)$$

Firms' net operating cash flows ($\pi_{i,t}^*$) is given by profits minus variation of inventories. When the net operating cash flows is positive, firms pay taxes ($T_{i,t}^\pi$) to the government and dividends ($Div_{i,t}^\pi$) to shareholders, thus firms may retain part of the profit in the form of deposits:

$$T_{i,t}^\pi = \begin{cases} \tau_{k,t}\pi_{i,t}^*, & \text{if } \pi_{i,t}^* > 0 \\ 0, & \text{if } \pi_{i,t}^* \leq 0 \end{cases} \quad (11)$$

$$Div_{i,t}^\pi = \begin{cases} \rho(\pi_{i,t}^* - T_{i,t}^\pi), & \text{if } \pi_{i,t}^* > 0 \\ 0, & \text{if } \pi_{i,t}^* \leq 0 \end{cases} \quad (12)$$

3.3 Households

Households play different roles in this model: first of all, they are workers. The labor supply of each household l^S is inelastically equal to one. Each households can be employed by different firms (until a maximum of ψ firms).⁴ We define $l_{hi,t}$ as the quantity of labor sold by household h to firm i , and $w_{hi,t}$ the wage received, thus the labor effectively supplied by the household h is: $l_{h,t} = \sum_{i,l_{hi,t}>0} l_{hi,t}$.

Workers do not accept jobs below a reservation level $w_{h,t}$ which depends on their past employment status and cannot be lower than zero by assumption:

$$w_{h,t} = \begin{cases} w_{h,t-1}(1 + U[0, \delta]), & \text{if } l^S - l_{h,t-1} = 0 \text{ with } Pr(w_{h,t}^+) = \nu_H \\ w_{h,t-1}(1 - U[0, \delta]), & \text{if } l^S - l_{h,t-1} > 0 \text{ with } Pr(w_{h,t}^-) = 1 - \nu_H \end{cases} \quad (13)$$

Moreover, workers receive an extra wage from innovation activities. R&D expenditure ($R\&D_{i,t}$, see section 3.2) is added on to workers' income according to the quantity of labor that each worker provides to the firm.

In addition, households also receive interests on deposits $D_{h,t}$ from banks, dividends from

⁴This is just a simplifying assumption which allows us avoiding to enlarge and reduce for each time period the size of the model; accordingly, we keep unchanged the number of agents in order to reduce the computational burden. We tested the model with the possibility for each worker of being employed just in one firm; the results are very similar compared to the baseline model. Simulation results are available upon request.

shared firms and banks ($Div_{h,t}$), and monetary transfer (G_t/H) from the government.

Given the tax rate on income (τ_t), households' gross and net income (respectively $y_{h,t}$ and $y_{h,t}^D$) are:

$$y_{h,t} = \sum_{i,l_{hi,t}>0}^{I_t} w_{hi,t} l_{hi,t} + r_{d,t} D_{h,t} + Div_{h,t} + \sum_{i,l_{hi,t}>0}^{I_t} R\&D_{i,t} \frac{l_{hit}}{l_{it}} \quad (14)$$

$$y_{h,t}^D = (1 - \tau_t) y_{h,t} + \frac{G_t}{H} \quad (15)$$

Households' desired consumption ($C_{i,t}^D$) is a function of current disposable income ($y_{h,t}^D$) and current wealth held in the form of deposits ($D_{h,t}$), given the marginal propensities c_y and c_d :

$$C_{h,t}^D = c_y y_{h,t}^D + c_d D_{h,t} \quad (16)$$

The final good market follows a monopolistic competition model à la Hotelling (Hotelling, 1929; Salop, 1979). Firms produce different varieties of a final goods and consumers have different preferences; each firm and each consumer is located on a circle and the distance between them represents the difference between the variety produced by a firm and a consumer's preference. We define d_{hi} as the distance between consumer h and a firm i . Consumers samples randomly ψ firms and compare the supply of those firms ranking them according to price and variety; household h prefers firm i to firm j if:

$$\frac{1}{d_{hi}^\beta} \frac{P_t}{p_{i,t}} > \frac{1}{d_{hj}^\beta} \frac{P_t}{p_{j,t}} \quad (17)$$

where $p_{i,t}$ and $p_{j,t}$ are the prices of two firm, P_t is the average price, and $\beta \geq 0$ is a parameter representing the importance of variety with respect to prices. Based on the ranking, each consumer buys her/his most preferred good according to her/his desired consumption. Both firms and consumers are randomly located in the circle that represents varieties of goods and preferences; agents cannot change their position. When a new firm enters the market its position is chosen randomly.

Households' income is mainly due to labor, but also by interest on deposits and dividends proportionally to the share that they own in banks and firms $A_{h,t}$. In each period, after the consumption decision, households have to decide how to allocate their savings between deposits or and equities that, in turn, will contribute to the creation of new firms and/or banks. The higher the past dividends the higher the investment in equities. Indicating by $lp_{h,t}$ the share or

wealth that households desire to hold as deposits, we have:

$$lp_{h,t} = \begin{cases} \lambda e^{-\left(\frac{Div_{h,t-1}}{A_{h,t-1}}(1-Pr_t^{default})-r_{d,t}\right)} & \text{if } \frac{Div_{h,t-1}}{A_{h,t-1}} \geq r_{d,t} \text{ and } A_{h,t-1} \geq 0 \\ \lambda & \text{if } \frac{Div_{h,t-1}}{A_{h,t-1}} < r_{d,t} \text{ or } A_{h,t-1} = 0 \end{cases} \quad (18)$$

with $0 < \lambda < 1$.

Thus, $NW_{h,t}^D = NW_{h,t-1} + y_{h,t}^D - C_{h,t}^D$ is the households' expected level of net worth; the desired level of equity and deposits can be then expressed as:

$$A_{h,t}^D = \max \{A_{h,t-1}, (1 - lp_{h,t})NW_{h,t}^D\} \quad (19)$$

$$D_{h,t}^D = NW_{h,t}^D - (A_{h,t}^D - A_{h,t-1}) \quad (20)$$

where $A_{h,t}^D - A_{h,t-1}$ is the desired investment in equity.

Households having a positive desired investment act together as an investment fund to create new firms or new banks. If funds collected are sufficient (i.e. see section 3.6.1), new enterprises are created.

3.4 Banks

Banks create money endogenously providing credit to firms. In order to avoid taking excessive risks, banks have to respect the capital requirement μ_1 : the maximum amount of credit that banks are willing to supply in any given period is a multiple μ of their equity $A_{z,t}$, thus $L_{z,t}^{DS} = \mu A_{z,t}$

Banks lend credit to firms with a probability $Pr(Loan_{i,t})$ of providing credit to a firm; the corresponding interest rate ($r_{i,t}$) depends on the desired level of firms' leverage ($L_{i,t}^D/A_{i,t}$):

$$Pr(Loan_{i,t}) = e^{-\chi \frac{L_{i,t}^D}{A_{i,t}}} \quad (21)$$

$$r_{i,t} = \chi \frac{L_{i,t}^D}{A_{i,t}} + r_t \quad (22)$$

Households and firms have deposit accounts into banks, paying an interest $r_{d,t}$ equal to a fraction ζ of the discount rate r_t set by the Central Bank. Banks have to respect minimal reserve requirements, computed as a share μ_2 of their deposits: $R_{z,t}^M = \mu_2 D_{z,t}$.

If reserves $R_{z,t}^M$ are larger than the minimum level, banks ask for advances to the Central Bank, receiving an amount $L_{zCB,t}$ at the discount rate r_t . Instead, if banks have more reserves

than needed, the excess is invested in public bonds ($B_{z,t}$) at the interest rate $r_{b,t}$. Banks' probability of purchasing each tranche of public debt follow equation (21), using as a measure of leverage the public debt-to-GDP ratio. Thus, banks' profits are:

$$\pi_{z,t} = \sum_{i, L_{iz,t} > 0}^{I_t} r_{i,t} L_{iz,t} + \sum_{k, B_{z,t} > 0}^{I_t} r_{b,t} B_{z,t} + r_{re} R_{z,t} - BD_{iz,t} - r_{d,t} D_{z,t} - r_t L_{zCB,t} \quad (23)$$

where $BD_{iz,t}$ is the bad debt, namely non-performing loans from defaulted firms.

Banks pay taxes on profits and distribute a fraction ρ of their net profits to shareholders.

$$T_{z,t}^\pi = \begin{cases} \tau_{k,t} \pi_{z,t}, & \text{if } \pi_{z,t} > 0 \\ 0, & \text{if } \pi_{z,t} \leq 0 \end{cases} \quad (24)$$

$$Div_{z,t}^\pi = \begin{cases} \rho(\pi_{z,t} - T_{z,t}^\pi), & \text{if } \pi_{z,t} > 0 \\ 0, & \text{if } \pi_{z,t} \leq 0 \end{cases} \quad (25)$$

3.5 Central Bank

The Central Bank sets the discount interest rate according to a Taylor rule (Taylor, 1993; Smets and Wouters, 2007; Gerali et al., 2010):

$$r_t = \bar{r}(1 - \xi) + \xi r_{t-1} + (1 - \xi) \xi^{\hat{P}} (\hat{P}_{t-1} - \hat{P}^*) \quad (26)$$

where \bar{r} is the parameter that represents the long-run interest rate, ξ is the adjustment rate, $\xi^{\hat{P}}$ is the sensitivity to inflation, \hat{P}_{t-1} is the past period inflation, and \hat{P}^* is the inflation target. For the sake of simplicity, following Caiani et al. (2018, 2019a), this Taylor rule only considers price variation (and not output gap).

The Central Bank holds the commercial banks' reserves ($R_{CB,t}$), provides advances to banks ($L_{CB,t}$), and buys bonds issued by the government ($B_{CB,t}$) which were not bought by banks. Central Bank' profit ($\pi_{CB,t} = r_{b,t} B_{CB,t} + r_t L_{CB,t} - r_{re} R_{CB,t}$) is passed to the government.

3.6 Government

Government collects taxes and provides money transfers to households. Households (h) pay income taxes, while firms (i) and banks (z) pay taxes on profits. Total taxes T_t are:

$$T_t = \sum_{h, y_{h,t} > 0}^H \tau_t y_{h,t} + \sum_{i, \pi^* > 0}^I \tau_t \pi_{i,t-1} + \sum_{z, \pi > 0}^Z \tau_t \pi_{z,t-1} \quad (27)$$

Public spending G_t takes the form of a monetary transfer that is equally distributed to households (G_t/H). The government intervention through fiscal policy then also acts as a consumption smoothing policy. Indeed, monetary transfers allow unemployed workers, and in general the worse off, to sustain consumption. Moreover, government pays interest on public debt. Public deficit is financed through the emission of bonds that last for one period only, while in case of surplus it is used in the following period to reduce the deficit. Government sets public spending (G_t) and the tax rate (τ_t) according to a given target deficit d^{max} .⁵

$$\text{if } d_{t-1} \geq d^{max} : \begin{cases} G_t = G_{t-1}(1 - U[0, \delta]) \\ \tau_{t+1} = \tau_t(1 + U[0, \delta]) \end{cases} \quad (28)$$

$$\text{if } d_{t-1} \leq d^{max} : \begin{cases} G_t = G_{t-1}(1 + U[0, \delta]) \\ \tau_{t+1} = \tau_t(1 - U[0, \delta]) \end{cases} \quad (29)$$

$$(30)$$

To avoid unreasonable values of public debt, both the tax rate and the public expenditure on GDP are bound within a given range, respectively $\{\tau_{min}, \tau_{max}\}$ and $\{g_{min}Y_t, g_{max}Y_t\}$.

The interest rate on bonds depends on the debt-to-GDP ratio:

$$r_{b,t} = \chi B_t / Y_t \quad (31)$$

Bonds are issued in 100 tranches ($b_{k,t} = B_{k,t}/100$) and sold on the bond market where they can be bought by banks. The Central Bank buys the part that is not bought by private banks.

Finally, government guarantees depositors in case of bank default emitting bonds that are directly purchased by the Central Bank.

3.6.1 Firms' and banks' endogenous entry and exit

A part of each household's saving is invested in the creation of new firms and new banks (see section 3.3). Thus, the number of banks and firms is determined endogenously in reason of

⁵Therefore, while fiscal policy tends to smooth households' consumption, due to the role of taxes and transfers, this kind of fiscal rule on the level of public deficit, reminiscent of one of the Eurozone's fiscal rules, can act both as a pro- and counter-cyclical force, depending on the phase of the business cycle.

households' investment. In other words, depending on the amount of resources households want to invest in the creation of new firms and/or banks, the number of agents in the economy changes along time. In order to avoid excessive imbalances in the relative dimension of the banking and firm sectors, the new entrant will be a bank when either the ratio between the number of banks and firms, or the ratio between total net-worth of banks and firms, are below a given percentage η .⁶ On the contrary, the new entrant will be a firm.

The equity level of the new agent is sampled in a range between the net worth of the smallest and the net worth of the larger incumbents. If funds invested by households are lower than the required net worth, no firm (bank) is created.

New firms' initial price ($p_{i,t}$), productivity (ϕ), offered wage ($w_{i,t}$) and sale expectations ($q_{i,t}^e$) are sampled within a range going from the lowest to the highest values of existing firms. Moreover, sales expectations cannot be lower than the minimum production level that can be fund with inner resources ($\frac{A_{i,t}}{w_{i,t}}\phi_{i,t}$).

Firms whose net worth is below the wage they offer to workers $F_t = w_{i,t}$ go bankrupt; in effect, firms that are not able to employ at least one worker exit the market. At the same time, banks having a net worth below the average wage default. In this way, we remove both firms having a negative net-worth and very small firms whose contribution to the dynamics of the economy is negligible: anyway, this mechanism does not impact on the concentration of sales. A firm default may generate a non-performing loan for lending banks, thus it may impact on banks' profit.

4 Expectations on sales variation

Agents interact in a complex environment. In the goods market, each firm faces the competition of the others with prices and sales changing in each period. Moreover, each incumbent firm competes with new entrants. Besides, sales are determined by aggregate demand fluctuations and production is influenced by both productivity and the variation of wages.

Forecasting is a difficult task in a complex system in which heterogeneous agents operate and interact under strong uncertainty. For this reason, we assumed that firms in our model try to forecast only the direction of change in trend growth in order to have a qualitative indication for production planning. This is in line with a large part of the literature on macro ABM in which

⁶We tested the effects of changing this share, both reducing and increasing η (from 0.1 to 0.05 and 0.15, respectively), and this does not hamper the predictive capabilities of the GA specification; moreover, the macro effects are still in line with the presented results. Simulation results regarding this sensitivity analysis are available upon request.

typically agents follow simple rules of behavior consisting in looking at past values of a variable, e.g. past sales, and then deciding the *direction* of change, e.g. increase production, whereas the *size* of the change is randomly determined by an idiosyncratic shock. Some works, as for instance Leitch and Tunner (1991), suggest that forecasting the direction of change may produce more profits compared to standard forecasting techniques. However, the performance of directional predictions depend on a number of factors and may vary across markets. There are applications of directional forecasting in various contexts: GDP (Pons, 2000), interest rates (Greer, 2003), exchange rates (Qi and Wu, 2003), and oil prices (Knetsch, 2007). Moreover, this kind of reasoning can be found in the literature on forecasting stock market movements: for instance, according to Qiu and Song (2016), in the business sector, predicting the exact daily price of the stock market index has always been considered a difficult task; for this reason, there is a great deal of research regarding the prediction of the direction of stock price index movement. Leung et al. (2000) point out that trading could be made profitable by an accurate prediction of the direction of the stock index movement; they suggest that financial forecasters and traders should focus on predicting the direction of movement so as to minimize the deviations of estimates from actual observed values. Complex systems like the stock market, as well as the macroeconomy, are dynamic and exhibit wide variations; therefore, the prediction of such complex systems becomes a highly challenging task because of the highly non-linear nature of emergent patterns deriving from the dispersed interaction of a multitude of heterogeneous agents. In our view, directional forecasting is an approach that can be experimented also to predict firms' sales. Indeed, in a complex environment like our simulated economy, making sales forecast is a quite difficult task. This is the main reason why we choose to focus only on forecasting the sign of the sales variation in each period. This kind of reasoning can be also found in the management literature that considers ABM as an effective tool to forecast variables. For instance, in an algorithm describing the price setting of a firm proposed in North and Macal (2007), a binary choice is faced by the firm when the price is revised up or down based on the sales volume being high or low (though more complicated schemes can be considered depending on the complexity of firms' choices). In a sense, the learning algorithm can be thought as a mechanism resembling the reasoning of managers that, based on the information perceived (regarding the economic environment in which the firm competes), provide a *qualitative forecasting* on what should be the next period behavior of the firm. As we will see, the choice of the forecasting technique has a relevant impact on firms' performance. Indeed, firms tend to make higher profits when they are able to correctly forecast the sign of sales' growth.

Accordingly, agents adapt gradually their behavior, thus a point prediction of sales growth would not be so useful in determining their choices. Instead, knowing if sales will increase or decrease is an information that may easily be incorporated in production and sales rules. In the baseline model, when sales are higher than the desired production, firms increase prices and production by a random amount. While, when sales are lower than the desired ones, firms reduce them. In this section, we define a new sales rule that improves firm profitability incorporating sales expectations.

With sales expectation $\Delta\hat{q}_{i,t+1}^E$, the decision rule regarding the setting of production and prices becomes:

$$\text{if } \hat{q}_{i,t-1} \geq q_{i,t-1}^S : \begin{cases} q_{i,t}^S = q_{i,t-1}^S(1 + U[0, \delta]) \\ p_{i,t} = p_{i,t-1}(1 + U[0, \delta]) \end{cases} \quad (32)$$

$$\text{if } \hat{q}_{i,t-1} < q_{i,t-1}^S : \begin{cases} \text{if } \Delta\hat{q}_{i,t+1}^E \geq 0 : \begin{cases} q_{i,t}^S = q_{i,t-1}^S \\ p_{i,t} = p_{i,t-1} \end{cases} \\ \text{if } \Delta\hat{q}_{i,t+1}^E < 0 : \begin{cases} q_{i,t}^S = q_{i,t-1}^S(1 - U[0, 2\delta]) \\ p_{i,t} = p_{i,t-1}(1 - U[0, 2\delta]) \end{cases} \end{cases} \quad (33)$$

As in the baseline scenario, if the previous period sales were larger than desired production, then firms increase the desired sales and prices. Indeed, when the previous period sales were larger than desired production, even if growth prediction is negative, it is still convenient for firms to increase production in order to reestablish their inventories. While, when past sales are lower than the desired production but the expected variation of sales is positive, firms do not decrease production and prices. On the contrary, when past sales are lower than the desired production and sales expectations are negative, firms reduce prices and production faster than in the baseline rule (by $U[0, 2\delta]$).

Associated with this new production rule, we implemented three forecasting methods: a naïve choice (N), a genetic algorithm (GA) and an autoregressive model (AR). The naïve choice (N) means that firm forecasts are equal to their past period result: if in the past period sales increased (decreased) firms predict a positive (negative) growth in the next period. The genetic algorithm (GA) uses an informative set of three variables: firm's past sales, the aggregate inflation growth rate and the aggregate unemployment growth rate. We use this set because after

several attempts is the smallest set of variables that allow both the *AR* and the *GA* to formulate unbiased expectations.⁷

The *GA* is nothing less than a map between these three input variables and sales prediction: it tries to individuate what will be the sign of sales given the sign of the input variables. We use the same variables to implement a vectorial autoregressive model (*AR*), which provides a punctual sales growth one-step-ahead predictions. The maximum length of the dataset used by the *AR* is of 100 periods.⁸

Both the output of the genetic algorithm (*GA*) and of the autoregressive model (*AR*) are translated in a simpler prediction of just the sign of the growth rate of sales. Moreover, in order to increase the precision of these two methods, which rely on the volume of data that they may process, new entering firms inherit data processed by a randomly chosen incumbent firm.⁹ In the next section, we provide simulation results on the model and in particular on the performance of different forecasting rules.

5 Simulations

5.1 Simulation setup

Simulation scheduling follows Caiani et al. (2018, 2019a). Even the parameters of the model are similar to the one in Caiani et al. (2018, 2019a) with some refinement due to the modifications necessary to adapt this model to a setup with only one country and that focuses on firm expectations (par. 3.1; see Table 1 in the Appendix). Moreover, following Caiani et al. (2018, 2019a) the model has been calibrated to provide realistic values of the macro variables in line with the average of Eurozone countries (see Table 2 in the Appendix). The only variable that is larger than the Eurozone average one is the public debt, which is compensated by a lower level of the private debt. This is due to the fact that following a procedure inspired by Godley and Lavoie (2007), the model starts with zero public and private resources and gradually money is created, through public expenditure, letting the economy grow.

⁷At the beginning of our research we allowed firms to use a larger set of variable to make forecast in addition to unemployment, past sales and inflation: the ratio between average and firm price, the number of firms, the average output level, aggregate income variation. However, we found that using only firms sales, unemployment and inflation was enough to allow firms to make correct prediction; therefore, we restricted the set of variables to just these three.

⁸A more detailed explanation of how algorithms are built and work can be found in Appendix B and C.

⁹Indeed, in order to work effectively, a learning mechanism needs a trial period, otherwise it is not able to make quite accurate predictions. For this reason, we assume that new entrants inherit the properties of the learning mechanism from a randomly chosen surviving firm. In a sense, this means that a firm, before entering the market, makes a kind of marketing study to understand some (local) properties of the economy, which in the specific case of the model corresponds to the features of a surviving firm picked at random, in order to have an initial direction to follow in a highly uncertain environment.

In line with the well-established macro AB validation approach (Dosi et al., 2010), we checked the consistency between the properties of the simulated times series and a set of key stylized facts;¹⁰ in particular:¹¹

- the volatility of consumption, investment, unemployment with respect to real GDP;
- the cyclical dynamics of consumption, public spending and public spending over GDP, and unemployment (Uribe and Schmitt-Grohé, 2017);
- the distribution of firm and bank size (Stanley et al., 1995).

5.2 Performances of forecasting rules

We compute the prediction mean squared error (MSE_t) and the average error (ER_t). Firms have only to try to predict if in their next period sales will grow or decline, thus we define a synthetic sales growth indicator x_{it} that is equal to 1 if sales are positive and, on the contrary, is equal to -1 if sales are negative. The mean squared error is given by the average square difference between the predicted and the effective value of the synthetic growth indicator of each firm i of the N firms in the market at time t :

$$MSE_t = \frac{\sum_i^N (x_{it} - x_{it}^E)^2}{N} \quad (34)$$

The average error (ER_t) is the sum of the prediction errors of firms:

$$ER_t = \frac{\sum_i^N (x_{it} - x_{it}^E)}{N} \quad (35)$$

We calibrated both the genetic algorithm (GA) and the autoregressive method (AR) in order to let their forecasting error converge to zero as shown in the left panel of figure 2. The negative bias of the naïve method (N) is due to the fact that the real output of the economy grows, thus on average also firm sales tend to increase through time.

The genetic algorithm (GA) outperforms all the other predictive methods (Right panel of figure 2). Indeed, even if the autoregressive method (AR) is unbiased it needs a more stable

¹⁰While our paper reproduces a reduced set of stylized facts, other works as Delli Gatti et al. (2011), Dawid et al. (2012), and Dosi et al. (2017) are able to replicate a larger list. Nevertheless, our model is a very simplified one representing a pure-labor closed economy. The motivation for such a simplified setting is that we aim at testing the effect of different learning capabilities, which can be a quite difficult task in more complex environment. In other words, we just aim at reproducing a plausible scenario in order to get relevant theoretical insights about learning dynamics. Therefore, we consider it as a first step in investigating how different degrees of sophistication in learning abilities can affect agents' behavior as well as the behavior of the system as a whole.

¹¹It is worth stressing that, as it is also the case for other ABMs, the model reproduces the stylized facts in a qualitative sense. In other words, we do not provide the results of statistical tests which could be applied in the case of a quantitative estimation.




Figure 2: Average error and Mean squared error. *GA* in red, *AR* in blue, *N* in orange. Solid lines represent averages over 100 simulation, dotted line are 95% confidence interval.

environment and a large set of data to be effective. Instead, the genetic algorithm (*GA*) is based on a simpler informative base: it just associates the sign of the variation of the input variables with the sign of the expected sales variation. Thus, the *GA* generates a map of associations between variables that results to be more effective in predicting the sign of sales variations in a complex environment.

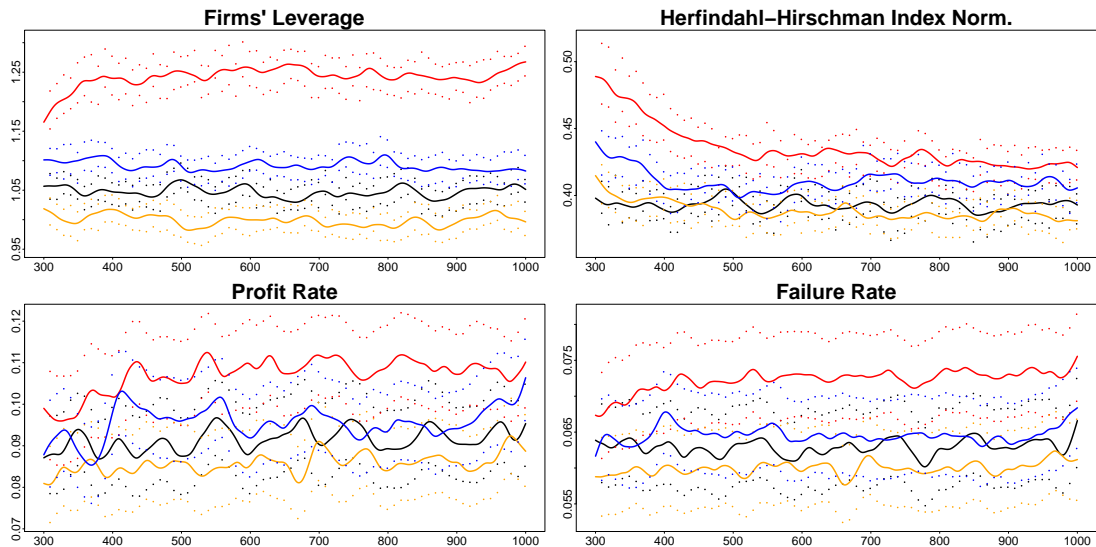


Figure 3: Firm Leverage, HI index, Profit and Failure Probability. *GA* in red, *AR* in blue, *N* in orange. The baseline scenario *B* in black. Solid lines represent averages over 100 simulation, dotted line are 95% confidence interval.

With respect to the baseline scenario, namely the one without expectations (B), using the predictive methods firms tend to increase their production, because they understand that the economy tends to grow. This lead to higher leverage with respect to the baseline scenario and generates also higher concentration in the market (fig. 3). In turn, higher concentration allows firms to increase their mark-up and, thus, to increase their profits. However, at the same time, higher leverage slightly increases failure probability, contributing to the reduction of the firm number and so to the higher market concentration.

Moreover, it is possible to notice that profitability and market concentration are higher when firms adopt the *GA* and the *AR* methods, which are more efficient than the naïve one (*N*). Indeed, the naïve method (*N*) is negatively biased, thus it is able to catch only partially the extent of the positive growth trend of the economy.

On aggregate level (fig. 4), the higher firm profitability of *GA* and *AR* reduces the wage share, that, in turn, depresses the aggregate demand. Indeed, only a part of the profit is reinvested or distributed as dividends, while the larger part of the worker income is translated in higher consumption. A weaker aggregate demand, thus, leads to a reduced dynamic of real wages. Moreover, firms' innovative effort is proportional to past sales, thus a slower dynamic of the demand reduces their innovativeness. Consequently, even if the demand is weak, the level of employment is sustained by lower real wages and lower productivity. In the long run, the economic system modeled is wage-led (Caiani et al., 2018, 2019a): higher profits come at the price of lower growth.

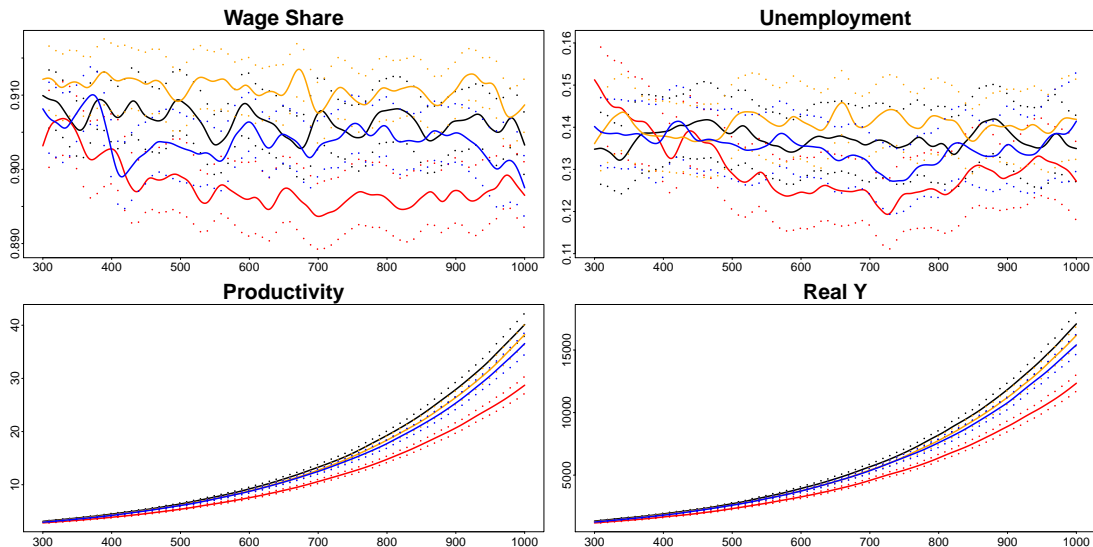


Figure 4: Macro variables. *GA* in red, *AR* in blue, *N* in orange. The baseline scenario *B* in black. Solid lines represent averages over 100 simulations, dotted line are 95% confidence interval.

6 Macprudential and Fiscal shocks

In this section we test the effect of shocks on the economy and, thus, on the predictive ability of firms.

We introduce only permanent shocks to better observe the reactivity of the predictive methods through time: a macroprudential shock in the form of a tighter capital requirement (μ_1) and a fiscal shock in the form of a cut of the public target deficit (d^{max}). Thus, we shock *permanently* the economy at half of the simulation (since period 500) and test the impact of shocks of different size.¹²

We implemented permanent contractions of the capital requirement and of the target deficit. Fig 5 shows that these shocks have a relevant contractionary impact on output. In particular, the effect of fiscal shocks are larger than the credit ones. For instance, a 50% reduction in the capital requirement (macroprudential shock) lowers the output by almost the 20% in the long run. While halving the target deficit (fiscal shock) reduces output by about 40% in the long run. In other words, the lower growth rate resulting from the fiscal shock leads to a difference of the 40% with respect to the normal scenario in the very long run (namely after 500 periods of simulation). Moreover, we have to take into account that we are considering a closed economy. Therefore, the fiscal shock reduces the demand and, thus, unemployment increases determining a reduction of wages. However, lower wages cannot increase country competitiveness and, thus, exports. Indeed, foreign demand cannot sustain the economy and the effect of the fiscal shock is larger than in an open economy.

In all the scenarios, therefore both when agents implement predictive methods and in the baseline setting, the shocks have qualitatively the same impact on the economy. Moreover, it is worth noting that shocks affect more the economy when agents adopt GA and AR as predictive methods. Indeed, when agents have more sophisticated expectations their leverage tends to be higher, thus the economic system is more vulnerable to exogenous shocks. Moreover, firms using GA and AR are more sensible to capital requirement variations because of their higher level of indebtedness. Therefore, the entity of the shocks is consistent and represents a good testing scenario for the different predictive methods we have implemented; indeed, they show the ability of these methods to re-adapt after huge perturbations of the economic environment that may

¹²In order to get rid of transient dynamics, there is a transition time, or training period, in the model, as it is typically done in agent based models. During this transition, the artificial economy evolves toward a kind of statistical equilibrium, after which we start analyzing simulation results. In order to be sure that we are not considering simulation data still hardly dependent on initial conditions, we set the threshold after which simulation data are analyzed at $t = 500$. Indeed, after this quite long transition firms have learned enough about the characteristics of the economy and then they should be able to make the predictions whose accuracy depend on the used forecasting rule.

have permanent effects on macro and micro dynamics.

Both the credit and the fiscal shock affect the predictive capacity of the agents (Fig. 6 and Fig. 7). We focus on the *GA* and *AR* methods that provide correct expectations. Both the macroprudential and fiscal shocks we take into exam have a restrictive impact on the economy. The negative impact of the fiscal shock tends to be larger but the macroprudential shocks have faster effects. In fact, the macroprudential shock that we analyze is a change in the capital requirement of banks that affect immediately the credit supply, while the fiscal shock consists in a variation of the public target deficit and current effective deficit only gradually realigns to the target one.

Considering the *GA* predictive method, until the shocks are not too strong, agents are still able to formulate unbiased expectations. Indeed, the *GA* allows agents to make correct expectations even after a fiscal shock of 75% and a macroprudential shock of 25%. When shocks are larger, expectations suddenly becomes biased and the *GA* needs time before starting to formulate unbiased expectations again. At the same time, at the beginning agents lose predictive capacity, the MSE increases sharply and later comes back to previous levels.

The *AR* method is less effective that the *GA* in dealing with fiscal shocks. When the fiscal shocks are too strong the *AR* method is not able to converge back to unbiased expectations. The worst performance of the *AR* is due to the fact that, as we have seen, *AR* expectations are less accurate than the *GA*. Moreover, the effect of fiscal shocks are stronger and they tend also to increase output volatility, while the credit shock reducing the leverage of banks impact negatively on output but increases the stability of the economic system. Indeed, the *AR* becomes biased with huge fiscal shocks. Before the shock, firms have developed positive expectations about the growth of the economy; the fiscal shock drastically reduces the growth rate of the economy and increases output volatility; therefore, *AR* firms face huge difficulties in reformulating their expectations. Instead, *GA* firms seem to be more effective in reacting to transformations of the economic environment.

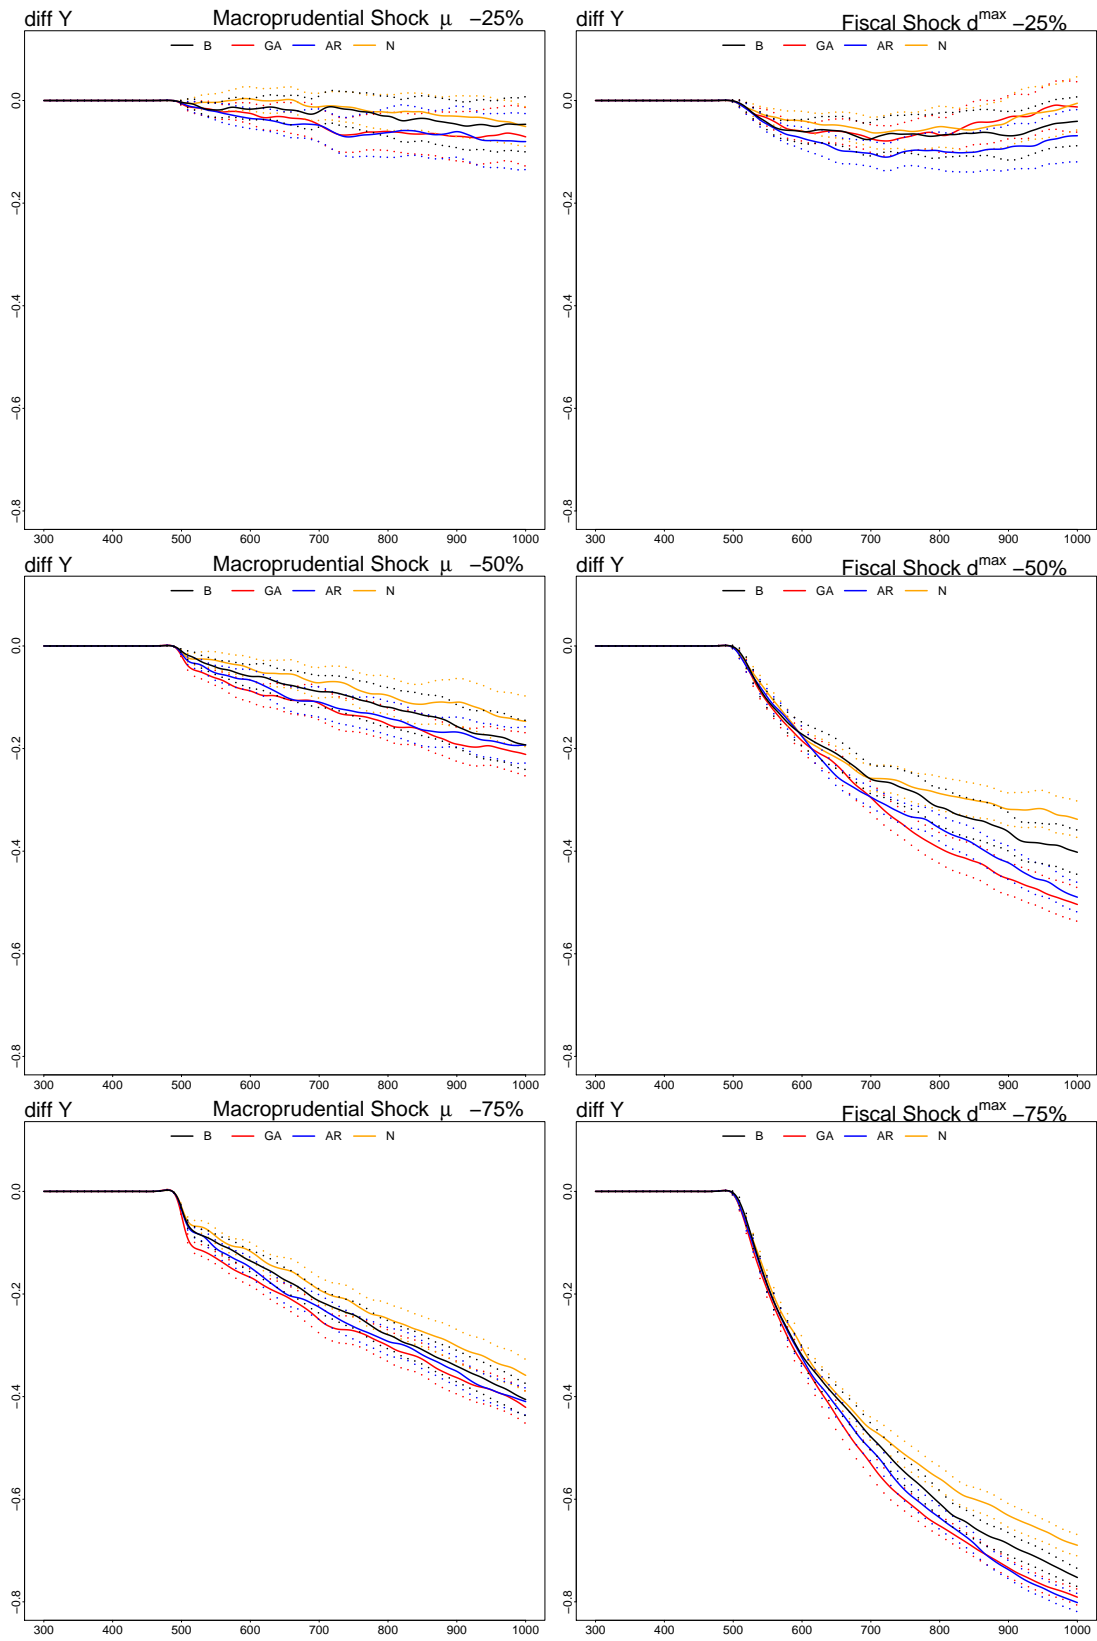


Figure 5: Average output percentage variation after credit and macroprudential shocks. *GA* in red, *AR* in blue, *N* in orange, the baseline scenario *B* in black. Solid lines represent averages over 50 simulation, dotted line are 95% confidence interval.




Figure 6: Average Error and Mean Squared Error with capital requirement variation (μ). *GA* in the baseline in red, *GA* after macroprudential shock in purple. *AR* in the baseline in blue, *AR* after macroprudential shock in brown. Solid lines represent averages over 50 simulation, dotted line are 95% confidence interval.

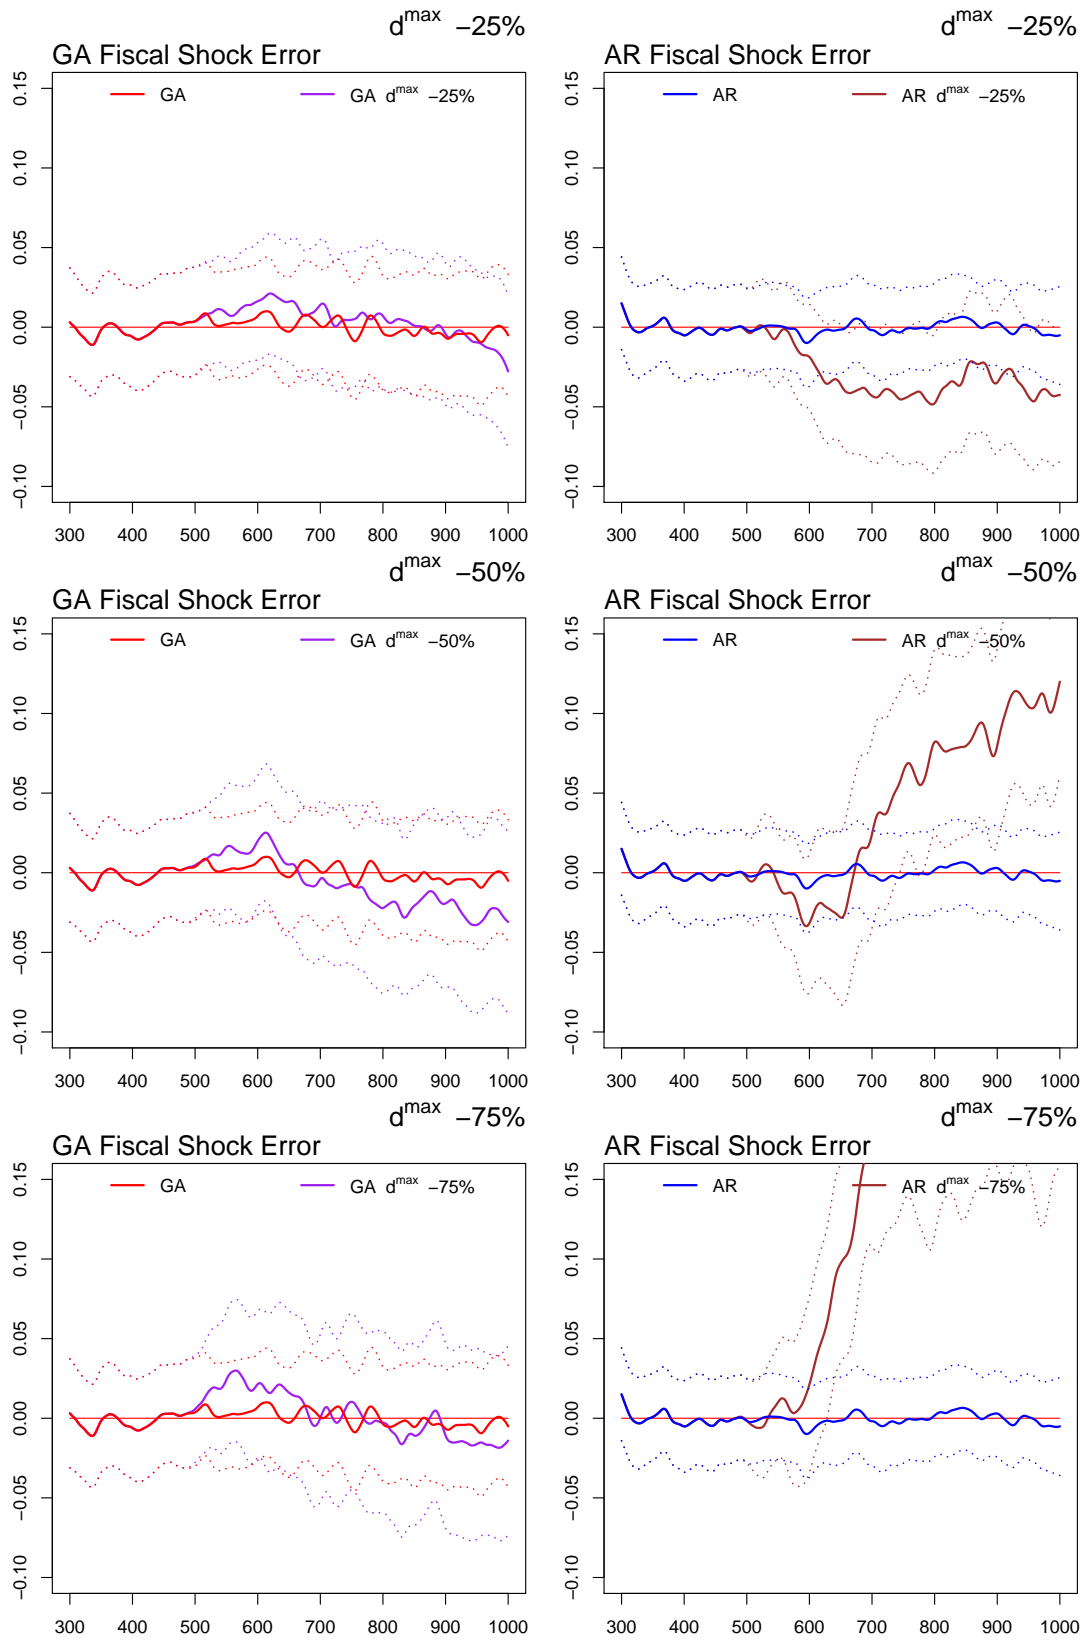


Figure 7: Average Error and Mean Squared Error with target public deficit variation (d^{\max}). *GA* in the baseline in red, *GA* after fiscal shock in purple. *AR* in the baseline in blue, *AR* after fiscal shock in brown. Solid lines represent averages over 50 simulation, dotted line are 95% confidence interval.

7 Market competitiveness and predictive capabilities

Market competitiveness is crucial in determining firms' performance and thus their predictive capabilities. In this section, we focus on the *GA*, that is the most efficient predictive method, to see the impact of market competitiveness in the formulation of expectation. In the model, a consumer's preference depends on both the price and the distance between their most preferred variety and the variety sold by firms. We modify the level of competitiveness, varying the parameter β , which represents the consumers' sensitivity to price. When we increase β the importance of prices decreased and thus the market becomes less competitive; the opposite happens when β decreases. In the baseline setting β is equal to 0.3, thus we firstly increase β to 0.4 and then we decrease it to 0.2. Changing the level of competitiveness gives the same results of the baseline scenario (section 4): *GA* firms make higher profits; this results in a lower growth due to a reduced demand (Figures 8 and 9).




Figure 8: Macro variables. *GA* in red, *AR* in blue, *N* in orange. The baseline scenario *B* in black. Solid lines represent averages over 50 simulations, dotted line are 95% confidence intervals. $\beta = 0.2$.

Fig. 10 shows the maximum distance of the error from zero after fiscal and macroprudential shocks of different entity. Changing the level of competitiveness does not impact on the predictive performance under fiscal shocks: expectation remain unbiased even after massive fiscal contractions. While, considering macroprudential shocks, when competitiveness raises ($\beta = 0.2$) the effectiveness of the predictive method slightly declines: prediction are no longer unbiased after a 25% shock. Indeed, in a more competitive environment even a small variation of prices and quantity sold by other firms may have a larger impact on sales of the other agents, thus

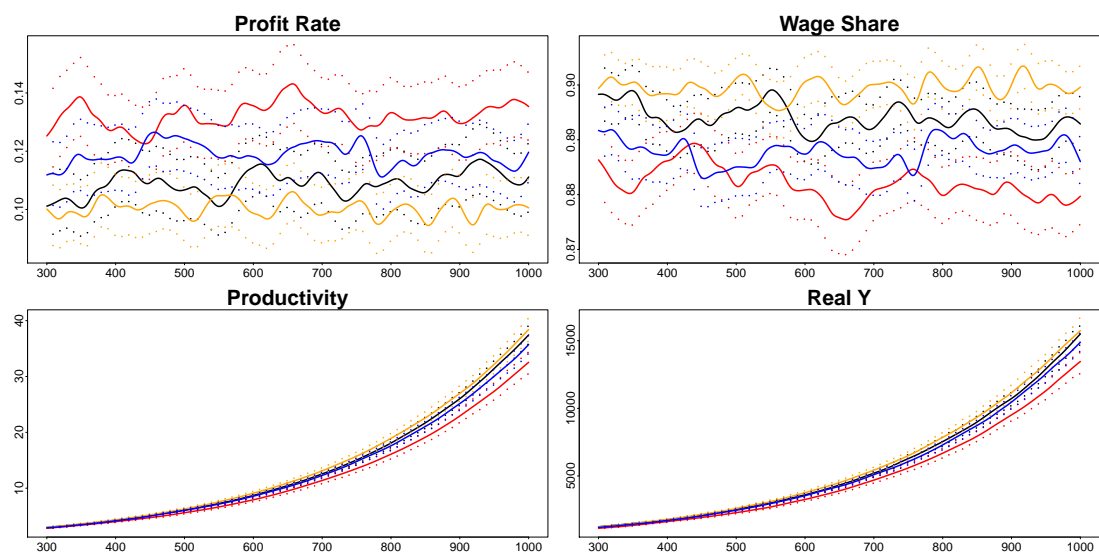


Figure 9: Macro variables. *GA* in red, *AR* in blue, *N* in orange. The baseline scenario *B* in black. Solid lines represent averages over 50 simulations, dotted lines are 95% confidence intervals. $\beta = 0.4$.

formulating correct expectation becomes harder.

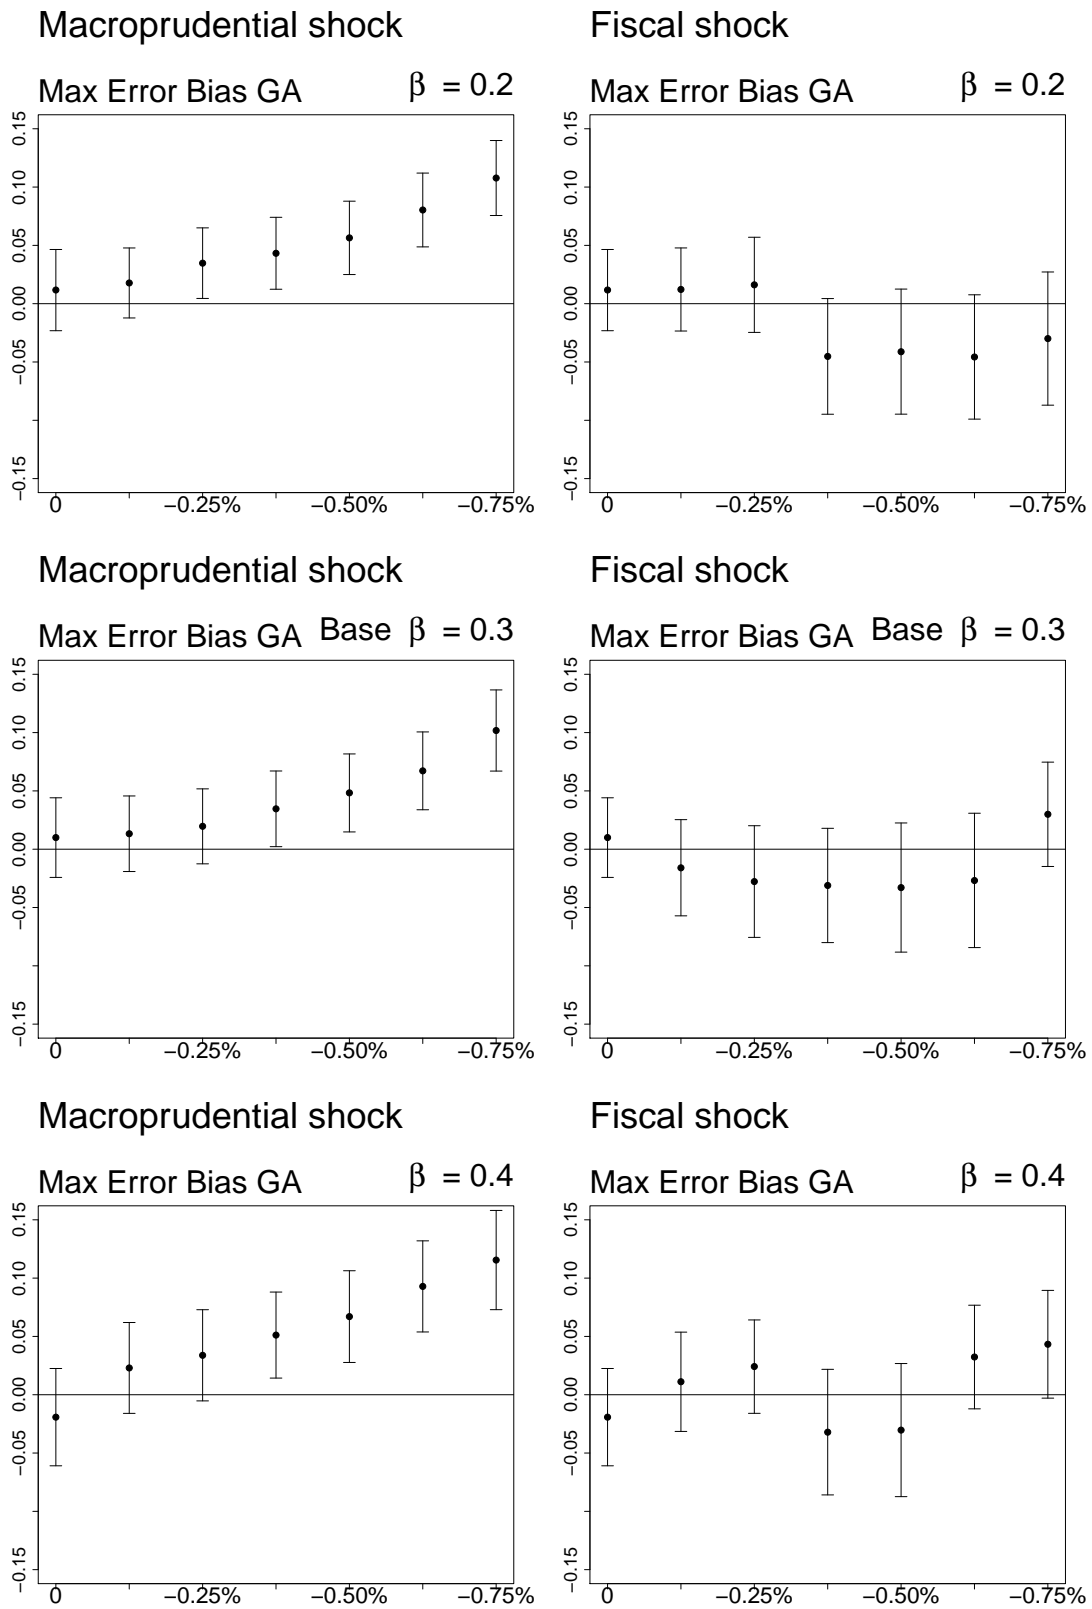


Figure 10: Error bias after fiscal and macroprudential shock with different level of market competition ($\beta = 0.2$, the baseline specification $\beta = 0.3$, $\beta = 0.4$). Bar line are 95% confidence intervals. Averages over 50 simulations. On the x axis shock entity and on the y axis the maximum error bias.

8 Conclusions

Adopting computational techniques, even in a complex economic system, it is possible to endow agents with expectations that are not biased and show a certain degree of accuracy. In particular, firms that implement a generic algorithm (*GA*) are able to formulate sales variation predictions with errors that converge to zero and with a relatively low mean squared error. In other words, firms are able to effectively adapt to a complex environment based on evolutionary principles as selection and experimentation.

When firms employ more sophisticated behavioral rules, rather than naïve expectations, they obtain higher profits. In more detail, better forecasting the increasing trend of a growing economy makes firms to expand production, resulting in a rising financial leverage, and in general a larger scale of activity. Larger size causes a reduction of the number of firms, thus competition among firms becomes weaker and firms see their profits increase with respect to wages. But a lower wage share tends to reduce the aggregate demand which, in turn, slows down economy growth (Caiani et al., 2018, 2019a). Therefore, while more sophisticated prediction methods allow firms to improve forecasting abilities and increasing their profits, the aggregate result is a decline of the long-run growth rate of the economy. Micro and macro implications are not going in the same direction as it is not surprising in a complex economic system. Moreover, short run gains in firms' profit are in opposition with a reduction of the growth potential of the economy in the long-run.

Fiscal and macroprudential shocks have a significant impact on economic dynamics: the effect of these shocks are similar in the scenario when agents have predictive capabilities and in the baseline one where firms do not formulate prediction and follow simple adaptive rules. Nevertheless, macroprudential shocks affect more the economy when agents are able make predictions. In this last case, agents tend to have higher leverage, so they are more affected by restrictive macroprudential measures. When shocks are not too strong the predictions formulated by the *GA* algorithm still remain unbiased, even if they suffer from a reduction of their accuracy. For shocks that are not massive, agents' predictions and, thus, their implied behaviors, remain consistent with the changed economic environment. Therefore, computational predictive methods may endow firms with models that, with a certain approximation, may be consistent with the Lucas' critique, in the sense that agents are able to adapt to a changing environment, even after unexpected shocks of non-negligible size. In other words, we show that our model, when agents are endowed with appropriate forecasting rules like a genetic algorithm, is able to provide unbiased expectations, even after a shock, though requiring some time to converge.

Firms in our model are then able to adapt to an evolving system, improving in this way their profits. For this reason, we refer to the Lucas Critique though we are not maintaining that our model is assimilated to more standard models with optimizing agents and rational expectations or that we are converging toward this kind of modeling approach. We are trying to assess the effectiveness of (elaborated) adaptive rules in getting unbiased expectations, thus allowing agents to successfully adapt to an ever changing environment. This result is rooted in the alternative microfoundations provided by an agent-based model.

In general, our findings show that agent-based models could be used as a computational laboratory to test different behavioral rules, from static and adaptive expectations to more sophisticated predictive methods that exploit selection mechanisms and experimentation along evolutionary lines. Considering agents' forecasting ability allows us to study the reaction of boundedly rational (but not naïve) agents to policy changes. This is a first step in this direction. The model can be extended in several directions. Indeed, the financial sector and the labor market are quite simplified in the present version. Basically, we aimed at keeping the macro model with multiple agents as simple as possible, by focusing on the real side of the economy, that is firms' production related to the performance of different forecasting rules. We showed that making accurate predictions, for instance based on a simple machine learning process like a GA, is crucial for firms and may have relevant consequences for the whole economy. Nevertheless, in future works it could be very interesting extending the model to improve the mechanism regarding the financial side of the economy, in order to test the effects of introducing different forecasting abilities for agents operating in the stock market, as well as for banks, in their relationship with firms and households and among them on the interbank market. A promising direction is to test various predictive methods for banks as well as for households, especially for investigating the response to macroprudential policy changes. Finally, experimental studies provide a source of inspiration for testing alternative rules of behavior with different degrees of rationality and forward-looking attitude. Indeed, heterogeneous agents can be characterized by differential abilities to forecast (within the same scenario) and this may have further implications for micro, meso and macro properties. This is something we would investigate in a next step of our analysis. The integration between computational and experimental approaches is a promising way that we want to follow for future research on the effects of heterogeneous degrees of agents' rationality.

References

- Anufriev, M. and Hommes, C. (2012). Evolutionary Selection of Individual Expectations and Aggregate Outcomes in Asset Pricing Experiments. *American Economic Journal: Microeconomics*, 4(4):35–64.
- Arifovic, J., Dawid, H., Deissenberg, C., and Kostyshyna, O. (2010). Learning benevolent leadership in a heterogenous agents economy. *Journal of Economic Dynamics and Control*, 34(9):1768–1790.
- Arifovic, J. and Petersen, L. (2017). Stabilizing expectations at the zero lower bound: Experimental evidence. *Journal of Economic Dynamics and Control*, 82:21–43.
- Bargigli, L., Gallegati, M., Riccetti, L., and Russo, A. (2014). Network analysis and calibration of the “leveraged network-based financial accelerator”. *Journal of Economic Behavior & Organization*, 99(C):109–125.
- Bargigli, L., Riccetti, L., Russo, A., and Gallegati, M. (2018). Network calibration and meta-modeling of a financial accelerator agent based model. *Journal of Economic Interaction and Coordination*, forthcoming.
- Bianchi, C., Cirillo, P., Gallegati, M., and Vagliasindi, P. (2007). Validating and Calibrating Agent-Based Models: A Case Study. *Computational Economics*, 30(3):245–264.
- Caiani, A., Catullo, E., and Gallegati, M. (2018). The effects of fiscal targets in a monetary union: a multi-country agent based-stock flow consistent model. *Industrial and Corporate Change*, Available online, printed version forthcoming:–.
- Caiani, A., Catullo, E., and Gallegati, M. (2019a). The effects of alternative wage regimes in a monetary union: A multi-country agent based-stock flow consistent model. *Journal of Economic Behavior and Organization*, 162:389–416.
- Caiani, A., Godin, A., Caverzasi, E., Gallegati, M., Kinsella, S., and Stiglitz, J. E. (2016). Agent Based-Stock Flow Consistent macroeconomics: Towards a benchmark model. *Journal of Economic Dynamics & Control*, 69:375–408.
- Caiani, A., Russo, A., and Gallegati, M. (2019b). Does inequality hamper innovation and growth? an ab-sfc analysis. *Journal of Evolutionary Economics*, 29:177–228.
- Caiani, A., Russo, A., and Gallegati, M. (2020). Are Higher Wages Good For Business? An Assessment Under Alternative Innovation And Investment Scenarios. *Macroeconomic Dynamics*, 24(1):191–230.

- Caverzasi, E. and Russo, A. (2018). Toward a new microfounded macroeconomics in the wake of the crisis. *Industrial and Corporate Change*, 27(6):999–1014.
- Cincotti, S., Raberto, M., and Teglio, A. (2010). Credit money and macroeconomic instability in the agent-based model and simulator Eurace. *Economics - The Open-Access, Open-Assessment E-Journal*, 4:1–32.
- Colasante, A., Palestrini, A., Russo, A., and Gallegati, M. (2017). Adaptive expectations versus rational expectations: Evidence from the lab. *International Journal of Forecasting*, 33(4):988–1006.
- Dawid, H. and Delli Gatti, D. (2018). Agent-based macroeconomics. In Hommes, C. and LeBaron, B., editors, *Handbook of Computational Economics: Volume 4 Heterogeneous Agents Modeling*, 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 design. Bielefeld Working Papers in Economics and Management 5, University of Bielefeld.
- Dawid, H., Harting, P., and Neugart, M. (2014). Economic convergence: Policy implications from a heterogeneous agent model. *Journal of Economic Dynamics and Control*, 44(C):54–80.
- Dawid, H., Harting, P., van der Hoog, S., and Neugart, M. (2018). A heterogeneous agent macroeconomic model for policy evaluation: Improving transparency and reproducibility. *Journal of Evolutionary Economics*, 29:467–538.
- Delli Gatti, D., Desiderio, S., Gaffeo, E., Cirillo, P., and Gallegati, M. (2011). *Macroeconomics from the Bottom-up*. Springer.
- Delli Gatti, D., Gallegati, M., Greenwald, B. C., Russo, A., and Stiglitz, J. E. (2010). The financial accelerator in an evolving credit network. *Journal of Economic Dynamics & Control*, 34:1627–1650.
- Dosi, G., Fagiolo, G., Napoletano, M., Roventini, A., and Treibich, T. (2015). Fiscal and monetary policies in complex evolving economies. *Journal of Economic Dynamics and Control*, 52(C):166–189.
- Dosi, G., Fagiolo, G., and Roventini, A. (2010). Schumpeter Meeting Keynes: A Policy-Friendly Model of Endogenous Growth and Business Cycles. *Journal of Economic Dynamics and Control*, 34(9):1748–1767.
- Dosi, G., Pereira, M., Roventini, A., and Virgillito, M. (2017). When more flexibility yields more fragility: The microfoundations of Keynesian aggregate unemployment. *Journal of Economic*

- Dynamics and Control*, 81(C):162–186.
- Evans, G. W. and Honkapohja, S. (1996). *Growing Artificial Societies: Social Science from the Bottom Up*. The MIT Press.
- Evans, G. W. and Honkapohja, S. (2012). *Learning and Expectations in Macroeconomics*. Princeton University Press.
- Gerali, A., Neri, S., Sessa, L., and Signoretti, F. (2010). Credit and banking in a dsge model of the euro area. *Journal of Money, Credit and Banking*, 42:107–141.
- Giri, F., Riccetti, L., Russo, A., and Gallegati, M. (2019). Monetary policy and large crises in a financial accelerator agent-based model. *Journal of Economic Behavior & Organization*, 157(C):42–58.
- Gode, D. K. and Sunder, S. (1993). Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, 101(1):119–37.
- Godley, W. and Lavoie, M. (2007). *Monetary Economics An Integrated Approach to Credit, Money, Income, Production and Wealth*. Palgrave MacMillan, New York.
- Grazzini, J., Richiardi, M. G., and Tsionas, M. (2017). Bayesian estimation of agent-based models. *Journal of Economic Dynamics and Control*, 77:26–47.
- Greer, M. (2003). Directional accuracy tests of long-term interest rate forecasts. *International Journal of Forecasting*, 19:291–298.
- Hendry, D. and Mizon, G. (2010). On the mathematical basis of inter-temporal optimization. *Discussion Paper, Department of Economics, University of Oxford*, (497).
- Hommes, C. (2011). The heterogeneous expectations hypothesis: some evidence from the lab. *Journal of Economic Dynamics and Control*, 35:1–24.
- Hommes, C. (2013). *Behavioural Rationality and Heterogenous Expectations in Complex Economic Systems*. Cambridge University Press.
- Hommes, C., Sonnemans, J., Tuinstra, J., and van de Velden, H. (2005). Coordination of Expectations in Asset Pricing Experiments. *Review of Financial Studies*, 18(3):955–980.
- Hotelling, H. (1929). Stability in competition. *Economic Journal*, 39-153:41–57.
- Keynes, J. M. (1937). The general theory of employment. *Quarterly Journal of Economics*, (51):209–223.
- Knetsch, T. (2007). Forecasting the price of crude oil via convenience yield predictions. *Journal*

- of *Forecasting*, 26:527–549.
- Knight, F. (1921). *Risk, Uncertainty, and Profits*. Chicago University Press, Chicago.
- Leitch, G. and Tunner, J. E. (1991). Economic forecast evaluation: profits versus the conventional error measures. *American Economic Review*, 81:580–590.
- Leung, M., Daouk, H., and Chen, A. (2000). Forecasting stock indices: a comparison of classification and level estimation models. *International Journal of Forecasting*, 16(2):173–190.
- Lucas, R. J. (1976). Econometric policy evaluation: A critique. *Carnegie-Rochester Conference Series on Public Policy*, 1(1):19–46.
- Marimon, R., Spear, S. E., and Sunder, S. (1993). Expectationally driven market volatility: an experimental study. *Journal of Economic Theory*, 61(1):74–103.
- Meyers, S. (1984). Capital Structure Puzzle. *Journal of Finance*, 39-3:575–592.
- Muth, J. F. (1961). Rational expectations and the theory of price movements. *Econometrica*, (29):315–335.
- North, M. J. and Macal, C. M. (2007). *Managing Business Complexity*. Oxford University Press.
- Onaran, O. and Galanis, G. (2012). Is aggregate demand wage-led or profit-led? National and global effects. Conditions of Work and Employment Series 40, International Labour Office, Geneva.
- Onaran, O. and Obst, T. (2016). Wage-led growth in the EU15 member-states: the effects of income distribution on growth, investment, trade balance and inflation. *Cambridge Journal of Economics*, 40(6):1517–1551.
- Pons, J. (2000). The accuracy of IMF and OECD forecasts for G7 countries. *Journal of Forecasting*, 19(1):53–63.
- Qi, M. and Wu, Y. (2003). Non-linear prediction of exchange rates with monetary fundamentals. *International Journal of Forecasting*, 10:623–640.
- Qiu, M. and Song, Y. (2016). Predicting the direction of stock market index movement using an optimized artificial neural network model. *PlosONE*, 11(5):e0155133.
- Riccetti, L., Russo, A., and Gallegati, M. (2015). An agent-based decentralized matching macroeconomic model. *Journal of Economic Interaction and Coordination*, 10-2:305–332.
- Riccetti, L., Russo, A., and Gallegati, M. (2016). Financialisation and crisis in an agent based macroeconomic model. *Economic Modelling*, 52(PA):162–172.
- Riccetti, L., Russo, A., and Gallegati, M. (2018). Financial Regulation And Endogenous Macroe-

- conomic Crises. *Macroeconomic Dynamics*, 22(4):896–930.
- Russo, A., Catalano, M., Gaffeo, E., Gallegati, M., and Napoletano, M. (2007). Industrial dynamics, fiscal policy and R&D: Evidence from a computational experiment. *Journal of Economic Behavior & Organization*, 64(3-4):426–447.
- Russo, A., Riccetti, L., and Gallegati, M. (2016). Increasing inequality, consumer credit and financial fragility in an agent based macroeconomic model. *Journal of Evolutionary Economics*, 26(1):25–47.
- Salle, I. L. (2015). Modeling expectations in agent-based models — An application to central bank’s communication and monetary policy. *Economic Modelling*, 46(C):130–141.
- Salop, S. C. (1979). Monopolistic competition with outside goods. *The Bell Journal of Economics*, 10-1:141–156.
- Simon, H. (1959). Theories of decision-making in economics and behavioral science. *American Economic Review*, (49):253–283.
- Simon, H. (1976). From Substantive to Procedural Rationality. In Latsis, S., editor, *Method and Appraisal in Economics*, pages 129–148. Cambridge University Press, Cambridge UK.
- Sims, C. (1980). Macroeconomics and reality. *Econometrica*, 48:1–48.
- Sinitetskaya, E. and Tesfatsion, L. (2015). Macroeconomies as constructively rational games. *Journal of Economic Dynamics and Control*, 61(C):152–182.
- Smets, F. and Wouters, R. (2003). An Estimated Dynamic Stochastic General Equilibrium Model of the Euro Area. *Journal of the European Economic Association*, 1(5):1123–1175.
- Smets, F. and Wouters, R. (2007). Shocks and frictions in us business cycles: A bayesian dsge approach. *American Economic Review*, 97:586–606.
- Stanley, M., Buldyrev, S., Havhn, S., Mantegna, R., Salinger, M., and Stanley, E. (1995). Zipf plots and the size distribution of firms. *Economic Letters*, 49:453–457.
- Stiglitz, J. E. (2016). Rethinking macroeconomics: What failed, and how to repair it. *Journal of the European Economic Association*, (9):591–645.
- Stockhammer, E., Onaran, O., and Ederer, S. (2009). Functional income distribution and aggregate demand in the Euro area. *Cambridge Journal of Economics*, 33(1):139–159.
- Stockhammer, E. and Sotiropoulos, D. P. (2014). Rebalancing the euro area: The costs of internal devaluation. *Review of Political Economy*, 26(2):210–233.
- Taylor, J. B. (1993). Discretion versus policy rules in practice. *Carnegie-Rochester Conference*

Series on Public Policy, 39-1:195–214.

Teglio, A., Mazzocchetti, A., Ponta, L., Raberto, M., and Cincotti, S. (2019). Budgetary rigour with stimulus in lean times: Policy advices from an agent-based model. *Journal of Economic Behavior and Organization*, 157:59–83.

Turrell, A. (2015). Agent-based models: understanding the economy from the bottom up. *Quarterly Bulletin, Bank of England*, Q4.

Tversky, A. and Kahneman, D. (1974). Judgment under Uncertainty: Heuristics and Biases. *Science*, 185(4157):1124–1131.

Uribe, M. and Schmitt-Grohé, S. (2017). *Open Economy Macroeconomics*. Princeton NJ: Princeton University Press.

Wilson, W. S. (1995). Classifier Fitness Based on Accuracy. *Evolutionary Computation*, 2(3):149–175.

Wilson, W. S. (2002). Classifiers that approximate functions. *Journal of Natural Computations*, 1(2-3):211–234.

Woodford, M. (2013). Macroeconomic analysis without the rational expectations hypothesis. *Annual Review of Economics*, (5):303–346.

A Appendix: Baseline Setup

Table 1: Parameters

H : Number of Households	500	ι_l : Loan probability parameter	0.4
l^S : Workers' labor supply	1.0	χ : Loan interest parameter	0.003
ψ : Matching parameter	10	ι_b : Bond probability parameter	0.1
v : Wage revision probability parameter	1.625	r_{re} : Interest paid on banks' reserves	0.0
v_H : Wage revision probability households	0.7	r_{b0} : Initial interest on bonds	0.001
v_F : Wage revision probability firms	1.0	w_0 : Initial wage	1.0
ϕ_0 : Initial productivity	1.0	\bar{r} : Taylor rule long run interest rate	0.0075
τ_0 : Initial tax rate	0.4	ξ : Taylor rule adjustment speed parameter	0.8
c_y : Propensity to consume out of income	0.9	$\xi^{\Delta P}$: Taylor rule sensitivity to inflation	2
c_D : Propensity to consume out of wealth	0.1	$\bar{\Delta P}$: Inflation Target (quarterly)	0.005
δ : Adaptive Parameter	0.04	d^{max} : Maximum deficit-GDP ratio	0.03
μ_1 : Bank capital requirement	30	τ_{min} : Minimum tax rate	0.46
β : Hotelling circle parameter	0.3	τ_{max} : Maximum tax rate	0.56
λ : Liquidity preference parameter	0.3	g_{min} : Minimum G/GDP	0.4
θ : Share of sales as inventories	0.2	g_{max} : Maximum G/GDP	0.5
γ : R&D expenditure parameter	0.04	η : Banks-firms minimum proportion	0.1
ν : R&D success probability parameter	0.8	ϖ : Minimum investment threshold parameter	0.1
ρ : Share of profits distributed	0.95	A^0 : First firms' initial net worth	10.0
ζ : Deposit interest-discount rate ratio	0.1	σ : Banks' minimum dimension relative to firms	4
μ_2 : Minimal reserve requirement parameter	0.1	G : Initial real value of public spending	200

Table 2: Baseline summary

Simulations			
Variable		Euro Area 1999-2007	Euro Area 2007-2016
Real GDP Growth	1.42 (0.111)	2.3	0.4
Real labour	1.42	Total Economy: 0.9	0.4
Productivity Growth	(0.111)	Manufacturing only: 3.3	1.8
Inflation	2.95 (0.148)	2.2***	1.4
Unemployment	14.0 (2.079)	8.8	10.4
Public Debt over GDP	130 (6.308)	68.3	86.0
Public Deficit over GDP	1.42 (0.132)	2.0	3.5
Private Debt over GDP	76.18 (4.022)	122.1**	142.2
Public Expenditure over GDP	49.3 (0.218)	46.6	49.1
R&D expenditure over GDP	1.47 (0.070)	1.8*	2.1
Household debt to GDP ratio	5.88 (0.186)	6.9	5.8

Table 2: Average simulated of 50 Monte Carlo simulation runs and empirical macro-variables. The standard errors of the country averages across Monte Carlo simulations are also displayed in brackets.

* data availability for the EA (19 country) starts from 2000; ** data for the EA (19 country) starts for the EA (19 country) from 2001; *** data for the EA (19 country) starts from 2002

B Appendix: Genetic Algorithm setup

Each firm in the GA specification model is provided with a genetic algorithm. The genetic algorithm is based on the XCSF methodology which combines a classifier system and a Q-learning approach (Wilson, 1995, 2002). It is necessary to adopt a Q-learning approach because the economy in which firms interact is continuously evolving in reason of the competition among firms, macroeconomic cycles and exogenous shocks.

Each classifier is made by three perceptor genes and one prediction gene. The three perceptor genes correspond to past sales, prices and unemployment variations. The prediction gene

represents next period sales variations. Each gene can take two values: 1,-1 (if one of those variable is increasing it takes the value 1, otherwise -1).

Given the current perceptions, namely the combination of the variations of the three variables that we take into account, the GA associates these three variables with prediction, thus producing a classifier. The algorithm uses a Q-learning method to compute a value that shows the prediction strength of each classifier. Therefore, given the perception genes that represent what firm observes, in each period the GA chooses the classifier with the higher prediction strength.

According to the simulation code, the payoff function of each classifier is given by:

```
PayOff=(1-(abs(cl.action-effectiveValue)/float(max(Laction)-min(Laction))))**xPayoff)
```

cl.action is the action associated with the classifier, the effective value is the effective realization, *Laction* is the set of prediction that our classifier can make; in our specification we restricted the prediction to only two values (1,-1). *xPayoff* is a parameter that influences the sensitivity to errors.

In order to update classifiers' prediction strength:

```
cl.p=cl.p+self.beta*(PayOff-cl.p)
```

where *cl.p* is the prediction strength and *self.beta* is the learning parameter. The prediction strength, weighted for the number of times the classifier was in the action set (the set of classifier that can be chosen given the perceptrons), is used to choose which rule is more effective to make forecasts.

The creation of a new classifier, thus a new association between perceptions and predictions, is based on crossover and mutation processes. Moreover, there is a small probability of error that allow firms to choose not the best classifier to make their predictions.

Every period new firms that enter the market take the classifier developed by a random chosen existing firm; this is assumed in order to allow firms to develop an effective set of classifiers.

C Appendix: AR setup

The autoregression that we use to make prediction is a simple VAR, with three variables, the same variables used by the GA algorithm to allow comparisons among the two prediction methods. We use one step ahead predictions:

$$\hat{y}_{T+k} = \hat{A}_1 \hat{y}_{T+k-1} + \dots + \hat{A}_p \hat{y}_{T+k-p} \quad (36)$$

Table 3: Genetic Algorithm Parameters

N^{GA} : maximum number of classifiers	5000	(all the possible classifiers in our case)	
β^{GA} : learning rate	0.45	α^{GA} : calculating classifier fitness (multiplier)	0.1
ν^{GA} : calculating classifier fitness (exponential)	5	ϵ_0^{GA} : calculating classifier fitness (minimum error)	0.00
γ^{GA} : multi step processes Q learning parameter	0.7	Pashtag: probability of using # as covering of an attribute	0.05
$theta_{Mna}^{GA}$: minimum number of action in a matching set	2	θ_{Del}^{GA} : fitness threshold	20
θ^{GA} : subsumption threshold	20	pExpl: probability of exploration	0.05
δ^{GA} : mean fitness minimum	10	θ^{GA} : waiting time before creating another classifier	30
χ^{GA} probability of applying crossover	1.0	μ^{GA} : probability of mutating an allene	0.05
pI: payoffs	0.05	ϵ_f^{GA} : error	0.0
FI: fitness	0.05	xPayOff: sensitivity to payoff	3.0

We use ten lags of the variable to make predictions and a time series of maximum 100 periods used to estimates the VAR coefficient matrices. We use this relatively large number of lags and this long data time span to allow firms to perform a good estimate of the VAR.

As in the GA, new entering firms take the time series that have been collected by a random chosen existing firm to estimate the VAR.

D Appendix: Sensitivity analysis

In the paper we test the effect of different shocks, moreover we verify the robustness of the model to variation of the degree of competitiveness (β). In this appendix we test the impact of other crucial variables such as: the adjustment parameter (δ),¹³ the probability of introducing an innovation (ν), the dividend rate (ρ).

D.1 Variation of the adjustment parameter, δ

In the baseline scenario the adjustment parameter δ is equal to 0.04. In additional experiments we evaluate the effects of both reducing it to 0.035 and increasing it to 0.045. As Figures 11 and 12 show, the results of simulations using different prediction methods are in line with the baseline specification. Fixing the parameter δ equal to 0.04 as in the baseline improves the model calibration. Indeed, when δ is lower the public debt tends to grow, while with higher values of δ unemployment increases. Moreover, figure 13 shows that the prediction error variation in relation to both a fiscal and a macroprudential shock are in line with the baseline specification of the model.

¹³This parameter is related to different behavioral rules and equations though it is always used as a parameter governing some adaptive process. In a sense, by experimenting on this parameter we want to test the sensitivity of the system to changes in the degree of adaptation in various mechanisms.




Figure 11: Macro variables. *GA* in red, *AR* in blue, *N* in orange. The baseline scenario *B* in black. Solid lines represent averages over 100 simulations, dotted lines are 95% confidence intervals. $\delta = 0.035$.




Figure 12: Macro variables. *GA* in red, *AR* in blue, *N* in orange. The baseline scenario *B* in black. Solid lines represent averages over 100 simulations, dotted lines are 95% confidence intervals. $\delta = 0.045$.




Figure 13: Error bias after fiscal and macroprudential shock. Bars are 95% confidence intervals. $\delta = 0.035$ and $\delta = 0.045$. On the x axis shock entity and on the y axis the maximum error bias.

D.2 Variation of the innovation parameter, ν

The innovation parameter ν determines the probability of translating the expenditure in R&D in a successful innovation. Increasing ν augments the probability of increasing firms' productivity, while the opposite occurs when the parameter is reduced. Figure 14 and 15 show that modifying the parameter ν leads to the same results as in the baseline scenario (with $\nu = 1.0$). When ν is reduced to 0.9, the growth rate of the economy is reduced in reason of the slowest productivity dynamics, while when we increase ν to 1.1, productivity growth augments. Figure 16 shows that the results in terms of error variations are in line with the baseline scenario.

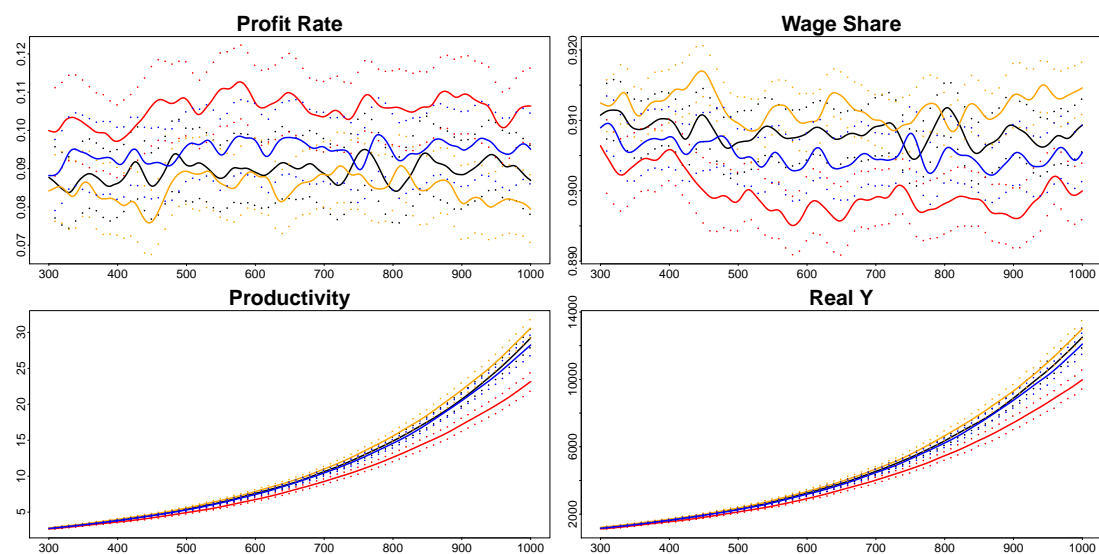


Figure 14: $\nu=0.9$. *GA* in red, *AR* in blue, *N* in orange. The baseline scenario *B* in black. Solid lines represent averages over 100 simulations, dotted line are 95% confidence intervals.




Figure 15: $\nu=1.1$. *GA* in red, *AR* in blue, *N* in orange. The baseline scenario *B* in black. Solid lines represent averages over 100 simulations, dotted line are 95% confidence intervals.

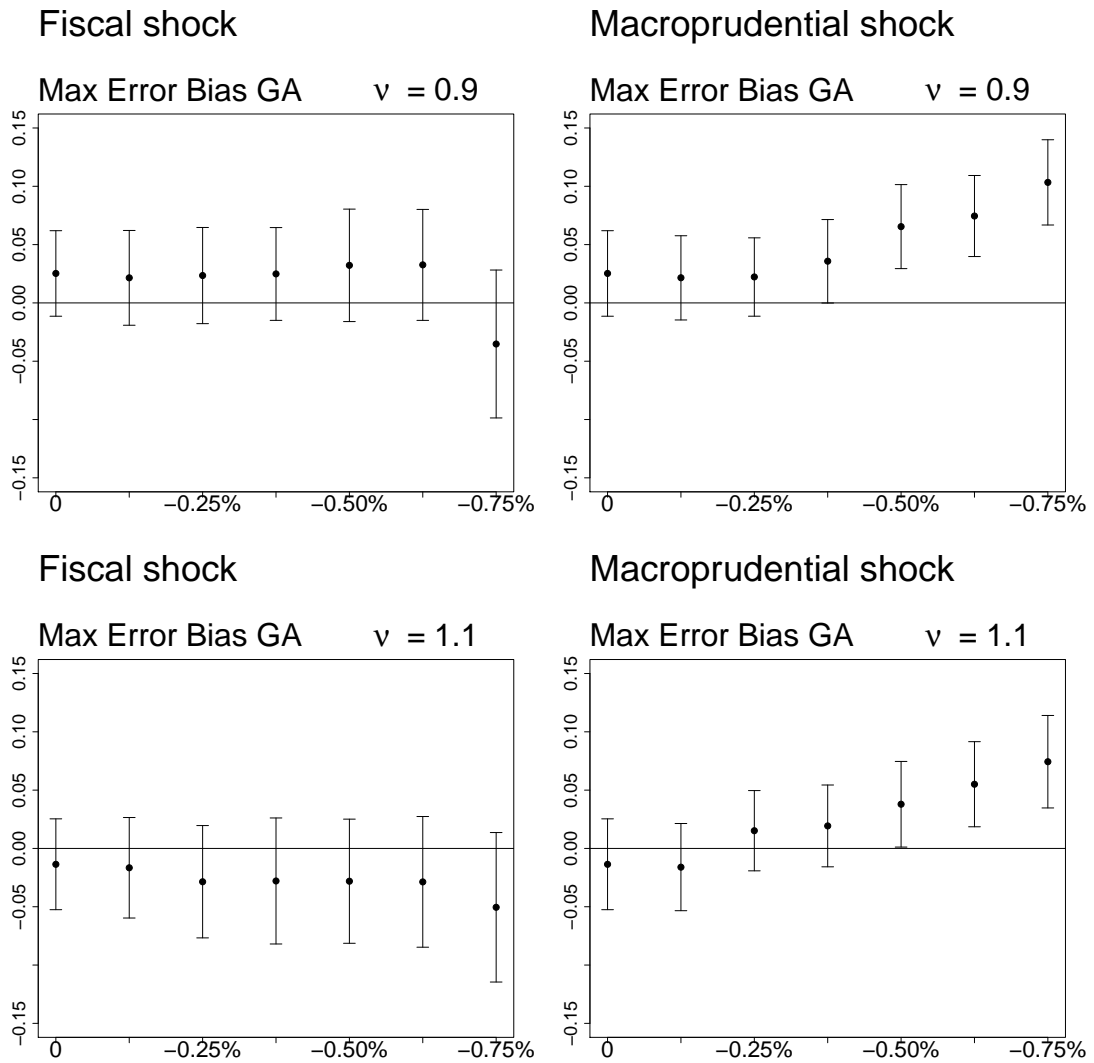


Figure 16: Error bias after fiscal and macroprudential shock. Bars are 95% confidence intervals. $\nu = 0.9$ and $\nu = 1.1$. On the x axis shock entity and on the y axis the maximum error bias.

D.3 Variations of the dividend rate, ρ

The dividend rate (ρ) determines the capacity of firms of increasing its net-worth and, thus, the need of external funding through bank lending. In the baseline scenario ρ is equal to 0.9. We firstly increased ρ to 0.98 and then we decreased it to 0.8. Figure 17 and 17 show that varying ρ does not change the main results of the baseline scenario regarding the different prediction methods that we test. As expected, when the ρ parameter increases, the need for external funds increases, thus firm leverage goes up. The opposite happens when ρ decreases. Moreover, when ρ increases and firms are more indebted the economy is more sensible to macroprudential shocks.

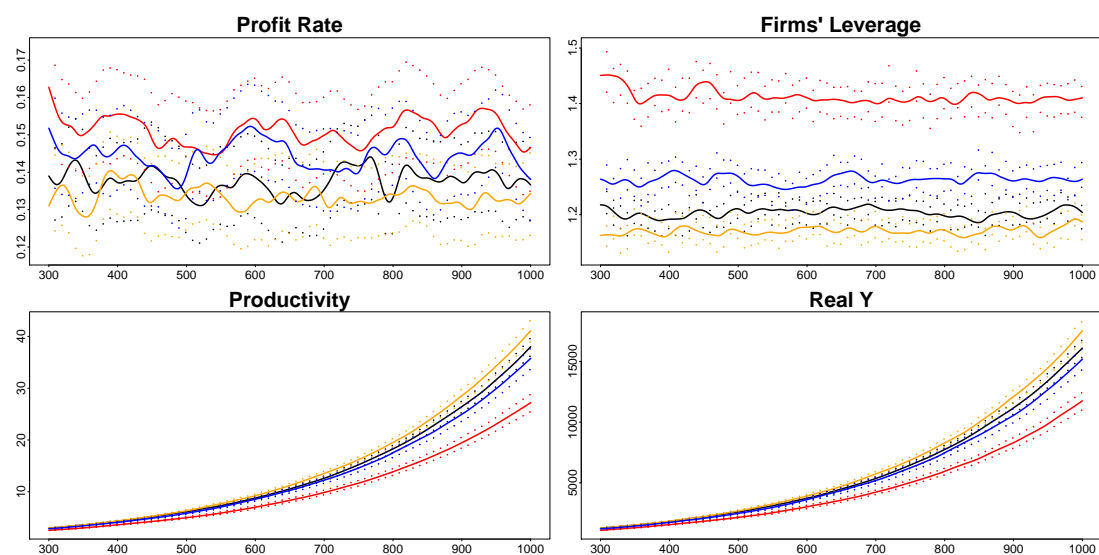


Figure 17: $\rho = 0.98$. *GA* in red, *AR* in blue, *N* in orange. The baseline scenario *B* in black. Solid lines represent averages over 50 simulations, dotted line are 95% confidence intervals.




Figure 18: $\rho = 0.8$. *GA* in red, *AR* in blue, *N* in orange. The baseline scenario *B* in black. Solid lines represent averages over 50 simulations, dotted line are 95% confidence intervals.

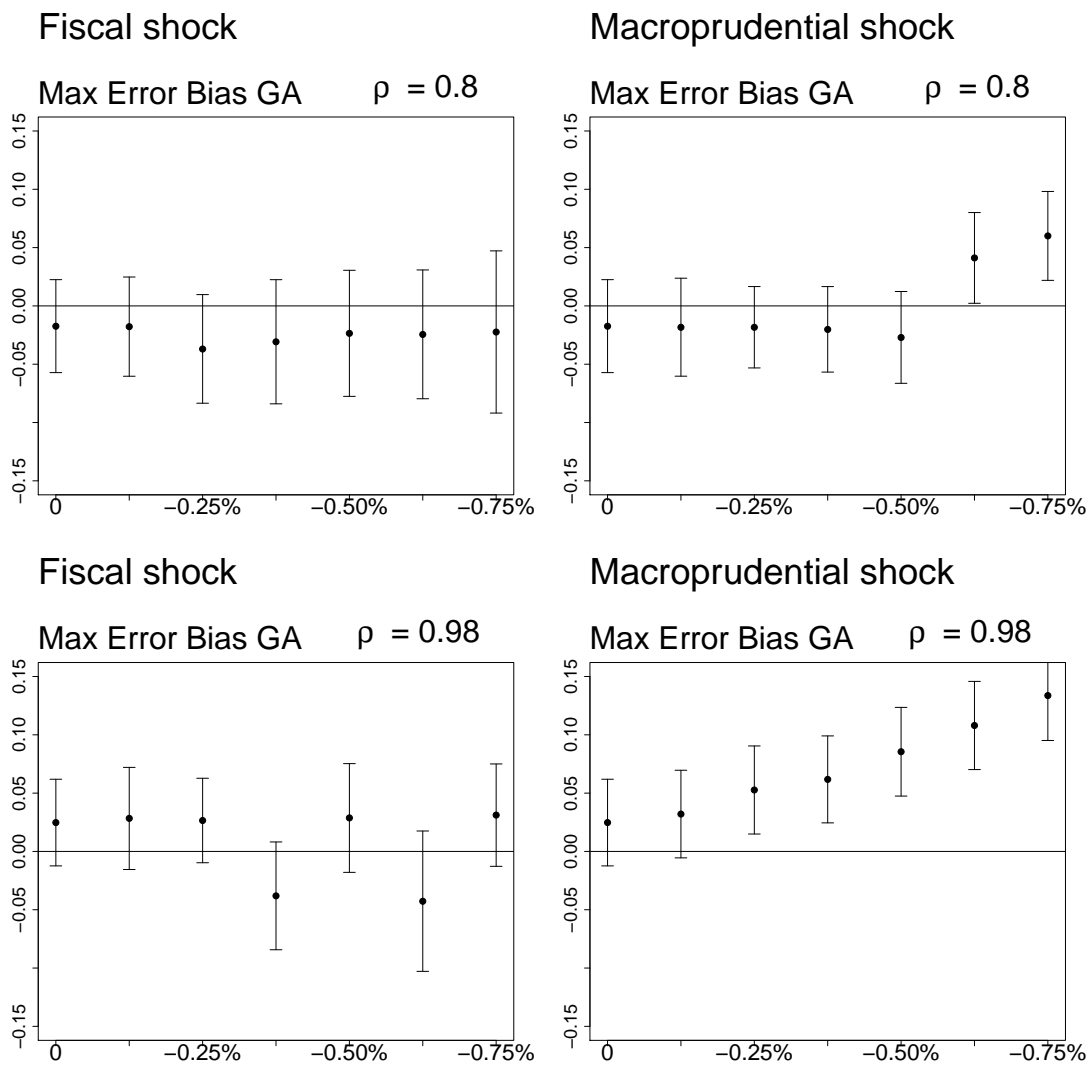


Figure 19: Error bias after fiscal and macroprudential shock. Bars are 95% confidence intervals. $\rho = 0.98$ and $\rho = 0.8$. On the x axis shock entity and on the y axis the maximum error bias.

D.4 Zero Intelligence Test

We tested the effect of random choices on sales expectation: growth or decline expectations are chosen randomly with the same probability of 50%. Figure 20 shows that random choice are strongly negative biased, this is because firms are not able to understand that in the simulated economy thanks to productivity growth sales tends to increase. As a consequence the mean square error is larger than in the GA scenario. The negative bias is larger with respect to the naive expectation (fig. 2), indeed with naive expectations firms at least can base their action on previous results. In conclusion, random agents perform worst than naive ones in terms of predictive capabilities.




Figure 20: Simulations with Zero intelligence, Average error and Mean squared error. *GA* in red, *Random* in green. Solid lines represent averages over 50 simulation, dotted line are 95% confidence interval.

The macro results are in line with the baseline scenario the one without firm predictive capabilities (Figures 21 and 22).

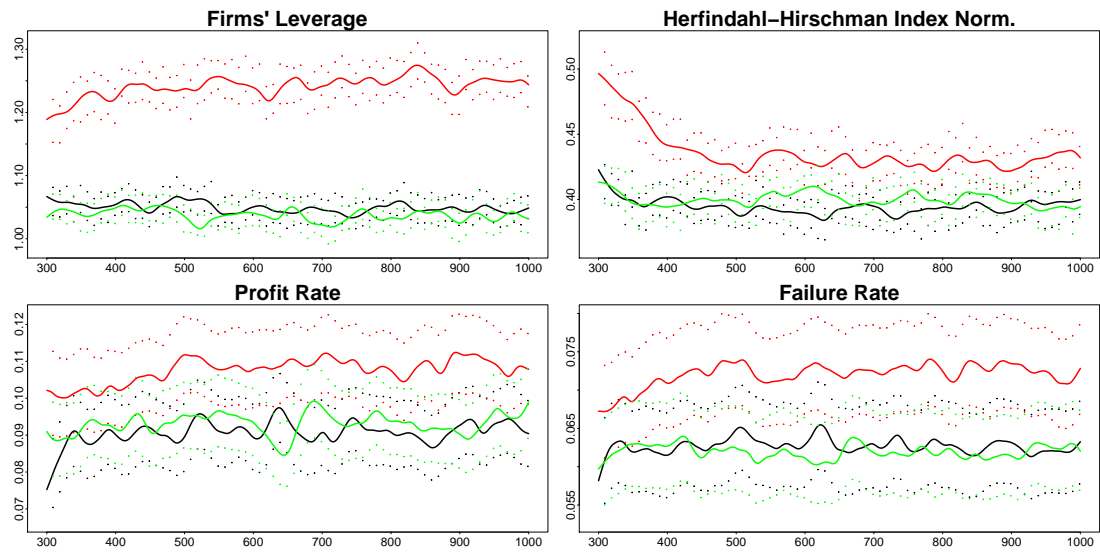


Figure 21: Simulations with zero intelligence. Firm Leverage, HI index, Profit and Failure Probability. *GA* in red, *Random* in green. The baseline scenario *B* in black. Solid lines represent averages over 50 simulation, dotted line are 95% confidence interval.

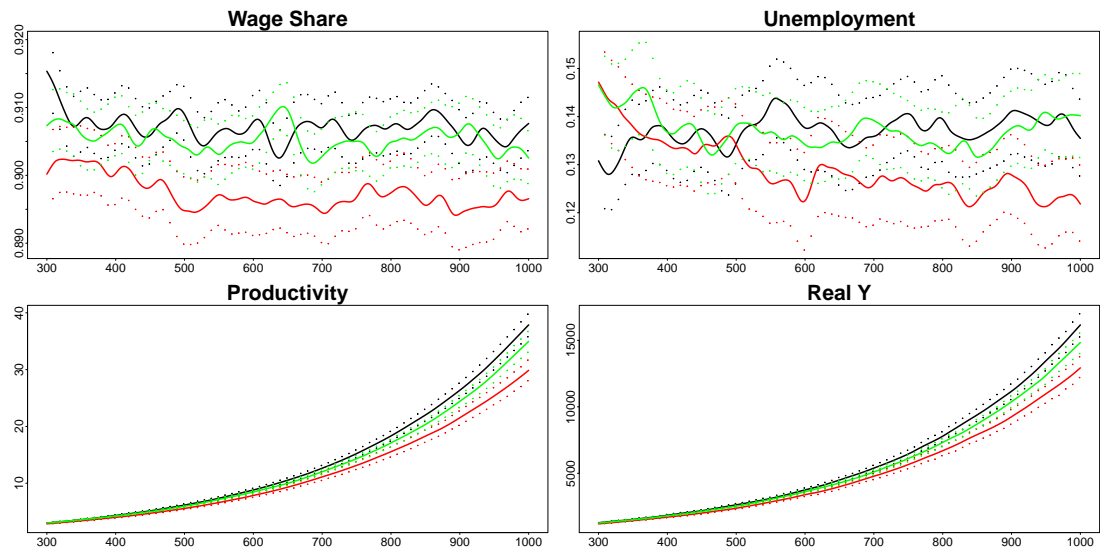


Figure 22: Simulations with zero intelligence. Macro variables. *GA* in red, *Random* in green, *N* in orange. The baseline scenario *B* in black. Solid lines represent averages over 100 simulations, dotted line are 95% confidence interval.