UNIVERSITÀ POLITECNICA DELLE MARCHE
Facoltà di Economia G.Fuà

DIPARTIMENTO DI SCIENZE ECONOMICHE E SOCIALI (DISES)
Ph.D. degree in Economics - XXXII cycle

Doctoral Thesis

Three Essays in Agent-Based Macroeconomics

Author:
Luca Eduardo Fierro

Supervisors:
Professor Alberto Russo
Doctor Alessandro Caiani

List of Figures

List of Tables

Contents

Introduction

This dissertation seeks to contribute to the growing literature on Agent-Based Macroeconomics along three lines of research organised in three distinct chapters: (i) an ABM application to a specific macroeconomic problem; (ii) a methodological contribution; (iii) a critical review about open challenges still to be faced by ABM modellers;
Although the three chapters can be read independently from each other, I would argue that they show a certain degree of complementarity. Indeed, some insights and findings presented in chapter 2 and 3 support at least partially the modelling strategy implemented in chapter 1, whereas others can be used in the future to refine and improve the model presented in chapter 1.
The organisation of the chapters may seem unusual, in particular the choice to place the literature review at the end of the dissertation. This reflects the trajectory of my research: admittedly, when I started my PhD journey I was not fully aware of many critical issues affecting ABMs, however as my research advanced unresolved challenges stood in my way, some of which I decided to directly address in this dissertation, this is the case of chapter 2, others I decided to critically discuss, which is the case of chapter 3.

In chapter 1 I propose a macroeconomic model suitable for studying technological innovations and structural change, moreover I provide an application of such model which elucidates a plausible and empirically sound mechanism leading from *automation* to *job polarization*.
The model proposed in this chapter extends and complements models already well established in the literature, in particular it introduces heterogenous consumption goods, the possibility to model a non vertically integrated multi-sectors economy, and a novel production system in which different types of labor must be combined in the production process.
The model, in its general version, is intended to be flexible enough to accommodate for future extensions and therefore to address a variety of research questions dealing with technological change, structural change, and labor market outcomes in terms of: aggregate employment, labor flows across sectors and skill specific reactions to technological shocks.
In chapter 1 I also provide an application of the model to study how *automation* can lead to *job polarization*. The modelling strategy has been designed to be as close as possible to available empirical evidence on robots, sectorial workers skill distribution, and consumer preferences over differentiated goods. This allows to provide empirical grounding to some interesting and sometimes counterintuitive results. In particular, the model helps to understand a possible mechanism leading from *automation* to *job polarization*. As we will see, a sector specific labor-saving and skill-biased technological shock, which *per se* should depress the employment share of low skilled workers and increase the employment share of high skilled workers, can actually set in motion a chain of causal events leading to structural change at first and eventually to job polarization. In the chapter such mechanism will be explained in

details, moreover it will be justified in light of available empirical evidence and its robustness proved by means of extensive sensitivity analysis.

The second chapter is a methodological contribution, which tries to understand how to model expectations in ABMs. To begin with, the chapter tries to clarify which type of rationality is best suited in the ABM framework, maintaining that: (i) rational expectations *a lá* Muth are neither applicable, nor needed in the ABM framework; (ii) what is sufficient to achieve in ABMs is *collective* rationality, which simply implies that the aggregate mean forecasting error is on average zero, i.e. the economy as a whole is not systematically mistaken in making predictions.
Therefore, under the assumption that *collective* rationality is sufficient in the ABM framework, the chapter studies the performances of different expectations formation mechanism within two agent based models. Moreover, I introduce a learning algorithm which combined to "static" expectations allows to update the otherwise fixed parameters contained in the expectation rules. Thus, the goal is to study whether it is possible to obtain aggregate unbiased expectations in an ABM framework. Since typically a macro ABM does not have a closed form solution, I rely on extensive computer simulations in order to assess the performances of different expectation formation mechanisms in different contexts. I do so in two macro environments: (i) a very simple and stylised model in which agents try to forecast a stationary variable and (ii) a full fledged macro ABM in which agents try to forecast a trended variable. In case (i) I designed a simple model in which a central bank set the interest rate and households try to forecast the one-step ahed inflation rate. In this context I will assess the performances of different expectation formation mechanisms across policy regimes and policy shocks. In case (ii) I will employ the model put forward in Caiani et al. (2016) augmented by technological innovation as in Caiani et al. (2019). This is a full fledged macro ABM which I use as a laboratory to assess different expectation formation mechanisms applied to firms trying to forecast future sales. The expectation rules employed in the following exercise are: naive expectations, where the expected value of a variable equals its past realisation, adaptive expectations, trend following, and social learning in the form of a very simple genetic algorithm. Moreover, I also employ hybrid expectations in which adaptive expectations and trend following are combined with learning.

The last chapter discusses some of the main challenges for ABMs, in particular it deals with how to bridge models to real data and how to address the Lucas critique in ABMs. Overall, it is intended as a review of ongoing research on these issues, however I also suggest some possible ways to deal with the specific problems surveyed. The chapter concludes with a preliminary meta-analysis trying to assess the state of the art of currently widespread modelling practises, with respect to the challenges laid down in the chapter.

Chapter 1

Automation, Structural Change and Job Polarization in an ABM Framework

Abstract

In this paper I present an Agent-Based macroeconomic model suitable for studying the effects of technological innovations in a multi-sector economy where a combination of differentiated labor and machines are needed in production.
I therefore employ the model to study the aggregate effect of a sector specific, labor-saving, and skill-biased technological shock, which, consistently with available empirical evidence on robots, it is interpreted as an *automation* shock. Interestingly, the model shows that a skill-biased technological change can in principle trigger a job polarization dynamics, moreover it suggests a simple mechanism behind this results which logically depends on two empirically grounded assumptions: (i) the service sector disproportionally employs low skilled workers; (ii) high-income households feature higher preferences towards services relative to other households.
I performed an extensive sensitivity analysis on the parameters governing the strength of the technological shock, showing that the emergence of job polarization is robust to the kind of shock imposed within the model. Moreover, the most important parameters have been carefully calibrated using empirical data.

1.1 Introduction

In this paper I propose a macroeconomic model suitable for studying technological innovations and structural change, moreover I provide an application of such model which elucidates a plausible and empirically sound mechanism leading from *automation* to *job polarization*.
I therefore designed a multi-sector macroeconomic Agent-Based model (ABM, hereafter) allowing for heterogenous workers and capital items. The model extends the contribution of Caiani et al. (2019) in several ways, however the most important boil down to: (i) a new definition of capital, which requires different types of skilled labor in order to be operated. This mimics the fact that different tasks are needed in the production process and that each task requires particular skills. Moreover, the type of tasks/skills needed in production depends on the technology employed and it is therefore embedded in capital goods; (ii) a dual consumption sector, where households can consume *manufactory* goods and *personal* services. Moreover, different sectors involve different tasks in the production process and therefore require different skills.
The advantage of using ABMs in economic analysis rests on their ability to realistically model economies as complex adaptive systems in which heterogenous and bounded rationally agents interact without need to impose any strong notion of equilibrium at any level[1], i.e. micro, meso, or macro. Specifically for the purposes of this paper, I shall stress the flexibility guaranteed by ABMs which accommodates for a variety of features, making the model fit to tackle a wide range of research questions. For example, the very structure of the economy, that is the number of productive sectors and the links among them, can be easily adjusted to fit the scope of the research. A similar claim holds true for skills, the model indeed allows for simple skill definitions as "hierarchical" skills, i.e. from higher to lower skills, to more complex skill definitions, where each worker is endowed with a bundle of abilities. Moreover, whatever the skill definition chosen, skills can evolve endogenously according to a variety of mechanisms, therefore making the model suitable to study issues related to education and technological change. Another attractive feature is the flexibility in modelling technological evolution, which can impact on productivity, capital-labor ratio, and the composition of tasks/skills needed in production. Moreover, it can be modelled as an exogenous technological shock or as an endogenous pattern shaped by the wide economic environment. Finally, macro ABMs provide a rich policy laboratory, where a range of policies spamming from monetary to fiscal or innovation related interventions can be tested and assessed.

The application presented in this paper can help to rationalise the chain of causal relations leading to job polarization in the face of an *automation* shock occurring in the manufactory sector. Where job polarization is defined as a tendency according to the employment shares of jobs located at the two poles of the skill distribution increase, therefore depressing the shares of middle-skilled employment; and automation is defined as a labor-saving skill-biased technological shock. Let me also clarify that throughout the paper a skill-biased technological shock is assumed to complement high skilled workers, substitute for low-skilled workers, and leave unaffected middle-skilled workers. The reason to define automation as such has to be found in the empirical findings of Graetz and Michaels (2018), who document a seizable

[1]For an in-depth discussion on Agent-Based modelling and its advantages the reader should refer to Epstein (2006), LeBaron and Tesfatsion (2008), Delli Gatti et al. (2010a), Gallegati and Kirman (2012), Fagiolo and Roventini (2016), Haldane and Turrell (2018), and Caverzasi and Russo (2018).

adverse effect on low-skilled employment share caused by robots. Whereas the choice of leaving middle-skilled employment unaffected is simply out of convenience but without loss of generality, as it will be clarified later on.
Clearly, automation, as defined above, does not *per se* lead to job polarization, as one of its direct effects is to reduce the low-skilled employment share. Indeed, the full mechanism behind the emergence of job polarization is slightly more complicated than a simple direct effect of the technological shock and it depends on two empirically grounded assumptions embedded in the model. Assumption (1): non-homothetic preferences, i.e. high-income, high-skilled households have stronger preference towards *personal* services relatively to other households. Such hypothesis was originally formulated by Manning (2004) who maintains that because of the high opportunity costs to produce home-substitutes goods by high-income households, they would be more willing to buy them in the form of *personal* services produced by low-income, low-skilled households. The hypothesis were confirmed empirically by Manning (2004) e more recently by Mazzolari and Ragusa (2013) and Lee and Clarke (2019) who find evidence of high-skilled consumption spill-over on low-skilled employment particularly concentrated in the *personal* service sector; Assumption (2): the *personal* service sector disproportionally employ low-skilled workers. This is a well known empirical fact, which will be confirmed by the calibration exercise presented later on in the paper.
When the model integrate said assumptions, the response to an automation shock in the manufactory sector can be disentangled in two steps: in step 1 the shock hits the manufactory sector, therefore increasing the share of high-skill employment, reducing the share of low-skilled employment and leaving unaffected the share of middle-skill employment *within* the manufactory sector. At this point high-skilled workers enjoy a gain in terms of employment, which reduces high-skilled labor supply and determines an upward pressure on high-skilled wages. As high-skilled wages grow relative to middle and low-skilled wages, aggregate demand lean towards personal services because of assumption (2). Such change in the aggregate demand composition brings about a demand-led structural change, causing growth of the service sector. Service sector growth generates growth in the share of low-skilled employment because of assumption (1). The net effect can be job polarization and in fact, as it will be shown later on, job polarization occurs for a wide range of shock intensities.

The rest of the paper is organised as follows: section 2 discusses the related literature; section 3 describes the model; section 4 presents the simulation procedure, the calibration strategy and the baseline dynamics; section 5 shows the main result and sensitivity analysis; section 6 concludes.

1.2 Related Literature

The paper contributes to three strands of literature: (i) ABMs for macroeconomics; (ii) automation, and (iii) job polarization.
The model presented in this paper extends the model put forward by Caiani et al. (2019), which in turn builds on the benchmark *stock-flow-consistent*-ABM (SFC, hereafter) of Caiani et al. (2016). Other examples of SFC-ABM are the EURACE model, see Deissenberg et al. (2008), Kinsella et al. (2011), Riccetti et al. (2015), and the JAMEL model, see Seppecher (2012), just to cite some early attempts to integrate the SFC and the ABM methodologies.
The model deals with workers heterogeneity, which is not a new topic in the wider ABM

literature. For example Dosi et al. (2018a) proposes a model in which a worker's individual skill level determines her individual labor productivity, moreover the skill level endogenously evolve according to the worker's employment status in a *learning-by-doing* fashion. A similar idea is contained in the model proposed by Dawid et al. (2008), although with an additional level of complexity: each worker is endowed with an exogenous level of *general* skill and with a set of endogenous skills each of them specific to a particular type of technology. A specific skill evolves depending on the general skill level embedded into the worker and the time the worker spend working with the relative technology, again, in a learning-by-doing fashion. On the other hand, Ciarli et al. (2010) and Caiani et al. (2019) model workers heterogeneity in a purely hierarchical fashion, where blu collars are directly employed in production, whereas white collar are hired for organisational reasons only. The approach pursued in this paper is somehow different, first of all productivity is a technological characteristics only and therefore embedded in capital items and not in workers. Moreover, I took a sort of stylised task approach, where different technologies requires different tasks, or different compositions of the same tasks, and different tasks requires different skills. Therefore capital items *command* skills, which in this version of the model are assumed to be exogenously given to workers. Also in this version of the model, skills are assumed to be somehow hierarchical in the sense that there are higher and lower skills, although differently form Ciarli et al. (2010) and Caiani et al. (2019) those are directly needed in the production process, i.e. there is no blue/white collar distinction.
Finally, multi-sector ABMs often model vertically integrated sectors, usually a capital and consumption good sector[2], in this paper I introduce a dual-consumption good sector where households can consume *manufactory* good and *personal* services, although the model can accommodate for any type of consumption sectors.

The second contribution the paper makes is to provide a new model to study the effects of automation on the labor market and the economy at large. A renewed interest in this topic was prompted by recent empirical findings about the potential disruptive effects that robotization may have for workers. The debate was initiated by Frey and Osborne (2017) who estimate that 47% of U.S. jobs are likely to be automated in the near future. Others have tried to replicate the study, finding rather different figures, for example Arntz et al. (2016) found that across OECD countries only 9% of jobs are likely to be automated, or Pajarinen and Rouvinen (2014), who instead confirm Frey and Osborne's finding for US using a more updated dataset and estimate for Finland a 35% share of jobs at high risk of automation. These studies try to estimate the future effect of robotization on employment, yet robotization is an ongoing process and therefore it is possible to investigate its already occurred impact. Using IFRs data Graetz and Michaels (2018) estimate no seizable effect on aggregate employment, but a negative effect for low-skilled workers. On the other hand, Acemoglu and Restrepo (2017) and Chiacchio et al. (2018) estimate the spatial general equilibrium model put forward in Acemoglu and Restrepo (2017) for US and Europe respectively. They also use IFRs data, but unlike Graetz and Michaels (2018) they both find that the displacement effect dominates.
Such conflicting empirical evidences signal a lack of complete understanding of the forces at work, their interactions and relative strengths. To fill this gap some models have been

[2]See for example Dosi et al. (2010), Ciarli et al. (2010), Assenza et al. (2015), Dawid et al. (2016), Caiani et al. (2016), and Seppecher et al. (2018).

proposed for studying automation. Sachs and Kotlikoff (2012) focus on the generational conflicts of automation, they assume automation to complement for old skilled-labor and substitute for young and relatively unskilled labor. They therefore predicts wage depression for young workers, hampering their ability to invest in physical and human capital. Under certain parametrizations the model predicts "long-term misery" where young workers suffer low wages both in the short and long term, Sachs and Kotlikoff (2012) advice for an inter-generational redistributive tax policy allowing for more efficient distribution of the productivity gains provided by automation. DeCanio (2016) using a production function involving labor, robots, and ordinary capital finds that the effect of robots on capital depends on the elasticity of substitution between labor and robots. His estimates suggest that with an elasticity larger than 1.9, robots diffusion is likely to depress wages. Finally, Berg et al. (2018) devise a model in which capitalists own robots and traditional capital, which are combined with low and high skilled labor in production. They use a CES production function and assume robots to complement for high-skilled labor and substitute for low-skilled labor. They find that automation exerts a positive effect on growth, but it exacerbates inequality.

The last strand of literature the paper contributes to is job polarization, which originated from the seminal contribution of Goos and Manning (2007), who starting from the task-based approach of Autor et al. (2003) show that non-routine cognitive task jobs are located at the top of the wage distribution, whereas non-routine manual jobs are located at the bottom of the wage distribution. Therefore, they show how the UK labor market polarised between 1975 and 1999 proxying labor quality by median wage. Following Goos and Manning (2007) an abundant empirical literature has confirmed job polarization to be a common feature of virtually all developed country, see for for example Autor et al. (2015), Ciarli et al. (2018), Goos et al. (2009), Goos et al. (2014), Michaels et al. (2014), Foote and Ryan (2015) and Jaimovich and Siu (2012).
Closely related to this paper are Autor and Dorn (2013) and Bárány and Siegel (2018), who point out that the growth in low skilled jobs has been concentrated in *personal* service jobs. Bárány and Siegel (2018) are probably the first to show neatly that a considerable part of job polarization happened between sectors, suggesting that structural change plays a pivotal role in explaining job polarization.

1.3 The Model

The model integrates the Stock-Flow Consistent macro modelling approach in an ABM framework, extending the SFC-ABM benchmark presented in Caiani et al. (2016) and enriched with workers heterogeneity in Caiani et al. (2019). The SFC macro structure ultimately imposes accounting discipline in the model: indeed it requires that (i) except for physical capital, any asset owned by an agent must have its liability counterpart entering the balance sheet of another agent; (ii) every flow is a vector moving chunks of stocks from one balance sheet to another. This implies that any expenditure of one agent has to generate income for another agent and that this transaction must be recorded as a variation in the respective stocks. SFC therefore provides a realistic picture of the macroeconomy, moreover as discussed in Caiani et al. (2016), even small accounting mistakes violating stock-flow consistency tend to build up as the simulation unfolds, potentially affecting model dynamics and biasing results.
As already mentioned, this paper extends the model presented in Caiani et al. (2019) in two

main respects, i.e. it introduces a task-based production process and allows for heterogenous consumption goods. Other novelties introduced are non-homothetic preferences and a slightly modified mark-up updating process.
The model is composed of different classes of heterogenous interacting agents which can be summarised as:

- A set Φ_H of households consuming, selling labor to firms and the government, paying taxes on net income, saving in the form of deposit and owning firms and banks proportionally to their share of net wealth. Moreover, households are further divided in three skill groups: low, middle, and high skilled.

- Three sets of firms: service firms, Φ_S, consumption good firms, Φ_C, and capital good firms, Φ_K. Service firms produce an homogenous good using labor only. Consumption good firms produce an homogenous good combining labor and machines. Capital good firms produce machines using labor only.
  Firms demand loans to banks in order to finance production and investment, retain profits in forms of bank deposits, pay taxes on profit, and distribute profits to their owners.

- A set Φ_B of banks, collecting deposits from households and firms, providing credit for firms, buying government bonds, and distributing profits to their owners.

- A government hiring public workers, paying unemployment benefits, collecting taxes, and issuing bonds.

- A central bank, holding banks' and government's reserve accounts, accommodating banks' demand for cash advances, and buying public bonds when government supply exceeds banks demand.

Agents interact on seven markets:

- consumption good market
- service market
- capital good market
- three differentiated labor market
- deposit market
- credit market
- bond market

Agents interact on markets following a common matching protocol: at every step of the simulation each demander is endowed with the supplier chosen in the previous period. Moreover, she can survey the price offered by a sample ϕ_s of the entire suppliers population. The size of ϕ_s defines the degree of competition within the market and it is set exogenously with a parameter χ. Then, the demander picks the lowest price among those offered by the suppliers belonging ϕ_s, call it P_n, and compares it to the price offered by the her original supplier, call it P_o. If $P_n < P_o$, then, following Delli Gatti et al. (2010b) the demander switch to the new supplier with a probability given by:

$$Pr_s = \begin{cases} 1 - e^{\epsilon\left(\frac{P_n - P_o}{P_n}\right)} & \text{if } \; P_n < P_o \\ 0 & \text{Otherwise} \end{cases} \tag{1.1}$$

Where $\epsilon > 0$ is an intensity choice parameters exogenously set. Therefore, the probability of switching supplier is an increasing non-linear function of the difference between the old and the new price[3].

1.3.1 Sequence of events

In each period of the simulations agents' decisions/actions and market interactions take place in the following order:

1. *Production planning*: consumption, service, and capital firms set their desired output in order to match expected demand plus planned inventories.
2. *Labor demand*: Given desired output and technology, firms calculate their labor demand for each skill group.
3. *Prices, interest and wage settings:* Firms set prices, banks set interest rates on deposit and loans, and workers updates their reservation wage.
4. *Expanding capacity*: consumption firms determine their desired production capacity growth and therefore their capital demand.
5. *Credit demand*: firms compute their credit demand
6. *Credit supply*: banks gather loans application and grant credit to firms.
7. *Labour markets*: each unemployed household posts her demanded wage on the labour market relative to her skill group (skill mismatch is not allowed).
   Firms try to satisfy the labour demand for each skill group hiring available workers with the lowest reservation wage.
8. *Production* Once workers are hired, firms can produce.
9. *Capital goods market* Consumption firms buy machines in order to match their desired capacity growth.
10. *Consumption markets* Consumption good and service markets open simultaneously. Households buy goods and services form suppliers.

[3]Note that in case of the deposit market the price is the interest on deposit offered by the bank, therefore demander are actually seeking the largest possible price. So, equation (1.1) is adjusted accordingly:

$$Pr_s = \begin{cases} 1 - e^{\epsilon\left(\frac{P_o - P_n}{P_o}\right)} & \text{if } \; P_o < P_n \\ 0 & \text{Otherwise} \end{cases} \tag{1.2}$$

11. *Interests payment* Banks pay interest on deposits, firms pay interests on loans, and the government pays interests on bonds.

12. *Wages and dole*: firms pay wages and government pays wages and unemployment benefits.

13. *Taxes*: government collects profit taxes from firms and banks and income taxes from households.

14. *Dividends* banks and firms distribute dividends to households when profits are positive.

15. *Deposit market* firms and households select banks to deposit savings.

16. *Bonds market*: government emits new bonds if needed, banks buy them, and the central bank buys the difference between supply and demand.

1.3.2 Agents

In this section I will give full account of agent's behavioural rules, which are largely drawn from the parent model presented in Caiani et al. (2016). To begin with, let me anticipate that anywhere in the paper when the expectation operator is invoked, I will always refer to adaptive expectations, expressed as:

$$x_t^e = x_{t-1}^e + \lambda \left(x_{t-1} - x_{t-1}^e \right) \tag{1.3}$$

Also, to clarify notation, I will call x a generic firm. In case I seek to specify whether a firm belong to the consumption good, capital, or service sector I will use respectively c, k, and s. Similarly, a worker of generic skill is identified as σ, when I seek to specify which particular skill group she belongs to, I will use l, m, and h meaning low, middle, and high skilled.

1.3.2.1 Firms

1.3.2.1.1 Production and labor demand

In the spirit of Steindl (1976) and Lavoie (1992), consumption good and capital firms accumulate planned inventories in order to cope with unforeseen demand and to avoid frustrating potential costumers with supply constraints. Therefore, firms set desired output in order to match expected demand plus planned inventories, where planned inventories are defined as a constant share of expected sales:

$$y_{x,t}^D = (1+v)s_{x,t}^e - inv_{x,t-1} \quad \text{with} \quad x = \{c, k\} \tag{1.4}$$

Where $y_{x,t}^D$ is desired output at time t for firm x, v is the constant inventories/sales target ratio, $s_{x,t}^e$ are expected sales, and $inv_{x,t-1}$ are accumulated inventories up to the previous period.
Service firms do not accumulated inventories, as they provide an intangible good which cannot be stored. This implies that equation (1.4) for service firms must be modified as:

$$y_{s,t}^D = (1+v)s_{s,t}^e \tag{1.5}$$

Where, with a slight abuse of notation, v represents the proportion of expected sales that firms s is willing to supply in excess of its own expectations so to cope with unforeseen demand upswings.

The way firms produce, and consequently define labour demand, depends on whether they use capital in production. In case they do not, as for capital and service firms, I assume that each type of labor must enter the production process as a fixed share of the total labor employed. Which amounts to assume a fixed coefficient production function, where inputs are differentiated types of labor. Indeed, assume a single input production function:

$$Y = \mu N$$

Where Y is output, μ labor productivity, and N is labor. The amount of labor needed to produce one unit of output is clearly given by $\frac{1}{\mu}$. If we further assume three types of labor, low, middle, and high skilled, and impose that each must be employed as a fixed share of total labor we have:

$$N = N^{ls} + N^{ms} + N^{hs}$$

with:

$$\begin{cases} N^{ls} = \alpha^{ls} N \\ N^{ms} = \alpha^{ms} N \\ N^{hs} = \alpha^{hs} N \end{cases}$$

where N^{ls} is the number of low-skilled workers, α^{ls} is the required share low-skilled workers and so on. It follows that the number of σ-skilled workers required to produce one unit of output is given by $\frac{\alpha^{\sigma}}{\mu}$

Thus, we can express the production function of the service and capital firms as:

$$Y_{x,t} = \mu_x min\left(\frac{N^{ls}_{x,t}}{\alpha^{ls}_x}, \frac{N^{ms}_{x,t}}{\alpha^{ms}_x}, \frac{N^{hs}_{x,t}}{\alpha^{hs}_x}\right) \quad \text{with} \quad x = \{s, k\} \tag{1.6}$$

Where $Y_{x,t}$ is output of firm x at time t, μ_x is labor productivity of firm x, $N^{ls}_{x,t}$ is the number of low skilled workers hired by firm x at time t, α^{ls}_x is the share of low skilled workers required for production in firm x, and so on.

Note that μ_x and α_x's are firm specific and time independent. Also, I am going to assume that firms belonging to the same sector are homogenous as far as production technology is concerned, i.e. firms belonging to the same sector share the same μ_x and $\alpha' s$.

Finally, let me anticipate that the values of α's are an important driver of results, therefore I calibrated them using real data. The exact procedure is provided later on.

Finally, labour demand for each skill group is given by:

$$N^{D,\sigma}_{x,t} = Y^{D}_{x,t} \frac{\alpha^{\sigma}_x}{\mu_x} \quad \text{with} \quad \sigma = (ls, ms, hs) \tag{1.7}$$

Consumption firms combine labor and capital in production, therefore the procedure for computing labor demand turns out to be slightly different. First of all, let us define capital: any capital items is defined by five parameters: $\Omega = \{\mu_{\kappa}, l_{\kappa}, \alpha^{ls}_{\kappa}, \alpha^{ms}_{\kappa}, \alpha^{hs}_{\kappa}\}$ which respectively define capital productivity, capital-labor ratio, low-skilled share, and so on. This implies that each capital item embeds a certain level of productivity and requires a specific numbers

of workers, which is defined as the inverse of l_κ. Moreover, given the number of workers needed to operate a machine of type κ, a share α^{ls} of such workers must be of the low skilled type, and so on. The logic is the same as the one featured in equation (1.6), with the only difference that here the α's are specific to the capital item and not to the sector.
In principle we could have many κ-types which differentiate from each other along any dimension of Ω, however to keep things as simple as possible I am going to assume one type of capital κ only. This simplifies production sensibly,
When firm c has set its desired output level it can derive its desired capacity utilisation as:

$$u_{c,t}^D = \frac{y_{c,t}^D}{y_c^{tot}} \tag{1.8}$$

Where y_c^{tot} is the maximum output c can produce given its capital stock. Therefore labor demand for a generic skill is simply defined as:

$$N_{c,t}^{\sigma,D} = min(1, u_{c,t}^D) K_{c,t} \frac{\alpha_\kappa^\sigma}{l_\kappa} \quad \text{with} \quad \sigma = (ls, ms, hs) \tag{1.9}$$

Where $K_{c,t}$ is c's capital stock.
Note that since only one type of capital is allowed, the consumption firm consumption function turns out to be the same as equation (1.6) where the $\alpha's$ are κ-specific.

$$Y_{c,t} = \mu_\kappa min\left(\frac{N_{c,t}^{ls}}{\alpha_\kappa^{ls}}, \frac{N_{c,t}^{ms}}{\alpha_\kappa^{ms}}, \frac{N_{c,t}^{hs}}{\alpha_\kappa^{hs}}\right) \tag{1.10}$$

To be precise this is true if and only if there are enough machines to employ all the workers available to the firm, otherwise production should be scaled down accordingly. However, consumption firms never hire more workers that they can employ given the number of machines available, see the term $min(1, u_{c,t}^D)$ in equation (1.9).
From now on I will refer to the parameters α's describing skill requirements as *technical parameters.*

1.3.2.1.2 Pricing

Firms set price as a mark-up over unit costs of production:

$$p_{x,t} = (1 + \mu_x)\left(\frac{\sum_\sigma w_{x,t}^{e,\sigma} N_{x,t}^{D,\sigma}}{Y_{x,t}^D}\right) \quad \text{with} \quad \sigma = (ls, ms, hs) \tag{1.11}$$

Moreover, firm x's mark-up evolve endogenously following a simple adaptive rule: when real sales exceed expected sales the mark-up is revised up by a random amount, vice-versa is reduced by the same token. Moreover, the scale of the adjustment is given by the sales forecasting mistake occurred in the previous period:

$$\mu_{x,t} = \begin{cases} \mu_{x,t-1}(1 + FN_{x,t}^1)^{\frac{|s_{x,t-1}^e - s_{x,t-1}|}{s_{x,t-1}^e}\varkappa} & \text{if} \quad s_{x,t-1} > s_{x,t-1}^e \\ \mu_{x,t-1}(1 - FN_{x,t}^1)^{\frac{|s_{x,t-1}^e - s_{x,t-1}|}{s_{x,t-1}^e}\varkappa} & \text{if} \quad s_{x,t-1} < s_{x,t-1}^e \end{cases} \tag{1.12}$$

Where $FN_{x,t}^1$ is a random draw from the folded normal distribution FN^1 defined over the parameters $(\mu_{FN^1}, \sigma_{FN^1}^2)$, and $\varkappa$ is a scaling factor.

1.3.2.1.3 Investment

I assume that consumption firms invest in capital items in order to reach a given capacity utilization target. I therefore define firm c's desired capacity growth as:

$$g_{c,t}^{D} = \gamma_u \frac{u_{c,t}^{D} - \bar{u}}{\bar{u}} \tag{1.13}$$

Where $u_{c,t}^{D}$ is the desired capacity utilization, defined as desired output over total capacity, γ_u an exogenous parameter, and $\bar{u}$ is target capacity utilization, which is exogenous and equal across firms.
Once consumption firms have set desired capacity growth, they can go on the capital market and buy the required number of machines to reach the capacity target.

1.3.2.1.4 Profits, taxes, dividends, and credit demand

Consumption pre-tax profits are given by the difference between cash inflows plus investment in inventories and disbursements plus capital amortization:

$$\begin{aligned}\pi_{c,t} = s_{c,t}p_{c,t} + i_{b,t-1}^{d}D_{c,t-1} + (inv_{c,t}uc_{c,t} - inv_{c,t-1}uc_{c,t-1}) + \\ - \textstyle\sum_{n\in N_{c,t}} w_{n,t} - \sum_{j=t-\eta}^{t-1} i_j^l L_{c,j} \frac{\eta-[(t-1)-j]}{\eta} - \sum_{k_{c,\kappa}\in K_{c,t-1}} k_{c,\kappa}p_{k_{c,\kappa}} \frac{1}{\delta}\end{aligned} \tag{1.14}$$

Where $i_{b,t-1}^{d}$ is the interest rate on deposits granted by bank b and $D_{c,t-1}$ is the total amount of deposits. i_j^l is the interest on loans charged on $L_{c,j}$. η is the loan duration and therefore the second last block of equation (1.14) represents the sum of disbursements relative to each firm c's loan. The last block of equation (1.14) is capital amortization. Finally note that investment in inventories is evaluated at unit costs of production.
Capital firms evaluate profits as in equation (1.14), with the only difference that capital items stored as inventories do not depreciate, so no amortization is needed. Service firms evaluate profits as capital firms, with the difference that they do not store inventories.

Corporate tax rate, τ_π, is set by the government so that taxes are simply given by:

$$T_{x,t} = max\left\{\tau_\pi \pi_{x,t},\ 0\right\} \tag{1.15}$$

I assume τ_π to be constant across time and homogenous across firms.

Dividends are defined as a share of after-tax profits and distributed to households:

$$div_{x,t} = max\left\{\rho_x\left(1-\tau_\pi\right)\pi_{x,t},\ 0\right\} \tag{1.16}$$

Where ρ_x is the dividend share of profits, which is assumed to be constant over time, moreover firms belonging to the same sector share the same ρ_x.

Firms follow a pecking-order approach in defining credit demand: they make an estimation of the deficit between cash in and out-flows based on their expectations and cover it using internal resources. They rely on debt only when internal resources are not enough. However, for precautionary reasons firms retain an amount equal to a proportion φ of the expected wage disbursement.

1.3.2.2 Banks

1.3.2.2.1 Credit

Each bank offers the same interest rate to any potential costumer, but it can credit crunch firms based on a case by case credit worthiness assessment. Loan interest rate evolves endogenously and depends on the bank's capital ratio. In particular, we assume that if the capital ratio $CR_{b,t}$ of bank b is lower than a target capital ratio CR^T_t determined at the sector level and equal for all banks, then bank b increases its loan interest rate relative to the average loan interest rate of the previous period. For simplicity we assume CR^T_t to be the average capital ratio of the banking sector realised in the previous period. Therefore the loan interest rate charged by bank b evolves as:

$$i^l_{b,t} = \begin{cases} \bar{i}^l_{t-1}(1 + FN^2_{b,t}) & \text{If} \quad CR_{b,t} < CR^T_t \\ \bar{i}^l_{t-1}(1 - FN^2_{b,t}) & \text{Otherwise} \end{cases} \tag{1.17}$$

Where $\bar{i}^l_{t-1}$ is the average loan rate in the previous period and $FN^2_{b,t}$ is a random draw from a folded normal distribution FN^2 defined over the parameters $(\mu_{FN^2}, \sigma^2_{FN^2})$.

1.3.2.2.2 Deposit and bonds markets

Banks compete on the deposit market to attract savings from households and firms. The interest rate on deposits follows a similar evolution as the interest on loans:

$$i^d_{b,t} = \begin{cases} \bar{i}^d_{t-1}(1 - FN^2_{b,t}) & \text{If} \quad LR_{b,t} \geq LR^T_t \\ \bar{i}^d_{t-1}(1 + FN^2_{b,t}) & \text{Otherwise} \end{cases} \tag{1.18}$$

Where $\bar{i}^d_{t-1}$ is the average deposit interest rate relative to the previous period, $LR_{b,t}$ is the liquidity ratio of bank b, and LR^T_t is the sector average liquidity ratio.

Finally, banks use reserve in excess to buy government bonds when available. If supply exceed demand the difference is absorbed by the central bank.

1.3.2.3 Households

Households essentially engage in two activities: working and consuming. Each household inelastically supplies one unit of labor in her specific labor market, however if she cannot find a full time job she can be hired part-time. Part-time labor is used by firms in order to make adjustments at the margin. Anyhow, each household updates her demanded wage following a simple heuristic: if in a given time window of length T she has been unemployed for less than t_u periods, she scales up her demanded wage by a random amount, vice-versa she scales it down by the same token.

$$w^{d,t}_i = \begin{cases} w^{d,t-1}_i(1 - FN^1) & \text{If} \quad \sum_{n=1}^{T} u_{h,t-n} > t_u \\ w^{d,t-1}_i(1 + FN^1) & \text{If} \quad \sum_{n=1}^{T} u_{h,t-n} \leq t_u \end{cases} \tag{1.19}$$

Where $u_{h,t}$ is dummy variable taking value 1 if h is employed in period t and 0 otherwise. In case the worker is hired part-time, say ω% of one full unit of labor, she receives a percentage of her demanded wage equal to $\omega\%$.

Since households consume two differentiated goods, I assume that each household at first set her own desired consumption budget, which is defined as a portion of net income plus a portion of net wealth:

$$C^D_{h,t} = \alpha_{NI} NI_{h,t} + \alpha_{NW} NW_{h,t} \tag{1.20}$$

Where α_{NI} is the propensity to consume out of income and α_{NW} is the propensity to consume out of net wealth.
Once the consumption budget has been set, households decide how much of it to spend for services and how much for consumption goods. I assume that the consumption budget is split in fixed proportion between the two markets. Moreover, I differentiate such proportion across skill group, assuming that higher skill groups spend higher proportion of consumption budget for services relative to lower skill groups:

$$\begin{cases} C^{D,s}_{h,t} = \gamma^{\sigma}_{s} C^D_{h,t} \\ C^{D,c}_{h,t} = (1-\gamma^{\sigma}_{s}) C^D_{h,t} \end{cases} \qquad \text{with} \qquad \gamma^{hs} > \gamma^{ms} > \gamma^{ls} \tag{1.21}$$

Finally, each household owns each firm and bank in proportion to her share of wealth. It follows that each household receives dividends and contributes to bankruptcies in proportion to her share of wealth.

1.3.2.4 Government

The government hires a fixed number of public workers from each skill group and pays a dole to unemployed workers. The dole is calculated in each period as a percentage of the average low skilled workers wage:

$$w^d_t = u^g \bar{w}_{ls} \tag{1.22}$$

The government taxes households' net income and profits realised by banks and firms. The tax rates are kept fixed throughout the simulation.

After having collected taxes and paid public servant wages and unemployment benefits, the government emits bonds to cover the deficit.
For simplicity, I assume that bonds pay a fixed interest rate and last for one period only. The evolution of public debt in each period is expressed as:

$$\bar{p}^b \Delta b_t = def_{g,t} = T_t + \pi_{CB,t} - \sum_{n \in N_{g,t}} w_n - w^d_t U_t - \bar{i}^b \bar{p}^b b_{t-1} \tag{1.23}$$

Where T_t are total taxes collected, $\pi_{CB,t}$ are central bank profit distributed to the government, $N_{g,t}$ are public workers, U_t is the number of employed workers, $\bar{i}^b$ is the bond interest rate, and $\bar{p}^s$ is the bond price.

1.4 Baseline Simulation

In this section I describe the calibration and validation procedures implemented for the baseline version of the model. I follow the calibration procedure designed by Caiani et al. (2016) to initialise stocks and flows, moreover an empirical motivation is provided for the key

parameters.
As it is common in the ABM literature, simulated data are the results of a Montecarlo experiment: the model has been run 25 times, each time for 500 periods. I allowed for a 50-periods burn-in during which the model stabilises, leaving 450 periods available for analysis. Results hereafter, always refer to averages across the 25 rounds of simulation.

1.4.1 Calibration

1.4.1.1 Initial stock, flows and interactions

Three entities must be initialised before the model can be simulated: (i) sectoral stocks, (ii) agent's stocks, and (iii) initial interactions. The strategy followed is to calibrate stocks at the sectoral level and afterwards distribute them across agents. To do so I can exploit the model SFC structure, following a three steps procedure: (i) I derive the SFC aggregate version of the model, which is a simple system of linear equations; (ii) I solve the system in steady state, which is defined as an equilibrium in which real stocks are stable and nominal stocks grow at a constant rate, i.e. the inflation rate; (iii) I record each and any stock aggregated at the sectorial level consistent with the steady state.
In order to solve the system for the initial stocks and flows I need to fix before hand the parameters. Some of the parameters have been precisely calibrated using real data. The others are set at ranges consistent with previous literature, and afterwards fine tuned in order to achieve a realistic distribution of initial stocks across sectors. After having computed stocks at the sectoral level, I simply distribute them homogeneously across agents within each sector. This implies that at the beginning of the simulation agents belonging to the same group are indistinguishable from a balance sheet point of view.
As anticipated above, interactions need to be initialised as well. The interactions I refer to are debt relationship between banks and firms, costumer-supplier and employer-worker relationships. I initialise them in a random-controlled way: it is random in the sense that each agent can be connected to another one randomly, but it is controlled in a way that, for example, each bank as the same amount of credit, or each firm has the same number of workers, or that each supplier has the same number of costumers, and so on.
The aim of this calibration procedure is to provide the most possible homogenous situation across agents belonging to the same group and let heterogeneity emerge as the simulation unfolds.

1.4.1.2 Technical parameters

Technical parameters, α's, play a pivotal role in the model, insofar they define the labor skill composition within sectors, therefore they are calibrated using empirical data. In this particular instance I will refer to US data.
In subsection 1.3.2.1.1 I provided a description of those parameters and how they enter production. In particular, I stressed that unlike the service and capital good sectors, from the consumption firms' point of view the α's are embedded in capital items. However, to keep things simple I assume only one type of machines in the baseline scenario, therefore I can treat the α_m's as the consumption sector technical parameters in the calibration exercise.
In order to calibrate the technical parameters, I need to precisely identify skills and sectors that I wish to model. As mentioned before, I proxy skills by educational attainment, roughly

speaking I define low-skilled as workers with no formal education, middled-skilled as workers with high-school diploma and high-skilled as workers with bachelor degree or above, see table 1.1.
The way sectors are identified depends on the calibration strategy, as we will see there is a tradeoff between the degree of precision according to which we can measure the α's and the degree of sectoral disaggregation. One way to go is to use BLS data which directly show workers educational attainment by sector. The positive side of this approach is that data are provided with a good level of disaggregation, 5-digit in NAICS classification. The drawback is that this measure does not match exactly the technical parameters. Indeed, the α's indicate the share of skills *needed* by each sector. However, in real world there is a considerable mismatch between skills and jobs, so the shares of skills actually observed in an industry does not necessarily represent the *needed* ones. To overcome this issue I decided to use BLS data on occupations disaggregated by sector. The positive side of this strategy is that I can match the minimum education attainment required for each occupation and therefore I can infer the skill shares proxied by educational attainment required in each sector. The drawback is that data are available only at 2-digits disaggregation, making much more difficult to identify the industries I wish to model. To simplify the matter I assumed the capital good sector and consumption good sector to be both broad manufactory, details are presented in table (1.2) .
Table (1.3) presents the values for the technical parameters we picked for manufactory, service and the government. Notice that, the service sector employs the larger share of low-skilled jobs in our sample of industries.

Table 1.1: Skills definition by educational attainement

Skill level	Qualification
high	Bachelor's degree","Master's degree", "Doctoral or professional degree"
medium	"High school diploma or equivalent", "Associate's degree", "Some college, no degree","Postsecondary nondegree award"
low	"No formal educational credential"

Table 1.2: Sectors definition in NAICS classification

Sectors	NAICS classification
Manufactory/Final Goods	31-33, 42, 44-45
Service	72, 81
Government	99

Table 1.3: Estimated technical parameters

	ls	ms	hs
Manufactory/Capital	0.348287	0.5279909	0.1237221
Service	0.6792776	0.2712603	0.04946206
Government	0.06438509	0.6031763	0.3324386

1.4.1.3 Skill groups size

In principle, it is possible to calibrate the relative size of households populations sorted by skill group. However, because of the way I calibrated the technical parameters this would determine a constant shortage of labour supply, particularly of low skilled workers, unless I allow for skill mismatch in the labor market. However, at this stage of analysis I prefer to keep things as simple as possible. Also because, in this paper I am mostly concerned with how, where and which type of jobs are generated, and not how workers flow from one type of occupation to another, in response to a skill-biased technological shock.
Therefore, in this particular case I decided to ensure internal consistency at the expenses of external validity and I set the number of workers for each skill group so to ensure the same initial within group unemployment rate. In this way I make sure that there are enough workers for each skill group in order to initiate the artificial economy and that there is a similar wage pressure across groups (see equation, 1.19).

1.4.1.4 Wage distribution across skill groups

I calibrated relative wages across skill groups following a similar procedure as the one for the technical parameters. The starting point are occupations matched by their minimum educational attainment. Then occupations are sorted in low, middle, high skill group according to the classification of table (1.1). Finally, we defined the σ-skill wage as the average wage across occupations belonging to group σ.
Table (1.4) presents results of our calibration exercise, where we reported relative wages across skill groups. Values reported in table 1.4 are used to initialise wages.

Table 1.4: Relative wages across skill groups

	percent of hs wage
ls	0.28
ms	0.46
hs	1

1.5 Results

1.5.1 Baseline Dynamics

Before presenting the main results, I will discuss some properties of the baseline version of the model. Figures in appendix shows that the model stabilises at fairly reasonable levels of unemployment, 4%, aggregate capacity utilization, 0.8 and public debt to GDP ratio, which remains steadily below the 60% threshold. Real sales in the final good sectors are stable as well as skill employment shares. Note that this is rather important, because it shows that in absence of external shocks the model does not generate neither structural change nor job polarization.
Overall the model resembles the aggregate dynamics featured in the parent models: the simulated autocorrelation plots are similar to the real autocorrelations derived from US data. Also the cross-correlation structure is in line with empirical observations: real investment and real consumption are pro-cyclical and unemployment anti-cyclical, as they are expected. However, real consumption is slightly lagged with respect to GDP, as in Caiani et al. (2019). Also investment is lagged, this is probably due to the fact that in this specification investment only depends on capacity utilization, which is driven by aggregate demand. Finally, also the cyclical components match stylised facts about business cycle: unemployment is more volatile than consumption and GDP, whereas investment only moderately so. Again, this may be due to the investment function specification employed in this model.

1.5.2 A simple shock

I model a permanent technological shock by exogenously imposing a new type of capital, r, which, with respect to the old type m, is assumed to be skill-biased and labor saving.

$$\begin{cases} l_r > l_m \\ \alpha_r^{ls} < \alpha_m^{ls} \\ \alpha_r^{ms} = \alpha_m^{ms} \\ \alpha_r^{hs} > \alpha_m^{hs} \end{cases} \tag{1.24}$$

Note that increasing the capital-labor ratio[4] is needed in order to ensure the new type of capital to be more efficient, i.e. to bear lower unit cost of production, with respect to the old one.
The new type of capital is introduced at a specific point in time, period 150, and it affects all the capital firms, which suddenly produce and supply it only.
Shocking the model at a specific point in time allows to use the baseline version of the model as a counterfactual experiment. Indeed, before the shock is implemented the two configurations are exactly alike.
In table (1.5) are reported the parameter values of the shock and those used for the baseline scenario.

Figure (1.1) shows how the model reacts to the shock. To improve comparability between the shock and the baseline scenario I plotted the relevant time series as shock-baseline ratios. As expected, previous to the shock all the ratios are 1. Panel (a) of figure (1.1) shows the

[4] Alternatively, we could have raised capital productivity.

Table 1.5: Simple shock

	Baseline	Shock
α^{ls}	0.35	0.335
α^{ms}	0.53	0.53
α^{hs}	0.12	0.135
l	8	9.5

Figure 1.1: Simple Shock Vs Baseline Configuration

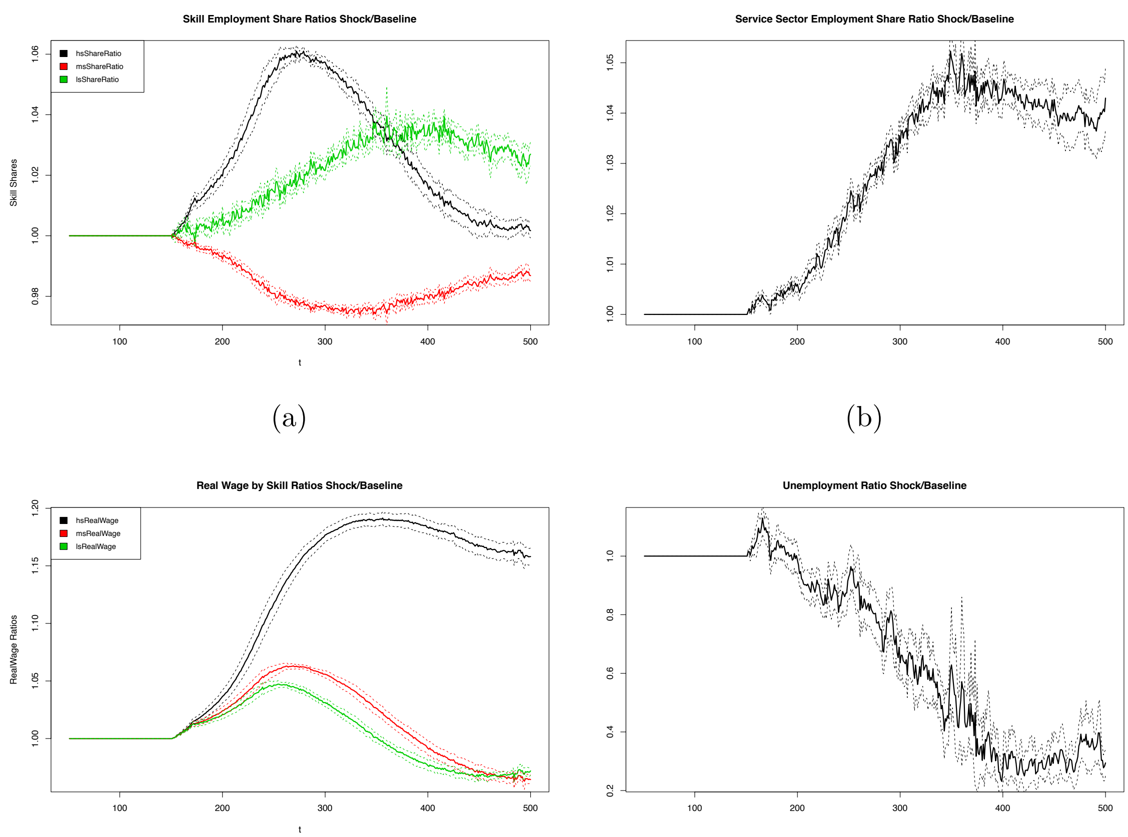

(a) (b)

(c) (d)

evolution of the employment shares relative to the three skill groups. After the shock we observe an increase in the high skilled share relative to the baseline scenario. This is not very surprising, as I imposed a skill biased technological shock favouring high skill employment. What is more puzzling is the growth in the low-skill employment share, since the shock reduces the need of low skill workers in the consumption good sector.

In panel (b) I plotted the employment share relative to the service sector, which features a strong upward trend. Such trend in the relative size of the service sector signal a structural change process which incidentally is able to explain the growth in the low-skill employment share: indeed, as suggested by the technical parameters calibration exercise, the service sector disproportionately employs low skilled workers. Therefore, as the service sector expands it creates more low-skill jobs relative to middle and high-skill jobs.

At this point, it easy to understand the engine behind the growth in low-skill jobs, what I am left to explain is the relation between the skill biased technological shock affecting the consumption good sector and the structural change dynamics. The answer is given by panel (c), showing the evolution of the average real wage within skill groups. Right after the shock there is a surge in the real wages of all skill groups. This is explained by a drop in the price for consumption good, leading to a decrease in the general price level. Indeed, the labor saving character of the new technology allows consumption firms to reduce unit costs of production and since prices are set as a mark-up over unit costs it follows a drop in prices. Clearly, the rise in real wages is not as strong and prolonged for all skill groups, high-skill wages are shown to grow much faster and steadily then low and middle-skill ones. The growth in high skill wages, coupled with the growth in high skill employment, explains structural change: indeed, I assumed that high skilled households have a higher preference to consume services relatively to low and middle skill households. When high skill relative wages increase, high skill relative consumption increase as well, moreover because of the high skilled households preferences also the aggregate service share of consumption increases. The increasing share of service consumption trigger a demand-led structural change, which eventually explains the growth in the service employment share.

The last thing I want to analyse is the effect of the shock on aggregate unemployment. By inspection of panel (d) we observe a short lived negative effect on employment occurring right after the shock, which is quickly absorbed and even reversed in the long-run. Therefore, our model, at least in this instance, does not generate technological unemployment in response of a labor saving technological shock. However, a word of caution is needed: the present configuration of the model is probably not well suited to deal with issues related to technological unemployment. The reason is that any gain in productive efficiency is fully translated onto prices. In this case, a labor saving technology allows to produce the same amount with less labor, therefore reducing the unit cost of production. Assuming fix mark-up, an $x\%$ drop in the unit cost of production reduces prices by the same percentage. If the downward adjustment of nominal wages due to higher unemployment is slower than the drop in prices, we must observe a surge in real wages which sustains aggregate demand and reduces unemployment. The mechanism described is logically sensible, but it rests on the very assumption on price adjustment in response to a technological shock. In this paper I will not tackle such issue, however, it is important to highlight that the result on employment directly depends on the price adjustment mechanism which, admittedly, may not reflect reality.

1.5.3 Sensitivity analysis

Results shown above may depend on the particular shock imposed. To investigate the robustness of the results I perform an extensive sensitivity analyses on the capital-labor ratio and technical parameters embedded in the new capital r.

In order to perform the parameters exploration I employ a kriging algorithm which allows to estimate a continuous response-surface starting from finite number of (l_r, α_r^{hs}) couples, 25 in this case. Each of the 25 shocks configurations have been run 25 times, ending up with a total 625 simulation rounds.

The kriging algorithm was originally developed in Geostatistics, however it is well suited to perform extensive sensitivity analysis of computationally heavy models of any kind. In particular, It is important to remark the seminal work by Salle and Yıldızoğlu (2014), who

proposed the kriging methodology for ABMs.
In order to produce the 25 shock combinations I varied the capital-labor ratio in the (8.5 : 10.5) interval with steps of 0.5 length and the technical parameter α^{hs} in the interval (0.125 : 0.145) with steps of 0.005 length. Results are shown in figure 1.2, they always refer to averages of shock-baseline ratios calculated in the period following the shock.
The first row of figure 1.2 shows the responses of the skill employment shares to the shocks. We observe that the middle-skill employment share is always below 1, whereas the low and high-skill shares are always above 1 independently on the particular shock implemented. We can therefore conclude that the emergence of job polarization is robust to the intensity of the shock both in the labor-saving and skill-bias dimensions. Moreover, we observe that the growth in low and high-skill employment share is more pronounced the more skill-biased and the more labor-saving is the shock. By inspection of panel (c) we can also confirm the mechanism driving the growth in the low-skill employment share. Indeed, we observe that the service share of employment is constantly above 1 and that follow the same pattern as the low-skilled employment share.
Panel (f) depicts the response of the relative price, defined as the ratio of the average consumption good price with respect to the average service price. The pattern reflects the unit cost of production response to the technological shock. The shock-baseline ratio is consistently below 1, meaning that the efficiency gains are translated into consumption good prices. Moreover, the extent of such variation is increasing with the degree of labor-saving and decreasing in the degree of skill-bias of the shock. This reflects that unit costs of production are obviously decreasing in the capital-labor ratio and increasing in the high-skill technical parameter.
The sensitivity analysis also confirms the positive effect of the technological shock on employment, as it is observable in figure (e).
Finally, we observe a positive effect on real wages for all class of workers, but more pronounced for high skilled ones. Interestingly middle and low wages are decreasing the skill-bias dimension of the shock, but increasing in the capital ratio. This suggest that the skill bias tends to favour high skill workers, whereas the efficiency gains brought about by the labor saving character of the shock are distributed across skill groups.

Figure 1.2: Sensitivity

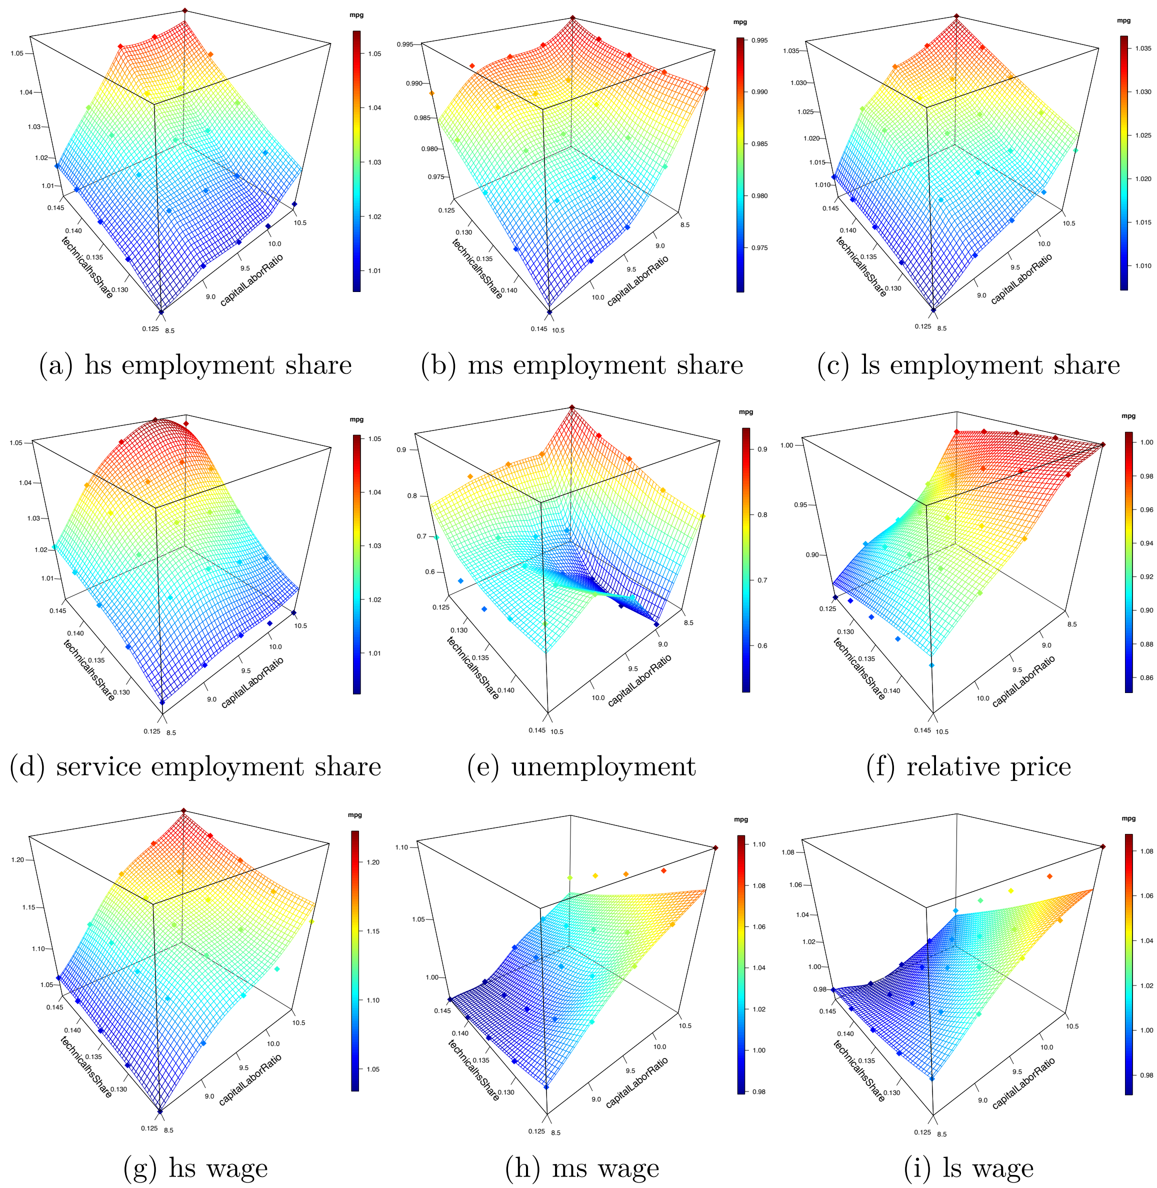

(a) hs employment share (b) ms employment share (c) ls employment share

(d) service employment share (e) unemployment (f) relative price

(g) hs wage (h) ms wage (i) ls wage

1.6 Conclusions

In this paper I laid down a macroeconomic framework intended to study issues related to technological change, structural change, and labor market adjustments in terms of skills demand, employment flows across productive sectors, and aggregate employment. The model presented here is intended to be simple and flexible so to accommodate for future extensions: a potentially interesting extension may be to model a more realistic network of productive sectors. This would allow to better investigate how a technological shock affecting one sector propagate throughout the economy and in particular how labour flows across sectors in response to the shock. Another possible extension is a more complex definition of workers skills and skills content of productive tasks. A modelling strategy of this type may be suitable to answer research questions more specific in terms of which skills and how are affected by

technological change, moreover it may allow for a more precise empirical calibration of the technical parameters.
The model could also prove to be useful to study education and investment policies, such cases would probably require to devise endogenous mechanisms for individual skills evolution and technological innovations affecting productivity and skills demand.

In this paper I also provide an application of the model intended to study the relation between automation, job polarization, and structural change. The model shows a simple mechanism according to a labor-saving and skill-biased technological shock occurring in the manufactory sector prompts a change in the income distribution, which, in turns trigger a demand-led structural change favouring growth in the service sector. The "general equilibrium" effect of the automation shock is job polarization, where the growth of high-skilled employment share is concentrated in the manufactory sector, while the growth in the low-skilled employment share is concentrated in the service sector. This is the central results of the application proposed in this paper, it is robust to a wide range of shock intensities, both in the labor saving and skill-bias dimension, and it is consistent with empirical findings gathered from different strands of literature.
Another results is the overall positive effect on aggregate employment of automation, this is also robust to a wide range of shocks, even though its empirical underpinning is less clear, since at the moment empirical studies on technological unemployment and robots are not conclusive.

1.7 Appendix

Figure 1.3: Auto-correlations

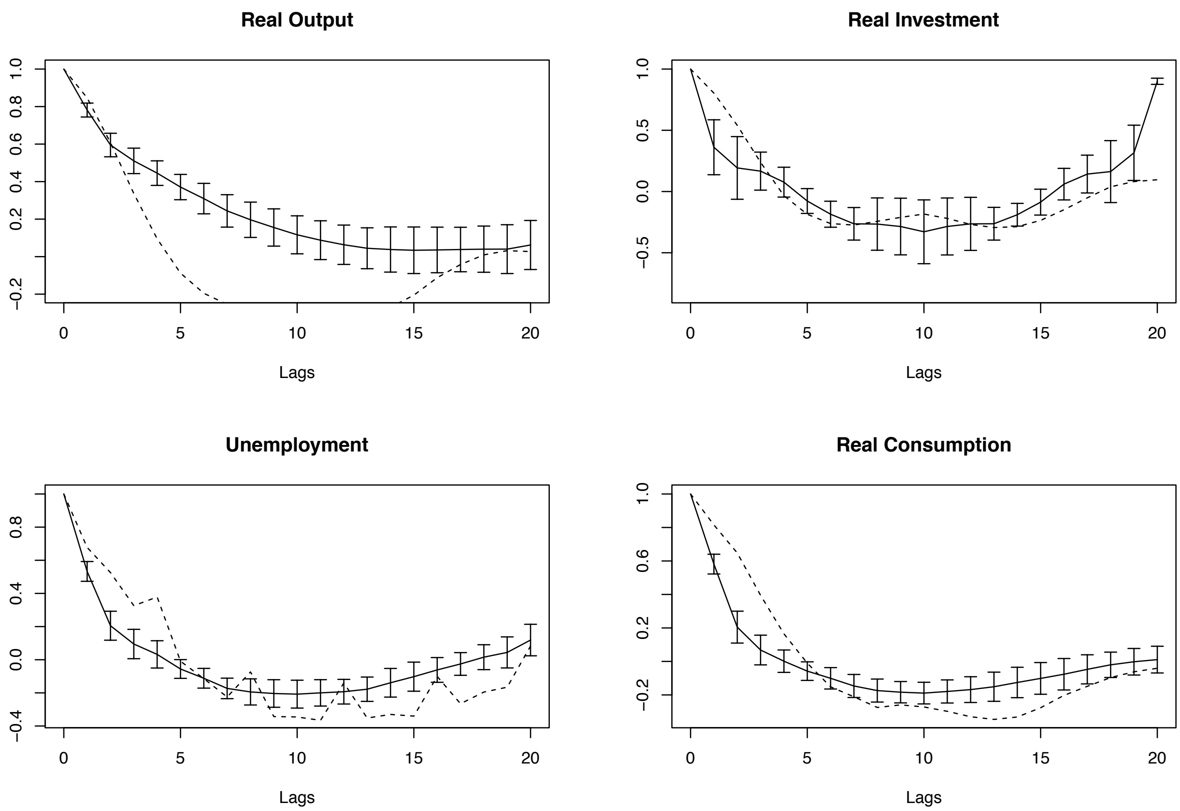

Dotted line: US data
Continuous line: simulated data

Figure 1.4: Cross-correlations

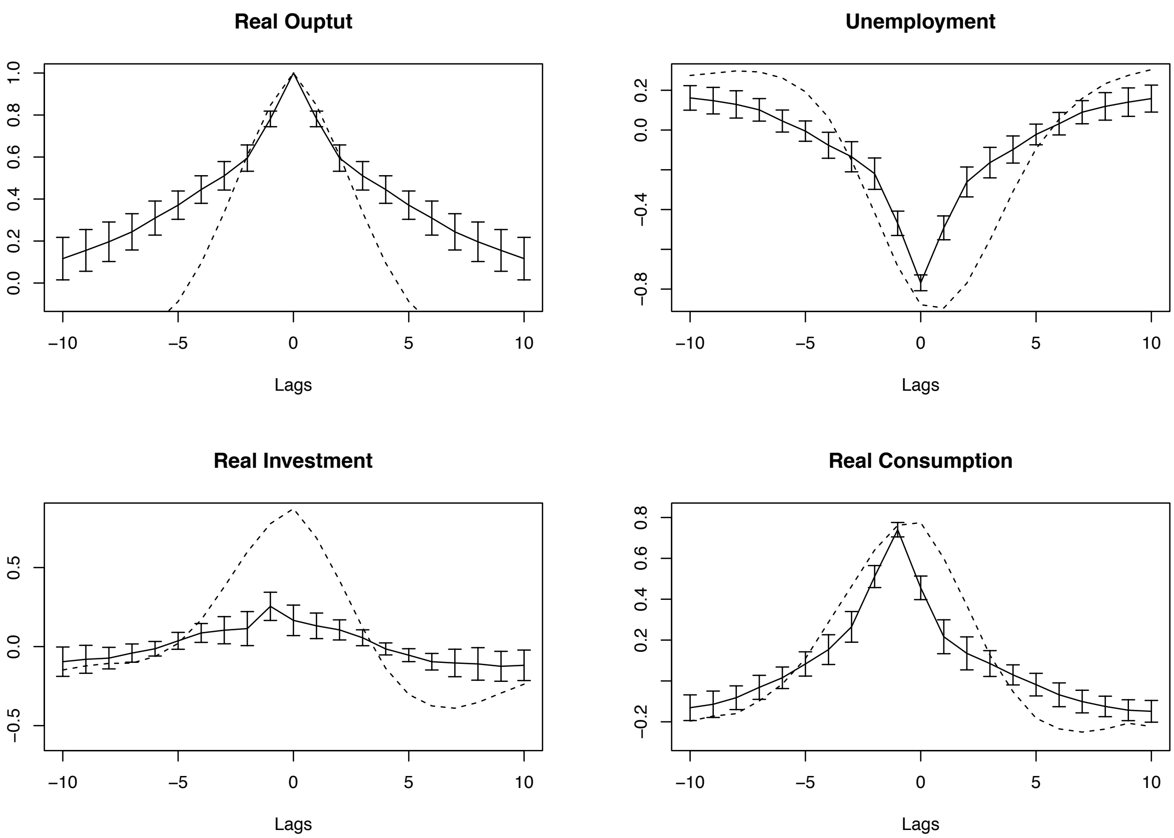

Dotted line: US data
Continuous line: simulated data

Figure 1.5: Cyclical components

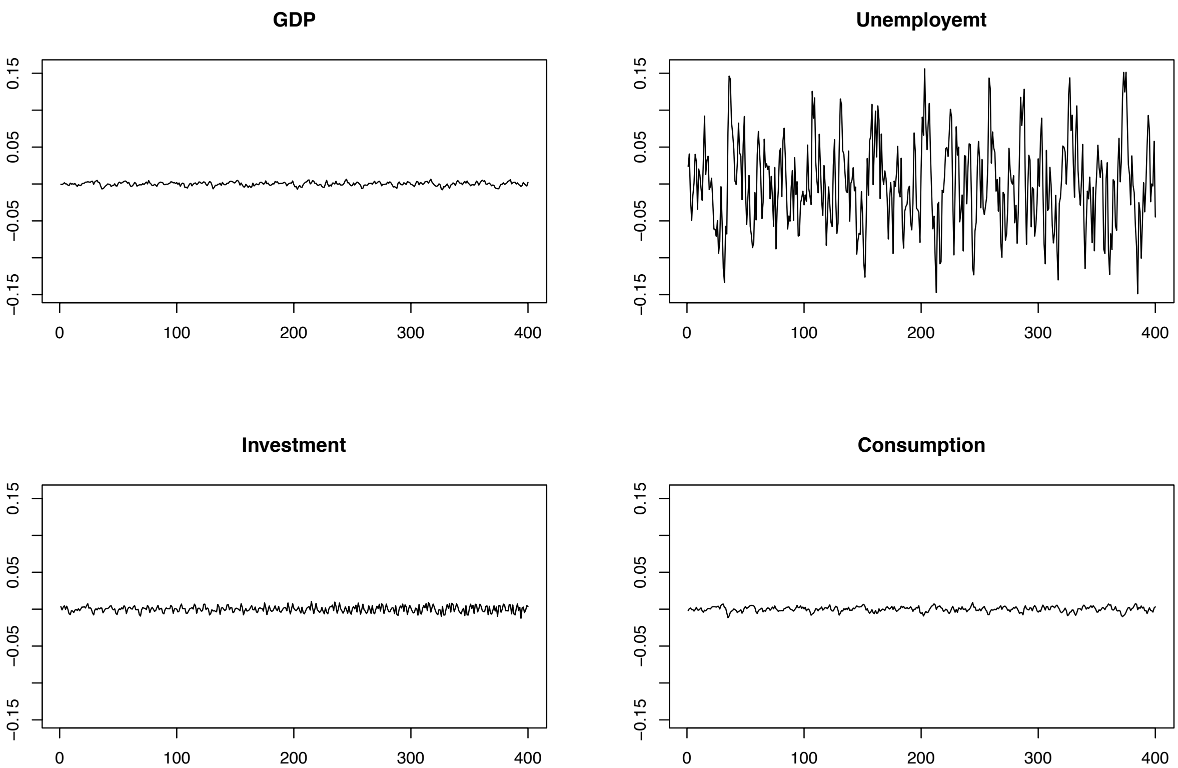

Figure 1.6: Baseline Dynamics

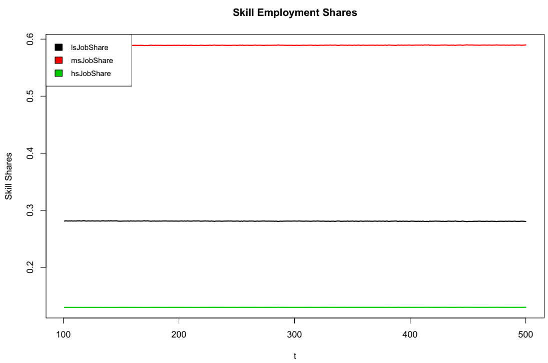

(a)

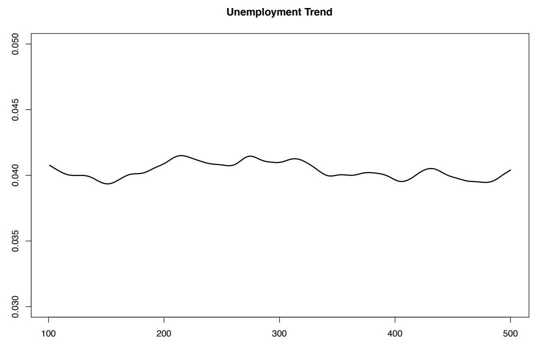

(b)

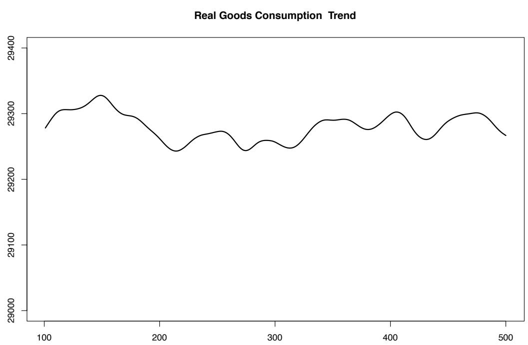

(c)

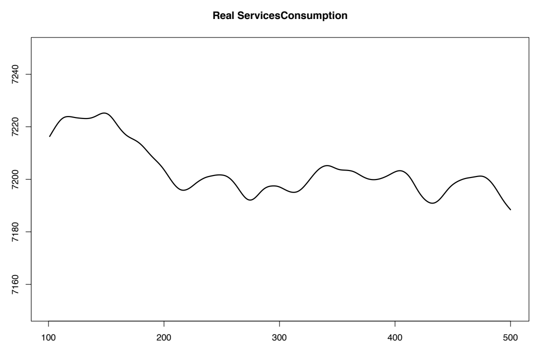

(d)

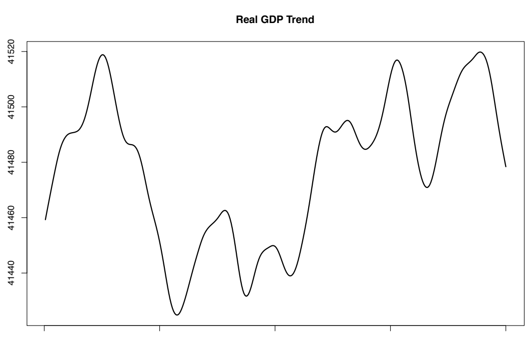

(e)

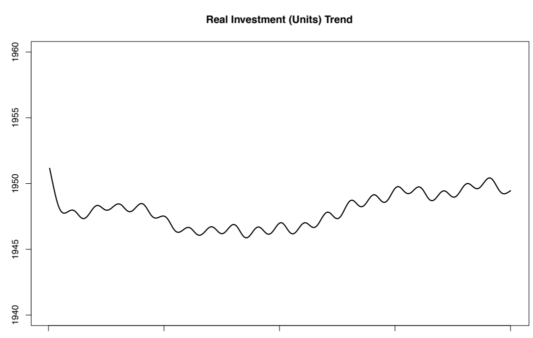

(f)

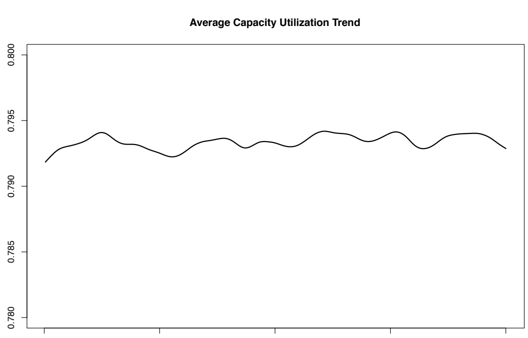

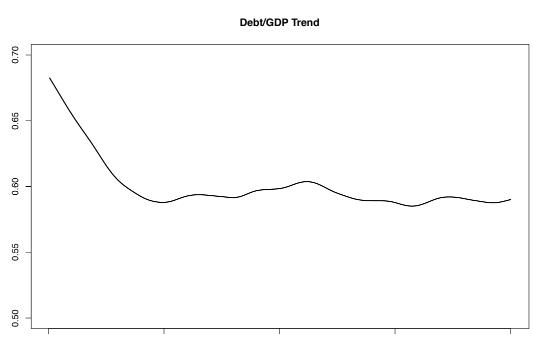

(g)

(h)

Chapter 2

Modelling Expectations and Learning in ABMs

Abstract

This is a methodological paper investigating how expectations should be modelled in ABMs.
I will argue that in ABMs rational expectations *a lá* Muth are neither applicable, nor needed. Moreover, what is sufficient to achieve in ABMs is *collective* rationality, which simply implies that the aggregate mean forecasting error is on average zero, i.e. the economy as a whole is not systematically mistaken in making predictions.
Therefore I will study if in ABMs it is possible to achieve *collective* rationality, i.e. aggregate unbiased expectations. In order to do so, I will experiment various expectation formation mechanisms coupled with a learning algorithm. Results suggest that, under certain condition, it is possible to achieve unbiased expectations at the aggregate level, moreover a simple learning algorithm is enough to sensibly improve agents' forecasting performances.

2.1 Introduction

In the last twenty years a considerable number of Agent Based models (ABMs) have been designed for the purpose of macroeconomic policy analysis. Alongside their proliferation, ABMs attracted significant interest in academic circles and among policy makers to the point that, today, some regard this new tool as a potential complement to the workhorse of modern macroeconomics, i.e. DSGE models[1]. To understand the reasons behind this trend one should appreciate the advantages embedded in ABMs *vis-a-vis* DSGEs, which are many, or at least as many as the disadvantages they bear, and although a thorough comparison of the two methodologies falls beyond the scope of this paper, three aspects in which ABMs seem to be better suited are worth mentioning: *heterogeneity*, *bounded rationality*, and *interaction*.

It is widely acknowledged that heterogeneity and bounded rationality are relevant features in macroeconomics, indeed DSGE practitioners had already moved forward from the perfect rational representative agent framework, for example integrating households heterogeneity and some sort of limited rationality (see Kaplan et al. (2018) and Gabaix (2016)). However, DSGE models are not quite flexible and any departure from the core framework comes with many caveats and limitations. For example, at the best of my knowledge the two aforementioned research programs are kept distinct given the difficulties of integrating both features in the same model. Moreover, the degree of heterogeneity as well as departures from full rationality are still quite modest. Keeping in mind that such limitations may be overcome in the future, it is fair to say that - at the current stage - ABMs do entail a comparative advantage in both respects. In fact, ABMs easily accommodate for a wide span of heterogeneities for example in expectations, wealth, income, skills, firm size, and broadly speaking behavioural rules. These heterogeneities can be imposed by the modeller, or can emerge endogenously as the model is simulated, and they can all coexist in the very same model.

A similar discussion holds true for what concerns bounded rationality. The main difference between the two approaches is that in DSGEs any limitation to rationality is in fact a narrow departure from perfect rationality. On the other hand, ABMs are flexible enough to allow for bounded rationality in a more genuine sense, that is to say they allow to follow closely the notion of *procedural* rationality introduced by Simon (1976).

As far as interaction is concerned instead, the gap between ABMs and DSGEs is more fundamental and seemingly impossible to close. Kirman (1992) in his famous critique of the representative agents highlighted the pivotal role of interactions occurring among heterogenous agents and in particular its local character, as local phenomena can propagate globally and can therefore explain aggregate dynamics. The point has been recently further discussed by Caverzasi and Russo (2018) who point out that the main difference between DSGEs and ABMs is that the former assumes a complete network of agents, whereas the latter allows for a range of different topologies. Caverzasi and Russo (2018) focus in particular on the implications that networked interaction among agents have for studying financial crises and bankruptcy cascades[2], but, as we will see later on, agents direct interactions have interesting implications in many other applications, for example strategies selection.

[1] For an interesting discussion on this point see Fagiolo and Roventini (2016) and Haldane and Turrell (2018)

[2] see Delli Gatti et al. (2010b) for an example

Despite the many advantages ABMs brings to the field of macroeconomics, there are still many aspects that need to be improved and further investigated. One of such aspects is how expectations should be modelled in ABMs and according to which rationality criterion. Let me start by saying that rational expectations (RE, hereafter) as envisaged by Muth (1961) are very difficult, if not impossible, to implement in ABMs. This is because RE imply that agents internalise the *true* underling model of the economy and therefore they know the distributions of economic variables conditional to some information set. But ABMs are complex adaptive systems, characterised by non-linearities, endogenous dynamics, and fundamental uncertainty. The very notion of "underling model" in this a context is very difficult to formalise, let alone to model as part of agents' information sets[3].
Thus, RE are not well suited in an ABM type of environment, neither strictly needed in more general terms. Indeed, I would argue, what we really need in modelling expectations are mechanism consistent with the notion of *collective* rationality, which simply requires *aggregate* unbiasedness. This implies that the across agents average forecasting error must have zero mean, and therefore that the economy as a whole does not make systematic mistakes in forecasting, even though its micro entities might. I would also argue that forecasting rules consistent with the notion of *collective* rationality are to some extent robust to the Lucas critique[4], indeed if expectations are on average correct, especially in the face of policy shocks, than it means that the economy as a whole is able to adjust its forecasting in order to accommodate to changes in the environment.
Therefore, my goal is to study whether it is possible to obtain *aggregate* unbiased expectations in an ABM framework. Since typically a macro ABM does not have a closed form solution, I will rely on extensive computer simulations in order to assess the performances of different expectation formation mechanisms in different contexts. I will do so in two macro environments: (i) a very simple and stylised model in which agents try to forecast a stationary variable and (ii) a full fledged macro ABM in which agents try to forecast a trended variable. In case (i) I designed a simple model in which a central bank set the interest rate and households try to forecast the one-step ahed inflation rate. In this context I will assess the performances of different expectation formation mechanisms across policy regimes and policy shocks. In case (ii) I will employ the model put forward in Caiani et al. (2016) augmented by technological innovation as in Caiani et al. (2019). This is a full fledged macro ABM which I will use as a laboratory to assess different expectation formation mechanisms applied to firms trying to forecast future sales.
The expectation rules employed in the following exercise are: naive expectations, where the expected value of a variable equals its past realisation, adaptive expectations, trend following, and social learning in the form of a very simple genetic algorithm. Moreover, I will also employ hybrid expectations in which adaptive expectations and trend following are combined with learning. In the spirit of Palestrini and Gallegati (2015), this allows to update the adaptive parameter in order to correct biases which might arise otherwise.

The rest of the paper is organised as follow: section 2 will revise the relevant literature; section 3 will present in detail the expectations formation mechanisms employed in the paper; section 4 will host the first experiment. It will describe the model and discuss the results; section 5 will host the second experiment. It will briefly describe the model and present the

[3] An interesting attempt in this direction is provided by Salle (2015), who allows agents to form mental models of the economy using neural network algorithms.

[4] See Lucas Jr (1976)

results; section 6 will suggest future developments and conclude.

2.2 Related Literature

The paper touches upon the Lucas' critique debate initiated by the classical contributions of Lucas Jr (1976) and chiefly summarised to the present days by Sergi (2018). As well known, the debate quickly took an empirical turn[5], however a handful of theoretical papers can be found, in particular those criticising the way modern macroeconomics has been dealt with the critique, see for example Kirman (1992), Altissimo et al. (2002), and Marcellino and Salmon (2002). This paper is exclusively concerned with the theoretical side of the debate and given the technical difficulties of implementing rational expectations in ABMs, this paper naturally refers to the vast literature on learning in macroeconomics. Modern macroeconomics has used different types of learning essentially in two contexts: selecting the best equilibrium in multiple equilibria settings and studying the learning convergence to rational expectation equilibrium[6]. On the other hand, within the ABM research program learning has been used as a necessary tool to model agents' forecasting with an agnostic view about the equilibrium being attained. Applications of learning in ABMs are many, for example social learning in the form of Genetic Algorithms (GA hereafter) has found fruitful applications: Delli Gatti et al. (2005) model a central bank trying to learn the optimal Taylor rule parameters by means of a GA, also in the context of monetary policy Salle et al. (2013) design a GA applied to both households and firms trying to learn optimal decision rules. Social learning has also been applied to ABM in a General Equilibrium setting as in Gintis (2007) and Salle et al. (2017). A simpler learning algorithm, but with some similarities to GA, is the *blanketing shotgun process* implemented by Seppecher et al. (2019), where firms are assigned behavioural rules randomly and the selection of successful rules is left to market forces. Another mechanism which has drawn attention in the ABM community is reinforcement learning: Catullo et al. (2015) apply it in a financial accelerator model where banks try to learn the optimal leverage ratio; Dosi et al. (2017a) model switching among a fixed set of expectation formation mechanism using a sort of replicator dynamics, moreover they employ *recursive least square* in order to update adaptive parameters in otherwise fixed expectation formation mechanisms. Finally, Catullo et al. (2019) use a more sophisticated machine learning algorithm in a macro model where firms try to forecast future sales.
The type of learning this paper is concerned with is GA, which since the seminal contribution of Holland (1970) and Arifovic (1991) has found many application in economics (for an exhaustive survey see Arifovic (2000)): common application of GA can be found in the context of the Cobweb model, where firms try to learn optimal prices and quantities, see Arifovic (1994), Franke (1998), and Dawid and Kopel (1998). More recently, models dealing with monetary economics has made use of GA algorithms, both in OLG and NK frameworks, as in Arifovic (1995) and Arifovic et al. (2012).
Finally, the paper which mostly relates to this one is certainly Dosi et al. (2017a) who also experiment and assess different expectation formation mechanisms. The main difference be-

[5]For example see Alogoskoufis and Smith (1991), Blanchard (1984), Chang et al. (2010), Cuthbertson and Taylor (1990), Engle and Hendry (1993), Favero and Hendry (1992), Fischer (1989), Hendry (1988), and Taylor (1984)

[6]For example see Evans and McGough (2005), Giannitsarou (2003), Honkapohja and Mitra (2004), and Honkapohja and Mitra (2006)

tween this contribution and Dosi et al. (2017a) are essentially three: (i) in this paper I am not only concerned with the relative performances of various expectation formation mechanisms, but also with the absolute performance of each expectation formation mechanism I experiment with; (ii) I use a different learning algorithm, i.e. a GA; (iii) most importantly, my results differs substantially form those of Dosi et al. (2017a). In particular, I find that learning improves forecasting performances with respect to less dynamic rules.

2.3 Expectation Formation Mechanisms

This section gives full account of the expectations formation mechanisms employed in the following two macro experiments. Four forecasting strategies are described: Naive, Adaptive, and Trend Following expectations. Moreover, the learning algorithm employed throughout the paper is discussed. I shall already mention that not each and any expectation formation mechanism is employed in both macro experiments. This is because some rules are not well suited in some specific environments, for example, it would not be very interesting to experiment with trend following expectations when agents try to forecast a stationary variable.

2.3.1 Naive Expectations

Naive expectations are the simplest: they assume that the expected future value of a variable x equals its past realisation:

$$x_t^e = x_{t-1} \tag{2.1}$$

2.3.2 Adaptive Expectations

Adaptive expectations represent a step forward in terms of rationality, they use past forecasting errors in order to adjust past expectations:

$$x_t^e = x_{t-1}^e + \lambda \left(x_{t-1} - x_{t-1}^e \right) \tag{2.2}$$

Where λ is the exogenous time invariant adjustment parameter.
Therefore, expectations at t equal expectations at t-1 adjusted by the weighted past forecasting error.

2.3.3 Trend Following Expectations

In this case agents impose some more structure in their mental model of the economy. They assume that past trends will be observed in the future and therefore they form expectations as follows:

$$x_t^e = x_{t-1} + \lambda(x_{t-1} - x_{t-2}) \tag{2.3}$$

Where λ is the exogenous time invariant adjustment parameter.

2.3.4 Social Learning

Social learning involves exchanging ideas among agents through communication or imitation and it has a twofold effects in our framework: (i)it is inherently dynamic, so it allows agents to continuously adapt in an ever-evolving environment; (ii) it allows for heterogeneity among agents.

Social learning is implemented in different fashions throughout the paper, but it always follows the same logic, which is the one of an extremely simple genetic algorithm, using two genetic operators: *tournament* and *mutation*. *Tournament* allows for the spreading of successful rules: in each period of the simulation two agents are randomly paired and their forecasting strategies compared using a fitness function. Thereafter, the agent endowed with the relative less performing expectation rule copies the relative more successful one, whereas the other agent is left unaffected. *Mutation* allows for new rules to be discovered and later on compared with rules already present in the genetic pool. At each step of the simulation an agent is drawn from the population with probability Pr_m. Once the agent is effectively drawn, she randomly picks a new rule from the population of all conceivable rules and adopts it as her new forecasting strategy.

2.3.4.1 Augmented Adaptive Expectations

As previously stated, adaptive expectations allow to internalise past forecasting errors in future expectations. However, the weight assigned to past forecasting errors is arbitrarily chosen. To overcome this shortcoming we employ a mixed expectation formation mechanism, which integrates social learning in an otherwise simple adaptive expectations framework. The object of learning is the adaptive parameter λ, i.e. the weight assigned to past forecasting errors, and the procedure the same as the one presented above: expectation rules are subject to tournament and mutation operators. The population of possible rules is $\lambda \in [0, 1]$. In case of mutation the new λ is a random draw from a uniform distribution bounded between 0 and 1. Therefore, the rule can be written as:

$$x_t^e = x_{t-1}^e + \lambda_t^h \left(x_{t-1} - x_{t-1}^e\right) \tag{2.4}$$

Where λ_t^h now specifically refers to time t and agent h.

2.3.4.2 Augmented Trend Following Expectations

The same procedure implemented for the augmented adaptive expectation is applied in the case of trend following. Also in this case agents try to learn the best possible adaptive parameter λ. The only difference is that the population of conceivable rules is defined over the range $\lambda \in [0, 1.5]$. Clearly, augmented trend following expectations are defined as:

$$x_t^e = x_{t-1} + \lambda_t^h (x_{t-1} - x_{t-2}) \tag{2.5}$$

2.4 Case I: Forecasting a Stationary Variable

For this first experiment I designed a bare-bone ABM in which households try to forecast the one-step-ahed inflation rate, which is obviously a stationary variable. The forecasting rules I will consider in this experiment are: (i) naive expectations, (ii) adaptive expectations, (iii)

augmented adaptive expectations, and (iv) augmented anchored expectations.
My primary goal is to understand whether these forecasting rules are consistent with the notion of *collective* rationality, which in turn requires *aggregate unbiasedness*. I will therefore assess each forecasting rule based on its observed *aggregate* mean forecasting error: the closer to zero, the more collectively rational the rule will be evaluated.

Before describing the model and presenting the results, two preliminary remarks are in order: (i) the inflation rate is influenced by the overall economic activity and the central bank, which implements a single-mandate Taylor rule aiming at stabilising inflation. Modelling the central bank allows to study the forecasting rules performances across a variety of policy regimes and shocks, enriching the analysis and providing robustness checks to the baseline results; (ii) the aim of the simple model implemented hereafter is exclusively to provide a laboratory to perform controlled experiments with different expectation formation mechanisms. The model indeed, does not provide any meaningful economic insights and it is not intended for anything different than the purpose already made clear. The only model requirement is stability, so to assure that results are not biased by extreme dynamics.

2.4.1 The model

The model is an extremely simple stock-flow-consistent (SFC) ABM [7]and it is composed of five types of agents:

- households
- consumption firm
- commercial bank
- central bank
- government

Only households are modelled as a multitude of heterogeneous interactive agents, making effectively the model an agent-based.
Agents interact on 5 markets:

- *consumption market* where households buy goods from the consumption firm
- *labor market* where the consumption firm hires households
- *credit market* where the firm demands credit to the bank
- *deposit market* where the bank collects deposits form households and the firm
- *bond market* where the government sells public bond to the bank.

Beside market interactions, the government taxes the private sector, pays unemployment benefits and emits bonds to cover possible deficits. The central bank sets the base interest rate following an inflation stabilisation single-mandate Taylor rule and buys bonds which are not absorbed by the bank.

[7]See Caiani et al. (2016) for a general discussion about the importance to impose a SFC structure to ABMs

2.4.1.1 Sequence of Events

The model is simulated step-by-step within each simulation period following a precise sequence of events:

1. *Policy rate setting* The central bank recovers past inflation and sets the current period policy rate through a Taylor Rule.

2. *Production planning* The firm sets its desired production, computed as expected sales plus planned inventories. Thereafter, it computes labour demand, given its technology and desired production.

3. *Labour market* Households update their reservation wage and supply labour on the labour market. Consumption firm hires needed labour if available.

4. *Production, pricing and credit demand* Once workers are hired, the consumption firm can produce and compute its unit cost of production. It therefore sets its price as a mark-up over unit costs and it is also able to set its credit demand if internal resources are not enough to pay any disbursement due.

5. *Consumption, credit market, wages* Households set their consumption demand, the bank grant loans to the consumption firm, and wages are paid. Consumption market opens and households try to satisfy their desired consumption subject to product availability and their own resource constraint.

6. *Taxes, dole and bonds* Government collects taxes, pays unemployment benefits, and pays interests on bonds. If deficit is positive it emits new bonds. The Bank demands bonds or cash advances depending on its resources. Bonds market opens and bonds are sold to the bank. If bond supply exceeds demand, the central bank steps in to buy the difference.

7. *Dividends* If profits are positive, bank and consumption firm pay dividends to households.

2.4.1.2 Agents

2.4.1.2.1 Households

Households engage in two activities, working and consuming, and participate in three markets: consumption, labour, and deposit market. Moreover, they own the bank and the firm. Household h ownership share is given by her share of total wealth defined as:

$$WS_t^h = \frac{D_t^h}{\sum_{i=1}^{H} D_t^i} \quad \text{with} \quad i \in \Phi_H \tag{2.6}$$

Where WS_t^h is the wealth share of households h at time t, D_t^i is the deposit amount of household i at time t, Φ_H is the set of households, and H is the dimension of Φ_H.
Each household receive dividends proportionally to her ownership share.

wage setting

The wage setting equation is borrowed from Caiani et al. (2016). It is a simple heuristic which tries to catch bargaining power swinging between workers and firms as labour market conditions mutate: when the economy is strong and job opportunities are abundant, workers revise up their demanded wage. On the contrary, when workers encounter difficulties in finding a job, the reservation wage is reduced.
Each worker proxies labour market conditions by her own employment status over a given time span. The wage setting mechanism is expressed as:

$$w_t^{h,d} = \begin{cases} w_{t-1}^{h,d}(1 - FN_{h,t}) & if \quad \sum_{n=1}^{4} u_{t-n}^h \leq 2 \\ \\ w_{t-1}^{h,d}(1 + FN_{h,t}) & if \quad \sum_{n=1}^{4} u_{t-n}^h > 2 \end{cases} \tag{2.7}$$

Where $w_t^{h,d}$ is the demanded wage of household h at time t, FN is a random draw from a folded normal distribution defined over the parameters (μ_{FN}, σ_{FN}), u_t^h is a dummy variable taking value 1 if household h is unemployed at time t and zero otherwise.
Therefore, if in the last four periods household h has been unemployed for at least two periods, the reservation wage is scaled down by a random amount. Otherwise, the reservation wage is revised up by the same token.

consumption function

The consumption function is a simplified version of the one employed in Bouchaud et al. (2017). It internalises an Euler equation logic, insofar it picks up the inter-temporal substitution effect at work as the real interest rate changes. This is achieved by defining the propensity to consume out of income and wealth as a function of the difference between nominal interest rate and inflation. The formulation employed in this paper is given by:

$$C_t^{h,b} = c_t^h(NI_t^h + D_t^h) \tag{2.8}$$

Where NI_t^h is household h's net income at time t, D_t^h is h's deposit amount at time t, and c_t^h is h's propensity to consume, which is in turn defined as:

$$c_t^h = c_0 \left[1 + \alpha_c \left(\pi_t^e - i_t^d\right)\right] \quad \text{with} \quad c_t^h \in [0, 1] \tag{2.9}$$

Where c_0 is the normal propensity to consume, i.e. where the interest rate and inflation expectations exactly cancels out, α_c is the sensitivity with respect to the real interest rate, π_t^e is expected inflation[8], and i_t^d is the deposit interest rate. Note that c_t^h is not naturally bounded between 0 and 1, however anytime it exceeds 1 it is set equal to 1, whereas in case in turns out to be negative it is set to 0. The economic interpretation for such formulation is straightforward: when the deposit rate grows relative to the expected inflation rate, it is rational to delay consumption and increase savings. On the other hand, if inflation is expected to be high relative to the deposit rate, it is rational to anticipate consumption in order to minimise the impact of future high prices on individual welfare..
Finally, equation (2.9) defines the main (direct) transmission channel for monetary policy.

[8]Inflation expectations are defined according to different rules, see subsection (2.4.1.3).

Net Income and taxes

At the end of each period households calculate their gross income, which is given by the sum of received wage, dividends and interests on deposits:

$$GI_t^h = w_t^h + div_t^h + i_t^D D_t^h \tag{2.10}$$

Where GI_t^h is h's gross income at time t, w_t^h is h's wage at time t if h is employed or unemployment benefit otherwise, div^h are h's received dividends at time t, i_t^D is the interest rate on deposit at time t, and D_t^h is h's deposit amount at time t.
Households pay taxes on gross income. The tax rate is an exogenous and time independent $\tau \in [0, 1]$, therefore net income is given by:

$$NI_t^h = (1 - \tau)(w_t^h + div_t^h + i_t^D D_t^h) \tag{2.11}$$

2.4.1.2.2 Consumption Firm

Firm c plans production in order to meet actual demand, sets price as a mark-up over unit costs of production, pays dividends to households when profits, and pays taxes on profits. Firm c also interacts with workers and the bank on three markets: it sells good to households on the good market, hires workers on the labour market, and asks for loans to the bank on the credit market.

planned production, labor demand, and price setting

Firm's desired production is the sum of expected sales[9] plus planned inventories. Inventories serve as a buffer-stock against unforeseen demand and are set as a fixed proportion of expected sales. Therefore, desired production is given by:

$$y_t^{d,c} = (1 + v)s_t^{e,c} - inv_{t-1}^c \tag{2.12}$$

Where $y_t^{d,c}$ is desired output at time t, $s_t^{e,c}$ are expected sales at time t, v is the fixed planned inventories/expected sales ratio, and inv_{t-1}^c are inventories left from the previous period.

I assume a single factor, constant return to scale production function:

$$y_t^c = \gamma N_t^c \tag{2.13}$$

Where y_t^c is actual production at time t, γ is the time invariant labour productivity, and N_t^c is the number of workers employed. Therefore, once desired production is calculated, firm's labour demand is simply given by:

$$N_t^{d,c} = \frac{y_t^{d,c}}{\gamma} \tag{2.14}$$

Finally, price is set as a fixed mark-up over actual unit costs of production:

$$p_t^c = (1 + \mu^p)\frac{\sum_{i=1}^{N_t^c} w_i}{y_t^c} \tag{2.15}$$

[9]We assume that expected sales follow and adaptive expectations scheme throughout the paper, i.e.

$$s_t^{e,c} = s_{t-1}^{e,c} + \lambda_s(s_{t-1}^c - s_{t-1}^{e,c})$$

Where μ^p is the exogenous time invariant mark-up[10], L_t^c is the number of workers hired by the firm, and w_i is worker i's wage.

credit demand

It is assumed a pecking order approach to financial requirements, which amounts to say that internal resources are always preferred to costly debt when a disbursement needs to be covered. Pecking order implies that firm c relies on debt only when its deposits are exhausted, or to put it in another way, when it faces a liquidity constraint. Therefore, firm's credit demand is given by:

$$L_t^{d,c} = NPD_t^{e,c} - D_t^c \tag{2.16}$$

Where, $L_t^{d,c}$ is the loan demand of firm c at time t, $NPD_t^{e,c}$ are the expected net payments due, and D_t^c is the total amount of deposits owned by the firm.
In defining $NPD_t^{e,c}$, it is pivotal to consider the sequence of events: wages are set to be paid before the market for consumption goods opens, i.e. before revenues are cashed in by the firm. Such lag between the moment wage payments are due and profits are realised, it is likely to push the firm in a liquidity constraint situation. To escape it, I am going to assume that the expected payments due are defined as:

$$NPD_t^{e,c} = \sum_{i=1}^{L_t} w_i + max\left\{0, db_t - ir_t + div_t^e - rev_t^e + tax_t^e\right\} \tag{2.17}$$

Where db_t is the debt burden, i.e. interest plus capital payments due, ir_t are received interests on deposit, div_t^e are expected dividends to pay, rev_t^e are expected revenues, and tax_t^e are expected taxes.
Note that the wage bill enters equation (2.14) regardless whether net profits are expected to cover it, even partially, or not.

Dividends, taxes, and bankruptcy

Firm c gross profits are given by nominal sales and interest received, minus labor costs and debt interest payments, plus investments in inventories:

$$GP_t^c = p_t^c s_t^c + i_t^D D_t^c - \sum_{i \in N_t^c} w_t^i - DebtInt_t^c + \Delta NomInventories_t^c \tag{2.18}$$

Where GP_t^c are c's gross profits at time t, s_t^c are c's realised sales in t, D_t^c are c's deposits at time t, $DebtInt_t^c$ are debt interest payment due at time t[11], and $\Delta NomInventories_t^c$ is the change in nominal inventories between t and $t-1$.
The government lays taxes on profits equal to a share $\tau \in [0, 1]$ over gross profits, so that net profits turns out to be:

$$NP_t^c = (1-\tau)\left(p_t^c s_t^c + i_t^D D_t^c - \sum_{i \in N_t^c} w_t^i - DebtInt_t^c + \Delta NomInventories_t^c\right) \tag{2.19}$$

[10] It is assumed a fixed and exogenous mark-up, which in general does not have to be the case and in fact it is not in the second experiment, where a full fledged model is used. Here, allowing for an endogenous mark-up would introduce a further monetary policy transmission channel, which we do not wish to include in the present model.

[11] See paragraph (2.4.1.2.3)

Note that it is assumed the same τ for income and profits taxes.
Profits are redistributed as a fixed proportion of positive profits, i.e.:

$$Div_t^c = max(0, \beta^c NP_t^c) \tag{2.20}$$

Where Div_t^c are c's total redistributed profits and β^c is the exogenous time invariant rate of c's redistributed profits.

Firms c's maybe be unable to meet wages, debt, or taxes payments. In such occasions c is not force to bankruptcy, but:

- In case c is not able to meet its wage payments obligations, each worker h's wage is reduced as it follows:

$$w_t^{A,h} = \frac{D_t^c}{\sum_{i \in N_t^c} w_t^i} w_t^h \tag{2.21}$$

  Where $w_t^{A,h}$ is h's actual waged received at time t and w_t^i is i's wage bargained at the beginning of time t.

- In case c is not able to meet it debt payments obligations, all c's debt is simply rolled one period ahed

- In case c is not able to pay taxes in full, it only pays an amount of taxes equal to its residual internal resources.

labour and consumption goods market

Since the consumption good sector has been treated as one aggregated entity, labour and consumption goods market turn out to be quite simple: for what concerns the labour market, unemployed workers simply post their reservation wages. The firm sorts unemployed workers according to their reservation wages and starts hiring from the cheaper worker onwards, until its demand is satisfied or there are no unemployed workers available.
In the consumption markets households arrive in a random order and demand a number of goods equal to their consumption budget, i.e. their nominal demand specified in equation (7), divided by the price offered by the consumption firm. The market closes when either firm's inventories are exhausted or when all consumers are being served.

2.4.1.2.3 Bank

The bank provides deposits and loans, purchases government's bonds, demands cash advances to the central bank, and redistribute profits.

Deposits

The bank accept any amount of deposits from households and firm, and it sets the interest rate on deposits as a mark-up over the policy rate:

$$i_t^D = (1 - \mu^D) i_t \tag{2.22}$$

Where i_t^D is the interest rate paid on deposits at time t, μ^D is the exogenous time invariant deposit interest mark-up, and i_t is the policy rate at time t.

Credit

The bank always accommodates firm's credit demand in full. It charges an interest on loan defined as a mark-up over the policy rate:

$$i_t^L = (1 + \mu^L)i_t \tag{2.23}$$

Where i_t^L is the interest rate charged on loans at time t, μ^L is the exogenous time invariant credit interest mark-up, and i_t is the policy rate at time t.
Each loan is issued with an original duration LL and at each point in time the creditor pays the interest on principle plus a share of the principle equal to the inverse of LL. Therefore, creditor payments at each point in time relative to a loan of amount L issued at time t is given by $i_t^L L + \frac{L}{LL}$. Clearly, the debt is extinguished when actual duration reaches zero.
Whenever the firm is not able to meet its loan payments obligations, firm's loans are simply rolled over one period ahed bearing their original interest rate.

Bonds and advances demand

I assume a fixed and exogenous bank liquidity ratio, therefore the bank ask for cash advancement to the central bank only in case it is not able to to meet the liquidity ratio requirement. Bank's cash advancement demand is therefore given by:

$$A_t^{d,b} = max(0, LR * D_t^b - R_t^b) \tag{2.24}$$

Where $A^{d,b}$ is the cash advancement demand at time t, LR is the liquidity ratio, D_t^b is the total amount of deposits detained by the bank, and R_t^b are reserves owned by the the bank. Bank's cash advancement demand is always satisfied by the central bank.

Similarly, the bank is willing to buy as many bonds as the liquidity-ratio constraint allows. Banks' bonds demand is therefore given by:

$$B_t^{d,b} = \frac{max(0, R_t^b - LR * D_t^b)}{P^b} \tag{2.25}$$

Where $B_t^{d,b}$ is the bank's bonds demand at time t and P^b is the bonds' price, which for simplicity is exogenously set to 1 throughout the simulation.

Dividends

Bank's total dividends are calculated as a fixed portion of after tax profits:

$$Div_t^b = max(0, \beta^b(1 - \tau)\Pi_t^b) \tag{2.26}$$

Where Div_t^b are total dividends paid by the bank at time t, β^b the exogenous time invariant bank's dividend rate, τ is time invariant exogenous the tax rate, and Π_t^b are bank's realised profits at time t.

Bankruptcy

The bank declares bankruptcy whenever its net wealth turns out to be negative, in such case it is alway bailed out by the government and it continues its operations.

2.4.1.2.4 Government and Central Bank

The government collects income and profit taxes from households, consumption firm and bank. The total amount due is simply calculated as the tax rate multiplied by gross income, in case of households, or by profits, when positive, in case of consumption firm and bank. The tax rate τ is assumed to be time invariant, exogenous, and equal for income and profits. The government also provides unemployment benefits, which are calculated as a fixed proportion θ of the average wage. When public deficit is positive, the government emits one period duration bonds up to the deficit value. Finally, the government bails-out the bank in case of bankruptcy.

The central bank provides advances to the bank when those are required, buys bonds supplied by the government exceeding the bank's demand and sets the interest rate following a single mandate Taylor rule defined as:

$$i_t = i^* + \rho(\pi_{t-1} - \pi^*) \tag{2.27}$$

where i^* is the "normal" interest rate exogenously defined and π^* is the central bank inflation target.
I define two policy regimes for the central bank, a *static* regime and a *dynamic* regime. In the former the Taylor rule parameter, ρ is set at the beginning of the simulation and never updated. In case of *dynamic* regime, I allow the central bank to experiment different ρ's as the simulation unfolds and to retain those that appears to be more effective in reaching the inflation target. The ρ-updating procedure follows: every four periods the central bank updates ρ with a mutation probability Pr_ρ. If actual updating takes place the new parameter is:

$$\rho^{new} = \rho^{old}(1 + N_{cb,t}) \tag{2.28}$$

Where $N_{cb,t}$ is a random draw from a normal distribution defined as $N_{cb} \sim N(\mu_{N_{cb}}, \sigma_{N_{cb}})$ Then, in the next four periods ρ is kept fixed and at the end of the this time window the new rules is assessed against the old one: if the average squared error, that is the distance between actual inflation and target inflation has been larger with the old rule, the old rule is definitely discarded, otherwise, the old rule is resumed and the new one discarded.

2.4.1.3 Expectations

In what follows, I will experiment with different specifications for the households one-step-ahed inflation expectations. Four forecasting strategies will be implemented: naive expectations, adaptive expectations, augmented adaptive expectations, and anchored augmented expectations. As already discussed in section (2.3) naive expectations are simply defined as:

$$\pi^e_t = \pi_{t-1} \tag{2.29}$$

Adaptive expectations are defined as:

$$\pi^e_t = \pi^e_{t-1} + \lambda^{ada}\left(\pi_{t-1} - \pi^e_{t-1}\right) \tag{2.30}$$

Augmented adaptive expectations are defined as:

$$\pi^e_t = \pi^e_{t-1} + \lambda^{ada,h}_t\left(\pi_{t-1} - \pi^e_{t-1}\right) \quad \text{with} \quad \lambda^{ada,h}_t \in [0,1], \forall h, t \tag{2.31}$$

Where $\lambda_t^{ada,h}$ is agent specific and endogenously evolve as time elapses, see paragraph (2.4.1.3.1). Anchored augmented expectations are a modified version of inflation expectations anchored to the central bank inflation target, which in turn are defined as:

$$\pi_t^e = \beta^{\pi}\pi_{t-1} + \beta^{\pi^*}\pi^* \tag{2.32}$$

Where for simplicity I assume: where:

$$\beta^{\pi^*} = 1 - \beta^{\pi} \quad \text{with} \quad \beta^{\pi} \in [0,1]$$

So, *augmented* simply signals that the parameter β^{π} is agent specific end endogenously evolve as time elapses following the algorithm described in paragraph (2.4.1.3.1).
Therefore, augmented anchored expectations are defined as:

$$\pi_t^e = \beta_t^{\pi,h}\pi_{t-1} + (1 - \beta_t^{\pi,h})\pi^* \tag{2.33}$$

2.4.1.3.1 GA strategy to forecast inflation

The implemented GA is defined as in section 2.3: two genetic operators are borrowed from the genetic algorithm literature, *tournament* and *mutation*. Tournament allows the spread of successful rules, whereas mutation allows for exploration of new rules.
Note in both cases of augmented adaptive expectations and augmented anchored expectations the rules are defined by a single parameter only, $\lambda_t^{ada,h}$ and $\beta_t^{\pi,h}$ respectively. This implies that in case of tournament the aforementioned parameters swap between agents, whereas in case case of mutation are simply replaced by a random draw.
In each period of the simulation any two households are randomly paired and their respective expectation rules compared using the fitness function:

$$FIT = \frac{\sum_{i=1}^{T}\left(\pi_{t-i}^e - \pi_{t-i}\right)^2}{T} \tag{2.34}$$

Which is the average squared error, calculated over a fixed time window of length T representing the fitness memory. The rule providing the better fitness gets copied by the agent who previously used the less performing rule, while the "successful" agent retains her own rule.
Moreover, each household undergoes mutation with probability Pr_m. In such case the parameter is updated as:

$$\begin{gathered} \beta_t^{\pi,h} = U_{h,t} \\ \text{or} \\ \lambda_t^{ada,h} = U_{h,t} \end{gathered} \tag{2.35}$$

Where $U_{h,t}$ is a random draw from a uniform distribution defined in the [0,1] interval. This is also how $\beta^{\pi,h}$ and $\lambda_t^{ada,h}$ initialised across agents, i.e.:

$$\begin{gathered} \beta_0^{\pi,h} = U_{h,t} \\ \text{or} \\ \lambda_0^{ada,h} = U_{h,t} \end{gathered} \tag{2.36}$$

2.5 Baseline Dynamics

The model is solved by means of computer simulations. Each model configuration has been run fifty times and each run is 600 periods long. As it is common practice in the ABM literature, we allowed for a burn-in period during which the model stabilises. In this case the burn-in is of 100 periods, leaving 500 periods available for analysis. Results presented hereafter always refer to the average across the fifty runs.
Before turning to the bulk of our analysis, I will show some aggregate time series generated by the model in the two baseline cases: static and dynamic Taylor rule regimes. This should not be interpreted as some sort of model validation exercise, since the model at hand is solely intended for forecasting strategies assessment and not for general economic analysis. Interest in aggregate time series is rather motivated by the need of working with a stable model. If that was not the case, then results might be driven by extreme model dynamics instead of being rooted in a business as usual situation, which is ultimately my interest at this stage of analysis.
Table 2.1 lists parameters values, while figures (2.1) and (2.2) shows time series plot of inflation, nominal GDP, real GDP, policy rate, and unemployment. Moreover, it is shown the unemployment-inflation scatterplot.

Figures show an overall stability across forecasting rules and policy regimes. The model features nominal growth, whereas real GDP, unemployment, and the policy rate are rather stable. Inflation appears to be stable as well, although showing different degrees of volatility across expectation formation mechanisms. Finally, the unemployment-inflation scatterplot shows the emergence of a Phillips-curve kind of relationship.
Since a comfortable degree of model stability is achieved, I can now move on to assess how the different forecasting strategies perform in the two different policy regimes. Let me remind again that my main interest is to detect which expectation rule achieves *collective* rationality which implies an *aggregate* zero mean forecasting error.

2.6 Results

Figure (2.3) shows the time series of the aggregate forecasting error for the eight configurations under consideration, four for each policy regime. It is interesting to notice that all of them seem to fluctuate around zero, suggesting that all the expectation formation rules employed are likely to be unbiased and therefore collectively rational. It also seems that volatility changes quite visibly across expectation formation mechanisms. This is an important point, since, given unbiasedness, the higher the volatility the more sever are the errors at each point in time. Therefore the less volatile the more "collectively rational" the rule is.

Table (2.2) confirms that all four expectation mechanisms provide mean errors very close to zero, I can therefore conclude that, at least in those two baseline scenarios, they are all unbiased. Another interesting finding is that volatility varies across different specification with naive expectations being much more volatile with respect to the others and therefore being subject to more severe errors. Interestingly, social learning does not perform better than adaptive expectations, however augmented adaptive expectations reduces volatility by a discrete amount with respect to simple adaptive expectation.
It is worth noticing that in case of dynamic central bank, the system is overall more volatile,

Table 2.1: Model Parametrization

Parameter	**Description**	**Value**
H	households #	100
c_0	normal propensity to consume	0.85
α_c	propensity to consume speed of adjustment	4
υ	desired inventories to output ratio	0.1
γ	labor productivity	1
μ^p	price mark-up	0.045
μ^D	deposits mark-up	0.1
μ^L	loans mark-up	0.2
β^c	consumption firm's dividend rate	0.9
β^b	bank's dividend rate	0.7
LL	loans duration	20
LR	minimum compulsory liquidity-ratio	0.08
p^b	bond price	1
θ	unemployment benefit	0.4
τ	tax rate	0.02
i^*	normal interest rate	0.0075
π^*	inflation target	0.0075
ρ	Taylor rule parameter	1.5
Pr_ρ	Central bank mutation probability	0.04
μ_{FN}	folded normal distribution mean	0
σ_{FN}	folded normal distribution standard deviation	0.01
μ_n	normal distribution mean	0
σ_n	normal distribution standard deviation	0.25
λ^{ada}	adaptive expectations parameter	0.25
T	GA memory	4
Pr_m	GA mutation probability	0.04

Table 2.2: Average Forecasting Error: Baseline Cases

	Static Taylor rule		Dynamic Taylor rule	
Expectation Mechanism	Mean	SD	Mean	SD
Adaptive	-4.36E-07 (0.9732)	0.0002897669	-2.44E-07 (0.9884)	0.000374252
Augmented Adaptive	6.096E-06 (0.5691)	0.0002392173	-4.04E-07 (0.9765)	0.000306617
Social Learning	0.000419004 (2.2e-16)	0.0003374981	0.000408832 (2.2e-16)	0.0003743263
Naive	1.516E-06 (0.9738)	0.001029729	1.18E-06 (0.9795)	0.001028179

p-values in brackets refers to t-tests under the 0 null

this is also reflected in the standard deviations associated to mean forecasting errors, signalling that the performance of each forecasting strategy deteriorates.

Figure 2.1: Static Taylor Rule

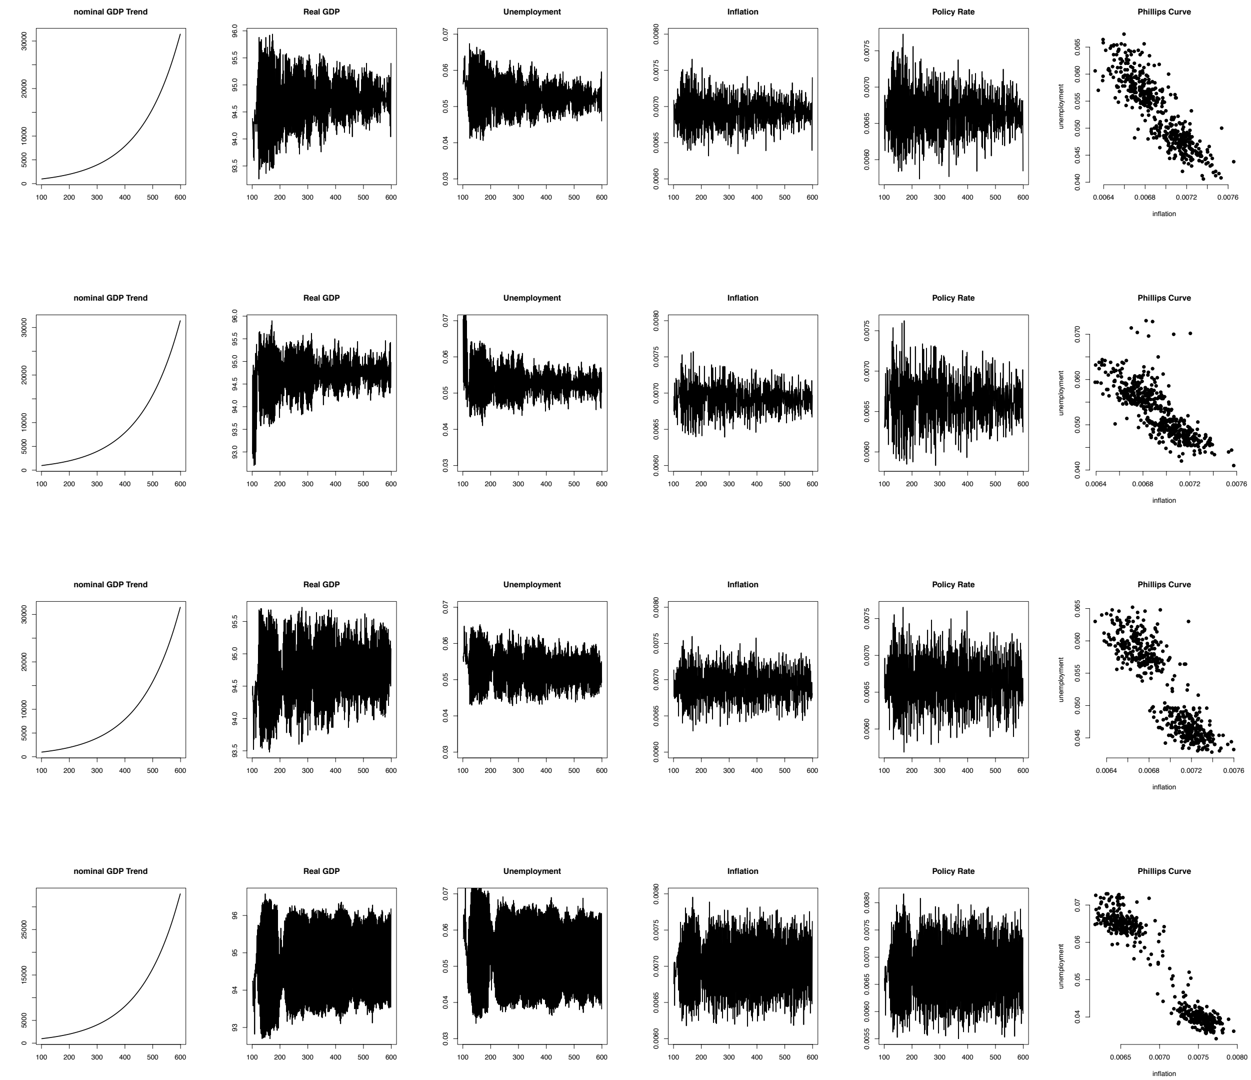

First row, adaptive expectations; second row, augmented adaptive expectations; third row, GA; fourth row, naive expectations

The baseline results therefore suggest that, although not being rational in a strict Muth's sense, those expectation formation mechanisms are collectively rational. Also, results show that adaptive expectations works reasonably well and comparatively better than other specifications. I finally notice, that combining social learning and adaptive expectations reduces volatility without introducing any bias, therefore at this stage it can be concluded that augmented adaptive expectations are the preferred specification.

2.6.1 Robustness Checks

Results obtained from the baseline scenario may be very much sensitive to the particular model configuration employed. Natural candidates to be critical parameters in this respects

Figure 2.2: Dynamic Taylor rule

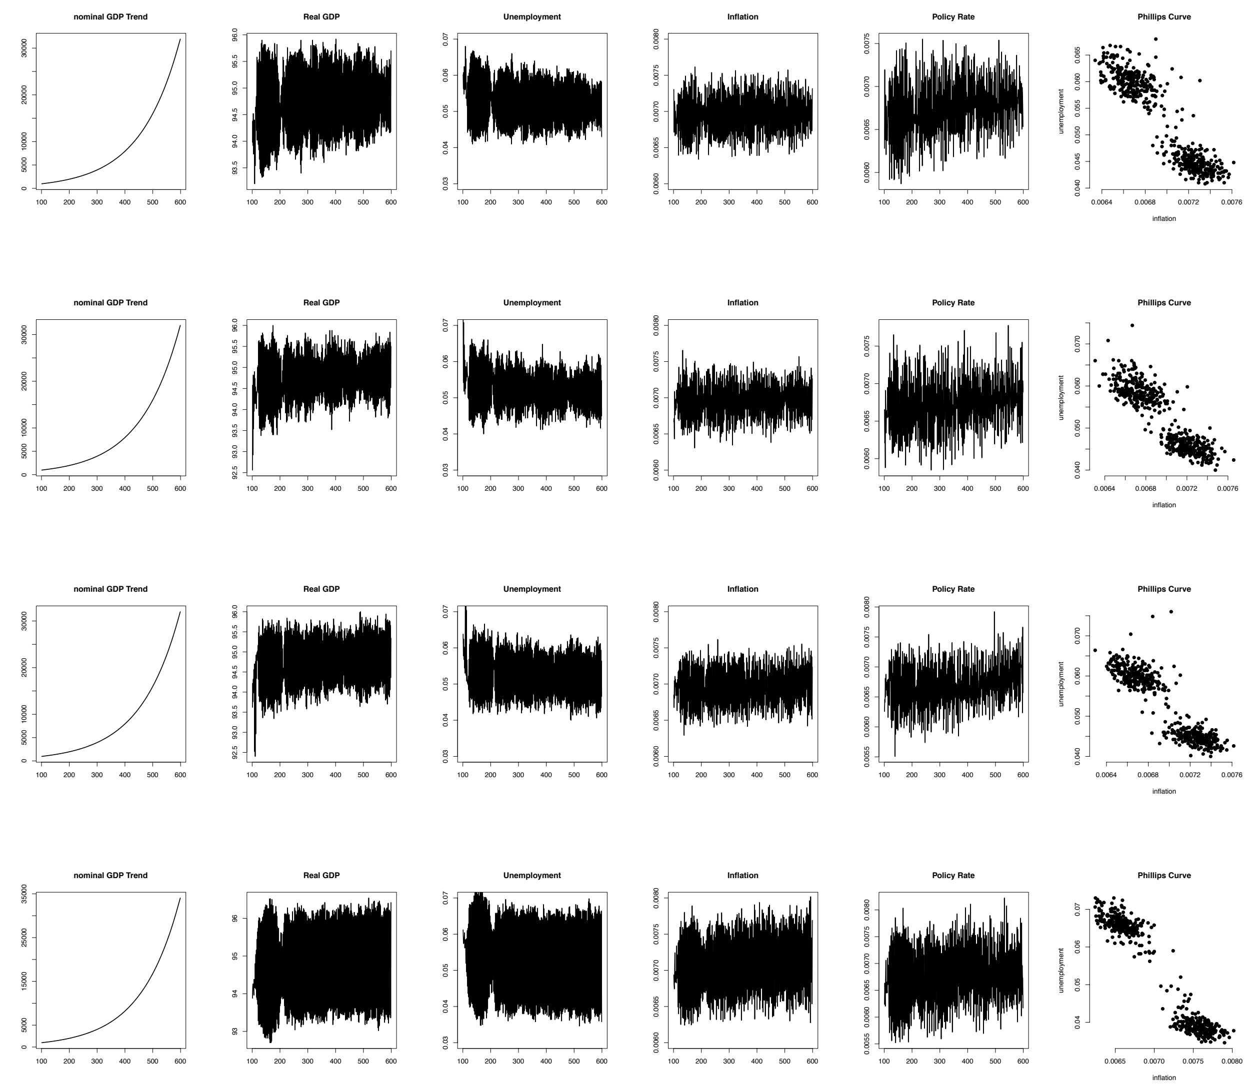

First row, adaptive expectations; second row, augmented adaptive expectations; third row, GA; fourth row, naive expectations

are: (i) parameters governing the feedback effects between inflation and consumption; (ii) parameters governing learning; (iii) Taylor rule parameters.
In this subsection point (i) and (ii) will be analysed, whereas to point (iii) is devoted the following subsection, dealing with policy shocks.

2.6.1.1 Abstracting from feedbacks effects

Forecasting performances are not independent from the economic environment they are tested in. It follows that, to be able to properly compare different expectation formation mechanisms, those must be applied in the exact same economic environment. Which, at first glance, it is what I have done in the baseline scenario: where indeed, the four expectation formation

Figure 2.3: Aggregate Forecasting Errors

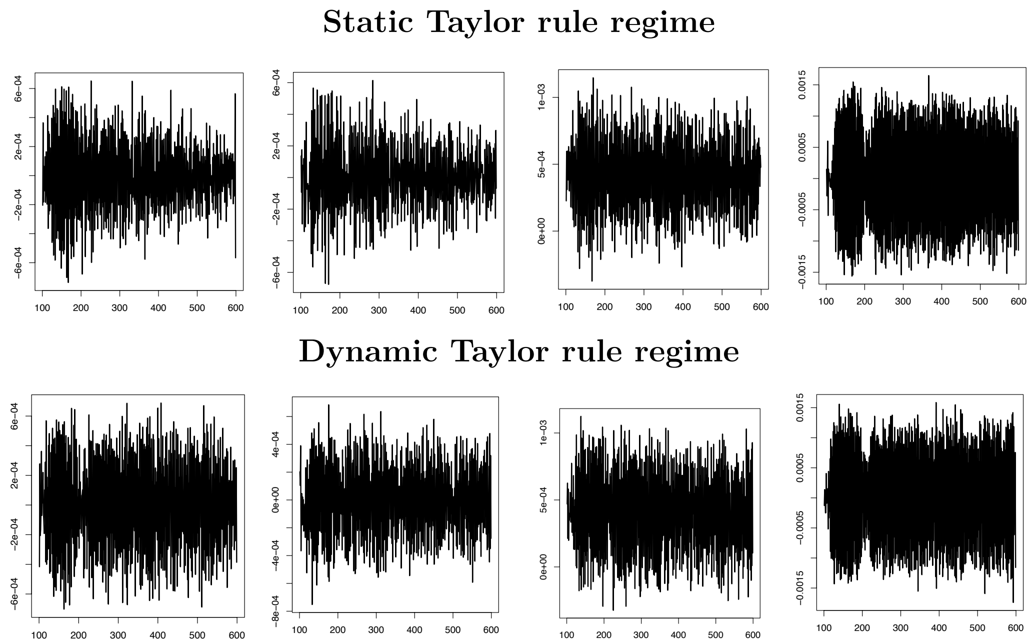

First row, adaptive expectations; second row, augmented adaptive expectations; third row, GA; fourth row, naive expectations

mechanisms have been tested within the same model. However, recall equation (2.9), which defines the marginal propensity to consume:

$$c_t^h = c_0 \left[1 + \alpha_c \left(\pi_t^e - i_t^d\right)\right]$$

Clearly, there is a feedback effect which can be summarised as follows: actual inflation affects inflation expectations, which affect consumption, which feedbacks into inflation. Therefore, as it is evident from inspection of figure (2.1) , changing the expectation formation mechanism affects the overall model dynamics.

It is therefore interesting to understand what is the impact of such feedback effect on forecasting performances. To polish from such feedback effect, I simply get rid of equation (2.9) and set exogenously the propensity to consume.

Table (2.3) presents means and standard deviations of the aggregate forecasting error across the eight model configurations. Comparing these results with those presented in the first two columns of table (2.2) shows that: (i) the average (aggregate) forecasting error is close to 0, so expectations turn out be *collectively* rational also when the feedback effect is absent; (ii) In the no feedback case the forecasting error is overall considerably more volatile. This is due to higher volatility in the inflation rate, implying a stabilising effect of the feedback mechanism; (iii) Augmented adaptive expectations provide the lower standard deviations, remaining the preferred specification; (iv) GA improves its relative performance in terms of volatility. This is due by the fact that in the no feedback scenario the central bank is more effective in reaching its target, however at the cost of higher volatility. Therefore, using the central bank target in forming inflation expectation has a positive impact on the GA performance.

Table 2.3: Average Forecasting Error: No Feedback

	Mean	SD
Adaptive	5.72E-07	0.0008765545
	(0.9884)	
Adaptive GA	1.68E-07	0.0007901013
	(0.9962)	
GA	0.000365344	0.0008097736
	(2.2e-16)	
Naive	2.032E-06	0.001525916
	(0.9763)	

p-values in brackets refers to t-tests under the 0 null

This experiment shows that the feedback effect is not the main driver of the results achieved in the baseline scenario. Most importantly, the four forecasting strategies are unbiased even when the feedback effect is absent. Moreover, augmented adaptive expectations remain the best performing strategy as it provides the lowest volatility.

2.6.1.2 Exploring learning parameters

The learning algorithm is governed by two parameters: (i) the fitness memory, i.e. the parameter T in equation (2.34) and (ii) the mutation parameter, Pr_m. In the baseline scenario they were exogenously fixed respectively at 4 and 0.04, meaning that: the tournament operator compare different forecasting performances over the last 4 periods of time, and that at each point in time each forecasting strategy mutates with a 4% probability.
It is in principle possible that varying these two parameters seriously affects our results, so further exploration is needed in order to assess such potential issue.
I consider the augmented adaptive expectation case and rerun the baseline static central bank scenario varying the parameter T and the parameter Pr_m in the [1,8] and [0.01,0.08] interval respectively. Pr_m is varied by 0.01 steps, ending up with 8 values for each parameter and therefore 64 combinations. I run each of this 64 for combinations 25 times and took averages. Detailed results are provided in appendix, however I can conclude that results obtained in the baseline scenario are robust to changes in the fitness memory and mutation probability, indeed neither the mean average forecasting error, nor its variance are seriously affected: augmented adaptive expectations turn out to be unbiased throughout the sensitivity exercise, moreover they keep providing low volatility even in the most extreme cases.

2.6.1.3 Policy Shocks

As already pointed out, it is good news that our expectations formation mechanism are unbiased in the baseline scenarios, nevertheless this is not enough to rule out systematic mistake in the face of policy shocks. The dynamic central bank scenario already gives us a positive indication about it, but to further investigate this point we take the static central bank case and we impose different policy shocks at a specific point in time. We analyse the collective forecasting performances in the vicinity of the shock, to understand whether shocks deteriorates forecasting performances and if there are forecasting rules which are more

susceptible than others.
For every shock under considerations we proceed as follows: we impose the shock at period 200 and we analyse the forecasting performances for the subsequent 50, 100 and 400 periods. In this way we are able to see if the forecasting performance deteriorates and if so, whether it recovers as time goes.
Shocks are specified in terms of Taylor parameter and inflation target shifts, as reported in table 2.3. Results are reported in Appendix and they show overall unbiasedness also in the

Table 2.4: Policy Shocks

	Taylor Parameter	Inflation Target
shock1	0.5	0.0035
shock2	1	0.0055
Baseline	1.5	0.0075
shock3	2	0.0095
shock4	2.5	0.0115

vicinity of the policy shocks, showing a satisfactorily degree of robustness in particular for adaptive expectations. We can observe that augmenting adaptive expectation with learning it proves beneficial to simple adaptive expectations reducing considerably volatility. We can also confirm higher volatility in the case of naive expectations. However, just a pinch of extra rationality with respect to naive expectations deliver a more than satisfactorily result: our expectations formation mechanism are unbiased and with relatively low degree of volatility even in case of a policy shocks. I should mention that, although not so systematically, volatility tend to be higher in the vicinity of the shock and dampen as the simulation unfolds. Overall, augmented adaptive expectations remain the preferred specification, although simple adaptive expectations prove to perform well enough and in a few cases even to outperform their augmented counterpart.

2.7 Case II: Forecasting a Trended Variable

In this setting I am going to analyse the performances of consumption firms trying to forecast their future real demand. Unlike the previous case, here a full fledged SFC-ABM model is used for the purpose. In particular I will use the Caiani et al. (2016) model augmented by technological innovation as in Caiani et al. (2019) and Caiani et al. (2018b). Technological innovation is needed in order to generate real growth so to introduce a trend in consumption firms' real demand. It is important to stress that the individual real demand is determined by two main components: (i) technological innovation determines the long run trend; (ii) competition among firms determines short term fluctuation for individual demand, which maybe very severe.
In this context I will experiment with 6 expectation formation mechanisms: (i) naive; (ii) adaptive; (iii) weak trend following; (iv) strong trend following; (v) augmented adaptive; (vi) augmented trend following.

2.7.1 The Model

Hereafter the exact model presented in Caiani et al. (2019) will be used, including its parametrisation and initial conditions. A full description of the model would be too cumbersome, therefore in the following I am going to give a superficial overview of the model and focus only on the blocks which are relevant for this paper.

The model is fully agent-based, i.e. each and every sector is disaggregated, and is composed of: (i) capital good sector; (ii) consumption good sector; (iii) household sector, which is further disaggregated in three different skill groups; (iv) bank sector; (v)government; (vi) central bank.

Consumption firms buy capital goods from capital firms and hire workers on a skill-differentiated labour market. Once capital goods are bought and workers are hired, consumption firms combine these two factors of production in a fixed coefficient production function.

Capital firms continuously engage in R&D activities aiming at improving capital productivity Consumption firms sell goods to households on the consumption market and ask for loans to banks in the credit market.

The government taxes the private sector, provides unemployment benefits and emits bonds when public expenditures exceeds taxes. Banks buy bonds on the bond market and the central bank absorbs the difference between supply and demand.

2.7.1.1 Focus 1: consumption firms

Since I am going to experiment on consumption firms forecasting performances, I am going to describe consumption firms behaviours in details: Consumption firms in this model behave similarly to the single consumption firm modelled above. The main difference is that in this model consumption firms needs machines and labour in order to produce

Desired output is chosen as:

$$y_t^{d,c} = (1+v)s_t^{e,c} - inv_{t-1}^c$$

Once $y_t^{d,c}$ is set, the consumption firm can calculate its labour demand. In order to do so, the consumption firm has to decide which machineries to employ in production. Each machine is characterised by two parameters: capital productivity, μ_κ, and capital labour ratio, l_k, which is assumed to be constant across capital vintages. Therefore firm c will try to reach its target $y_t^{d,c}$ using the most productive capital at her disposal.

Call the number of machines needed to reach desired production $k_t^{c,eff}$, then labour demand for firm c at time t is given by:

$$l_t^{d,c} = \frac{k_t^{c,eff}}{\mu_\kappa} \tag{2.37}$$

Given labour demand, firm c can compute its price as a mark over expected unit cost of production:

$$p_t^c = (1+\mu_t^c) = \frac{W_t^{e,c} l_t^{d,c}}{y_t^{d,c}} \tag{2.38}$$

Note that: (i) the mark-up is time and firm dependent, meaning that it endogenously evolves; (ii) in this equation enters the expected wage bill, instead of the actual wage bill. This is due by the fact that the sequence of events in this model requires firms to set their prices before hiring.

Firms adjust mark-up according to market condition, if actual sales exceeds expected sales they increase the mark-up and lower it vice-versa following the rule:

$$\mu_t^c = \begin{cases} \mu_{t-1}^c(1 + FN_{c,t}^1) & \text{if} \quad \frac{inv^c t-1}{s_{t-1}^c} \geq v \\ \mu_{t-1}^c(1 - FN_{c,t}^1) & \text{if} \quad \frac{inv^c t-1}{s_{t-1}^c} < v \end{cases} \tag{2.39}$$

Where $FN_{c,t}^1$ is a random draw from a folded normal distribution FN^1 defined over the parameters $(\mu_{FN^1}, \sigma_{FN^1})$.

Finally, firms can increase their productive capacity by investing in physical capital. Firm c desired capacity growth is determined by the investment function:

$$g_{c,t}^D = \gamma_1 \frac{r_{c,t-1} - \bar{r}}{\bar{r}} + \gamma_2 \frac{u_{c,t}^D - \bar{u}}{\bar{u}} \tag{2.40}$$

Where γ_1 and γ_2 are exogenous and time invariant parameters, $r_{c,t-1}$ is the profit rate realised by firm c at time t, $u_{c,t}^D$ is desired capacity utilization, $\bar{r}$ and $\bar{u}$ are respectively the exogenous and time invariant "normal" levels of profit rate and capacity utilisation.
Note that firms invest in response to above normal profitability and to keep a certain level of capacity utilization, but not in order to catch up with the technological frontier.
Once $g_{c,t}^D$ has been set, firm c goes on the capital market and try to buy enough machines to reach its capacity target. The exact procedure can be found in Caiani et al. (2019), here will be enough to say that firm c is only able to survey a subset of capital suppliers and it chooses the most productive machine available.

2.7.1.2 Focus 2: competition

Consumers interact with consumption firms following a 2-stages matching protocol: in stage 1 each demander is endowed with one supplier. Moreover, she can survey the price offered by a number of alternative suppliers. The number of alternative suppliers defines the degree of competition within the market and it is set exogenously with a parameter χ. The demander compares all the prices surveyed and picks the lowest among them. In stage 2 the demander compares the price of her old supplier with the price offered by the new potential supplier. If the new price is lower than the old one, then, following Delli Gatti et al. (2010a), the demander switches supplier with a probability given by:

$$Pr_s = \begin{cases} 1 - e^{\epsilon\left(\frac{P_n - P_o}{P_n}\right)} & \text{if } P_n < P_o \\ 0 & \text{Otherwise} \end{cases} \tag{2.41}$$

Where P_n and P_o are respectively the prices offered by the potential new supplier and the old one. Therefore, the probability of switching supplier is a non linear function of the difference between the new and the old price.
Note that sudden changes in relative prices may result in severe swings in real sales. Of course the parameters governing market competition define the extent of the swing for a given change in relative prices. In the following I will analyse the role played by the number of potential suppliers χ, which in the baseline simulation is set to 5.

2.7.1.3 Focus 3: R&D

Technological innovation is modelled as in Dosi et al. (2010): it is a 2-stages process updating capital productivity. In the first stage each capital firm k performs a Bernoulli experiment to determine whether innovation has been successful. The probability of innovate for firm k is given by:

$$Pr^{inn}_{k,t} = 1 - e^{\xi^{inn} N_{k,r,t}} \tag{2.42}$$

Where $N_{k,r,t}$ is the number of researchers hired by firm k at time t and ξ^{inn} is an exogenous time invariant parameter.
If innovation has been successful, the productivity of machines produced by firm k evolves as:

$$\mu_{k,t} = (1 + FN^2_{k,t})\mu_{k,t-1} \tag{2.43}$$

Where $\mu_{k,t}$ is the productivity of machines produced by firm k at time t and $FN^2_{k,t}$ is a random draw form a folded normal distribution FN^2 defined over the parameters $(\mu_{FN^2}, \sigma^2_{FN^2})$. Firm k also engages in a process of imitation, where it can copy the technology of a competitor. Firm k succeeds in imitation with probability:

$$Pr^{im}_{k,t} = 1 - e^{\xi^{im} N_{k,r,t}} \tag{2.44}$$

If firm k is successful in imitating, it is allowed to survey a subset of competitors' technologies and copy the best technology surveyed if better than its own.

2.7.1.4 Expectations

The forecasting strategies implemented hereafter are the same as those used in Dosi et al. (2017a), except for the learning algorithm. As mentioned above, I am going to experiment with: (i) naive, (ii) adaptive, (iii) augmented adaptive, (iv) weak trend following, (v) strong trend following, and (vi) augmented trend following expectations.
Augmented adaptive and augmented trend following expectations implement the GA algorithm already presented in section 2.3. The fitness function adopted in this framework is defined as:

$$FIT = \left[\sum_{i=1}^{T} \frac{\frac{|s^{e,c}_{t-i} - s^{c}_{t-i}|}{s^{c}_{t-i}}}{T} \right] \tag{2.45}$$

Learning parameters are exogenously set as: $T = 4$ and $Pr_m = 0.04$.
The adaptive parameter in the simple adaptive expectation scheme is set to 0.25, as in Caiani et al. (2019). Whereas the adaptive parameters regarding the weak and strong trend following schemes are set respectively to 0.4 and 1.3, as in Dosi et al. (2017a).

2.7.2 Results

I run 6 model configurations, one for each expectation formation mechanism. Each configuration has been run 25 times and each run is of 700 periods. Following Caiani et al. (2019) we used 500 periods for the burn-in process, leaving 200 periods available for analysis. Results hereafter always refer to the average across the 25 runs.
I shall highlight that the model produces sometimes extreme outliers in the individual forecasting errors. This is due to violent swings in individual demand, which may, for example

drive, sales expectations close to 0 very quickly and therefore pushing the forecasting error to unreasonable high values. Those outliers, although not so frequent can be so extreme to seriously bias the results. I will address this issue directly later on showing a somehow interesting result, however at this stage of analysis I decided to exclude individual errors exceeding the value of 100% relative to expectation in the computation of the aggregate mean forecasting error.

2.7.2.1 Baseline

Figures in appendix show real GDP, unemployment, and aggregate forecasting error time series for the six configurations under consideration. Aggregate time series are not of much interests, however it is important to notice a relatively stable unemployment, signalling that results are not biased by extreme model dynamics, and the upward trend in real GDP. Real growth is in fact needed to have trended real sales.

Table 2.5: Summary Baseline

	MEAN	SD
Adaptive	-0.013	0.0061
	(0.000)	
Adaptive GA	-0.040	0.0057
	(0.000)	
Naive	-0.045	0.0050
	(0.00)	
S. Trend Following	-0.124	0.0076
	(0.000)	
W. Trend Following	-0.066	0.0054
	(0.000)	
Trend Following GA	-0.105	0.0062
	(0.000)	

p-values in brackets under the 0 null

Table (2.5) summarises the forecasting strategies performances, showing two main facts: (i) in this context none of the forecasting strategies turns out to be unbiased. Simple adaptive expectations provide the best performance with an average aggregate forecasting error equal to 1.3% relative to the expected value; (ii) unlike the previous case, learning deteriorates forecasting performances. Indeed, augmented schemes provide higher volatilities than the simpler schemes. Moreover, the augmented adaptive scheme provide a much larger error than the simple adaptive scheme.

Result (ii) is similar to the one obtained in Dosi et al. (2017a), who put forward the following explanation: fundamental uncertainty brought about by technological innovations favours simple and frugal forecasting strategies over more elaborated rules. Innovation waves occur in an unpredictable way, producing continuous structural breaks which seriously bias learning algorithms. However, attentive scrutiny would suggest some caution with respect to such explanation: let us assume that those unpredictable structural breaks *alone* are responsible for the overall deterioration of forecasting strategies performances and in particular for those

using learning. Then, any forecasting strategy should underestimate realised real sales, because, in this context, a wave of technological innovation would trigger an unexpected surge in real demand, resulting in a serious and generalised underestimation of real sales. However, negative values on the first column of table 2.4 suggests the opposite, i.e. expectations on average tend to overestimate actual real sales.
This observation does not reject the hypothesis according to unpredictable technological changes play a relevant role in determining forecasting errors, nevertheless it calls for further analysis which I am going to carry out in the following sections.

2.7.2.2 The role of competition: a preliminary analysis

As discussed above individual real sales follow a trend determined by technological innovations, but they also turn out to be quite noisy because of competition among firms which try to gain market shares at the expense of competitors.
Here, I will analyse if these two elements are somehow linked and, if so, whether the interplay between technological innovation and competition can explain forecasting strategies performances.
First of all, let me revise the functioning of the goods market: when the consumption good market opens, each household is endowed by her own consumption good supplier. However, before buying she is allowed to survey a certain number of competitors' price and compare them with her own supplier's price. It follows that each households, if able to find more convenient prices, can change supplier with a probability increasing in the difference between her own old supplier's price and the potentially new supplier's price. Clearly, relative firm's prices and the degree of competition play a fundamental role in determining individual demand swings.
Recall that price is set as a mark-up over unit cost of production, which is determined by the technology embedded in the capital stock of each firm. Let us abstract from the role of mark-up and focus on the role technology asking the question: is it possible to have heterogeneity in firms' unit cost of production because of different technologies employed across firms? The answer is positive and directly depend on two features of the model: (i) as households, consumption firms only observe a subset of the entire capital firms population. It follows that if a particular consumption firm is unlucky in the sense that it is not able to observe the capital firms offering the most productive capital, than firm c will lack behind the technological frontier, at least for some time, and will experience a loss in competitiveness. (ii) Firms do not invest in capital items with the scope of catching up to the technological frontier, but only in response to fluctuation in capacity utilization and profit rate. Consider a situation in which a technological wave occurs, but according to its own investment function firm c does not seek to invest in productive capital. By not investing, firm c will lack behind with respect to the technological frontier therefore loosing in terms of competitiveness.
I claim that, at least in principle, if those dynamics occur too violently, they produce excessive volatility in firms' individual good demand which in turn is responsible for the degradation in forecasting strategies performances. To test my claim I re-run the model as it is, but with one modification: I will reduce the number of new potential suppliers each households is allowed to survey, in order to reduce competition.
In the baseline scenario each households was allowed to survey 5 different consumption firms, results presented hereafter refer to a setting in which households only see 1 potential new supplier.

Table 2.6: Summary Reduced Competition

	MEAN	SD
Adaptive	0.004	0.0119
	(0.000)	
Adaptive GA	0.006	0.0061
	(0.000)	
Naive	0.0000	0.0057
	(0.885)	
S. Trend Following	-0.022	0.0092
	(0.000)	
W. Trend Following	-0.006	0.0059
	(0.000)	
Trend Following GA	-0.006	0.0083
	(0.000)	

p-values in brackets

Comparison of table 2.4 and table 2.5 shows relevant differences, suggesting that competition indeed plays a relevant role in determining agents forecasting performances: (i) the aggregate mean forecasting error is very close to 0 across specifications, suggesting that in this context expectations are *collectively* rational; (ii) in case of adaptive expectation, learning reduces volatility by a discrete amount therefore improving the forecasting performances.

2.7.2.3 Sensitivity Analysis

2.7.2.3.1 Competition Vs Innovation

The local sensitivity analysis performed by reducing the number of potential suppliers to 1 indicates that competition may affect the performances of forecasting strategies. Here I am going to perform a more detailed analysis in order to better understand the role of competition *per se* and relatively to the innovation dynamics. In particular, the goal is to understand what is the main driver behind the deterioration of forecasting strategies. In order to do so, I performed a sensitivity analysis over two parameters governing innovation dynamics and competition. For what concerns the innovation dynamics I vary the parameter $\sigma^2_{FN^2}$, which is the variance of the folded normal distribution from which productivity updates are drawn (see equation, 2.26). The larger $\sigma^2_{FN^2}$, the larger the jumps in productivity and therefore the structural breaks. For what concerns competition I varied the parameter χ which gives the number of potential suppliers an household is able to survey on the consumption market. Clearly, the larger is χ the more competitive is the market. $\sigma^2_{FN^2}$ is varied in the range $(0.005 : 0.015)$ with steps equal to 0.0025, whereas the parameter χ is varied in the range $(1 : 9)$ with steps equal to 2. I end up with 5 values for each parameter and therefore with 25 configurations. Each configuration has been run 10 times and averages across rounds are hereby considered. We interpolate the 25 configurations to obtain a continuous surface using a kriging algorithm as suggested by Salle and Yıldızoğlu (2014).

Figure 2.4 shows a few interesting patterns: (i) reducing innovation dynamics has only a limited impact. Reducing $\sigma^2_{FN^2}$ mildly improves the performances as the mean forecasting

Figure 2.4: Competition-Innovation kriging

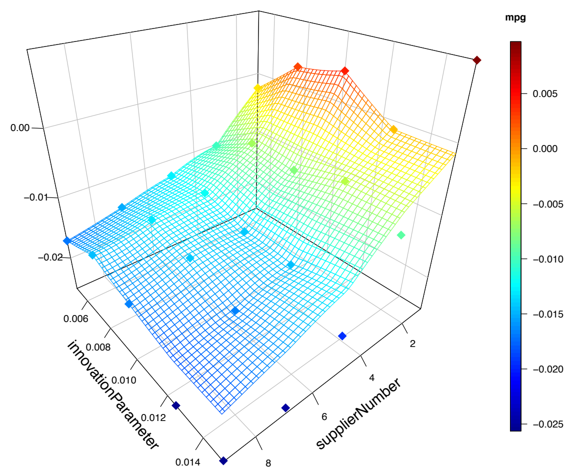

(a) mean forecasting error adaptive case

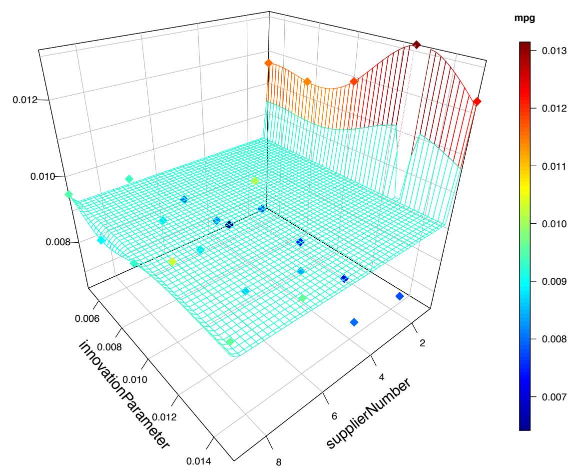

(b) s.e. forecasting error adaptive case

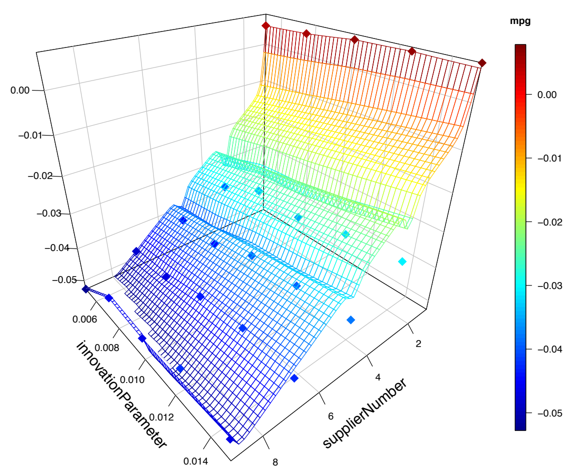

(c) mean forecasting error adaptive GA case

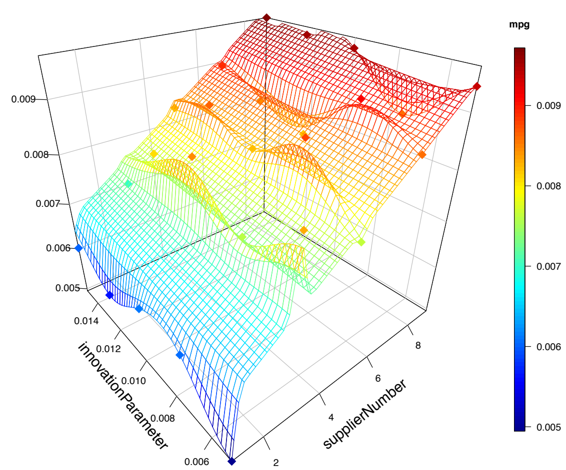

(d) s.e. forecasting error adaptive GA case

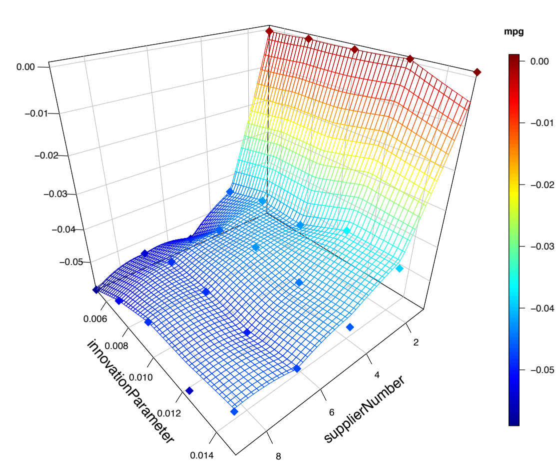

(e) mean forecasting error naive case

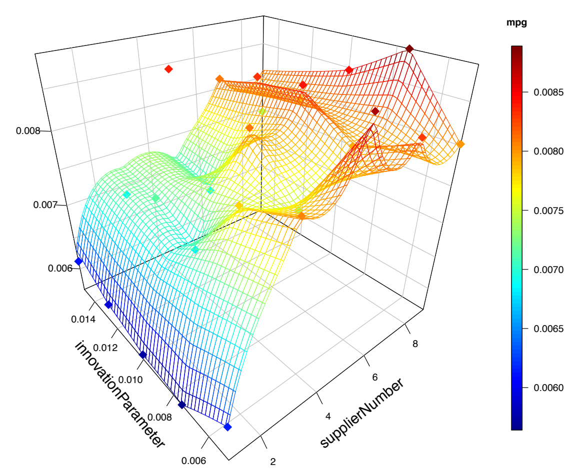

(f) s.e. forecasting error naive case

error seems to getting closer to zero, but only in the adaptive case; (ii) on the other hand,

competition seems to have a much important impact. All the surfaces describing the mean forecasting error pattern are quite steep in the χ dimension. In particular, augmented adaptive expectations quickly approach a zero mean error as χ is reduced; (iii) competition also seems to influence the volatility of the mean forecasting error. In the naive and augmented adaptive cases volatility decreases with χ. Suggesting that as competition is tamed the mean forecasting error approaches zero as well as volatility. On the contrary, the standard error surface relative to simple adaptive expectations is flat, which is somehow puzzling, when confronted with the other two.
This sensitivity experiment suggests the insight gained in the previous section. Technological innovation and competition explain the deterioration of the forecasting performances. However, it is competition to be the main driver through the channel outlined above: differences in prices due to gaps in technology determine sudden switches, which in turn determine sudden swing in individual firms' demand.
This is particularly true for learning, that as competition is tamed, improves its performance relative to the other forecasting strategies.

2.8 Conclusions

The paper shows that it is possible to achieve aggregate unbiased expectations in ABMs. Such result addresses an important caveat posed by the Lucas' critique: even if rational expectations are not applicable in an agent-based framework, there are alternatives which ensure *collective* rationality. Moreover, it shows that a simple learning algorithm can improve the performances of static expectation schemes, like adaptive or trend following expectations. The result is robust across different economic environments, variables following different dynamics, i.e. stationary and trended, and different parameterisations of the learning algorithm. A somehow novel insights has been given in the context of a full fledged model, where forecasting performances appear to deteriorate and learning ceases to be effective. In order to understand the main driver of such pattern, I performed a sensitivity analysis, suggesting that excessive competition on the good market is far more likely to determine it than innovation dynamics. Although the two are most likely linked.
This last insight calls for a reconsideration of market design in ABM, where an excessive degree of competition may be a source of excess volatility in firms individual demand and therefore affecting the overall model dynamics.

2.9 Appendix

2.9.1 Experiment 1: learning sensitivity

Table 2.7: Mean Aggregate Errors Across Learning Configurations

	1	**2**	**3**	**4**	**5**	**6**	**7**	**8**
0.01	-1.296E-06	4.312E-06	8.936E-06	1.256E-06	-2E-07	3.136E-06	-4.4E-07	-2.96E-06
0.02	3.28E-06	6.344E-06	4.6E-06	5.144E-06	1.0272E-05	3.048E-06	3.672E-06	3.68E-07
0.03	3.856E-06	9.032E-06	7.176E-06	3.832E-06	5.52E-06	-9.006E-08	2.344E-06	1.064E-06
0.04	-1.464E-06	2.912E-06	-3.408E-06	4.784E-06	6E-06	6.272E-06	1E-06	-3.76E-07
0.05	3.584E-06	6.8E-06	2.584E-06	4.928E-06	6.248E-06	5.28E-06	-1.28E-07	2.688E-06
0.06	1.768E-06	5.568E-06	3.192E-06	7.6E-06	5.32E-06	-3.44E-07	-1.392E-06	2.144E-06
0.07	1.2000E-07	-3.840E-07	6.32E-06	1.1568E-05	3.12E-06	5.944E-06	1.384E-06	7.28E-07
0.08	3.36E-06	4.544E-06	4.144E-06	7.096E-06	2.568E-06	2.8E-06	4.92E-06	-9.6E-08

On the first row: T values

On the first column: Pr_m values

2.9.2 Experiment 1: policy shock

2.9.2.1 Taylor Parameter shock

Table 2.8: ρ=0.5

	MEAN50	SD50	MEAN100	SD100	MEAN	SD
Adaptive	-1.216E-05	0.0005991703	-4.58E-06	0.0006766183	-4.85E-07	0.0007018976
AdaptiveGA	-2.228E-05	0.0004646909	-1.102E-05	0.0005176399	-5.2E-07	0.0005594697
GA	0.00038456	0.0006221426	0.00038768	0.0006596317	0.000391445	0.0006609009
Naive	3.76E-06	0.0008718723	2.54E-06	0.001068527	-1.5E-08	0.001316011

Table 2.9: ρ=1

	MEAN50	SD50	MEAN100	SD100	MEAN	SD
Adaptive	-1.088E-05	0.0004098006	-1.58E-06	0.0004493224	-1.38E-06	0.0004160927
AdaptiveGA	-4.24E-06	0.0003460956	-1.14E-06	0.0003777404	4.135E-06	0.0003248124
GA	0.00041404	0.0004550623	0.0004186	0.0004735153	0.00042109	0.000458323
Naive	7.84E-06	0.001158699	5.78E-06	0.001403276	7.55E-07	0.001628533

Table 2.10: ρ=2

	MEAN50	SD50	MEAN100	SD100	MEAN	SD
Adaptive	-5.28E-06	0.0002751729	-2.16E-06	0.0002639545	-1.265E-06	0.0002189036
AdaptiveGA	-4.76E-06	0.0002292636	-1.54E-06	0.0002273323	1.4E-06	0.0001851678
GA	0.00040924	0.0002599582	0.00042408	0.0002508576	0.000415855	0.0002087152
Naive	-9.6E-07	0.0006828342	2.7E-06	0.0006466851	4.15E-07	0.0005832353

Table 2.11: ρ=2.5

	MEAN50	SD50	MEAN100	SD100	MEAN	SD
Adaptive	-7.44E-06	0.000225912	-7.2E-07	0.0002061167	-1.3E-07	0.0001765817
AdaptiveGA	-1.608E-05	0.0002000767	-9.82E-06	0.0001941151	-2.19E-06	0.0001729705
GA	0.00041296	0.0002331408	0.00042092	0.0002239174	0.000416815	0.0001875215
Naive	2.84E-06	0.0004905183	3.24E-06	0.0005098484	3.65E-07	0.0003831442

2.9.2.2 Inflation target shock

Table 2.12: π^{*}=0.0035

	MEAN50	SD50	MEAN100	SD100	MEAN	SD
Adaptive	-2.36E-05	0.0005097974	-1.072E-05	0.0005611713	-2.92E-06	0.0005275701
AdaptiveGA	-6.5E-05	0.0004026213	-2.528E-05	0.0004307025	1.115E-06	0.0004119144
GA	6.884E-05	0.001204417	5.192E-05	0.00139866	4.3345E-05	0.001495167
Naive	9.4E-06	0.001514799	3.86E-06	0.001713174	1.55E-06	0.00183845

Table 2.13: π^{*}=0.0055

	MEAN50	SD50	MEAN100	SD100	MEAN	SD
Adaptive	-7.96E-06	0.0003932507	-5.94E-06	0.0004205856	-1.16E-06	0.0003968069
AdaptiveGA	-1.436E-05	0.0003306942	-7.32E-06	0.0003529415	1.76E-06	0.0003135034
GA	0.00018704	0.0007794354	0.00017108	0.0008204448	0.00016476	0.000801965
Naive	5.72E-06	0.001221315	3.1E-06	0.00136008	9.65E-07	0.00145383

Table 2.14: π^{*}=0.0095

	MEAN50	SD50	MEAN100	SD100	MEAN	SD
Adaptive	2.2E-06	0.0002518047	3.52E-06	0.000204981	3.05E-07	0.0001763827
AdaptiveGA	1.596E-05	0.0002160954	1.33E-05	0.0001860803	7.845E-06	0.0001608133
GA	0.00012048	0.0004803287	0.00010394	0.0005087538	9.154E-05	0.0004256358
Naive	8.64E-06	0.0006401138	1.22E-06	0.0006495089	7.6E-07	0.0006113241

Table 2.15: π^{*}=0.0115

	MEAN50	SD50	MEAN100	SD100	MEAN	SD
Adaptive	5.64E-06	0.0002827069	4.02E-06	0.0002264759	1.12E-06	0.0001761751
AdaptiveGA	3.032E-05	0.0002377007	1.784E-05	0.0002030741	9.295E-06	0.0001558868
GA	7.228E-05	0.0003036265	6.196E-05	0.000282854	5.1405E-05	0.0003063952
Naive	3.32E-06	0.0005035837	1.54E-06	0.000405043	2.3E-07	0.0003145736

2.9.3 Full Fledged Model Dynamics

Figure 2.5: Baseline Dynamics

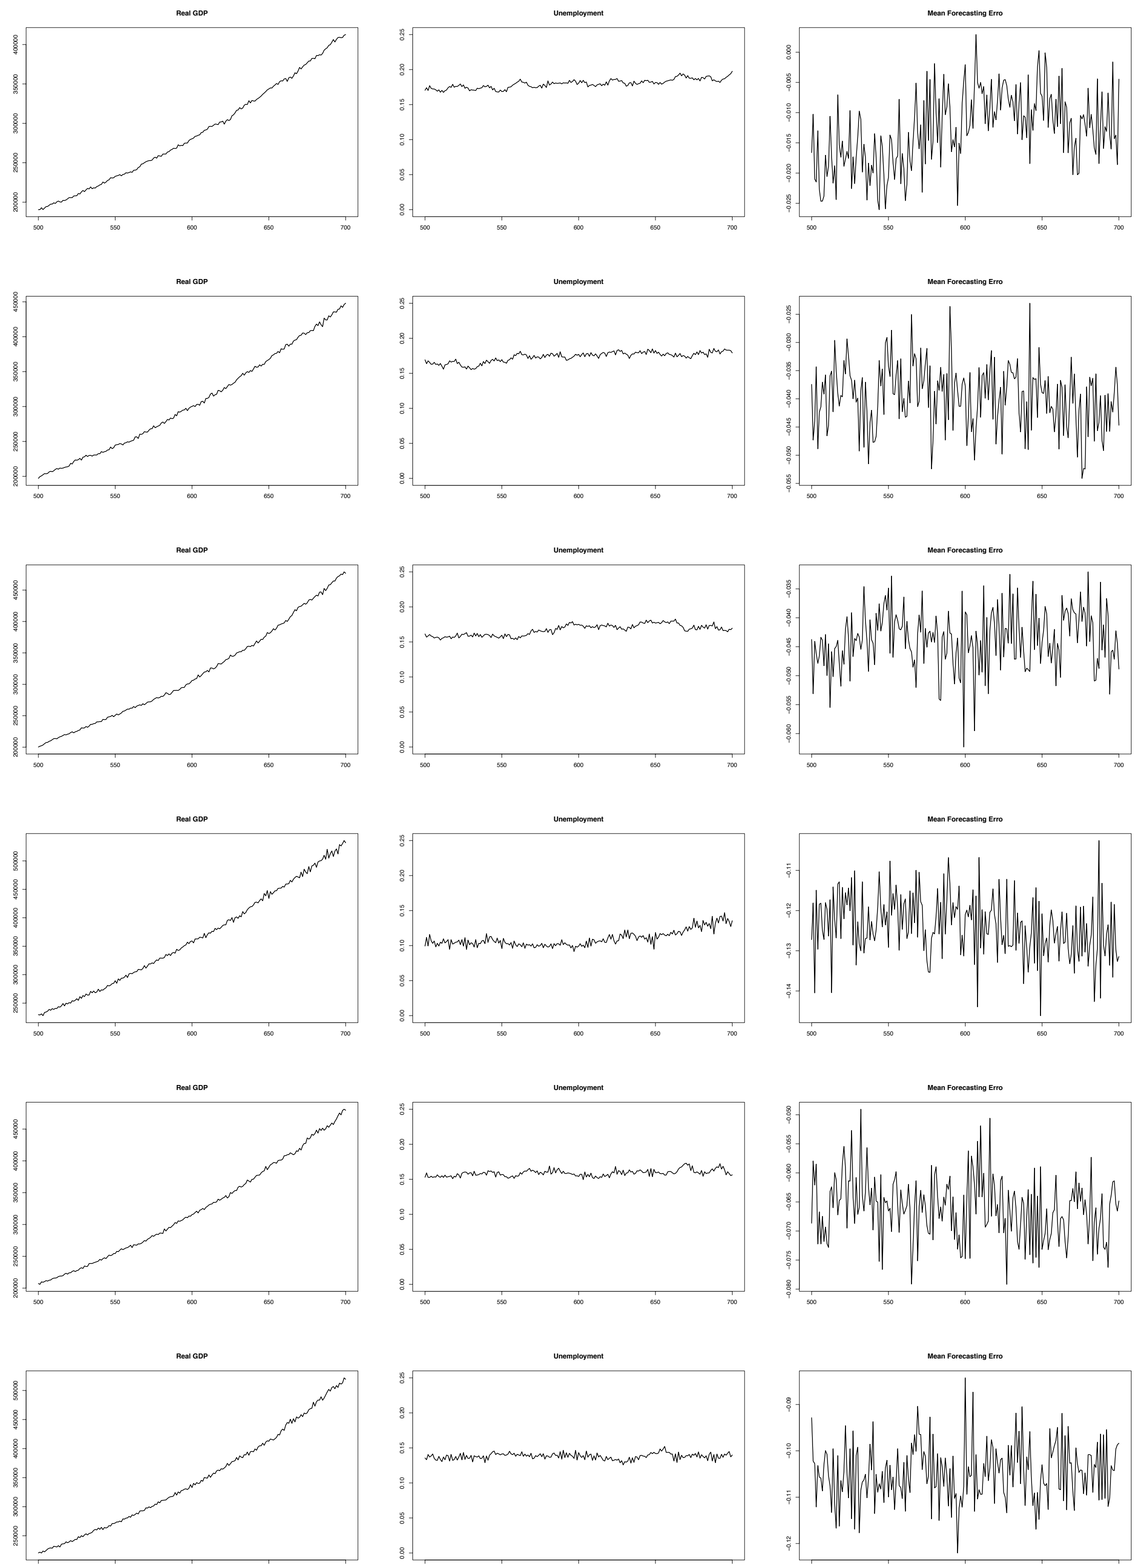

First row, Adaptive; Second row, Augmented Adaptive; Third row, Naive; Fourth row, Strong Trend Following; Fifth row, Weak Trend Following; Sixth row, Augmented Trend Following expectations

Figure 2.6: Reduced Competition: χ=1

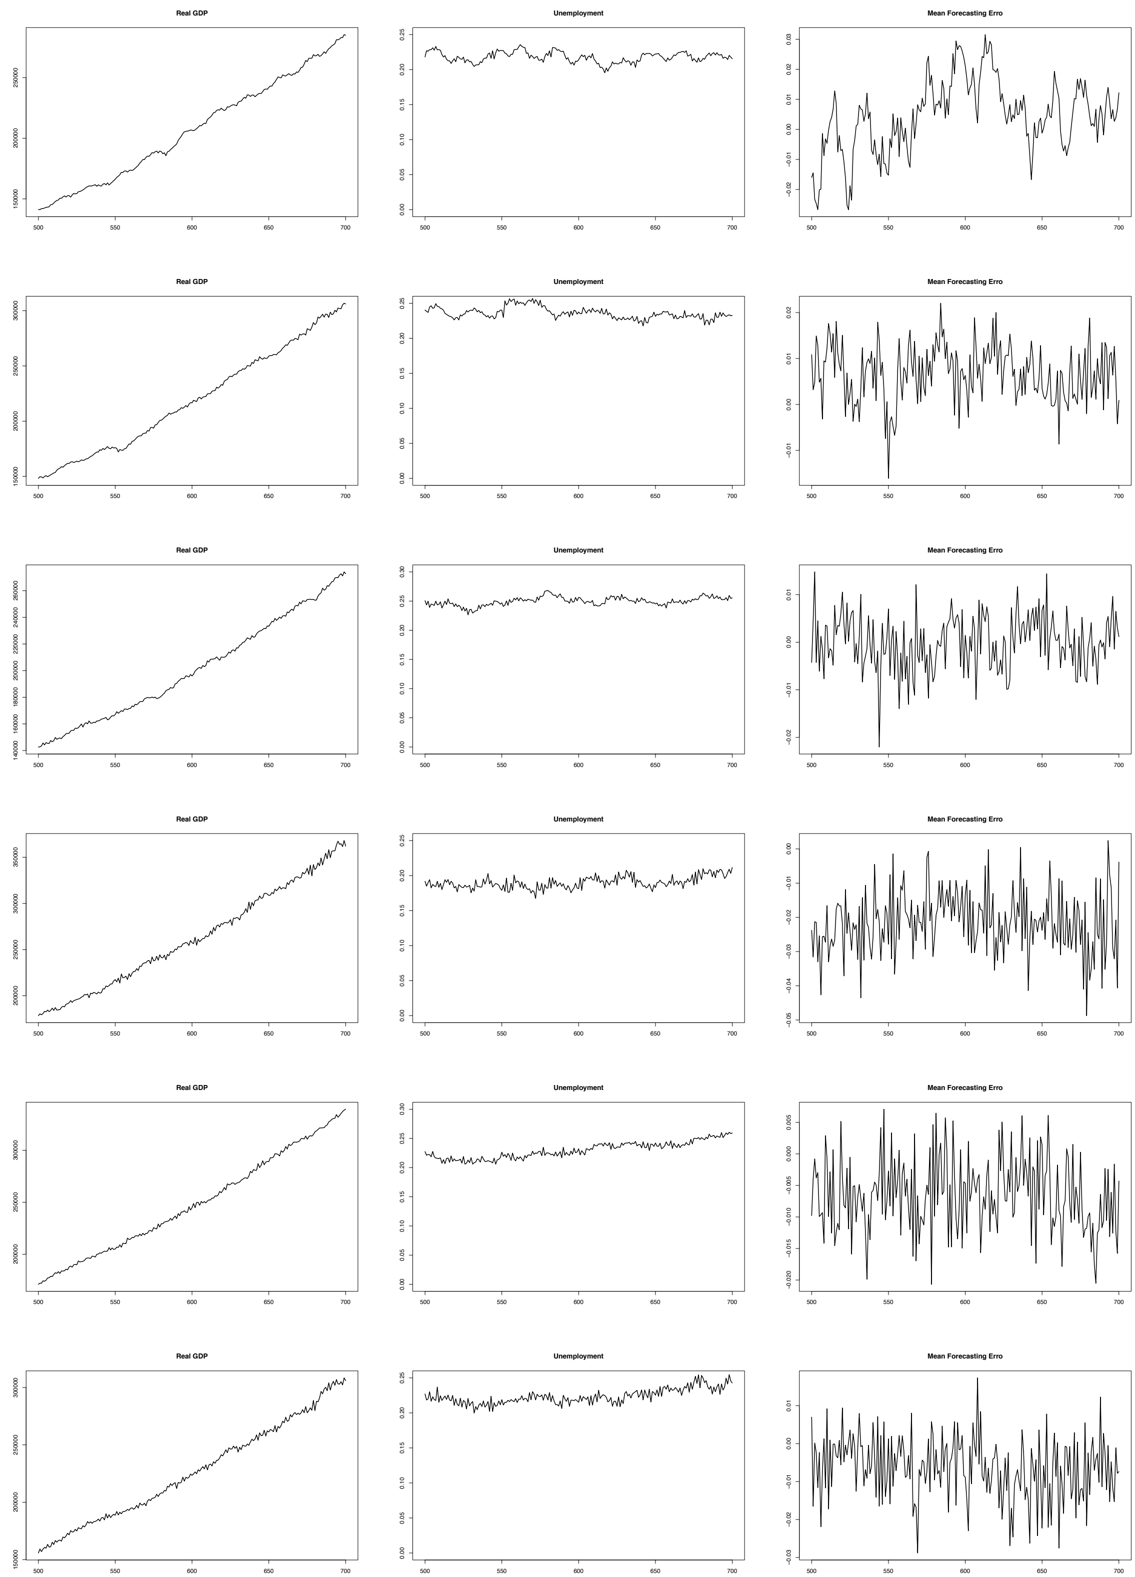

First row, Adaptive; Second row, Augmented Adaptive; Third row, Naive; Fourth row, Strong Trend Following; Fifth row, Weak Trend Following; Sixth row, Augmented Trend Following expectations

Chapter 3

Challenges for Macroeconomic ABMs

Abstract

In this paper I discuss some of the main open challenges faced by ABM modellers, specifically how to bridge models to data and how to deal with the Lucas critique.
The problem of model estimation/validation has attracted much interest within the ABM community and ongoing research on the topic is producing interesting insights and techniques. The paper summarises such new developments, moreover it proposes a taxonomy in order to match modelling strategies with appropriate estimation/validation techniques and proposes a possible strategy in order to validate model results.
The paper also tries to clarify the particular challenges posed by the Lucas critique for ABMs and suggests possible ways to overcome them, in particular it discusses how *stock-flow-consistency* can be interpreted as a necessary, although not sufficient condition, to address the critique.
The paper concludes by presenting a preliminary meta-analysis in which it tries to assess the state of research in ABM with respect to the two aforementioned issues.

3.1 Introduction

Agent-Based models (ABMs, hereafter) represent a relatively young approach in macroeconomics, as such, it brings interesting methodological novelties, unconventional insights, internal issues, and sometimes it attracts casual criticisms. The goal of this paper is to provide a stylised description of the methodology and to discuss some critical aspects. Let me start defining ABMs, which are *complex adaptive systems in which heterogenous, bounded rational agents interact locally and directly.* Being a complex system implies that aggregate dynamics cannot be understood by *ex-ante* aggregation, that is to say aggregating individuals in a representative agent, see Kirman (1992). On the contrary, an ABM can only be aggregated *ex-post*, that is to say: once agents have taken decisions and accordingly taken actions, the aggregate output can be computed as the sum of individual outputs. Therefore, in ABMs we model the individual entities composing the economic system, their behaviours, interaction structures, etc. and observe where the system lead in terms of aggregate dynamics. In such framework equilibrium is not exogenously imposed, it rather is a possible *emergent* outcome, see Delli Gatti et al. (2010a), as they are out-of-equilibrium dynamics. Therefore, in such system dominated by non-linearities and fundamental uncertainty the strict notion of perfect rationality is not applicable, which by no means imply that agent are necessarily irrational. Indeed, let us take the perfect rational agent commonly modelled in economics, she has two distinctive features: she maximises her utility under various constraint and has rational expectations about the future. In order to maximise utility, at the very least, one has to know the entire set of possible strategies and be able to construct a mapping between the set of possible strategies and the set of *objective* outcomes, that is outcomes not *subjectively* evaluated in light of the utility function. First of all it is not always true that individuals are aware of all the possible strategies at their disposal, although let us assume they are: in such a system it is impossible to know/learn a precise and global mapping between strategies and outcome. Similarly, in system dominated by non-linearities and fundamental uncertainty it makes little sense to express future events in terms of probability distributions. Therefore, being bounded rational is not a second best, in fact in such an environment *procedural* rationality *a lá* Simon (1976) allows agent to locally and adaptively search for *better* performing strategies, see Epstein (2006) and Arthur (1994), which is the best one can do given the environmental constraints.

Heterogeneity is sometimes perceived as *the* distinctive feature of agent-based models, but this is not entirely correct. Indeed, what makes ABMs unique in the realm of macroeconomic modelling is the direct and local *interactions* among heterogenous agent, or as defined by Caverzasi and Russo (2018) a *strong* type of heterogeneity. This is not just a nuance, in fact the real difference between a DSGE, for example, and an ABM is that the former implicitly assume a complete network of agents, whereas the latter allow for any topology, moreover the network topology can be exogenously imposed or evolve endogenously. This feature is central in studying how local shocks are transmitted to the wider economic system, indeed if shocks are amplified or tamed is likely to depend on the particular topology in place. This has major consequences for policies, indeed so called prudential policies are aimed at designing topologies which are more likely to tame unexpected shocks, while policies directed at counteracting a shock in place have to take into account the particular network in order to be effective, as the same shock may have different implications depending on the system topology. This

discussion obviously hints at the problem of financial contagion and bankruptcy cascade, as described in the financial accelerator model of Delli Gatti et al. (2010b). In this cases the nodes of the network are heterogenous and competing, however it is sometimes interesting to analyse network relations among *quasi-homogenous* agents, i.e. heterogenous agents belonging to the same class, not necessarily competing. A case in point is how information spread across agents and how this affects aggregate dynamics. For example in the second chapter of this dissertation I used a genetic algorithm in order to allow the updating of expectations across agents who are not in any kind of competition.
Thus, direct and local interactions are key to study a number of interesting macroeconomic questions, and ABMs provide a suitable framework in such respect.

There is a clear trend in the popularity of macroeconomics ABMs, this can be proxied by the number of papers, researchers and topics addressed using such models in the last 10 years or so. One reason, beside the methodological ones listed above, is probably the quality (micro, meso, and macro) and amount of simulated data ABMs can produce, and therefore the wide range of stylised facts ABMs can potentially match, see Fagiolo and Roventini (2016) and Haldane and Turrell (2018). However, there are many aspects to be improved and clarified, the aim of the paper is precisely to discuss some of those aspects. In particular, I will address a practical and a theoretical issue: how to bring ABMs to the data and how to address the Lucas critique.

3.2 Bridging ABMs To Data

Even though at the moment a solid and commonly accepted methodology to estimate and validate ABMs is missing, the necessity to provide empirical underpinning for macroeconomic models is not underestimated within the ABM community, see for example LeBaron and Tesfatsion (2008). Indeed, improvements in this respect are ongoing and if it is premature to provide a full assessment of this strand of research[1], it is at least fair to say that estimation/validation techniques for macro ABMs are neither as developed nor as frequently implemented as in other area of macro, like DSGEs[2], or other areas of ABMs applications, like finance[3].
The main goal of this section is therefore to provide a review of the techniques so far designed for the estimation/validation of macro ABMs and to suggest a mapping between modelling approaches and empirical strategies. Moreover, I will informally discuss some empirical strategies common in other area of economics which could potentially be adapted for ABMs. Before turning to the bulk of this section, let me introduce a broad distinction between two modelling approaches undertook by ABM modellers: the first I shall call the *strongly quantitatively* approach, the second *toy model* approach. As the name suggests, the former primary goal is to quantitatively match the real world data generating process (rdgp, hereafter) and to provide precise conditional predictions of the kind: *twisting a policy variable by x% affects some macroeconomic aggregate by y%.* On the other hand, the *toy model* approach aims at qualitatively approximating the rdgp and at providing scenario analysis and qualitative understanding of economic mechanisms. Both approaches require to be em-

[1]See Fagiolo et al. (2019) for a review.
[2]See Canova (2011) for an in-depth survey of the DSGEs estimation techniques.
[3]See Lux and Zwinkels (2018).

pirically validated, although different goals call for different strategies; to further elaborate on this point, let me use the classification proposed by Windrum et al. (2007) and Fagiolo et al. (2007) who distinguish between two[4] main approaches for ABMs empirical validation: *the indirect calibration approach* and *the Werker-Brenner approach*[5].

The indirect calibration approach follows a 4 steps procedure: (i) The modeller picks a set of empirical regularities she is interested to reproduce and explain. Such empirical regularities can be of any kind, for example it is possible to match macro auto-correlations, cross-correlations, volatilities, etc., or micro data, like firms size distribution, income distributions, wealth distributions, etc.; (ii) according to the type of empirical regularities the modeller wishes to reproduce, she models agents' behaviour and interactions. Unlike other macroeconomic models, ABMs do not force an *instrumentalist* approach, therefore the modeller tries to model behaviours and interactions in an - as realistic as possible - way, even relying on experimental evidence when available; (iii) the third step entails a more or less informal calibration exercise, where the parameter space is restricted in order to improve the match between model outputs and empirical regularities picked in stage (i). Sometimes, the modeller uses a mixed calibration strategy, where some parameters are directly estimated using real data, whereas the rests are indirectly calibrated; (iv) the final step is to use the model to learn causal mechanisms producing the empirical regularities picked in stage (i), or using the model as laboratory to produce new empirical regularities under particular conditions, i.e. the so called scenario analysis.

The indirect calibration approach is a rather informal way to choose model parameters. Indeed, under the hypothesis that the model is structural, there exists a parameter configuration which better than the others approximates the real world data generating process. However, the indirect calibration approach is not concerned to find it, it contents with finding a good approximation, not the best given the model at hand. Arguably, the goal of the indirect calibration procedure is to provide a qualitative approximation of the real world data generating process, and so it suites *toy models* for which is enough to provide a qualitative representation of the causal structure of the real world. So, indirect calibration is certainly a good starting point to validate a *toy* ABM, however later on I will argue that there are possible ways to go beyond such simple procedure in order to strengthen the empirical validity of *toy* ABMs.

The Werker-Brenner approach shares steps (i) and (ii) of the indirect calibration approach, which simply prescribe to choose the relevant empirical regularities and design the model accordingly. After having done that, the Werker-Brenner approach recommends to restrict the parameter space by choosing parameters ranges consistently with empirical observations. Then, Werker and Brenner suggest to explore the restricted parameter space in search of the configurations providing the *best* fit with the data.

This approach is better suited for *strongly quantitative* models, indeed it is not only concerned with providing a qualitative approximation of the real world dgp, its goal is to provide the best approximation to the dgp given the model at hand.

The main difficulties with the Werker-Brenner approach is how to choose among different model configurations and how to perform the parameter exploration. As we will see a few

[4]To be precise, Windrum et al. (2007) put forward a third validation approach called *the history friendly approach*. This approach is well suited for models dealing with industries or very specific economic entities, but it is not applicable for macro models.

[5]See Werker and Brenner (2004) for a deep in-depth explanation and Brenner (2006) for an early application

attempts have been made in order tackle both issues.

3.2.1 Validating a strongly quantitative model

Theoretically, the Werker-Brenner approach requires to explore the entire set of possible parameters configurations, or a subset of it in case empirical observations rule out *a priori* some parameters values. Such exploration can be in general very costly in terms of computational time, if not unfeasible given the size and computational requirements of many modern ABMs. Therefore, the first issue to face in order to devise an estimation technique for ABMs is how to reduce the computational time required for parameters exploration. The strategy undertook by Barde and Van Der Hoog (2017) and Lamperti et al. (2018) can be summarised in four steps: (i) define a measure of the distance between model output and empirical observations; (ii) run the model for a subset of the parameters space and calculate the measure defined in point (i) for each simulation run; (iii) use the points calculated in step (ii) to estimate the model response across the whole parameters space, that is to say, estimate the surface associating each point in the parameter space with a value of the measure defined in step (i); (iv) minimize (maximize) the surface obtained in point (iii) in order to single out the parameter configuration providing the best fit.

Barde and Van Der Hoog (2017) design an efficient algorithm to circumvent the issue of running lengthy simulation batteries. At the first stage Barde and Van Der Hoog suggest to reduce the parameter space using a Nearly-Orthogonal Latin Hypercube sampling (NOLH) method[6] and only run the parameter configurations surviving NOLH. Note that other strategies to reduce ex-ante the parameter space can be used, for example it would be possible to follow Werker-Brenner in reducing the range of parameters values using empirical observation and afterwards further reducing the subsample using NOLH. The novelty introduced by Barde and Van Der Hoog is the second step of the procedure, where they use the Markov Information Criterion (MIC) developed in Barde (2017) to score the simulated data against real data. So, to each configuration subsampled in stage 1 is attached a fitness indicator defined as the MIC score. At this point the MIC scores can be interpolated by means of stochastic kriging in order to produce a "MIC response surface", which can be subsequently minimised in order to find the best model configuration, i.e. the one bearing the lowest MIC score.

In the same spirit Lamperti et al. (2018) propose a machine learning approach to perform parameter space explorations. As in Barde and Van Der Hoog (2017) the first step is to define a criteria to measure the distance between model outputs and real data. For example, assume $v()$ is a function mapping a vector of model parameters x to β. Where β is a vector defining the distance between a set of moments generated by the model when is parametrised by x and the same set of moments measured observed in the real world. Then, we may want to single out all the x such that $\{x : v(x) < \alpha\}$ where α is some accuracy level defined by the researcher. As usual the problem is to infer the behaviour of the model across the entire parameter space from a finite number of points, which they address by employing an "extreme gradient boosted trees" (XGBoost). As a kriging, XGBoost allows to infer a surface model response starting from a predefined set of points in the parameter space. With the advantage that, unlike kriging, the XGBoost algorithm does not necessarily produce a smooth surface, as it is more sensitive to non-linearities in the mapping between parameter configurations

[6]See Cioppa and Lucas (2007) for an in-depth description of the NOLH methodology

and model output, typical of ABMs.
To conclude this section I should mention that beside techniques specifically designed for ABMs estimation, it is possible to adapt more traditional approaches to the ABM framework, this is for example the case of simulated minimum distance, see Grazzini and Richiardi (2015), or bayesian estimation, see Grazzini et al. (2017). The positive thing of such methodologies is that they are grounded in well established statistical theory and only require minor adjustments in order to be adapted to ABMs, however the main problem is that they rely on extensive model simulations which may or may not be feasible depending on the particular model one wish to estimate.

3.2.2 Validating a toy model

The bare goal of a toy model is to provide a logically consistent device to help devising theories, understanding transmission mechanisms, and evaluating the effects of policies and shocks of any nature. Of course, the devise used must resemble reality as much as possible, but it does not need to do so in a strictly quantitative sense. What is fundamental is that the *causal structure* of the model matches its real world counterpart.
The way most ABMs modellers have dealt so far with the issue of providing evidence that their model are indeed a good representation of the real world, is basically indirect calibration. Usually, the modeller defines a set of unconditional moments which she tries to match, by more or less informally twisting parameters' values. Such procedure is perfectly legitimate if we consider that ABM represents a very young methodology in macroeconomics, however there is probably space for improvements and I would like to use this space to highlight some interesting contributions in this direction.
At first, let us assume to be satisfied enough by matching unconditional moments. Then rather than manually twisting parameters in the attempt to match empirical regularities, it would be better to implement a more efficient parameter space exploration technique of the kind discussed in the previous section. For example the methodology put forward by Lamperti et al. (2018) may be well suited also for the calibration of toy models. Indeed, thanks to this methodology, it is possible to single out the model configurations which satisfy certain *qualitative* criteria specified as binary-outcomes. A binary-outcome can only be true or false, consider the following statement: *aggregate real investment must be more volatile than real GDP*. The model output can only be true of false in such respect. The modeller could therefore specify a set of qualitative criteria and implement a searching algorithm to find the configurations consistent with the criteria imposed.
However, matching unconditional moments does not ensure that the model is a good representation of the real world dgp. Moreover, what we really want is to model a good representation of the *causal structure* embedded in the real world. Guerini and Moneta (2017) provide a methodology to assess how close the model causal structure is to its real world counterpart, it does so by comparing a battery of SVAR estimated using simulated data against the exact same SVAR estimated using real data. In a nutshell, the procedure prescribe to estimate the reduced-form VAR real data and on simulated data, one for each round of simulation. In the second step, they estimate the causal structure for each VAR by means of causal search algorithm and finally they provide a metric to measure the distance between the model casual structure and the real world causal structure. They also provide an application of their methodology to the K+S model of Dosi et al. (2010), showing that according to their metric the model is able to match between the 65% and 85% of the estimated real world causal

structure.
Guerini and Moneta (2017) is probably the first validation methodology going in the direction of matching causal structures, however it is limited to the analysis of aggregate data and so does not exploit the full potential of an ABM, which is able to produce virtually any kind of micro, meso and macro data we can think of.

Finally, I would like to mention something which may seem obvious, that is the possibility of validating models *insights* rather than models *per se*. As we saw, proving that the model causal structure matches the real world causal structure at large is a very difficult matter. An alternative could be to test whether *local* model causal structures matches their real world counterpart. As already said toy models are designed to learn specific causal relations within the economic system. Such insights, sometimes, can be tested against real data irrespectively of the model. Of course, empirically unravelling causal mechanisms in macroeconomics can be tricky, however, as maintained by Nakamura and Steinsson (2018), applied macroeconomics have moved some considerable steps forward in such direction.
Although, such strategy seems to be seldom considered within the ABM community, it may prove a fruitful line of research.

3.3 Lucas Critique

A common criticism to ABMs is that these models supposedly fail to cope with the Lucas critique. The reason is that the use of heuristics and the technical difficulties to implement rational expectations in ABMs preempt the modeller to specify behaviours in a way that properly captures reactions to changes in the economic environment, like policy shifts or shocks of any kind.
I would argue that statements of this kind are rooted in a nearly ideological tradition, endorsing the view according to only models employing maximisation and rational expectations can be Lucas' critique robust[7]. But this is somehow arbitrary, because as far as behaviour is concerned the first and foremost goal should be modelling behaviours as realistically as possible, the second goal to express them as functions of deep structural parameters. In principle, there is no strong a priori reason to assume that parameters used in the maximisation framework are always *deeper* than any possible parameter describing any possible heuristic. This is ultimately an empirical question, which need to be addressed for each and any heuristic employed in ABMs.
For what concern expectations instead, the message of the critique, is that economic agents cannot be constantly wrong about their forecasting. Thus, rational expectations can be seen as a tool that provides unbiased expectations, but not necessarily the only one. Again, whether expectations employed in ABMs are biased is something that must be tested and cannot be assumed a priori.

In this section I will provide a more detailed discussion about modelling behaviour in ABMs in light of the Lucas' critique. Moreover I will try to stretch the boundaries of the critique a bit beyond agents' behaviour to discuss the relevance of stock-flow-consistency in this context. On the other hand, I will not discuss here issues related to expectations, as I already

[7] In truth maximization and rational expectations are not even theoretically sufficient conditions to achieve Lucas' critique robustness, see Kirman (1992) and Altissimo et al. (2002) for a discussion on this point.

did at length in chapter 2 of this dissertation.

3.3.1 Modelling behaviour

3.3.1.1 Realistic behaviour

A recognised ABMs feature is the high degree of flexibility they provide. Having many degree of freedoms is often regarded as a problem, because it favours a sort of expansionary realm of modelling practices, where every modeller employs the strategies that better suit her, impinging on model transparency and replicability. The positive side is however that the flexibility in modelling behaviours allows to reduce such degrees of freedom any time empirical evidence imposes it.

The question is rather how to acquire empirical evidence about human behaviours, which can be readily implemented in ABMs. According to Colasante (2017) cross-contamination between experimental economics and ABMs can be a fruitful line of research in order to provide a more sound model validation. In fact, evidence from experimental studies can restrict the range of behavioural rules ABM modellers choose from and also restrict parameter ranges which define such rules. Moreover, as suggested by Colasante (2017), by using the very same institutional framework to design a model and an experiment, it is possible to directly validate the model in terms of behavioural rules and parameters describing them, as long as model output of course. Indeed, thanks to the experimental design, we can observe how individuals behave in a controlled environment and readly compare them to agents in the model. Clearly, such methodology comes with all the drawbacks typical of experimental economics, small samples being the most obvious one.

Alternative empirical strategies to learn about agents' behaviour are survey studies or classical econometric analysis, especially those conducted using big-data[8], which can be very informative about micro relationships of the kind needed in this context.

Finally, Dawid and Harting (2012) suggest an interesting *pseudo* empirical strategy to model firms behaviour , that is to survey best practises endorsed by managerial sciences. This is not exactly an empirical analysis, however, under the assumption that management textbooks shape managers behaviour, it may provide sound indications about firms behaviours, difficult to grasp otherwise.

3.3.1.2 Testing heuristics

Are heuristics used in ABMs inconsistent with the Lucas' critique? We simply do not know, because they are not being tested. Whether a heuristic is robust to the critique or not depends on how deep the parameters defining it are, that is to say to what extent they are exogenous to shocks. This is ultimately an empirical question which can be addressed in two different ways. Sometimes heuristics can be directly estimated by means of standard regression analysis, which has a double advantage: it is possible to test for structural breaks in the relevant parameters in order to assess the degree of exogeneity. Moreover, estimating the relevant parameters gives an indication of the realistic values to be plugged into the model. Note that this strategy is common practice in the NK field. In fact, foundational elements of the NK-model are the so called rigidities, which are heuristics themselves and as such are being tested against empirical data. See for example the empirical literature trying

[8]See D'Orazio (2017) for a wider discussion about big-data and ABMs

to estimate the frequency of price changes, i.e. the degree of price stickiness, Bils and Klenow (2004), Nakamura and Steinsson (2008), and Klenow and Kryvtsov (2008).
However, it is not always possible to directly estimate heuristics. In such cases we must rely on indirect estimation of the kind I surveyed in the previous section. A case in point is Barde and Van Der Hoog (2017), who provide an application of their validation strategy to the Eurace@Unibi model presented in Dawid et al. (2016). Using their validation method they are able to estimate some of the structural parameters of the model. I would argue that it is in principle possible, although computationally demanding, to repeat such exercise at different points in time so to assess parameters stability across policy regimes and shocks historically occurred and temporally well identified.

3.3.2 Stock-flow consistency as a necessary antidote to the Lucas critique

Stock-flow consistency (SFC hereafter) is a methodology put forward for the first time in the ground-breaking book of Godley and Lavoie (2006), which prompted a vast macroeconomic literature (Caverzasi and Godin, 2014). SFC models developed independently from ABM, however as documented by Di Guilmi (2017) a growing number of researchers are working towards an integration of the two approaches.
In a nutshell, stock-flow consistency imposes accounting discipline in the model. It follows that, except for physical capital, each and any asset owned by each and any agent must have a liability counterpart recorded in an other agent's balance sheet. Also, each flow is intended as a vector moving chunks of stocks from one balance sheet to another. This implies that each expenditure of one agent is a source of income for another agent.
At first glance, SFC may seem to have little to do with the Lucas critique, however a closer inspection reveals an obvious nexus: agents' behaviours essentially depend on parameters *and* so called state variables. From the point of view of an agent, her own balance sheet is a matrix of state variables which heavily influence behaviour. For example, a household's consumption decision heavily depends on changes in her deposits amount or debt. Therefore, accounting inconsistencies, which fail to properly consider each any balance sheet effect shocks might have, impinge on the model ability to properly predict the economy reaction to shocks of any kind.
Therefore, the tendency of ABM modellers to impose a SFC macroeconomic structure can be also interpreted in light of the Lucas critique, insofar it provides a dimension of robustness to the critique sometimes overlooked in macroeconomics.

3.4 A preliminary meta-analysis

In the previous sections I discussed some controversial points and ongoing methodological developments within the ABM framework, specifically as far as empirical validation/estimation is concerned. The aim of this section is to preliminary assess whether current modelling practises have already internalised the aforementioned methodological advancements and directly addressed issues related to the Lucas' critique. In order to do so, I surveyed a sample of macro ABM papers published starting from 2010, focusing in particular on whether modellers provide (i) rigorous estimations; (ii) informal calibrations; (iii) insights validation; (iv) directly address the Lucas critique; (v) SFC macro structure.

Following the taxonomy laid down in the previous section, I will consider rigorously estimated any paper employing methodologies discussed in section (3.2) regardless being toy models or otherwise. On the other hand, I will consider informally calibrated any paper showing some sort of matching between simulated and real data, but not discussing how parameters has been singled out to obtain such matching. Insight validation refer to the last paragraph of section (3.2.2), i.e. the possibility to single out economic insights by the model, express them in terms of empirically testable equations, and estimating them by means of standard econometric analysis. Admittedly "directly address the Lucas critique" is quite vague and in fact it is intended to be a miscellaneous of either attempt to justify heuristics on the bases of empirical/experimental soundness or implementation of modelling strategies which somehow go in the direction of tackling the Lucas' critique. Finally, in this context I will consider SFC models, only those paper that state it clearly and with proper references.

Table 3.1: Meta Analysis

Paper	**Estimation**	**Calibration**	**Lucas Critique**	**SFC**
Bouchaud et al. (2017)			Euler Equation	x
Chen and Desiderio (2018)		x		x
Delli Gatti and Desiderio (2015)				x
Dosi et al. (2015)		x		
Giri et al. (2019)		x		
Gualdi et al. (2017)			Euler Equation	
Popoyan et al. (2017)		x		x
Salle (2015)			Learning	
Salle et al. (2013)			Learning	
Salle and Seppecher (2018)		x		x
Schasfoort et al. (2017)		x		x
van der Hoog (2018)		x	Survey	x
van der Hoog and Dawid (2017)		x		
Dosi et al. (2017b)		x		
Dosi et al. (2018b)				
Caiani et al. (2019)		x		x
Dosi et al. (2013)		x		
Russo et al. (2007)		x		
Caiani et al. (2018a)		x		x
Dawid et al. (2014)	x		Manag. Literature	
Caiani et al. (2016)		x		x
Dosi et al. (2010)		x		
Lengnick (2013)		x		x
Riccetti et al. (2015)		x		x
Cincotti et al. (2010)				x
Raberto et al. (2011)				x
Delli Gatti et al. (2010b)		x		
Gurgone et al. (2018)				x
Klimek et al. (2015)		x		x

Table (3.1) summarises the survey conducted over 29 papers published from 2010 onwards. Admittedly, the survey is just preliminary since many more papers should have been included in order to provide a more representative picture of recent developments in the macro ABM literature. Nevertheless, some interesting preliminary results emerge: the first one is that none of the papers surveyed tries to empirically test results independently from the model, that is by using standard econometric analysis. The second is that model estimation remains mostly in the methodology domain and still struggle to become an established standard for ABMs. On the other hand, most of the papers surveyed provide some sort of informal calibration, in which matching stylised facts is taken as a means for model validation. Interestingly, slightly more than half of the papers surveyed impose a macro SFC structure, suggesting a strong trend in the literature. Finally, I found Lucas critique to be somehow indirectly addressed in three ways: (i) using a consumption function which somehow tries to incorporate some Euler equation logic; (ii) learning, as a way to introduce flexibility in agents' behavioural rules and therefore allowing them to adapt to sudden changes in the environment, including policy shocks; (iii) gaining insights from surveys and managerial literature in order to model behaviour.
What emerges from this preliminary survey is an overall need to better integrate methodological advancements in agent-based macro modelling, especially when it comes to bridge models with data. On the other hand, good modelling practises like combining SFC with ABMs seems to have become common practise within the ABM community. This is a case in point where a methodological reflexion made its way through modelling and, as maintained in previous sections, it is a very positive development.
Finally, the issue of Lucas critique is much more subtle and will require time to be deeply investigated. Hopefully, a combination of theoretical analysis and empirical studies will clarify what it really means the critique in the ABM framework and possibly devise systematic modelling practises robust to it.

3.5 Conclusions

In this paper I highlighted some challenges affecting ABMs and surveyed ongoing research trying to address them. The survey shows that considerable effort has been put in devising techniques to estimate and validate ABMs. Moreover, I proposed a mapping between two broad model categories and empirical strategies, specifically for *strongly quantitative* and *toy* models. I also suggested a possible way to validate model *insights* when proper model estimation is not feasible or not interesting *per se*. Which I referred to as *insights* validation, and it simply suggests to gather insights learnt from the model, express them in terms of testable equations and conduct standard econometric analysis.
The second challenge I addressed is how the Lucas' critique relates to ABMs, I tried to argue that whether ABMs fail to cope with the critique is not obvious and to some extent difficult to assess, insofar it is not known wether heuristics commonly employed in ABMs are unstable in the face of policy shocks. This is ultimately an empirical question and should therefore be investigated using standard empirical techniques. Anyway, I also stressed that ABMs flexibility allows to change behavioural rules quite easily and therefore as empirical research selects better rules over worse ones, ABMs can accommodate for such new knowledge with relatively little effort and time. Finally, I argued that SFC may be seen as necessary, although not sufficient condition, to deal with the critique.

I also conducted a preliminary meta-analysis trying to assess the state of the arts in modelling practises, in particular I investigated whether methodological advancements has been in integrated in recent papers. Findings suggest that empirical techniques are still not much integrated in ABMs, despite advancements in the methodological literature. On the other hand, SFC seems to have become standard practise in the literature. Finally, ways to directly address the Lucas critique are various and seldom openly stated in ABM papers.

Bibliography

Acemoglu, D. and Restrepo, P. (2017). Robots and jobs: Evidence from us labor markets. *NBER working paper*, (w23285).

Alogoskoufis, G. S. and Smith, R. (1991). The phillips curve, the persistence of inflation, and the lucas critique: Evidence from exchange-rate regimes. *The American Economic Review*, pages 1254–1275.

Altissimo, F., Siviero, S., and Terlizzese, D. (2002). How deep are the deep parameters? *Annales d'Economie et de Statistique*, pages 207–226.

Arifovic, J. (1991). *Learning by genetic algorithms in economic environments*. PhD thesis, University of Chicago, Department of Economics.

Arifovic, J. (1994). Genetic algorithm learning and the cobweb model. *Journal of Economic dynamics and Control*, 18(1):3–28.

Arifovic, J. (1995). Genetic algorithms and inflationary economies. *Journal of Monetary Economics*, 36(1):219–243.

Arifovic, J. (2000). Evolutionary algorithms in macroeconomic models. *Macroeconomic Dynamics*, 4(3):373–414.

Arifovic, J., Bullard, J., and Kostyshyna, O. (2012). Social learning and monetary policy rules. *The Economic Journal*, 123(567):38–76.

Arntz, M., Gregory, T., and Zierahn, U. (2016). The risk of automation for jobs in oecd countries.

Arthur, W. B. (1994). Inductive reasoning and bounded rationality. *The American economic review*, 84(2):406–411.

Assenza, T., Gatti, D. D., and Grazzini, J. (2015). Emergent dynamics of a macroeconomic agent based model with capital and credit. *Journal of Economic Dynamics and Control*, 50:5–28.

Autor, D. H., Dorn, D., and Hanson, G. H. (2015). Untangling trade and technology: Evidence from local labour markets. *The Economic Journal*, 125(584):621–646.

Autor, D. H., Levy, F., and Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. *The Quarterly journal of economics*, 118(4):1279–1333.

Autor, H. and Dorn, D. (2013). The growth of low-skill service jobs and the polarization of the us labor market. *American Economic Review*, 103(5):1553–97.

Bárány, Z. L. and Siegel, C. (2018). Job polarization and structural change. *American Economic Journal: Macroeconomics*, 10(1):57–89.

Barde, S. (2017). A practical, accurate, information criterion for nth order markov processes. *Computational Economics*, 50(2):281–324.

Barde, S. and Van Der Hoog, S. (2017). An empirical validation protocol for large-scale agent-based models.

Berg, A., Buffie, E. F., and Zanna, L.-F. (2018). Should we fear the robot revolution?(the correct answer is yes). *Journal of Monetary Economics*, 97:117–148.

Bils, M. and Klenow, P. J. (2004). Some evidence on the importance of sticky prices. *Journal of political economy*, 112(5):947–985.

Blanchard, O. J. (1984). The lucas critique and the volcker deflation.

Bouchaud, J.-P., Gualdi, S., Tarzia, M., and Zamponi, F. (2017). Optimal inflation target: insights from an agent-based model. economics discussion papers, no 2017-64. *Kiel Institute for the World Economy. http://www. economics-ejournal. org/economics/discussionpapers/2017-64 Page*, 2(2).

Brenner, T. (2006). Agent learning representation: advice on modelling economic learning. *Handbook of computational economics*, 2:895–947.

Caiani, A., Catullo, E., and Gallegati, M. (2018a). The effects of fiscal targets in a monetary union: a multi-country agent-based stock flow consistent model. *Industrial and Corporate Change*, 27(6):1123–1154.

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 and Control*, 69:375–408.

Caiani, A., Russo, A., and Gallegati, M. (2018b). Are higher wages good for business? an assessment under alternative innovation and investment scenarios. *Macroeconomic Dynamics*, pages 1–40.

Caiani, A., Russo, A., and Gallegati, M. (2019). Does inequality hamper innovation and growth? an ab-sfc analysis. *Journal of Evolutionary Economics*, 29(1):177–228.

Canova, F. (2011). *Methods for applied macroeconomic research*. Princeton university press.

Catullo, E., Gallegati, M., and Palestrini, A. (2015). Towards a credit network based early warning indicator for crises. *Journal of Economic Dynamics and Control*, 50:78–97.

Catullo, E., Gallegati, M., and Russo, A. (2019). Forecasting in a complex environment: machine learning sales expectations in a sfc agent based simulation models. Unpublished.

Caverzasi, E. and Godin, A. (2014). Post-keynesian stock-flow-consistent modelling: a survey. *Cambridge Journal of Economics*, 39(1):157–187.

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.

Chang, Y., Kim, S.-B., and Schorfheide, F. (2010). Labor market heterogeneity and the lucas critique. *Unpublished manuscript, University of Rochester*.

Chen, S. and Desiderio, S. (2018). Computational evidence on the distributive properties of monetary policy. *Economics: The Open-Access, Open-Assessment E-Journal*, 12(2018-62):1–32.

Chiacchio, F., Petropoulos, G., and Pichler, D. (2018). The impact of industrial robots on eu employment and wages: A local labour market approach. bruegel working paper issue 02/18 april 2018.

Ciarli, T., Lorentz, A., Savona, M., and Valente, M. (2010). The effect of consumption and production structure on growth and distribution. a micro to macro model. *Metroeconomica*, 61(1):180–218.

Ciarli, T., Marzucchi, A., Salgado, E., and Savona, M. (2018). The effect of r&d growth on employment and self-employment in local labour markets.

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.

Cioppa, T. M. and Lucas, T. W. (2007). Efficient nearly orthogonal and space-filling latin hypercubes. *Technometrics*, 49(1):45–55.

Colasante, A. (2017). Experimental economics for abm validation. In *Introduction to Agent-Based Economics*, pages 143–162. Elsevier.

Cuthbertson, K. and Taylor, M. P. (1990). "the case of the missing money" and the lucas critique. *Journal of Macroeconomics*, 12(3):437–454.

Dawid, H., Gemkow, S., Harting, P., Kabus, K., Neugart, M., and Wersching, K. (2008). Skills, innovation, and growth: an agent-based policy analysis. *Jahrbücher für Nationalökonomie und Statistik*, 228(2-3):251–275.

Dawid, H. and Harting, P. (2012). Capturing firm behavior in agent-based models of industry evolution and macroeconomic dynamics. *Applied Evolutionary Economics, Behavior and Organizations, ed. by G. Bünsdorf, Edward-Elgar*, pages 103–130.

Dawid, H., Harting, P., and Neugart, M. (2014). Economic convergence: policy implications from a heterogeneous agent model. *Journal of Economic Dynamics and Control*, 44:54–80.

Dawid, H., Harting, P., van der Hoog, S., and Neugart, M. (2016). The use of simulations in developing robust knowledge about causal processes: Methodological considerations and an application to industrial evolution. *Bielefeld Working Papers in Economics and Management*.

Dawid, H. and Kopel, M. (1998). On economic applications of the genetic algorithm: a model of the cobweb type. *Journal of Evolutionary Economics*, 8(3):297–315.

DeCanio, S. J. (2016). Robots and humans–complements or substitutes? *Journal of Macroeconomics*, 49:280–291.

Deissenberg, C., Van Der Hoog, S., and Dawid, H. (2008). Eurace: A massively parallel agent-based model of the european economy. *Applied Mathematics and Computation*, 204(2):541–552.

Delli Gatti, D. D. and Desiderio, S. (2015). Monetary policy experiments in an agent-based model with financial frictions. *Journal of Economic Interaction and Coordination*, 10(2):265–286.

Delli Gatti, D. D., Gaffeo, E., and Gallegati, M. (2010a). Complex agent-based macroeconomics: a manifesto for a new paradigm. *Journal of Economic Interaction and Coordination*, 5(2):111–135.

Delli Gatti, D. D., Gaffeo, E., Gallegati, M., and Palestrini, A. (2005). The apprentice wizard: Montetary policy, complexity and learning. *New Mathematics and Natural Computation*, 1(01):109–128.

Delli Gatti, D. D., Gallegati, M., Greenwald, B., Russo, A., and Stiglitz, J. E. (2010b). The financial accelerator in an evolving credit network. *Journal of Economic Dynamics and Control*, 34(9):1627–1650.

Di Guilmi, C. (2017). The agent-based approach to post keynesian macro-modeling. *Journal of Economic Surveys*, 31(5):1183–1203.

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

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*, 52:166–189.

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., Napoletano, M., Roventini, A., Stiglitz, J. E., and Treibich, T. (2017a). Rational heuristics? expectations and behaviors in evolving economies with heterogeneous interacting agents. *Expectations and Behaviors in Evolving Economies with Heterogeneous Interacting Agents (December 8, 2017).*

Dosi, G., Pereira, M. C., Roventini, A., and Virgillito, M. E. (2017b). When more flexibility yields more fragility: the microfoundations of keynesian aggregate unemployment. *Journal of Economic Dynamics and Control*, 81:162–186.

Dosi, G., Pereira, M. C., Roventini, A., and Virgillito, M. E. (2018a). Causes and consequences of hysteresis: aggregate demand, productivity, and employment. *Industrial and Corporate Change*, 27(6):1015–1044.

Dosi, G., Pereira, M. C., Roventini, A., and Virgillito, M. E. (2018b). The effects of labour market reforms upon unemployment and income inequalities: an agent-based model. *Socio-Economic Review*, 16(4):687–720.

D'Orazio, P. (2017). Big data and complexity: Is macroeconomics heading toward a new paradigm? *Journal of Economic Methodology*, 24(4):410–429.

Engle, R. F. and Hendry, D. F. (1993). Testing superexogeneity and invariance in regression models. *Journal of Econometrics*, 56(1-2):119–139.

Epstein, J. M. (2006). Remarks on the foundations of agent-based generative social science. *Handbook of computational economics*, 2:1585–1604.

Evans, G. W. and McGough, B. (2005). Monetary policy, indeterminacy and learning. *Journal of Economic Dynamics and Control*, 29(11):1809–1840.

Fagiolo, G., Guerini, M., Lamperti, F., Moneta, A., and Roventini, A. (2019). Validation of agent-based models in economics and finance. In *Computer Simulation Validation*, pages 763–787. Springer.

Fagiolo, G., Moneta, A., and Windrum, P. (2007). A critical guide to empirical validation of agent-based models in economics: Methodologies, procedures, and open problems. *Computational Economics*, 30(3):195–226.

Fagiolo, G. and Roventini, A. (2016). Macroeconomic policy in dsge and agent-based models redux: New developments and challenges ahead. *Available at SSRN 2763735*.

Favero, C. and Hendry, F. D. (1992). Testing the lucas critique: A review. *Econometric Reviews*, 11(3):265–306.

Fischer, A. M. (1989). Policy regime changes and monetary expectations: testing for super exogeneity. *Journal of Monetary Economics*, 24(3):423–436.

Foote, C. L. and Ryan, R. W. (2015). Labor-market polarization over the business cycle. *NBER Macroeconomics Annual*, 29(1):371–413.

Franke, R. (1998). Coevolution and stable adjustments in the cobweb model. *Journal of Evolutionary Economics*, 8(4):383–406.

Frey, C. B. and Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerisation? *Technological forecasting and social change*, 114:254–280.

Gabaix, X. (2016). A behavioral new keynesian model.

Gallegati, M. and Kirman, A. (2012). Reconstructing economics: Agent based models and complexity. *Complexity Economics*, 1(1):5–31.

Giannitsarou, C. (2003). Heterogeneous learning. *Review of Economic Dynamics*, 6(4):885–906.

Gintis, H. (2007). The dynamics of general equilibrium. *The Economic Journal*, 117(523):1280–1309.

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*, 157:42–58.

Godley, W. and Lavoie, M. (2006). *Monetary economics: an integrated approach to credit, money, income, production and wealth*. Springer.

Goos, M. and Manning, A. (2007). Lousy and lovely jobs: The rising polarization of work in britain. *The review of economics and statistics*, 89(1):118–133.

Goos, M., Manning, A., and Salomons, A. (2009). Job polarization in europe. *American economic review*, 99(2):58–63.

Goos, M., Manning, A., and Salomons, A. (2014). Explaining job polarization: Routine-biased technological change and offshoring. *American Economic Review*, 104(8):2509–26.

Graetz, G. and Michaels, G. (2018). Robots at work. *Review of Economics and Statistics*, 100(5):753–768.

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

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

Gualdi, S., Tarzia, M., Zamponi, F., and Bouchaud, J.-P. (2017). Monetary policy and dark corners in a stylized agent-based model. *Journal of Economic Interaction and Coordination*, 12(3):507–537.

Guerini, M. and Moneta, A. (2017). A method for agent-based models validation. *Journal of Economic Dynamics and Control*, 82:125–141.

Gurgone, A., Iori, G., and Jafarey, S. (2018). The effects of interbank networks on efficiency and stability in a macroeconomic agent-based model. *Journal of Economic Dynamics and Control*, 91:257–288.

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

Hendry, D. F. (1988). The encompassing implications of feedback versus feedforward mechanisms in econometrics. *Oxford Economic Papers*, 40(1):132–149.

Holland, J. H. (1970). Robust algorithms for adaptation set in a general formal framework. In *1970 IEEE Symposium on Adaptive Processes (9th) Decision and Control*, pages 175–175. IEEE.

Honkapohja, S. and Mitra, K. (2004). Are non-fundamental equilibria learnable in models of monetary policy? *Journal of Monetary Economics*, 51(8):1743–1770.

Honkapohja, S. and Mitra, K. (2006). Learning stability in economies with heterogeneous agents. *Review of Economic Dynamics*, 9(2):284–309.

Jaimovich, N. and Siu, H. E. (2012). The trend is the cycle: Job polarization and jobless recoveries. Technical report, National Bureau of Economic Research.

Kaplan, G., Moll, B., and Violante, G. L. (2018). Monetary policy according to hank. *American Economic Review*, 108(3):697–743.

Kinsella, S., Greiff, M., and Nell, E. J. (2011). Income distribution in a stock-flow consistent model with education and technological change. *Eastern Economic Journal*, 37(1):134–149.

Kirman, A. P. (1992). Whom or what does the representative individual represent? *Journal of economic perspectives*, 6(2):117–136.

Klenow, P. J. and Kryvtsov, O. (2008). State-dependent or time-dependent pricing: Does it matter for recent us inflation? *The Quarterly Journal of Economics*, 123(3):863–904.

Klimek, P., Poledna, S., Farmer, J. D., and Thurner, S. (2015). To bail-out or to bail-in? answers from an agent-based model. *Journal of Economic Dynamics and Control*, 50:144–154.

Lamperti, F., Roventini, A., and Sani, A. (2018). Agent-based model calibration using machine learning surrogates. *Journal of Economic Dynamics and Control*, 90:366–389.

Lavoie, M. (1992). Foundations of post-keynesian economic analysis. *Books*.

LeBaron, B. and Tesfatsion, L. (2008). Modeling macroeconomies as open-ended dynamic systems of interacting agents. *American Economic Review*, 98(2):246–50.

Lee, N. and Clarke, S. (2019). Do low-skilled workers gain from high-tech employment growth? high-technology multipliers, employment and wages in britain. *Research Policy*, 48(9):103803.

Lengnick, M. (2013). Agent-based macroeconomics: A baseline model. *Journal of Economic Behavior & Organization*, 86:102–120.

Lucas Jr, R. E. (1976). Econometric policy evaluation: A critique. In *Carnegie-Rochester conference series on public policy*, volume 1, pages 19–46. North-Holland.

Lux, T. and Zwinkels, R. C. (2018). Empirical validation of agent-based models. In *Handbook of computational economics*, volume 4, pages 437–488. Elsevier.

Manning, A. (2004). We can work it out: the impact of technological change on the demand for low-skill workers. *Scottish Journal of Political Economy*, 51(5):581–608.

Marcellino, M. and Salmon, M. (2002). Robust decision theory and the lucas critique. *Macroeconomic Dynamics*, 6(1):167–185.

Mazzolari, F. and Ragusa, G. (2013). Spillovers from high-skill consumption to low-skill labor markets. *Review of Economics and Statistics*, 95(1):74–86.

Michaels, G., Natraj, A., and Van Reenen, J. (2014). Has ict polarized skill demand? evidence from eleven countries over twenty-five years. *Review of Economics and Statistics*, 96(1):60–77.

Muth, J. F. (1961). Rational expectations and the theory of price movements. *Econometrica: Journal of the Econometric Society*, pages 315–335.

Nakamura, E. and Steinsson, J. (2008). Five facts about prices: A reevaluation of menu cost models. *The Quarterly Journal of Economics*, 123(4):1415–1464.

Nakamura, E. and Steinsson, J. (2018). Identification in macroeconomics. *Journal of Economic Perspectives*, 32(3):59–86.

Pajarinen, M. and Rouvinen, P. (2014). Computerization threatens one third of finnish employment. *Etla Brief*, 22(13.1):2014.

Palestrini, A. and Gallegati, M. (2015). Unbiased adaptive expectation schemes. *Economics Bulletin*, 35(2):1185–1190.

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

Raberto, M., Teglio, A., and Cincotti, S. (2011). Debt deleveraging and business cycles: An agent-based perspective. *Economics Discussion Paper*, (2011-31).

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.

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.

Sachs, J. D. and Kotlikoff, L. J. (2012). Smart machines and long-term misery. Technical report, National Bureau of Economic Research.

Salle, I. and Seppecher, P. (2018). Stabilizing an unstable complex economy on the limitations of simple rules. *Journal of Economic Dynamics and Control*, 91:289–317.

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

Salle, I., Yıldızoğlu, M., and Sénégas, M.-A. (2013). Inflation targeting in a learning economy: An abm perspective. *Economic Modelling*, 34:114–128.

Salle, I., Yildizoglu, M., Zumpe, M., and Sénégas, M.-A. (2017). Coordination through social learning in a general equilibrium model. *Journal of Economic Behavior & Organization*, 141:64–82.

Salle, I. L. (2015). Modeling expectations in agent-based models—an application to central bank's communication and monetary policy. *Economic Modelling*, 46:130–141.

Schasfoort, J., Godin, A., Bezemer, D., Caiani, A., and Kinsella, S. (2017). Monetary policy transmission in a macroeconomic agent-based model. *Advances in Complex Systems*, 20(08):1850003.

Seppecher, P. (2012). Monnaie endogène et agents hétérogènes dans un modèle stock-flux cohérent. In *POLITICAL ECONOMY AND THE OUTLOOK FOR CAPITALISM (AHE, AFEP, IIPPE-JOINT CONFERENCE)*.

Seppecher, P., Salle, I., and Lang, D. (2019). Is the market really a good teacher? *Journal of Evolutionary Economics*, 29(1):299–335.

Seppecher, P., Salle, I. L., and Lavoie, M. (2018). What drives markups? evolutionary pricing in an agent-based stock-flow consistent macroeconomic model. *Industrial and Corporate Change*, 27(6):1045–1067.

Sergi, F. (2018). Dsge models and the lucas critique. a historical appraisal. *UWE Bristol, Economics Working Paper Series*.

Simon, H. A. (1976). From substantive to procedural rationality. In *25 years of economic theory*, pages 65–86. Springer.

Steindl, J. (1976). *Maturity and stagnation in American capitalism*. Number 4. NYU Press.

Taylor, J. B. (1984). Recent changes in macro policy and its effects: Some time-series evidence. *The American Economic Review*, 74(2):206–210.

van der Hoog, S. (2018). The limits to credit growth: Mitigation policies and macroprudential regulations to foster macrofinancial stability and sustainable debt. *Computational Economics*, 52(3):873–920.

van der Hoog, S. and Dawid, H. (2017). Bubbles, crashes, and the financial cycle: The impact of banking regulation on deep recessions. *Macroeconomic Dynamics*, pages 1–42.

Werker, C. and Brenner, T. (2004). Empirical calibration of simulation models. Technical report, Papers on Economics and Evolution.

Windrum, P., Fagiolo, G., and Moneta, A. (2007). Empirical validation of agent-based models: Alternatives and prospects. *Journal of Artificial Societies and Social Simulation*, 10(2):8.