

Coordination through Social Learning in a General Equilibrium Model

1st March 2017

1 Introduction

In the wake of the Lucas's critique, the General Equilibrium (GE) approach has been introduced as a way to provide microeconomic foundations to macroeconomic models. This has given rise to the so-called dynamic stochastic general equilibrium (DSGE) models. As suggested by their denomination, this class of models encapsulates two dimensions. The *intra-temporal* dimension draws the picture of the economy in every period: endogenous variables are determined by market interactions. The *inter-temporal* dimension describes the dynamics over an infinite horizon: endogenous variables depends on their expected future values.

The intertemporal dimension usually relies on rational expectations. Because this assumption has been increasingly regarded as too demanding, the literature has explored various ways of relaxing it, and assessing whether rational expectations equilibria can be viewed as the eventual outcome of a *learning* process.¹ The intratemporal aspects rely on substantive rationality and representative agents, which allow agents' behaviour to be optimized and mutually consistent, and markets to clear in every period. Those assumptions are equally strong and raise the conceptual question of what kind of learning or adaptation process has guided the economy towards equilibrium. However, the literature has remained rather silent about it. Hence, in this paper, we restrict our focus on intra-temporal aspects and addresses this question: can the simultaneous coordination on the optimal, welfare-maximizing state of a GE economy be learned by agents? or does relaxing those strong aforementioned assumptions necessarily imply deviations from

¹See especially the adaptive learning literature (Evans and Honkapohja, 2001), the heterogeneous agent literature and so-called learning-to-forecast laboratory experiments (Hommes, 2013).

equilibrium and related welfare losses? We envision 'top-down' equilibrium restrictions such as market clearing and the law of unique price not as starting points but, at best, as an (eventual) outcome of adaptation through learning. This question sounds interesting in light of the experimental literature as well, which has revealed the difficulty of the subjects to coordinate simultaneously on two interdependent markets.²

On the theoretical front, few attempts have been developed to address coordination in a GE model, e.g. the general disequilibrium theory (Bénassy, 1993) and the shadow price approach (Evans and McGough, 2015). However, these contributions do not really fit the purpose outlined above as they assume homogeneity of agents and operates on a centralized plan. We rather need a modelling approach that explicitly accounts for rationing of excess demand or supply and heterogeneity among boundedly rational agents interacting on a *decentralized* plan. We opt for an agent-based model (ABM) because agents' heterogeneity and bounded rationality are at the very heart of this modelling approach, which allows for fully individualised micro-foundations (Tesfatsion and LeBaron, 2008); emerging properties from individual local interactions can be observed at the macro level; and model analysis *via* computer simulations allows us to get rid of equilibrium constraints, thus paving the way for out-of-equilibrium economics. ABMs have first been developed for analysing the ability of agents to self-organize and the emergence of coordination problems in simple settings, inspired by game theory (Schelling, 1971; Arthur, 1994), in market games like the Cournot oligopoly (Vriend, 2000; Vallée and Yildizoglu, 2009), or simple macroeconomic models (Arifovic, 1995, 2000).³ The question posed in this line of research concerns the learnability of equilibria in these simple models, and the type of dynamics induced by learning.

We follow this literature by extending the learnability question to a *multi-market, multi-agent* model, and we use a social learning mechanism. Social learning has indeed

²Only few lab experiments have been run into a GE setting. Noussair et al (2015) run production-focused GE experiments and study firms' pricing behaviour. Fenig et al (2016) run semi-automated and entirely subject-driven experiments. With automatized firms and subject-households, labour supply is broadly in line with its equilibrium value, but output demand is consistently higher than equilibrium. In general, while coordination *towards* equilibrium is roughly observed in the lab, heterogeneity in individual behaviour leads to clashes, and sub-optimal aggregate outcomes persist. In a similar subject-driven experiments, Lian and Plott (1998) report that the level of aggregate production is consistently lower than the competitive GE equilibrium – by roughly 10%, but convergence is obtained in a weak sense (in their terminology) because the experimental economies significantly evolve towards the GE values.

³Some models also occupy a middle ground (Howitt and Clower, 2000; Gintis, 2007); see Gintis (2007) for a review of this literature.

several features that are particularly attractive for our purpose: it is intrinsically linked to agents' heterogeneity, and it is neither based on equilibrium conditions nor dependent on high-levelled rationality. What is more, social learning modelled by a genetic algorithm–type of evolutionary process is certainly the most widely used learning process in evolutionary and agent-based computational economics.⁴ Finally, genetic algorithms are well-known for fairly good convergence properties with respect to optimal equilibria (Holland et al, 1989). Due to this rather special status, it appears to be a somehow natural starting point for our investigation. To the best of our knowledge, (social) learning in a GE context has been suggested as an interesting research avenue (Arifovic, 2000), but never undertaken.⁵ Consequently, this paper investigates whether social learning could induce convergence to equilibrium in a rather non-restrictive class of general (dis)equilibrium frameworks that are simple enough to be transposed into an ABM, but encompass several key elements common to a rather wide range of macroeconomic schools of thought. We intend to characterize the type of macroeconomic dynamics induced by social learning in that kind of framework.

We proceed as follows. We develop a double specification of the framework: a *general equilibrium setting* based on homogeneity and complete information, and resulting market-clearing assumptions, and an *agent-based general disequilibrium setting* built on agents' heterogeneity, local information and resulting market rationing. Thanks to this double specification, the benchmark, welfare maximising equilibrium is perfectly known (*via* analytical resolution of the first specification), allowing for measurements (*via* simulations in the second specification) of the gap between this benchmark and the economic outcomes induced by individual interactions under learning.

The main results of our paper read as follows. Social learning gives rise to substantial coordination failures, especially between firms, which lead almost systematically to a

⁴See among others Arifovic (1995, 2000); Arifovic et al (2013); Vriend (2000); Vallée and Yildizoglu (2009).

⁵The macroeconomic frameworks in Arifovic (1995, 2000) differ substantially from agent-based principles. Arifovic et al (2013) introduce social learning in the log-linearised form of the New Keynesian model which already contains numerous market-clearing and equilibrium restrictions. Gintis (2007) is the closest in spirit to our investigation in this paper: Gintis uses an ABM of a decentralized Walrasian economy with producers and consumers engaged in repeated market interactions and evolving through imitation driven by a replicator dynamics. The author shows that the system often converges towards a steady state, in the sense of a stable attractor of the highly non-linear and complex dynamics of the ABM. However, learning only applies to the firms' operating characteristics, his ABM uses a rather specific design of agents' decision making and first order and market clearing conditions on parts of the model.

below-equilibrium state and persistent social welfare losses. Our finding echoes the conclusions drawn from the experimental literature cited above. Given that we implement a rather sophisticated form of social learning, which combines strong selectivity with minimum individual rationality principles, the observed coordination failures are likely to result from difficulties associated to the economic environment, rather than from a too restrictive design of the learning process. We therefore successively dig into the economic mechanisms that are inherent to the class of models under study, and could explain the observed miscoordination. Increasing market concentration (especially under decreasing rates of return) and the challenging simultaneous price-quantity learning problem of the firms play some role. However, the main mechanism behind coordination failures and related welfare losses is a self-reinforcing uneven distribution of income among households, resulting in a lack of aggregate demand and a stagnation of the economic system below the optimal equilibrium. This self-reinforcing mechanism echoes the amplification processes at work in the liquidity trap analysis of Eggertsson and Krugman (2012), and we draw some parallel with the recent discussions about raising income inequality as a source of economic stagnation after a financial crisis.

The remainder of this paper is organised as follows. The two model specifications are presented in Sections 2 and 3. Section 4 provides the simulation protocol. Simulation results are outlined in Section 5 and discussed in Section 6. Section 7 concludes.

2 A general equilibrium setting

This section defines a textbook GE setting which is embodied in most modern macroeconomic models (see e.g. Walsh (2003)).

2.1 Agents

We assume that the economy is populated by n households, indexed by $i = 1, \dots, n$ and m firms, indexed by $j = 1, \dots, m$. Lower-case letters stand for individual variables and capital letters for aggregate ones.

2.1.1 Households

Each household i has two choice variables: its labour supply h_i^s and its goods demand (i.e. its desired level of real consumption) c_i^d . It chooses levels of h_i^s and c_i^d so as to maximise its utility:

$$u(c_i, h_i) = \frac{c_i^{1-\sigma}}{1-\sigma} - \frac{h_i^{1+\phi}}{1+\phi}, \quad (1)$$

where h_i is the amount of labour actually worked by household i , and c_i its actual goods consumption level. The utility function $u(h_i, c_i)$ is supposed to be concave, increasing in c_i , and decreasing in h_i , implying $\sigma \geq 0$ and $\phi \geq 0$. Any agent i faces the flow (nominal) budget constraint:

$$P \cdot c_i \leq W \cdot h_i + \frac{\Pi}{n} \quad (2)$$

with W the nominal wage in the economy and P the goods price.⁶ It is usually assumed that each household i holds an equal ratio $1/n$ of the shares of each firm j , so that it receives a dividend flow equal to $1/n$ of the overall sum of nominal profits $\Pi = \sum_{j=1}^m \pi_j$ of the firms $j = 1, \dots, m$, with π_j being firm j 's individual profit.

2.1.2 Firms

Modern macroeconomics typically mobilizes monopolistic competition based on product differentiation à la Dixit and Stiglitz (1977) to model firms' market power. However, direct implementation of a monopolistically competitive environment is not feasible in the agent-based framework that we develop from Section 3 on, because it requires an infinite number of goods, which in turn requires an infinite number of firms. Therefore, we choose to model market power in a Cournot oligopoly setting that is easier to transpose to the ABM. However, it should be stressed that the differences between these two approaches are less pronounced than can appear at first sight. In fact, both modelling devices – product differentiation and Cournot competition in quantities – end up with a pricing scheme where firms apply some mark-up factor on their marginal costs as in Equation 9 below (see also Woodford (2003, p. 151) and Galí (2008, p. 44)).⁷

⁶We implicitly already assume that the law of unique prices applies, which holds at the symmetric equilibrium of the model. An alternative, less restrictive interpretation is to consider P and W as composite prices and wages.

⁷Noussair et al (2015) use an alternative approach for an implementation of the GE model in a laboratory experiment. They impose a different utility valuation of goods across consumers by using

For all these reasons, we assume that each firm j produces the same homogeneous good according to a labour-driven Cobb-Douglas production function, so that j 's goods supply is given by:

$$y_j^s = h_j^{1-\alpha}. \quad (3)$$

This production function is supposed to be increasing and concave, implying $0 \leq \alpha \leq 1$. Each firm j maximises its profit:

$$\pi_j = P \cdot y_j - W \cdot h_j, \quad (4)$$

with y_j and h_j the amounts of goods actually sold, and of labour actually hired by j .

The number of firms m is supposed to be finite, which implies imperfect competition: each firm j regards the production levels of the $m - 1$ other firms as given, and chooses the production level y_j^s , or equivalently the amount of labour input h_j^d , which provides a maximum profit. Recall that in the Cournot setting, the goods price is a function of the aggregate supply $Y^s = \sum_{j=1}^m y_j^s$:

$$P = P(Y^s). \quad (5)$$

2.2 General equilibrium

The resolution of the model requires the homogeneity of all agents of each type (*symmetry*) and market clearing (*equilibrium*): agents' *ex ante* plans are *ex post* fulfilled, and they all have the same plans.

2.2.1 Market clearing

Labour market Market clearing implies no unemployment and the firms hire the amount of labour consistent with their intended production levels, given by $h_j^d = (y_j^s)^{\frac{1}{1-\alpha}}$. As a consequence, we have perfect equality between the aggregate labour supply $H^s = \sum_{i=1}^n h_i^s$ and the aggregate actually worked labour $H = \sum_{i=1}^n h_i$. Identically, market clearing implies perfect equality between the aggregate labour demand $H^d = \sum_{j=1}^m h_j^d$ and hired labour $H = \sum_{j=1}^m h_j$, which trivially implies $H^d = H^s = H$.

individual taste shocks with different drifts for each consumer. This option would complicate the agent-based model quite substantively, while leading to the same equilibrium mark-up behaviour obtained in the Cournot market.

Goods market In the exact same way as on the labour market, the goods market must clear: all firms are able to sell their entire production, and all households end up buying the quantity of goods they desire to consume. We then have perfect equality between the aggregate production level Y^s and the aggregate goods sales $Y = \sum_{j=1}^m y_j$. Aggregate *desired* consumption $C^d = \sum_{i=1}^n c_i^d$ is also equal to aggregate consumption $C = \sum_{i=1}^n c_i$. Since consumption is the only component of demand, market clearing trivially implies $Y^s = Y = C = C^d$.

2.2.2 Resolution of the model

Symmetry of agent behaviour Agents of each type are supposed to be perfectly identical. All households necessarily choose the same quantities of labour supply and desired real consumption, denoted in what follows by h^s and c^d . In equilibrium, symmetry implies:

$$h_i^s = h^s \quad \text{and} \quad c_i^d = c^d \quad \forall i. \quad (6)$$

Similarly, we obtain for all firms the same production quantity y^s , and the same associated labour quantity $h^d = (y^s)^{\frac{1}{1-\alpha}}$:

$$y_j^s = y^s \quad \text{and} \quad h_j^d = h^d \quad \forall j. \quad (7)$$

Household first order condition Maximization of utility (1) under the budget constraint (2) gives the following first order condition:

$$(h^s)^\phi (c^d)^\sigma = \frac{W}{P}. \quad (8)$$

Note that we assume here that households are not aware that their dividend flows partially depend on their own consumption expenditures and wage earnings.

Firms' first order conditions Firms' competition in quantities implies that the goods price $P(Y^s)$ is higher than marginal costs, which means that firms have some degree of marker power. We have indeed:

$$P(Y^s) = \frac{m}{m - \sigma} \times \frac{W (h^d)^\alpha}{1 - \alpha}, \quad (9)$$

i.e. the goods price is the product of each firm's marginal cost $\frac{W(h^d)^\alpha}{1-\alpha}$ and the mark-up factor $\mu \equiv \frac{m}{m-\sigma} > 1$.

The exact expressions of the equilibrium values of endogenous variables are computed in Appendix A. Despite its simplicity, this model provides several key elements of modern macroeconomics which are common to a rather wide range of schools of thought: new classical economics, real business cycle, general disequilibrium theory and new Keynesian economics.⁸

3 Towards an agent-based general (dis)equilibrium setting

This section transposes the GE setting of Section 2 to a *learning/decentralized* framework. Such a transposition requires to introduce the heterogeneity and learning of agents, as well as disequilibrium dynamics.

3.1 Sequence of events

As ABMs are simulated in a strictly sequential way, we have to make explicit the sequence of events. In our model, economic activities follow six main sequential steps within each time period t , and the model is run for T periods. To be concrete, we can interpret a period as a day, households as daily workers who go every morning in the labour market to sell their workforce and shop every afternoon. They can transfer unspent wages from one period to the next, but we assume that these periods are sufficiently short to neglect interest incomes. The model runs then for T days.

1. Households and firms meet in the *labour market*. When all feasible transactions have taken place, the quantity of labour hired by each firm, as well as the amount of labour actually supplied by each household, and the associated wage bills of firms and labour incomes of households are determined.

⁸For instance, the utility function (1) meets all critical mathematical properties supposed in the highly influential works of Lucas (1972), Prescott (1986), Blanchard and Kiyotaki (1987) and Woodford (2003, p. 64 and 144) (to name but a few), and is often adopted as such (see e.g. Galí (2008, p. 42)). The same remark applies to the production function (3) (see in particular Lucas (1972), Bénassy (1993), Woodford (2003, p. 148) and Galí (2008, p. 43)).

2. Households compute their total incomes to be used to finance their consumption.
3. Each firm uses the quantity of hired labour to produce the consumption goods, and sets its selling price.
4. Firms sell their production to households in the goods market. At the end of this matching process, each firm computes its profits. The consumption of each household, together with its labour supply, determines its utility level. Remaining cash-on-hand is transferred to the next period.
5. Firms distribute their profits to households.
6. The strategies of firms (labour demand and wage offer) and of households (labour supply and goods demand) are updated through a social learning process.
7. The story starts all over again with step 1 (T repetitions of Steps 1 to 6).

We now detail the presentation of our agent-based version of the GE setting. For the sake of clarity, we start with the structural assumptions about the overall architecture of the economy (Steps 1 to 5), then we present the social learning process (Step 6).

3.2 Structural assumptions

Agents' strategies Each household has to decide upon its labour supply and its goods demand: $(h_{i,t}^s, c_{i,t}^d)$ in each period t . Symmetrically, each firm j has the two-component strategy $(h_{j,t}^d, w_{j,t})$, where $w_{j,t}$ is firm j 's wage offer. This latter point is due to the absence of a presupposed law of unique price: each firm has then to fix its own individual wage offer. Through the assumed mark-up behaviour, labour being the only input, those two firms' strategies determine the selling price of the goods.

Labour market The labour market is not Walrasian, and operates in a *decentralized* way. The ABM explicitly represents the market interactions between individual agents, and in particular the matching process of individual supplies and demands. Different ABMs adopt different interaction schemes, depending on their focus: pure random matching, (price or quantity) oriented matching, tournament type interactions, etc. In most

of the agent-based literature, buyers sample sellers according to price-driven rules, and start buying from the cheapest ones (see e.g. Riccetti et al (2015), Dosi et al (2013)). This rationing is said to be efficient – in the sense of consistent with the maximization of households’ surplus and full employment in the case of successful agents’ coordination, and we use it in our ABM as well. This rationing mechanism presents also the advantage of complying with some key properties of centralized rationing schemes (Bénassy, 1993): exchanges are voluntary, and the short side of the market is not rationed (frictionless rationing).

We assume that the firms having the highest wage offer deal first with the households proposing the highest labour supplies. The first ranked firm starts hiring labour supplied by the first-ranked household. Whenever a firm succeeds in hiring as labour as it demanded for, it exits the market (the same is true for any household having sold its entire labour supply). Rationing may then occur on either side of the market: firm j is rationed whenever $h_{j,t}^d > h_{j,t}$, and household i is rationed whenever $h_{i,t}^s > h_{i,t}$. When either all firms or all households have exited the matching process, labour market transactions cease: the quantity of labour that has been hired by each firm, and the quantity of labour that has been actually supplied by each household are fully determined. The average nominal wage in the economy is given by $W_t \equiv \frac{1}{H_t} \sum_{j=1}^m h_{j,t} w_{j,t}$, with H_t the aggregate labour. Firm j ’s labour costs are given by its wage bill $h_{j,t} \cdot w_{j,t}$. Household i ’s labour income reads similarly.

Households’ incomes In each period t , household i receives a (nominal) income, denoted by $\tilde{y}_{i,t}$:

$$\tilde{y}_{i,t} = w_{i,t} \cdot h_{i,t} + \frac{\Pi_{t-1}}{n} + m_{i,t-1}. \quad (10)$$

Household income is composed of three components: current labour income, dividend payments (corresponding to $1/n$ of last period’s aggregate profits Π_{t-1}), and cash-in-hand from the previous period, $m_{i,t-1}$. Several points related to income equation (10) may need further clarification.

First, GE models simultaneously compute the equilibrium values of all endogenous variables. This simultaneity allows firms to distribute current profits to households. Simultaneous feedbacks of that kind are inconsistent with the sequential structure of ABMs:

profits must first be collected by firms (at the end of period $t-1$), before being distributed to households (at the beginning of t).

Second, we may observe negative profits. As our model deals with out-of-equilibrium situations, some firms may not be able to sell their entire production for which they already have hired (and paid) the corresponding labour force. In this case, we simply assume that these negative profits are distributed to households in equal shares (just as the positive ones). This assumption is consistent with the supposed distribution of firms' shares, and makes unnecessary the introduction of more complex assumptions about the entry and exit of firms into the economy. In fact, we mainly observe such negative profits in the earlier periods of the simulations, before the settling down of the learning process, and they do not seem to play a determining role in our final results.

Third, some cash-on-hand may remain from the previous period if the household has not spent its entire income. This may happen for two reasons: *i*) the household may be unable to spend all its cash-on-hand, because it is rationed in the goods market⁹; *ii*) the household may not wish to spend its entire cash-on-hand (discrepancy between its income $\tilde{y}_{i,t}$ and goods demand $c_{i,t}^d$). These kinds of individual and/or aggregate coordination failures should precisely incite the agents to learn and modify their strategies towards the coordination on the utility-maximising equilibrium.

Production of the good, and pricing behaviour Labour is the only input, all firms use the same Cobb-Douglass production function (3). Each firm j then sets its price $p_{j,t}$ according to a (gross) mark-up $\mu \geq 1$ over its marginal cost as follows:¹⁰

$$p_{j,t} = \mu \left[\frac{w_{j,t} \cdot (y_{j,t}^s)^{\frac{\alpha}{1-\alpha}}}{1-\alpha} \right]. \quad (11)$$

Goods market Each firm enters the goods market with $y_{j,t}^s$ goods supplied at the selling price $p_{j,t}$, and the households are demanding $c_{i,t}^d$ goods and disposing of nominal income

⁹However, such a rationing is a rare event in the simulations.

¹⁰As we study a stationary environment without structural shocks and focus on the coordination on a static equilibrium, a fixed mark-up does not seem as such a restrictive assumption. In the simulations, the value of the mark-up depends on σ and m , see Section 2.1.2. In fact, a GE exists for any $\mu \geq 1$ value, so that the role of μ is only to pin down the equilibrium values of the wage rate and the other quantities that agents are trying to learn, but it does not affect the convergence and coordination properties of the evolutionary learning algorithm. However, additional simulations that we have conducted show that if the mark-up is a strategy of firms and is therefore time-varying, the GE equilibrium values are constantly modified and learning dynamics are much less convincing.

$\tilde{y}_{i,t}$ (see Equation (10)). Households are subject to the budget constraint

$$\tilde{c}_{i,t} \leq \tilde{y}_{i,t}, \quad (12)$$

where $\tilde{c}_{i,t}$ is household i 's actual nominal consumption expenditure. So households cannot spend more than their current income on consumption.

We introduce a similar decentralized matching process as in the labour market: the firms offering the lowest prices first interact with the households characterized by the highest level of desired (real) consumption. Note that a firm–household transaction may be truncated at the point where the household's remaining income becomes zero. When all feasible transactions have taken place, the goods market closes. At this step, the prices corresponding to each transaction, as well as the quantity of goods sold by each firm ($y_{j,t}$) and consumed by each household ($c_{i,t}$) are fully determined. The average price level is given by $P_t = \frac{1}{Y_t} \sum_{j=1}^m p_{j,t} y_{j,t}$, with Y_t the aggregate output.

Each firm j then computes its nominal profit:

$$\pi_{j,t} = p_{j,t} \cdot y_{j,t} - w_{j,t} \cdot h_{j,t}. \quad (13)$$

Each household i obtains the corresponding utility:

$$u(h_{i,t}, c_{i,t}) = \frac{c_{i,t}^{1-\sigma}}{1-\sigma} - \frac{h_{i,t}^{1+\phi}}{1+\phi}, \quad (14)$$

and computes its remaining cash-on-hand:

$$m_{i,t} = \tilde{y}_{i,t} - \tilde{c}_{i,t} \geq 0, \quad (15)$$

where we use the same utility and profit functions as in Section 2.

We now introduce a social learning process for the adaptation of households' and firms' strategies.

3.3 Social Learning

The agents' initial strategies $(h_{i,0}^s, c_{i,0}^d)$ and $(h_{j,0}^d, w_{j,0})$, are drawn from (independent) uniform distributions (see Appendix A). For any further period, these strategies are updated through a *social learning mechanism*. It is based on a very simple idea: each agent tries to copy the most successful strategies among their peers, while sometimes randomly experimenting with some new ones in order to discover new possibilities, which also ensures the ergodicity of the learning process. We use two operators: selection (imitation) and mutation (random experimentation). These two operators are also commonly used when modelling learning in evolutionary games (Fudenberg and Levine, 1998; Weibull, 1995).

We now detail the operator specifications adopted in our model. This presentation mainly focuses on the evolution of households' strategies because exactly the same scheme also applies to firms.

Imitation Imitation favours the diffusion of the most successful strategies among the population of agents. Utility is the performance indicator of the households, and firms' performance results from their profit level. With a given probability of imitation, P_{imit} , each household i has the opportunity to imitate another household's strategy. In this case, it copies both the labour supply and the consumption level of the other household. Strategies corresponding to higher performances have a higher probability of being imitated. We model the selection of the imitated agent using a standard roulette-wheel mechanism. The probability of being imitated for any household k depends on his *relative* performance, and is given by:

$$P[\textit{imitated} = k] = \frac{f_k}{\sum_{i=1}^n f_i}, \quad (16)$$

where $f_i \geq 0$ represents the current performance of the household i . It should be noted that we cannot directly use the utilities (and profits) as performance indicators, since they are potentially negative and their values may be very close over the domain of variation of strategies. As a consequence, we project them onto a performance domain with the following transformation of the utility:

$$f_i = \exp(\beta \bar{u}_i). \quad (17)$$

\bar{u} gives the average performance obtained with a strategy during its use since its adoption by the household. The exponential function is also used to correct possibly negative profits. $\beta > 0$ is a scale parameter. Depending on the supposed learning process, β allows several useful interpretations. In the heterogeneous agent literature, β is called the *discrimination rate* or the *intensity of choice*, and it is interpreted as the agents' degree of individual rationality (see e.g. Hommes (2013)). If $\beta \rightarrow 0$, the choice between strategies in the roulette-wheel becomes purely random. In the limiting case $\beta \rightarrow +\infty$, agents always imitate the currently best performing strategy. In the reinforcement learning literature (Sutton and Barto, 1998), this parameter has been interpreted in terms of *greediness*: the higher β , the more likely agents are to imitate strategies with the currently highest performance. Both β -interpretations may apply in our agent-based framework.

The use of the average past performance of a strategy (like \bar{u} in Equation (17), and similarly for profits) is also a frequent assumption in the learning literature (see e.g. Arifovic et al (2013)). This modelling device allows agents to use their whole experience with a given strategy, and not only the current performance level. In our set-up, this assumption is of particular relevance as the performance of any given strategy largely depends on the behaviour of all other agents (via aggregate quantities and market matching outcomes). These factors induce randomness and changes in strategy performances. Consequently, strategy evaluation over a longer horizon allows the agents to better discriminate between "good luck" and good-performing strategies.

Formulated as such, the imitation operator leaves little leeway for individual rationality: agents imitate strategies according to the roulette-wheel procedure, even if these strategies have given rise to performance levels that are lower than those obtained with their current strategy. Behavioural rules of this kind seem at odds with minimum requirements of rationality. We therefore introduce a filtering procedure called *selective replication* (Arifovic and Karaivanov, 2010): agents imitate another agent's strategy only if it results in a fitness that is higher than its own. If not, the agent keeps its current strategy. Selective replication has much in common with the *election* operator that is usually necessary for getting acceptable economic behaviour from genetic algorithm-based learning processes. For instance, Arifovic (1995) shows that this additional operator is necessary to discover the rational expectation equilibrium of the Cobweb model, and to obtain convergence in overlapping generation economies. Minimal individual rational-

ity of this kind is also systematically present in evolutionary games with social learning (Weibull, 1995; Dawid, 2007). Consequently, we adopt a social learning mechanism with selective replication.

Once established that household i imitates k , i 's two-component strategy is updated as follows:

$$(h_{i,t+1}^s, c_{i,t+1}^d) = (h_{k,t}, c_{k,t}). \quad (18)$$

Firm j imitating firm l updates its strategy in a similar way:

$$(h_{j,t+1}^d, w_{j,t+1}) = (h_{l,t}, w_{l,t}). \quad (19)$$

It should be underlined that according to Equations (18) and (19), agents imitate the *actual* strategies of other agents (potentially affected by market rationing) observed at the end of period t . They do not imitate the *desired* levels of those variables, because only the actual levels are responsible for the relative performances that are at the heart of the imitation process. Moreover, the desired levels can be considered as private information, and hard to observe for other agents.

Random experimenting (mutation) The mutation operator constantly introduces new strategies in the population, so that potentially better strategies than existing ones may be reached, and further diffuse among agents. We introduce these random experiments in a very standard and simple way (see, for example, Lux and Schornstein (2005) and the references therein).

With a given probability of mutation, P_{mut} , household i 's strategy is randomly modified. A second equiprobable random draw determines the element of i 's two-component strategy $(h_{i,t}^s, c_{i,t}^d)$ to be affected by the mutation process. In the case of a labour supply mutation, household i draws a new labour supply out of a normal distribution with a mean equal to i 's current labour supply, and with a given variance σ_{mut}^2 . An analogous scheme applies to the mutation of i 's goods demand, and for mutations of the firms' strategies $h_{j,t}^d$ and $w_{j,t}$.¹¹

If household i and firm j neither imitate nor perform a mutation, they keep their

¹¹Negative draws from these normal distributions are converted into 0.01 in order to avoid negative supplies, demands and prices. In our numerical simulations, this quasi-zero bound is rarely binding, because σ_{mut}^2 is sufficiently low.

current strategy for the next period. For the sake of parsimony, the learning parameters P_{imit} , β , P_{mut} and σ_{mut}^2 are assumed to be the same for households and firms, and are the only free parameters in the ABM.

3.4 Potential coordination failures with social learning

Even if social learning is able to solve important coordination problems in different setups, our multi-market framework, incorporating also a labour market, may present new challenges for this learning mechanism. To see this possibility, let us consider that the economy starts at the optimal symmetric equilibrium, and introduce small, unilateral changes in agents' strategies (like in the case of a mutation, for example). Even if they concern only one agent and one strategy variable at a time, we can show that those changes systematically pull the economy below the equilibrium in terms of aggregate output. Indeed, aggregate output, Y_t , is constrained by rationing on both markets:

- i)* by the level of aggregate demand: we have necessarily $Y_t \leq C_t^d$;
- ii)* by the level of aggregate goods supply: $Y_t \leq Y_t^s$, which encompasses two sub-bounds:
 - ii-a)* induced by aggregate labour demand H_t^d : $Y_t^s \leq \sum_{j=1}^m (h_{j,t}^d)^{1-\alpha}$;
 - ii-b)* induced by aggregate labour supply H_t^s : $Y_t^s \leq \sum_{i=1}^n (h_{i,t}^s)^{1-\alpha}$;
- iii)* by the level of aggregate purchasing power: $Y_t \leq \frac{\tilde{Y}_t}{P_t}$ (with $\tilde{Y}_t = W \cdot H_t + \Pi_{t-1} + M_{t-1}$).
- iv)* by an uneven distribution of purchasing power, i.e. purchasing power is highly concentrated in households with relatively low goods demand, and vice versa: $Y_t \leq \sum_{i=1}^n \text{Min} \left[c_{i,t}^d; \frac{\tilde{y}_{i,t}}{p_{i,t}} \right]$, where $p_{i,t}$ is the goods' price paid by household i .¹²

The crucial point is that at the optimal equilibrium, all these bounds to aggregate output are *simultaneously* binding. Now suppose isolated upward or downward changes in agents' strategy x , denoted by $\Delta x_{i,t} = x_{i,t} - x_{i,t-1}$. As summarized in Table 1, *all these changes result in lower aggregate output*. While decreasing supply or demand strategies will obviously lower aggregate output, the effects of positive changes in strategies may be less straightforward, as they operate through biases in the distribution of purchasing power. If $\Delta h_{i,t}^s > 0$, *ceteris paribus*, household i may work and earn more, without increasing its consumption and, as H_t^d remains unchanged, another household $k \neq i$ may

¹²Whenever i buys goods from firms applying different selling prices, $p_{i,t}$ is a composite price.

Change	Main mechanisms	Aggregate output	ΔY_t
$\Delta c_{i,t}^d > 0$	stability of goods' supply and income	$= Y_t^s = \frac{\tilde{Y}_t}{P_t}$	$= 0$
$\Delta c_{i,t}^d < 0$	decrease in goods' demand	$= C_t^d$	< 0
$\Delta h_{i,t}^s > 0$	household(s) rationed on labour market	$= \sum_i \text{Min} \left[c_{i,t}^d \cdot \frac{\tilde{y}_{i,t}}{p_{i,t}} \right]$	< 0
$\Delta h_{i,t}^s < 0$	decrease of goods' supply and income	$= \text{Min} \left[Y_t^s; \frac{\tilde{Y}_t}{P_t} \right]$	< 0
$\Delta h_{j,t}^d > 0$	skewed distribution of labour over firms	$= Y_t^s$	< 0
$\Delta h_{j,t}^d < 0$	decrease of goods' supply and income	$= \text{Min} \left[Y_t^s; \frac{\tilde{Y}_t}{P_t} \right]$	< 0
$\Delta w_{j,t} > 0$	skewed distribution of income over households	$= \sum_i \text{Min} \left[c_{i,t}^d \cdot \frac{\tilde{y}_{i,t}}{p_{i,t}} \right]$	< 0
$\Delta w_{j,t} < 0$	skewed distribution of income over households	$= \sum_i \text{Min} \left[c_{i,t}^d \cdot \frac{\tilde{y}_{i,t}}{p_{i,t}} \right]$	< 0

Table 1: Impact of strategy changes on aggregate output

be rationed on the labour market, lowering k 's purchasing power and consumption. If $\Delta h_{j,t}^d > 0$, firm j may hire more labour and produce more goods, while firm $l \neq j$ may be rationed on the labour market, and hence produce less. Importantly, due to decreasing returns ($1 - \alpha < 1$), l 's production decrease will be higher than j 's increase. If $\Delta w_{j,t} > 0$, one household, hired by j , will earn more than its consumption expenditure, while firm j , who is now asking for a higher price, may be rationed on the good market. A symmetrical situation will occur in the case of $\Delta w_{j,t} < 0$. Even if the economy is not perfectly positioned at the optimal equilibrium, those mechanisms also work in the neighbourhood of the equilibrium. We will show later that those mechanisms structure our main results.

4 Simulation protocol

We first focus on a *baseline setting* defined in Section 4.1. The robustness of the results with respect to a number of alternative structural and learning assumptions are analysed in Appendix B.

4.1 Parameter values

We use a very common calibration in the DSGE literature (see e.g. Galí (2008)), and set decreasing rate of returns $\alpha = 0.33$ and preferences $\sigma = 1$ (hence, we use $u(c) = \log(c)$) and $\phi = 1$. In the ABM, those parameters influence the shape of the fitness landscape for the social learning process, but our results are robust to alternative calibrations,

taken from the DSGE literature (Rotemberg and Woodford (1998): $\phi = 0.6, \sigma = 0.16$), or from the microeconomic estimates of the Frisch elasticity of labour supply (Whalen and Reichling (2017): $\phi = 2.5$). We retain an equal number of firms and consumers $n = m = 40$ to keep the learning problem of firms and households comparable, where the size of the population is chosen in line with the GA literature (Holland et al, 1989). We cover many combinations of learning parameters: $P_{imit} = \{0.01, 0.05, 0.1, 0.2, 0.5, 1\}$, $P_{mut} = \{0.01, 0.05, 0.1\}$ and $\sigma_{mut}^2 = \{0.01, 0.05, 0.1\}$, i.e. we consider $6 \times 3 \times 3 = 54$ learning configurations. Those parameter values are in line with standard calibrations of GAs. In the social learning mechanism, higher β values amplify the differences in relative utility, and make the roulette wheel more selective. This could be particularly relevant in ABMs derived from GE models, where performance functions are mainly chosen for their analytical properties in optimization problems, rather than for their ability to facilitate learning by agents with bounded rationality – see also Salle and Seppacher (2015) for a related discussion. Consequently, we adopt $\beta = 4$ in the simulations. Indeed, this case is more favourable to coordination than lower values of β , while we did not observe a better convergence with higher values.

The computer simulations are run over $T = 1,000$ periods, which should be sufficient for revealing the main dynamics and intrinsic properties. In order to cope with the inherently non-deterministic nature of ABMs, we generate 30 replications of each possible parameter combination, 30 being the minimum number of observations to apply normal distribution assumptions. We consequently run $54 \times 30 = 1,620$ simulations.

Appendix A reports the values of the structural and the learning parameters used in the baseline setting and the corresponding optimal equilibrium values in the GE model.

4.2 Convergence and coordination indicators

We focus the analysis on two main outcomes: *aggregate convergence* to the equilibrium and *individual coordination*.

We compute the aggregate convergence index ACI_t which is defined as follows:

$$ACI_t \equiv \frac{1}{3} \left(\frac{|C_t - C^*|}{C^*} + \frac{|H_t - H^*|}{H^*} + \frac{|\omega_t - \omega^*|}{\omega^*} \right) \quad (20)$$

Recall here that the star sign $*$ refers to optimal equilibrium values given in Appendix A.

ACI_t gives the average percentage distance of actual aggregated consumption C_t , aggregated labour H_t and wage rate ω_t with respect to their corresponding optimal equilibrium levels C^* , H^* and ω^* . The lower the index, the closer the economy to its optimal state, and the better the aggregate convergence.¹³ For instance, if $ACI_t < 0.1$, we can say that in period t , on average (over agents and markets), the economy is in the 10%-neighbourhood of the optimal equilibrium state. In practical terms, we mainly use the average of the aggregate convergence index over the last 100 periods of each simulation in order to get a grip on the asymptotic behaviour of the model. We also assess the level of individual coordination by measuring the variances of individual variables: $Var(h_{i,t})$ and $Var(c_{i,t})$ for households, and $Var(h_{j,t})$ and $Var(w_{j,t})$ for firms.

4.3 Numerical implementation

The agent-based version of the model is programmed in NetLogo 5 (Wilensky, 1999), and the code is available on the corresponding author's website.¹⁴ Simulation outcomes are saved every 50 periods, and analysed with R-project 3 (R Development Core Team, 2003).

5 Main simulation result: The difficulty of coordination

Figure 1 reports the distribution of the convergence index over the 1620 simulations across all learning configurations. In most runs, the economy does not converge to the equilibrium but reaches its neighbourhood: roughly two thirds of the runs reach a 20%-neighbourhood of the equilibrium, but only one quarter reaches the 10%-neighbourhood. The size of the neighbourhood considered can appear wide at a first glance, but recall that we allow for relatively large and systematic mutations till the end of the simulations. As a consequence, the system can only, by construction, approach the equilibrium. A deviation from equilibrium by about 10 to 20% is also consistent with the experimental evidence cited in the introduction.

Figure 2 gives the evolution in time of our main indicators. Utility and profit levels show that social welfare is clearly increasing in time, and these indicators stabilize at a

¹³Note, however, that the index cannot stay at zero even in the case of convergence because we allow for mutations throughout the simulations.

¹⁴<http://www.uu.nl/staff/ILSalle/>.

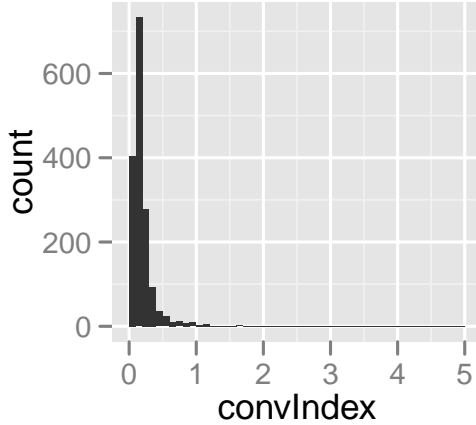


Figure 1: Distribution of the convergence index values, averaged over last 100 periods in the 1620 simulations

rather high level. The employment levels somewhat overshoot (see Figure 2c), but firms are not able to sell their entire production (compare goods supply and output levels in Figures 2e and 2f), implying a decrease in profits starting around period 250. Thus, aggregate dynamics converge *towards but below* the equilibrium, and some coordination problems remain. It should be noted that the results of Figure 1 do not depend on extreme values of the imitation probability, since we exclude them in Figure 2, and hereafter. Imitation probabilities of 50 or 100% imply a too frequent updating of the fitness of the strategies (every period or every two periods on average), which render more difficult for the agents to discriminate between a one-period "luck" and a sustained performance over a longer time frame.

As for individual coordination, we show in Figure 3 the evolution of the variances of the individual strategy components. Coordination between households' variables seems very clear, since the variances are decreasing towards zero (see Figures 3a and 3b).¹⁵ By contrast, coordination between firms is not obtained (see Figures 3c and 3d): the variance of their strategies instead steadily increases over time. Section 6.1 sheds light on this major regularity of our model.

In evolutionary terms, we can conclude that the optimal symmetric equilibrium appears *globally* unstable, and the model systematically reaches a sub-optimal state. If the

¹⁵Recall that mutation occurs throughout the simulations so that variability in strategies remain by construction till the end of the simulations, and the variance cannot drop to zero.

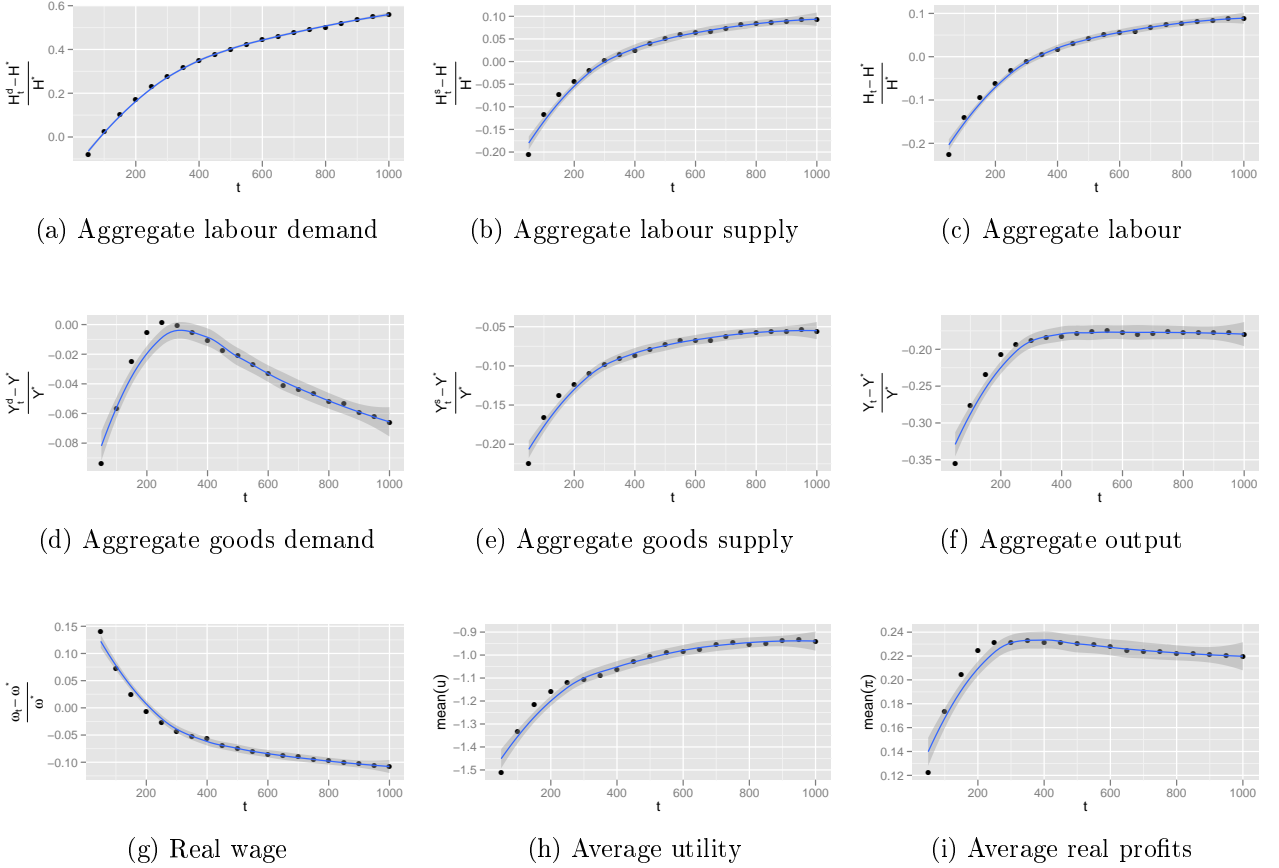




Figure 2: Assessment of *aggregate* convergence, $\beta = 4$, selective replication, $P_{imit} \in [1, 5, 10, 20]\%$, averaged over the 810 simulations. Relative deviations from the equilibrium values are given in figures a-g. Dots show period means, i.e. aggregate quantities averaged over all parameter configurations and replication runs. The line stands for a smoother, and the grey bands represent the standard error of this smooth (see the `ggplots2` graphical package of R-project).

optimal state were at least locally stable, the model, once located at that state, should converge back to it after a small departure, through the selection pressure of the learning process. To investigate this point, we initialize the simulations at the equilibrium. The mutation operator of the GA then gradually introduces deviations from this optimal state. As shown in Figure 4, the model quickly stabilizes at the sub-optimal state obtained in Figures 2h and 2i. This clearly demonstrates that this sub-optimal state seems to be the *globally* evolutionary stable attractor of the model.


Our results show *below-equilibrium* dynamics and *coordination between households* under a social learning process that allows for strong selectivity and some minimal individual rationality, but *coordination issues between individual firms* remain. Consequently, the




(a) Variance of households' labour



(b) Variance of households' consumption




(c) Variance of firms' labour




(d) Variance of firms' production

Figure 3: Assessment of *individual* coordination: Variances of individual indicators ($\beta = 4$, selective replication, $P_{imit} \in [1, 5, 10, 20]\%$, averaged over the 810 simulations)

model settles down in a near- but sub-optimal state that implies systematic welfare loss. In Appendix B, we consider alternative assumptions for the learning algorithm, and show that our results are robust to a decreasing size of mutation, changes in the number of agents (that corresponds to the size of the strategy pool in the GA) and different learning frequency. Since those different modalities of learning cannot prevent coordination failures, we now dig into the economic mechanisms that are inherent to the class of models under study, and may play a role in the observed miscoordination.



(a) Average utility



(b) Average profit

Figure 4: Drift of average utility and profit at sub-optimal levels when the model is initialized at the optimal state.

6 Economic sources of coordination problems

The main economic mechanisms that are potentially behind our results are: an increasing concentration of the production process over few firms combined with non-linearities in the production function, self-reinforcing local deviations from the equilibrium (as theoretically discussed in paragraph 3.4), and ultimately a self-reinforcing uneven distribution of purchasing power among the households.

6.1 Market concentration, non-linearities and firms' miscoordination)

The increasing variance between firms' quantities observed in Figures 3c and 3d indicates a increasing market concentration over time in the simulations. This concentration implies suboptimal aggregate output level due to the non-linearities induced by the decreasing rates of return of the production function (i.e. $\alpha > 0$ in Equation (3)). Under decreasing rates of return, we have:

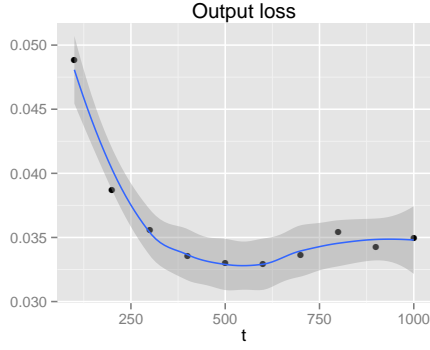
$$\frac{\sum_{j=1}^m f(h_j)}{m} \leq f\left(\frac{\sum_{j=1}^m h_j}{m}\right) \quad (21)$$

Consequently, even if aggregate convergence prevails in the labour market so that $\sum_{j=1}^m h_j = H^* = m \cdot h^*$, we will have, as soon as the quantity of labour is not evenly distributed across the m firms:


$$\sum_{j=1}^m h_j^{1-\alpha} \leq m \cdot (h^*)^{1-\alpha}$$

A given amount of labour allows to produce more goods under an even than under an uneven distribution over firms. This is the reason why the aggregate quantity of labour hired by firms is slightly above the optimal value (see Figure 2c), while the good supply stabilizes around 5% below the optimal output level (see Figure 2e). Figure 5a shows that the uneven distribution of labour among firms induces an average of 3.5% of output loss compared to a perfectly even distribution. This concentration loss explains therefore part of the discrepancy between aggregate output levels and the optimal one observed in the simulations.

Rationing may explain this concentration dynamics. As Figures 2a and 2b indicate clearly, aggregate labour supply is inferior to aggregate labour demand in the simulations.



(a) Normalized Herfindahl index of concentration



(b) Output loss (percentage of the equilibrium goods supply) due to concentration

In this case, all firms placed at the end of the firms' list in the labour market are completely rationed, i.e. they receive $h_j = 0$ (and at most one firm is partially rationed). Consequently, rationing induces concentration of aggregate labour over the firms placed at the top of the list. Figure 5b reports the evolution of the normalised Herfindahl index in the goods market. This index is comprised between 0 and 1, and the higher the index, the more concentrated the market. After an initial decrease (recall that firms are initialized randomly, and the initial situation may be quite asymmetric), we do observe that the average value of the index increases over time (starting around period 500), indicating an increase in concentration.

To illustrate this effect further, we run the simulations under constant rate of returns (i.e. setting $\alpha = 0$). In this case, the distribution of the production between firms in equilibrium is indeterminate, and does not weigh on aggregate outcomes. Results are reported in Appendix B.4. Removing non-linearities in the production process clearly allows for better *aggregate* convergence on the goods market, while the labour market overshoots the equilibrium level, just as in the case of the decreasing rate of returns. Additionally, we notice that stabilisation of the aggregate quantities is obtained rather quickly, after roughly 200 periods. However, coordination still only obtains among households.

One may think that coordination is actually more challenging among firms than among households because households are price takers, while firms have to learn both a quantity and a price on the labour market (and on the good market, given the fixed markup

pricing and production function).¹⁶ To investigate this intuition, we run an additional set of simulations, in which we simplify the firms' learning problem. We normalize the wage to unity, i.e. $w_{j,t} \equiv 1, \forall j, t$, so that firms are only left with the labour demand strategy. In this case, there is a one-to-one mapping between the quantity of labour hired by a firm j (or, conversely, goods produced) and its price.¹⁷ Results are reported in Appendix B.5. We observe a better aggregate convergence, and especially full convergence of the real wage, which results in a better coordination among firms. However, our qualitative results remain unchanged: output systematically lies below-equilibrium, and the variance between firms' quantity strategies steadily increases. This suggests that miscoordination between firms is an *asymptotic* pattern of our model, even with a simpler learning problem on their side. We then move to our second potential explanation: self-reinforcing local deviations from the equilibrium that we have theoretically discussed in Section 3.4.

6.2 Sources of deviations from the optimal state

In the baseline simulations, we now measure the relative contribution of the four bounds to aggregate output, operating through the self-reinforcing local deviations from the equilibrium that we have discussed in Section 3.4. In order to do so, we count the number of periods over an entire simulation duration ($T = 1000$) where the constraints *i*) to *iv*) in Section 3.4 are binding: the *demand constraint counter*, the *supply constraint counter*, the *budget constraint counter*, and the *distribution constraint counter*. Figure 5 displays boxplots and mean values (marked by red dots) of the distributions of the four constraint counters in the simulations. It is easily seen that the lack of aggregate demand exerts no effect (the mean counter value is almost zero: 0.015), nor do the supply and budget constraint (mean values are 42.5 and 42.6 and are mainly driven by outliers). According to the right-hand chart of figure 5, distribution problems are at the very core of the non-convergent dynamics: this constraint is almost always binding (on average 914.9 out of 1,000 periods). Aggregate output is thus frequently and massively constrained by an uneven distribution of purchasing power among households. This fact explains why, simultaneously, aggregate labour H stabilizes around 10% above equilibrium, goods supply

¹⁶We are grateful to an anonymous referee for bringing up this idea.

¹⁷We have also tried to fix both prices and wages, so that the real wage ω_t is always at its equilibrium value. In this case, the signalling role of prices disappears, which significantly hinders the profit-based imitation process.

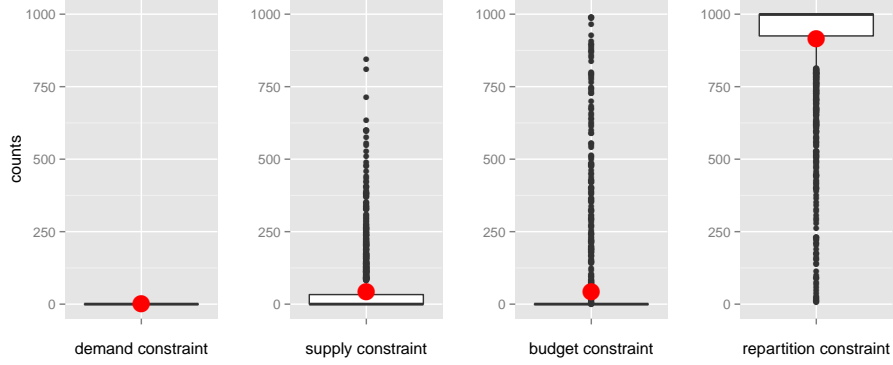


Figure 5: Distribution of constraint counters: numbers of periods where a given constraint is binding (distribution over all the simulations; the red dots indicate the distribution means).

Y^s lies only 5% below the equilibrium, but aggregate goods consumption Y settles down 15% below equilibrium (see Figure 2). Indeed, despite higher labour incomes, households are in general unable to buy more goods because of ill-suited distribution of those additional incomes. This effect is the main inherent driving force away from the optimal equilibrium. We therefore now focus the discussion on this effect.

6.3 Self-reinforcing uneven distribution of purchasing power

The importance of the distribution constraint just highlighted stems in large part from the existence of a feedback loop which reinforces any given uneven distribution of purchasing power between households. We provide a formal proof of this feature in Appendix C. We develop here the intuition. It should first be recalled that an uneven distribution implies that:

- i)* some households have oversized purchasing power with respect to their goods demand; this is the case of households verifying $wh_{i,t} + \frac{\Pi_{t-1}}{n} > p_{i,t}c_{i,t}$;
- ii)* while other households are marked by undersized purchasing power : $wh_{i,t} + \frac{\Pi_{t-1}}{n} < p_{i,t}c_{i,t}$.

For given agent-strategies, this initial gap necessarily results in worsening of the distribution of purchasing power. In fact, cash-on-hand $m_{i,t}^{over}$ of households with oversized

purchasing power increases:

$$\Delta m_{i,t}^{over} = m_{i,t}^{over} - m_{i,t-1}^{over} = wh_{i,t}^{over} + \frac{\Pi_{t-1}}{n} - p_{i,t}^{over} c_{i,t}^{over} > 0. \quad (22)$$

while cash-on-hand of households with undersized purchasing power decreases:

$$\Delta m_{i,t}^{under} = m_{i,t}^{under} - m_{i,t-1}^{under} = wh_{i,t}^{under} + \frac{\Pi_{t-1}}{n} - p_{i,t}^{under} c_{i,t}^{under} < 0. \quad (23)$$

As a consequence, purchasing power rises for households already endowed with "too much" purchasing power, and it falls for households marked by insufficient purchasing power. This will ultimately lead to strengthening of the distribution constraint and to contraction of aggregate output through a lack of aggregate demand.

This feedback loop is especially important at the beginning of the simulations. This is due to the fact that initial values are randomly drawn, implying that individual households' strategies are generally not *a priori* mutually consistent, i.e. purchasing power is either undersized or oversized. For the reasons just explained, this asymmetry is self-reinforcing, which amplifies the initial idiosyncratic income shocks. One could think that coordination failures due to an uneven income distribution may disappear in the simulations with a unique fixed nominal wage reported in Appendix B.5, because a common wage could play as a coordination device and an homogenization force among the population of initially heterogeneous households. However, unequal income distributions still systematically emerge, albeit at a lesser frequency (roughly in half of the periods, see Figure 8c), as a result of random experimenting and rationing on the markets. This shows the robustness and the importance of the role of an uneven distribution of purchasing power in the observed sub-optimal levels of output.

It may be interesting here to draw a parallel with the analysis of Eggertsson and Krugman (2012) that emphasises similar amplification mechanisms of uneven income distribution processes associated with debt deleveraging in the context of a liquidity trap. In their model, forced deleveraging by debt-constrained households leads them to reduce their consumption. When that reduction in spending cannot be offset by an increase in consumption on the part of unconstrained households, aggregate output, and therefore income, further depress. That depression feeds back into the current income of constrained

agents as that they consume further less. The amplification mechanisms differ in our two frameworks: in Eggertsson and Krugman (2012), the feedback force reinforcing the mismatch between aggregate income and consumption is the downward constraint on the interest rate while, in our model, such a mismatch arises from the interplay between income's heterogeneity, rationing, especially in the labour market, and the selection and reproduction process of behaviour under learning. However, in both analyses, the cause, i.e. an uneven distribution of purchasing power, and the consequence, i.e. suboptimal levels of output and aggregate consumption, are similar. More generally, the negative role of uneven income distribution that we uncover in this paper echoes the recent discussions about raising income inequality as a source of economic stagnation after a financial crisis.

7 Conclusion

This paper studies the question of the *intra-temporal* coordination of a collection of heterogeneous firms and households on the welfare-maximising state, in an *inter-dependent market* setting (labour and goods markets). The coordination device that we consider is social learning. Despite the use of a sophisticated learning mechanism, our first finding is a lack of coordination, especially between firms, and systematic welfare losses. This state of affairs calls for further research into learning algorithms, or even for careful handling of macroeconomic models relying on the equilibrium assumption, as they do not provide any explanation of how the economy has reached this equilibrium at the first place, or how it could reach it back after some disturbance. Recent experimental results in simple settings (see e.g. Bao et al (2013)) already suggest that learning to set an optimal quantity tends to be much more challenging for human subjects than learning to make accurate forecasts. It is worth pointing this out as almost all the learning literature focuses on non-rational expectations, assuming that the rest of the behaviour is fully optimal conditional on these expectations, while our results also tend to suggest that forecasting is the 'easy' face of the coin, and optimizing is a much more challenging task.

We then successively consider potential economic mechanisms that are inherent to the class of GE models under study and could explain the observed sub-optimal levels of output. After carefully highlighting the relative contribution of each of those economic mechanisms to our result, we conclude that a self-reinforcing uneven income distribution

among households, which results in a lack of aggregate demand, is the main culprit. The role of self-reinforcing income inequality in sub-optimal levels of output appears as a strong and robust finding of our analysis, and can be put in parallel with the recent discussions about raising income inequality as a source of economic stagnation after a financial crisis. On a technical plan, this major result stresses the importance of economic analysis within heterogeneous agent framework. On a more general plan, our analysis points out distribution problems in economic coordination.

Our finding opens up two interesting research questions. Our model abstracts from borrowing or insurance mechanisms against idiosyncratic income shocks. Whether or not allowing for credit markets could dampen the negative effect of an uneven income distribution on aggregate demand and output is a valuable related research question. While allowing for income and consumption smoothing, debt dynamics may also put the agents, and the system at a greater risk of financial instability. In the absence of credit markets, or equivalently, if their functioning have been impaired by a financial crisis and a subsequent credit crunch, a redistributive fiscal policy that transfers purchasing power from the wealthiest to the poorest households may dampen the self-reinforcing skewness in income distribution, and its downward effect on aggregate demand and output. In light of the recent experience after the financial crisis in 2008, this is also an important related research question.

References

- 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, Karaivanov A (2010) Learning by doing vs. learning from others in a principal-agent model. *Journal of Economic Dynamics and Control* 34(10):1967–1992
- Arifovic J, Bullard J, Kostyshyna O (2013) Social learning and monetary policy rules. *Economic Journal* 123:38–76
- Arthur BW (1994) Inductive reasoning and bounded rationality. *American Economic Review* 84:406–411

- Bao T, Duffy J, Hommes C (2013) Learning, forecasting and optimizing: An experimental study. *European Economic Review* 61(C):186–204
- Bénassy JP (1993) Nonclearing markets: Microeconomic concepts and macroeconomic applications. *Journal of Economic Literature* 31(2):732–761
- Blanchard OJ, Kiyotaki N (1987) Monopolistic competition and the effects of aggregate demand. *American Economic Review* 77:647–666
- Dawid H (2007) Evolutionary game dynamics and the analysis of agent-based imitation models: the long run, the medium run and the importance of global analysis. *Journal of Economic Dynamics and Control* 31(6):2108–2133
- Dixit A, Stiglitz J (1977) Monopolistic competition and optimum product diversity. *The American Economic Review* 67:297–308
- Dosi G, Fagiolo G, Napoletano M, 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
- Eggertsson GB, Krugman P (2012) Debt, Deleveraging, and the Liquidity Trap: A Fisher-Minsky-Koo Approach. *The Quarterly Journal of Economics* 127(3):1469–1513
- Evans GW, Honkapohja S (2001) *Learning and Expectations in Macroeconomics*. Princeton University Press
- Evans GW, McGough B (2015) *Learning to optimize*, university of Oregon, mimeo.
- Fenig G, Mileva M, Peterson L (2016) *Asset Trading and Monetary Policy in Production Economies*, simon Fraser University, BC, CA, mimeo
- Fudenberg D, Levine DK (1998) *The Theory of Learning in Games*. MIT Press
- Galí J (2008) *Monetary Policy, Inflation and the Business Cycle: An Introduction to the New Keynesian Framework*. Princeton University Press
- Gintis H (2007) The dynamics of general equilibrium. *Economic Journal* 117:1280–1309
- Holland J, Goldberg D, Booker L (1989) Classifier systems and genetic algorithms. *Artificial Intelligence* 40
- Hommes C (2013) *Behavioral rationality and heterogeneous expectations in complex economic systems*. Cambridge University Press, Cambridge (Massachusetts)
- Howitt P, Clower R (2000) The emergence of economic organization. *Journal of Economic Behavior & Organization* 41:55–84
- Lian P, Plott C (1998) General equilibrium, markets, macroeconomics and money in a laboratory experimental environment. *Economic Theory* 12(1):123–146
- Lucas RE (1972) Expectations and the neutrality of money. *Journal of Economic Theory* 4:103–124

- Lux T, Schornstein S (2005) Genetic learning as an explanation of stylized facts of foreign exchange markets. *Journal of Mathematical Economics* 41:169–196
- Noussair C, Pfajfar D, Zsiros J (2015) Pricing decisions in an experimental dynamic stochastic general equilibrium economy. *Journal of Economic Behavior & Organization* 109(C):188–202
- Prescott E (1986) Theory ahead of business cycle measurements. *Quarterly Review* 10:9–22
- R Development Core Team (2003) R: A Language and Environment for Statistical Computing, vol <http://www.r-project.org/>. R Foundation for Statistical Computing, Vienna
- Riccetti L, Russo A, Gallegati M (2015) An agent based decentralized matching macroeconomic model. *Journal of Economic Interaction and Coordination* 10(2):305–332
- Rotemberg JJ, Woodford M (1998) An optimization-based econometric framework for the evaluation of monetary policy : expanded version, nBER Technical Working Paper 233
- Salle I, Seppecher P (2015) Social Learning about Consumption. *Macroeconomic Dynamics* In press
- Schelling TC (1971) Dynamic models of segregation. *Journal of mathematical sociology* 1(2):143–186, URL <http://www.tandfonline.com/doi/abs/10.1080/0022250X.1971.9989794>
- Sutton RS, Barto AG (1998) Reinforcement Learning: an Introduction. MIT Press, Cambridge, MA
- Tesfatsion L, LeBaron B (2008) Modeling macroeconomies as open-ended dynamic systems of interacting agents. *American Economic Review* 98(2)
- Vallée T, Yildizoglu M (2009) Convergence in the finite cournot oligopoly with social and individual learning. *Journal of Economic Behaviour and Organization* (72):670–690
- Vriend N (2000) An illustration of the essential difference between individual and social learning, and its consequences for computational analyses. *Journal of Economic Dynamics and Control* 24(1):1–19
- Walsh CE (2003) *Monetary Theory and Policy*. the MIT Press
- Weibull JW (1995) *Evolutionary Game Theory*. MIT Press, Cambridge (Massachusetts)
- Whalen C, Reichling F (2017) Estimates of the Frisch Elasticity of Labor Supply: A Review. *Eastern Economic Journal* 43(1):37–42
- Wilensky U (1999) Netlogo. Center for Connected Learning and Computer-Based Modeling, Northwestern University. Evanston, Illinois, URL <http://ccl.northwestern.edu/netlogo/>
- Woodford M (2003) *Interest and Prices : Foundations of a Theory of Monetary Policy*. Princeton University Press

A Calibration and equilibrium values

Dividing (9) by W equalizes the left-hand side of (8) and the right-hand side of (9). Observing that in equilibrium $h_i = \frac{m}{n}h_j$ and $c_i = \frac{m}{n}y_j$, we obtain:

$$h_j^* = \left[\frac{m - \sigma}{m} (1 - \alpha) \left(\frac{m}{n} \right)^{-\sigma - \phi} \right]^{\frac{1}{\sigma(1-\alpha) + \phi + \alpha}} \quad (24)$$

Firm j 's optimal goods supply is therefore

$$y_j^* = (h_j^*)^{1-\alpha} \quad (25)$$

Now, we can express all real optimal quantities as functions of h_j^* and y_j^* :

$$h_i^* = \frac{m}{n} \times h_j^*, \quad (26)$$

$$c_i^* = \frac{m}{n} \times y_j^*, \quad (27)$$

$$u_i^* = u(c_i^*, h_i^*), \quad (28)$$

$$\omega^* \equiv \frac{W^*}{P^*} = \frac{1}{\mu} (1 - \alpha) (h_j^*)^{-\alpha} \quad (29)$$

$$\left(\frac{\pi_j}{P} \right)^* = y_j^* - \omega^* h_j^* \quad (30)$$

Aggregated quantities are simply given by: $C^* = n \times h_i^*$, $U^* = n \times u_i^*$, $H^* = n \times h_i^*$, $H^* = m \times h_j^*$, $Y^* = m \times y_j^*$ and $\Pi^* = m \times \pi_j^*$.

Structural parameters		ABM parameters		Equilibrium values		Initialization	
α	0.33	$m = n$	40	$h_i^* = h_j^*$	0.765	$h_{i,0}^s$	(0, 1)
μ	1.026	P_{imit}	{1, 5, 10, 20, 50, 100}%	$c_i^* = y_j^*$	0.836	$c_{i,0}^d$	(0, 1)
ϕ	1	P_{mut}	{1, 5, 10}%	ω^*	0.639	$m_{i,0}$	(0, 1)
σ	1	σ_{mut}^2	{0.01, 0.05, 0.1}	π_i^*	0.347	$w_{j,0}$	(0, 1)
T	1,000	β	4	u_i^*	-0.472	$h_{j,0}^d$	(0, 1)
				m_i^*	0		

Table 2: Parameter and equilibrium values

N.B: Households are endowed with an initial amount of cash-on-hand, randomly drawn between 0 and 1 (this initial endowment is heterogeneous across households). This is necessary as the distribution of current profits in the ABM is postponed to the next period. This initial endowment broadly corresponds to the realization of profits at period -1 .

B Additional simulation results

B.1 Controlling for decreasing size of mutations

One potential way of addressing persistent heterogeneity between agents' strategies is gradually reducing the variance of the mutations over time. Indeed, if the dispersion of mutations remains important, the strategies of the agents may continue to fluctuate quite significantly even at the end of the simulation, and perturb the coordination between agents, and the convergence to the optimal state. In order to check if the variance of the decisions are artificially maintained over zero by a constant mutation size, we test a case with decreasing mutation size, as in Arifovic et al (2013):

$$\sigma_{mut,t} = \sigma_{mut,t-1} \left(1 - 0.95 \times \frac{t}{T}\right) \quad (31)$$

Consequently random experiments are quite global at the start of the simulation, but they become more and more local towards its end. Note that σ_{mut} does not reduce to zero before the last period T .

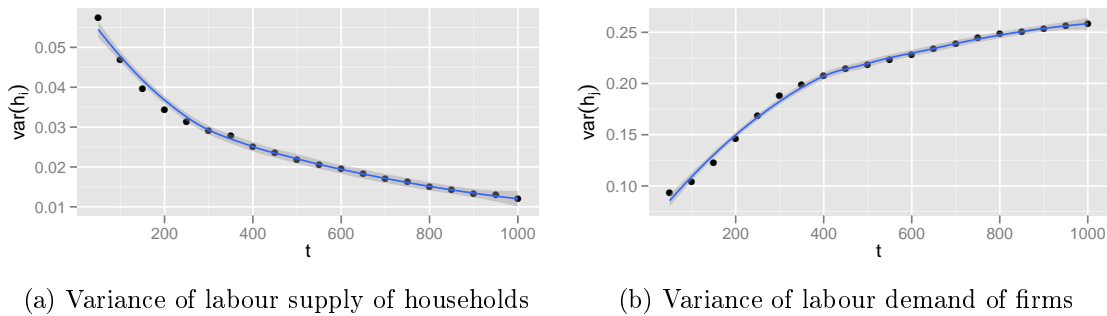


Figure 6: Assessment of individual coordination, $P_{imit} \in [1, 5, 10, 20]\%$, averaged over the 810 simulations, with decreasing mutation size.

Figure 6 shows that the decreasing of the mutation size does not affect the coordination among households and the divergence between firms remains important. A K-S test on the distribution of the average convergence index with and without decreasing mutation size leads to not reject the null hypothesis that convergence is the same in these two configurations. (K-S statistics 0.0148, with a p -value of 0.789). As a consequence, the main emergent features of our ABM are insensitive to the introduction of this mutation operator.

B.2 Robustness to the number of agents

Recall that under social learning, the number of agents corresponds to the number of strategies in the GA. More strategies imply more diversity and a higher exploratory strength of the GA. Consequently, convergence is indeed significantly weaker with only 20 agents compared to the case with 40 agents, but no further improvement is obtained with 60 agents of each type, see Table 3.

Distribution with $n = m = 20$ versus 40 agents $D^- = 0.0525$, p-value = 0.01156, convergence is better with 40 agents.
Distribution with $n = m = 40$ versus 60 agents $D^- = 0.0117$, p-value = 0.8002, convergence is not improved with 60 agents.

Table 3: K-S tests of robustness

B.3 Robustness to learning frequency (GARate)

The frequency of learning (the frequency with which imitation and mutation modify the strategy population) can play an important role (this parameter is called *GARate* in different papers, see for example Vriend (2000)). Vallée and Yildizoglu (2009) show, for example, that a difference in the values of this frequency can determine the type of equilibrium towards which firms' quantity strategies converge in a Cournot oligopoly. We also test the role of this frequency in our model. This is done by allowing imitation and mutation to occur only every *GARate* periods. We also increase the number of periods in order to keep constant the number of learning steps ($1000 \times \text{GARate}$). We conclude that one period is not enough for learning the performance of the strategies, but ten are sufficient for disciplining the learning process by limiting the over-valuation of opportunistic strategies. Taking into account several periods reveals the dynamic deficiency of these strategies: working less today has a positive effect on today's utility, but may reduce the purchasing power for tomorrow.

A common *GARate* has another side effect: it synchronizes the learning of agents, which is not a very realistic assumption even if it is adopted in articles using this variable. We can neutralize this effect by desynchronizing the learning of the agents: they use the same persistence of strategies (*GARate*) and memory, but they learn in different periods. With an individual frequency of 10, the learning process of agents becomes smoother and more efficient. Consequently, "obliging" agents to test the strategies for more than one period also considerably disciplines their learning process, with or without synchronization.

Overall, these refinements clearly ameliorate the global behaviour of the model, and favour better convergence of the model towards the optimal state. Nevertheless, the convergence below the optimal state persists.

Distribution with $\gamma_{GA} = 1$ (baseline scenario) versus $\gamma_{GA} = 10$ $D^- = 0.2272$, p-value = 0.0000, convergence is better with $\gamma_{GA} = 10$
Distribution with $\gamma_{GA} = 50$ versus $\gamma_{GA} = 10$ (two-sided test) $D = 0.0451$, p-value = 0.0745, convergence is similar with $\gamma_{GA} = 10$ and 50.
Distribution with $\gamma_{GA} = 10$ (common GARate) versus $\gamma_{GA} = 10$ (individual GARate) $D^- = 0.0525$, p-value = 0.0116, convergence is better with individual GARate
Distribution with $\gamma_{GA} = 50$ (individual GARate) versus $\gamma_{GA} = 10$ (individual GARate) $D^- = 0.0023$, p-value = 0.2116, convergence is not improved with an individual GARate of 50 periods

Table 4: K-S tests of robustness

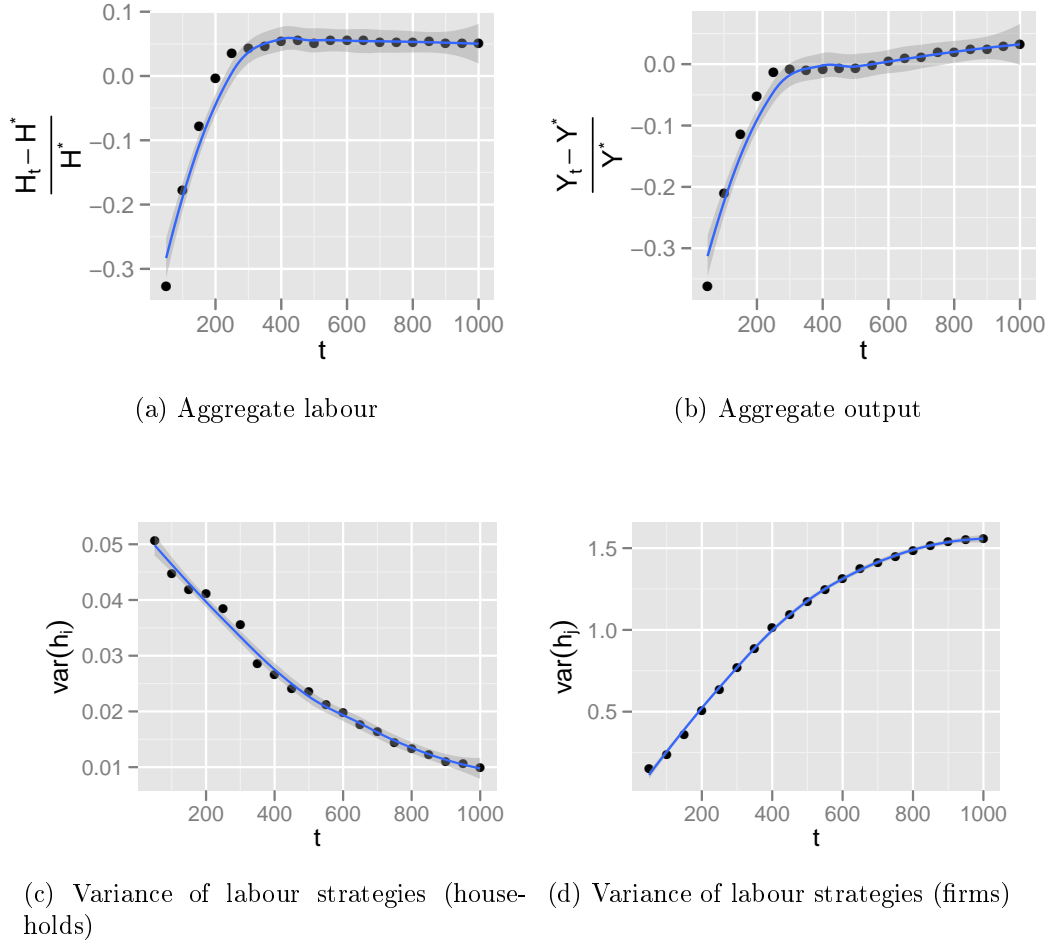


Figure 7: Behaviors under constant rate of returns ($\alpha = 0$), averaged over the 810 simulations.

B.4 Simulations under constant rates of return

B.5 Simulations with fixed nominal wages

Aggregate convergence is significantly better when the nominal wage is fixed (Figure 8a) than in the baseline (Figure 1, p-value= 0.0000). In particular, the median value of the convergence index is divided by two, from 0.149 in the baseline scenario to 0.068.

Coordination between firms is better than in the baseline (lower variance between firms' strategies), but steadily increasing, which denotes coordination failures and an increasing concentration on the goods market.

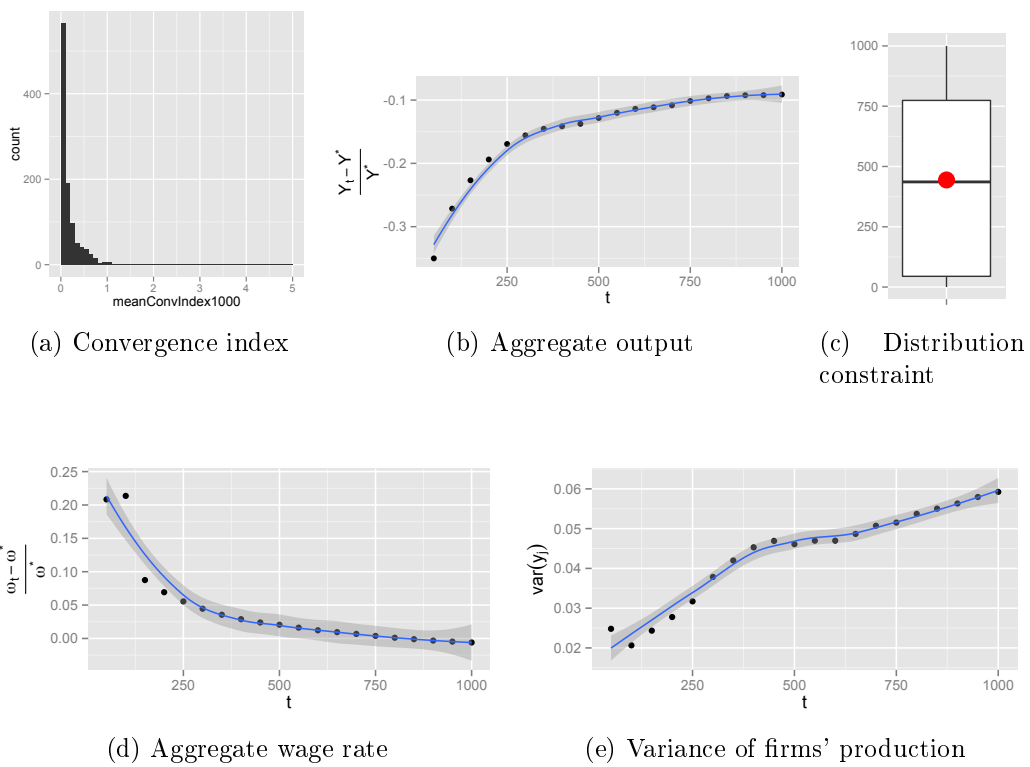


Figure 8: Simulations with fixed nominal wages, results over the 810 simulations.

C Output contraction induced by worsening of the distribution of purchasing power

Aggregate cash-on-hand M_t is given by

$$M_t = \tilde{Y}_t - P_t C_t. \quad (32)$$

We now distinguish between consumption expenditures of households with oversized and with undersized purchasing power:

$$P_t C_t = P_t^{over} C_t^{over} + P_t^{under} C_t^{under} \quad (33)$$

allowing rewriting equation (32):

$$C_t^{under} = \frac{\tilde{Y}_t - M_t - P_t^{over} C_t^{over}}{P_t^{under}}. \quad (34)$$

First difference of C_t^{under} now writes as

$$\Delta C_t^{under} = \frac{\Delta \tilde{Y}_t - \Delta M_t - P_t^{over} \Delta C_t^{over}}{P_t^{under}}. \quad (35)$$

Now take into account that for given agent strategies, $H_t = H_{t-1}$, implying $\tilde{Y}_t = \tilde{Y}_{t-1}$ (in fact, $\tilde{Y}_t = wH_t + \Pi_{t-1} + M_{t-1}$, $\Pi_{t-1} = P_{t-1}C_{t-1} - wH_{t-1}$ and $M_{t-1} = \tilde{Y}_{t-1} - P_{t-1}C_{t-1}$) and that consumption expenditure of households with oversized purchasing power is bounded by their goods demand: $C_t^{over} = C_{t-1}^{over} = C^{d,over}$. Consequently, we have

$$\Delta C_t^{under} = \frac{-\Delta M_t}{P_t^{under}}. \quad (36)$$

As soon as cash-on-hand of all households with undersized purchasing power has become zero, aggregate cash-on-hand M_t is bound to increase, resulting in contraction of C_t^{under} , C_t and Y_t .