Tipping points in macroeconomic agent-based models

Stanislao Gualdi [a,b,*], Marco Tarzia [a], Francesco Zamponi [c],
Jean-Philippe Bouchaud [d,e]

[a] Université Pierre et Marie Curie - Paris 6, Laboratoire de Physique, Théorique de la Matière Condensée, 4, Place Jussieu,
Tour 12, 75252 Paris Cedex 05, France
[b] IRAMIS, CEA-Saclay, 91191 Gif sur Yvette Cedex, France
[c] Laboratoire de Physique Théorique, École Normale Supérieure, UMR 8549 CNRS, 24 Rue Lhomond, 75231 Paris Cedex 05, France
[d] CFM, 23 rue de l'Université, 75007 Paris, France
[e] Ecole Polytechnique, 91120 Palaiseau, France

ARTICLE INFO

ABSTRACT

The aim of this work is to explore the possible types of *phenomena* that simple macro-economic Agent-Based models (ABMs) can reproduce. We propose a methodology, inspired by statistical physics, that characterizes a model through its "phase diagram" in the space of parameters. Our first motivation is to understand the large macro-economic fluctuations observed in the "Mark I" ABM devised by Delli Gatti and collaborators. In this regard, our major finding is the generic existence of a *phase transition* between a "good economy" where unemployment is low, and a "bad economy" where unemployment is high. We then introduce a simpler framework that allows us to show that this transition is *robust* against many modifications of the model, and is generically induced by an asymmetry between the rate of hiring and the rate of firing of the firms. The unemployment level remains small until a tipping point, beyond which the economy suddenly collapses. If the parameters are such that the system is close to this transition, any small fluctuation is amplified as the system jumps between the two equilibria. We have explored several natural extensions of the model. One is to introduce a bankruptcy threshold, limiting the firms maximum level of debt-to-sales ratio. This leads to a rich phase diagram with, in particular, a region where *acute endogenous crises* occur, during which the unemployment rate shoots up before the economy can recover. We also introduce simple wage policies. This leads to inflation (in the "good" phase) or deflation (in the "bad" phase), but leaves the overall phase diagram of the model essentially unchanged. We have also explored the effect of simple monetary policies that attempt to contain rising unemployment and defang crises. We end the paper with general comments on the usefulness of ABMs to model macroeconomic phenomena, in particular in view of the time needed to reach a steady state that raises the issue of *ergodicity* in these models.

© 2014 Elsevier B.V. All rights reserved.

It is human nature to think wisely and to act absurdly – Anatole France.

* Corresponding author.
  E-mail address: stanislao.gualdi@gmail.com (S. Gualdi).

1. Introduction

1.1. From micro-rules to macro-behaviour

Inferring the behaviour of large assemblies from the behaviour of its elementary constituents is arguably one of the most important problems in a host of different disciplines: physics, material sciences, biology, computer sciences, sociology and, of course, economics and finance. It is also a notoriously hard problem. Statistical physics has developed in the last 150 years essentially to understand this micro–macro link. Clearly, when interactions are absent or small enough, the system as a whole merely reflects the properties of individual entities. This is the canvas of traditional macro-economic approaches. Economic agents are assumed to be identical, non-interacting, rational agents, an idealization known as the "Representative Agent" (RA). In this framework, micro and macro trivially coincide. However, we know (in particular from physics) that discreteness, heterogeneities and/or interactions can lead to totally unexpected phenomena. Think for example of super-conductivity or super-fluidity[1]: before their experimental discovery, it was simply beyond human imagination that individual electrons or atoms could "conspire" to create a collective state that can flow without friction. Not only micro and macro behaviour do not coincide in general, but also genuinely *surprising* behaviour can emerge through aggregation. From the point of view of economic theory, this is interesting, because financial and economic history is strewn with bubbles, crashes, crises and upheavals of all sorts. These are very hard to fathom within a Representative Agent framework (Kirman, 1992), within which crises would require large aggregate shocks, when in fact small local shocks can trigger large systemic effects when heterogeneities, interactions and network effects are taken into account (see e.g. Brock and Durlauf, 2001; Ball, 2012, for a recent review, see Bouchaud, 2013; Bonart et al., 2014). The need to include these effects has spurred a large activity in "Agent-Based models" (ABMs) (for a concise overview of ABM for collective phenomena, see Goldstone and Janssen, 2005, for early papers on ABM for economics, see e.g. Axtell, 2000; Tesfatsion, 2002, More recent attempts can be found in Ashraf et al., 2011; Raberto et al., 2012; Dosi and et al., 2013, for some review papers on ABM for economics and finance, see e.g. Leijonhufvud, 2006; Hommes, 2006; LeBaron and Tesfatsion, 2008; Dawid and Neugart, 2011; Cristelli et al., 2011). These models need numerical simulations, but are extremely versatile because any possible behavioural rules, interactions, and heterogeneities can be taken into account.

In fact, these models are so versatile that they suffer from the "wilderness of high dimensional spaces" (paraphrasing Sims, 1980). The number of parameters and explicit or implicit choices of behavioural rules is so large (~ 10 or more, even in the simplest models, see below) that the results of the model may appear unreliable and arbitrary, and the calibration of the parameters is an hopeless (or highly unstable) task. Mainstream RA "Dynamic Stochastic General Equilibrium" (DSGE) models, on the other hand, are simple enough to lead to closed form analytical results, with simple narratives and well-trodden calibration avenues (see e.g. Gali, 2008). In spite of their unrealistic character, these models appear to perform satisfactorily in 'normal' times, when fluctuations are small. However, they become deeply flawed in times of economic instability (Buiter, 2009), suggesting that different assumptions are needed to understand what is observed in reality. But even after the 2008 crisis, these traditional models are still favoured by most economists, both in academia and in institutional and professional circles. ABMs are seen at best as a promising research direction and at worst as an unwarranted "black box" (see Fagiolo and Roventini, 2012 for an enlightening discussion on the debate between traditional DSGE models and ABMs, and Fagiolo et al., 2007; Pyka and Fagiolo, 2007; Kirman, 2010; Caballero, 2010 for further insights).

1.2. A methodological manifesto

At this stage, it seems to us that some clarifications are indeed needed, concerning both the objectives and the methodology of Agent-Based models. ABMs do indeed suffer from the wilderness of high dimensional spaces, and some guidance is necessary to put these models on a firm footing. In this respect, statistical physics offers a key concept: the *phase diagram* in parameter space (for a general, insightful introduction, see Goldenfeld, 1992). A classic example, shown in Fig. 1, is the phase diagram of usual substances as a function of two parameters, here temperature and pressure. The generic picture is that the number of distinct phases is usually small (e.g. three in the example of Fig. 1: solid, liquid, gas). Well within each phase, the properties are qualitatively similar and small changes of parameters have a small effect. Macroscopic (aggregate) properties do not fluctuate any more for very large systems and are robust against changes of microscopic details. This is the "nice" scenario, where the dynamics of the system can be described in terms of a small number of macroscopic (aggregate) variables, with some effective parameters that encode the microscopic details. But other scenarios are of course possible; for example, if one sits close to the boundary between two phases, fluctuations can remain large even for large systems and small changes of parameters can radically change the macroscopic behaviour of the system. There may be mechanisms naturally driving the system close to criticality (like in Self-Organized Criticality, Bak, 1996) or, alternatively, situations in which whole phases are critical, like for "spin-glasses" (on this point, see the interesting discussion in Kondor et al., 2014).

In any case, the very first step in exploring the properties of an Agent-Based model should be, we believe, to identify the different possible phases in parameter space and the location of the phase boundaries. In order to do this, numerical

---

[1] See e.g. Ref. Balibar (2007) for an history of the discovery of super-fluidity and a list of references.
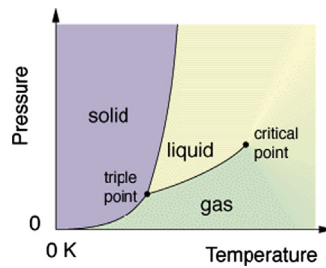
Fig. 1. A typical "phase diagram", here the solid–liquid–gas phases in the temperature-pressure plane. Far from phase transitions, within a given phase, the behaviour of the system is qualitatively similar for all values of the parameters. Close to phase boundaries, on the other hand, small changes of parameters can lead to dramatic macroscopic changes.

simulations turn out to be very helpful (Buchanan, 2008; Farmer and Foley, 2009): aggregate behaviour usually quickly sets in, even for small sizes. Some parameters usually turn out to be crucial, while others are found to play little role. This is useful to establish a qualitative *phenomenology* of the model – what kind of behaviour can the model reproduce, which basic mechanisms are important, which effects are potentially missing? This first, qualitative step allows one to unveil the "skeleton" of the ABM. Simplified models that retain most of the phenomenology can then be constructed and perhaps solved analytically, enhancing the understanding of the important mechanisms, and providing some narrative to make sense of the observed effects. In our opinion, calibration of an ABM using real data can only start to make sense after this initial qualitative investigation is in full command – which is in itself not trivial when the number of parameters is large. The phase diagram of the model allows one to restrict the region of parameters that lead to "reasonable" outcomes (see for example the discussion in Giardina and Bouchaud, 2003; Lye et al., 2012).

1.3. Outline, results and limitations of the paper

The aim of this paper is to put these ideas into practice in the context of a well-studied macroeconomic Agent-Based model (called "Mark I" below), devised by Delli Gatti and collaborators (The Mark I family of models was elaborated in a series of papers and books, in particular Gaffeo et al., 2008; Delli Gatti et al., 2008, 2011). This model is at the core of the European project "CRISIS", which partly justifies our attempt to shed some theoretical light on this framework.[2] In the first part of the paper, we briefly recall the main ingredients of the model and show that as one increases the baseline interest rate, there is a phase transition between a "good" state of the economy, where unemployment is low and a "bad" state of the economy where production and demand collapse. In the second part of the paper, we study the phase diagram of a highly simplified version of Mark I, dubbed 'Mark 0' that aims at capturing the main drivers of this phase transition. The model is a "hybrid" macro/ABM model where firms are treated individually but households are only described in aggregate. Mark 0 does not include any exogenous shock; crises can only be of endogenous origin. We show that the most important parameter in this regard is the asymmetry between the firms' propensity to hire (when business is good) or to fire (when business is bad). In Mark I, this asymmetry is induced by the reaction of firms to the level of interest rates, but other plausible mechanisms would lead to the same effect. The simplest version of the model is amenable to an analytic treatment and exhibits a "tipping point" (i.e. a discontinuous transition) between high employment and high unemployment. When a bankruptcy condition is introduced (in the form of a maximum level of debt-to-sales ratio), the model reveals an extremely rich phenomenology: the "good" phase of the economy is further split into three distinct phases: one of full employment, the second one with a substantial level of residual unemployment, and the third one, highly interesting region, where endogenous crises appear. We find that both the amount of credit available to firms (which in our model sets the bankrupcty threshold), and the way default costs are absorbed by the system, and are the most important aspects in shaping the qualitative behaviour of the economy. Finally, we allow wages to adapt (whereas they are kept fixed in Mark I) and allow inflation or deflation to set in. Still, the overall shape of the phase diagram is not modified. We investigate further enhancements of the models, in particular simple policy experiments. Open questions and future directions, in particular concerning macroeconomic ABMs in general, are discussed in our final section.

Before embarking to the core of our results, we want to clearly state what our ambition and objectives are, and what they are not. We do claim that the methodology proposed here is interesting and general, and could help improving the relevance of macroeconomic ABMs. We do believe that large aggregate volatility and crises (in particular those appearing in Mark I) can be understood through the lens of instabilities and phase transitions, as exemplified by our highly stylized Mark 0 model. We are also convinced that simple "skeleton" ABMs (for which most of the phenomenology can be fully dissected) must be developed and compared with traditional DSGE models before embarking into full-fledged models of the economy. On the other hand, we do *not* claim that our basic Mark 0 framework is necessarily the best starting point. Mark 0 was primarily set up as a simplification of Mark I. Still, we find that Mark 0 leads to a surprisingly rich and to some extent realistic set of possible behaviours, including business cycles and crises, inflation, and policy experiments. However, we do

---

[2] see www.crisis-economics.eu

not wish to claim that Mark 0 is able to reproduce all the known empirical stylized facts and could well be in contradiction with some of them.[3] In view of the simplicity of the model, this has to be expected. But we believe that Mark 0 can serve as a useful building block in the quest of a more comprehensive ABM, as more and more effects are progressively incorporated, in a controlled manner, to the model. Even if Mark 0 turns out to be little more than a methodological exercise, we hope that the ABM community (and perhaps beyond) will find it inspiring.

To conclude this (long) introduction, let us insist that all the claims made in this paper only refer to the studied *models*, but do not necessarily apply to economic reality. If fact, our central point is that a model has to be understood inside-out before even trying to match any empirical fact.

2. A phase transition in "Mark I"

2.1. Description of the model in a nutshell

The Mark I family of agent-based models was proposed by Delli Gatti and collaborators as a family of simple stylized macroeconomic models (Gaffeo et al., 2008; Delli Gatti et al., 2008, 2011). Note that several other macroeconomic Agent-Based models have been put forth in the recent years, see Deissenberg et al. (2008), Dawid et al. (2011), Dosi et al. (2005, 2010), Mandel et al. (2009), and Metzig and Gordon (2012). Mark I is particularly interesting because large fluctuations in unemployment and output seem to persist in the stationary state (a feature in fact shared by many ABMs cited above).

The Mark I economy (Gaffeo et al., 2008; Delli Gatti et al., 2008) is made up of a set of firms, households, firms owners and a bank. Firms produce a certain quantity of a single (and not storable) good, proportional to the number of their employees, that is sold at time-dependent and firm-dependent prices. Firms pay identical time-independent wages. When the cash owned by a firm is not enough to pay the wages, it asks banks for a loan. The bank provides loans to the firms at an interest rate that depends on the financial fragility of the firm. Households provide workforce in exchange of a salary and want to spend a fixed fraction of their savings (or wealth). The owners of the firms do not work but receive dividends if the firms make profits. Firms are adaptive, in the sense that they continuously update their production (i.e. they hire/fire workers) and their prices, in an attempt to match their production with the demand of goods issued by the households. They also choose how much extra loan they want to take on, as a function of the offered interest rate. This last feature, combined with the price and production update rules, will turn out to be crucial in the dynamics of the model.

The above description defines the basic structure of the Mark I family, but it is of course totally insufficient to code the model, since many additional choices have to be made, leading to several different possible implementations of the model. Here we will use as a baseline model one of the simplest implementation of Mark I, whose description can be found in Grazzini and Delli Gatti (2012) (the precise version of Mark I that we used as a reference); the total number of parameters in this version is 10 (but more parameters are actually implicitly fixed from the beginning). We have recoded this basic version and also a slightly different version that we call "Mark I+", which differs on minor details (some that we will specify below) but also on one major aspect: our version of the model *strictly conserves the amount of money in circulation*, i.e. the money in the bank + total firm assets + total households wealth (here savings), in order to avoid – at this stage – any effect due to uncontrolled money creation. A detailed pseudo-code of Mark I+ is provided in Appendix A.

2.2. State variables

In short (see Appendix A for a complete description), the dynamic evolution of the model is defined by the following state variables. The state of each firm $i = 1 \ldots N_F$ is specified by its price $p_i(t)$, the salary it offers $W_i(t)$, its production $Y_i(t)$, its target production $Y_i^T(t)$, its demand $D_i(t)$, its liquidity $\mathcal{L}_i(t)$, its total debt $\mathcal{D}_i^T(t)$. Moreover, each firm is owned by a household and has a list of employees that is dynamically updated. The state of each household $a = 1 \ldots N_H$ is specified by its wealth (in the form of savings) $S_a(t)$ and by the firm for which it works (if any).

2.3. Update rules for prices and production

Among all the micro-rules that any Agent-Based model has to postulate, some seem to be more crucial than others. An important item in Mark I is the behavioural rule for firms adaptation to their economic environment. Instead of the standard, infinite horizon, profit optimizing firm framework (that is both unrealistic and intractable), Mark I postulates a heuristic rule for production $Y_i(t)$ and price $p_i(t)$ update, which reads as follows:

$$Y_i(t) = D_i(t) \text{ and } p_i(t) > \overline{p}(t) \Rightarrow Y_i^T(t+1) = Y_i(t)[1 + \gamma_y \xi_i(t)]$$
$$Y_i(t) = D_i(t) \text{ and } p_i(t) < \overline{p}(t) \Rightarrow p_i(t+1) = p_i(t)[1 + \gamma_p \xi_i(t)]$$

---

[3] The idea of building *mathematical models of reality* that reproduce some phenomena, but might even be in contradiction with others, is at the heart of the development of physics. Most probably, its use was introduced in the Hellenistic period (Russo, 2004). A striking example Russo (2004) is Archimedes' *On Floating Bodies*. In the first of the two books, in fact, Archimedes provides a mathematical proof of the sphericity of Earth (assumed to be liquid and at rest). However, in the second book, he assumes the surface of Earth to be flat for the purpose of describing other phenomena, for which the sphericity of Earth is irrelevant.

$$Y_i(t) > D_i(t) \text{ and } p_i(t) < \overline{p}(t) \Rightarrow Y_i^T(t+1) = Y_i(t)[1 - \gamma_y \xi_i(t)]$$
$$Y_i(t) > D_i(t) \text{ and } p_i(t) > \overline{p}(t) \Rightarrow p_i(t+1) = p_i(t)[1 - \gamma_p \xi_i(t)], \tag{1}$$

where $D_i(t)$ is the total demand for the goods produced by firm i at time t, and

$$\overline{p}(t) = \frac{\sum_i p_i(t) D_i(t)}{\sum_i D_i(t)} \tag{2}$$

is the average price of sold goods at time t, $\xi_i(t)$ a $U[0, 1]$ random variable, independent across firms and across times, and γ_y, γ_p two parameters in [0, 1]. The quantity $Y_i^T(t)$ is the *target* production at time t, not necessarily the realized one, as described below. These heuristic rules can be interpreted as a plausible *tâtonnement* process of the firms, that attempt to guess their correct production level and price based on the information on the last time step. In spirit, each unit time step might correspond to a quarter, so the order of magnitude of the γ parameters should be a few percent. Note that in the version of Mark I that we consider, wages are fixed to a constant value $W_i(t) \equiv 1$, for all times and all firms.

As we shall see later, the adaptive price/production adjustments described in Eq. (1) lead to two stable attractors (full employment and full unemployment). Which of the two is be reached by the dynamics depends mainly on the level of asymmetry between the upward and downward production adjustments. In Eq. (1), the production adjustment depends on a single parameter γ_y and in this case the system evolves towards a full employment state in the absence of any other constraint. However, as it will become clear in the following section, financial constraints on loans may lead to an effective "weakening" of the upward adjustment, possibly driving the dynamics towards the full unemployment state. As long as such asymmetries between upward and downward production adjustments exist (together with some noise in the price dynamics), the scenario described above and in the following sections is very general and in fact do not depend on details of the update process.

2.4. Debt and loans

The model further assumes linear productivity, hence the target production corresponds to a target workforce $Y_i^T(t)/\alpha$, where α is a constant coefficient that can always be set to unity (gains in productivity are not considered at this stage). The financial need of the firm is $\max[0, Y_i^T(t)W_i(t) - \mathcal{L}_i(t)]$, where $\mathcal{L}_i(t)$ is the cash available. The total current debt of the firm is $\mathcal{D}_i^T(t)$. The financial fragility of the firm $\ell_i(t)$ is defined in Mark I as the ratio of debt over cash. The offered rate by the banks for the loan is given by

$$\rho_i(t) = \rho_0 G(\ell_i(t)) \times (1 + \xi_i'(t)), \tag{3}$$

where ρ_0 is the baseline (central bank) interest rate, G is an increasing function (taken to be $G(\ell) = 1 + \tanh(\ell)$ in the reference Mark I and $G(\ell) = 1 + \ln(1 + \ell)$ in Mark I+), and ξ' another noise term drawn from a uniform distribution $U[0, 1]$. Depending on the rate offered, firms decide to take the full loan they need or only a fraction $F(\rho)$ of it, where $F \leq 1$ is a decreasing function of ρ, called the "credit contraction". For example, in the reference Mark I, $F(\rho \leq 5\%) = 1$ and $F(\rho > 5\%) = 0.8$. We have played with the choice of the two functions F, G and the phase transition reported below is in fact robust whenever these functions are reasonable. In Mark I+, we chose a continuous function, such as to avoid built in discontinuities[4]:

$$F(x) = \begin{cases} 1 & \text{if } x < 5\% \\ 1 - \frac{x - 5\%}{5\%} & \text{if } 5\% < x < 10\% \\ 0 & \text{if } x > 10\%. \end{cases} \tag{4}$$

The important feature here is that when $F < 1$, the firm does not have enough money to hire the target workforce $Y_i^T(t)$ and is therefore obliged to hire less, or even to start firing in order to match its financial constraints. This financial constraint therefore induces an *asymmetry* in the hiring/firing process: when firms are indebted, hiring will be slowed down by the cost of further loans. As we will see later, this asymmetry is responsible for an abrupt change in the steady state of the economy.

2.5. Spending budget and bankruptcy

Firms pay salaries to workers and households determine their budget as a fraction c (constant in time and across households) of their total wealth (including the latest salary). Each household then selects M firms at random and sorts them according to their price; it then buys all it can buy from each firm sequentially, from the lower price to the highest price.[5] The budget left-over is added to the savings. Each firm sells a quantity $D_i(t) \leq Y_i(t)$, computes its profits (that includes interests paid on debt), and updates its cash and debt accordingly. Moreover, each firm pays back to the bank a fraction τ of

---

[4] Note however that the arbitrary thresholds (5% and 10%) in Eq. (4) are of little importance and only affect the precise location of the phase transition.
[5] In this sense, the good market is not efficient since the household demand is not necessarily satisfied when M is small. The job market instead, though not efficient due to the presence of unvoluntary unemployment, is characterized by perfect information since all the workers can contact all the firms until all the open positions are filled.

its total debt $\mathcal{D}_i^T(t)$. It also pays dividends to the firm owners if profits are positive. Firms with negative liquidity $\mathcal{L}_i(t) < 0$ go bankrupt. In Mark I+, the cost of the bankruptcy (i.e. $-\mathcal{L}_i(t)$) is spread over healthy firms and on households. Once a firm is bankrupt it is re-initialized in the next time step with the owner's money, to a firm with a price and production equal to their corresponding average values at that moment in time, and zero debt (see Appendix A for more precise statements).

2.6. Numerical results: a phase transition

When exploring the phase space of Mark I, it soon becomes clear that the baseline interest rate ρ_0 plays a major role. In order not to mix different effects, we remove altogether the noise term ξ in Eq. (3) that affects the actual rate offered to the firms. We find that as long as ρ_0 is smaller than a certain threshold ρ_c, firms are on average below the credit contraction threshold and always manage to have enough loans to pay wages. In this case the economy is stable and after few (~ 100) time steps reaches a stationary state where the unemployment rate is low. If on the other hand the baseline interest rate ρ_0 exceeds a critical value ρ_c, firms cannot afford to take as much loans as they would need to hire (or keep) the desired amount of workers. Surprisingly, this induces a sudden, catastrophic breakdown of the economy. Production collapses to very small values and unemployment sky-rockets. This transition between two states of the economy takes place in both the reference Mark I and the modified Mark I+; as we shall show in the next section, this transition is actually generic and occurs in simplified models as well. Note in particular that ρ_c is *different* from the value at which $F(x)$ starts decreasing.

The data we show in Fig. 2 corresponds to Mark I+ with parameters $\gamma_p = \gamma_y = 0.1$ and $M=3$ (see Appendix A for the general parameter setting of the model). While the qualitative behaviour of the model is robust, the details of the transition may change with other parameter settings. For example, smaller values of γ_p, γ_y lead to lower critical thresholds ρ_c (as well as smaller values of M) and to longer equilibration times (T_{eq} scales approximately as $1/\gamma_{y,p}$ for $\rho_0 < \rho_c$). Increasing the size of the economy only affects the magnitude of the fluctuations within one phase leaving the essential features of the transition unchanged. Interestingly, although it is not clear from Fig. 2, the model exhibits oscillatory patterns of the employment rate. The presence of these oscillations can be seen in the frequency domain of the employment rate time series (not shown here), which is essentially characterized by a white noise power spectrum with a well defined peak at intermediate frequencies. All these effects will become clearer within the reduced model described in the next section.

Anticipating the results of the next section, we have characterized, as a function of ρ_0, the asymmetry ratio R that characterizes how firms react to the need to hire or to fire. More precisely, we compute the ratio between the *target* number of job creations (resp. destructions) – as a reaction to excess demand (or supply) in the previous time step – to the *realized* number of hires (fires) after financial constraints are met. At each time step we can then extract the average target-to-realized ratios for firms in excess supply and demand. The ratio R (averaged over time) between these two numbers is a proxy of the asymmetry with which firms react to the need to hire or fire. It is clear that increasing ρ_0 hobbles the capacity of firms to hire when they need to, and therefore decreases the ratio R, as seen in Fig. 2, which suggests that the employment collapse is related to this ratio. The results of the next section will fully confirm this scenario.

To sum up, our most salient finding is that Mark I (or Mark I+) has essentially two stationary states, with a first order (discontinuous) transition line separating the two. If the parameters are such that the system lies close to this critical line, then any small modulation of these parameters – such as the noise term that appears in Eq. (3) – will be *amplified* by the proximity of the transition, and lead to interesting boom/bust oscillations, of the kind originally observed in Mark I (Grazzini
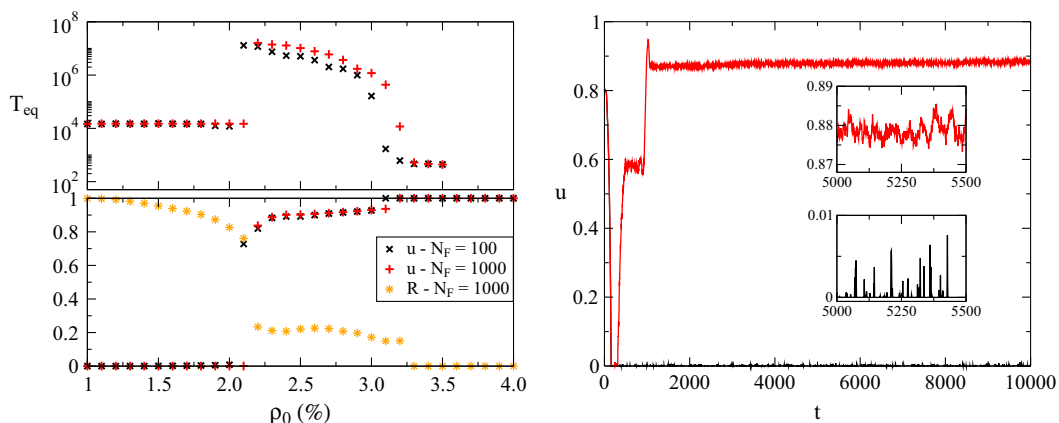

Fig. 2. Left: Average unemployment rate u as a function of the interest rate ρ_0 for two system sizes (with $N_H = 10N_F$ in both cases). The average is over 100 000 time steps discarding the first 50 000 time steps. The phase transition at $\rho_c = 2.1\%$ is of first order with u jumping discontinuously from small values to 1 (intermediate values obtained for $\rho_0 \in [2.1\%, 3.1\%]$ are only due to a much longer equilibration time near the critical point). Note that in the bottom graph we show averages for $T_{eq} = 50\,000$ regardless of ρ_0 while an estimate of the time needed to reach the steady state as a function of ρ_0 is plotted in logarithmic scale in the top graph. In the bottom graph we also show the average value of the asymmetry measure R (see main text) for $N_F = 1000$. As one can see R is decreasing with the interest rate up to a point where the asymmetry is strong enough to drive the economy into the high unemployment phase. Right: Two trajectories of the unemployment rate with $N_F = 1000$ at $\rho_0 = 1.9\%$ and $\rho_0 = 2.5\%$.

and Delli Gatti, 2012), and perhaps of economic relevance. The question, of course, is: how generic is this scenario? We will now show, by studying much simplified versions of Mark I, that this transition is generic, and can indeed be induced by the asymmetry between hiring and firing. We will then progressively enrich our watered-down model (call "Mark 0" below) and see how the qualitative picture that we propose is affected by additional features.

3. Hybrid ABM's: the minimal "Mark 0" model

Moving away from the RA framework, Agent-Based Modeling bites the bullet and attempts to represent in detail *all* the individual components of the economy (as, for example, in Deissenberg et al., 2008; Dawid et al., 2011). This might however be counter-productive, at least in the research stage we are still in: keeping too many details is not only computer-time consuming, but it may also hobble the understanding of the effects that ABMs attempt to capture. It may well be that some sectors of the economy can be adequately represented in terms of aggregate variables, while discreteness, heterogeneities and interactions are crucial in other sectors. In our attempt to simplify Mark I, we posit that the whole household sector can be represented by aggregate variables: total wealth (again entirely in the form of savings) $S(t)$, total income wage $W(t)$ and total consumption budget $C_B(t)$ (which, as we will emphasize below, is in general larger than the *actual* consumption $C(t)$). We also remove the banking sector and treat the loans in the simplest possible way – see below. While the interest rate is zero in the simplest version, the incentive to hire/fire provided by the interest rate, that was at play in Mark I, will be encoded in a phenomenological way in the update rule for production. The firms, on the other hand, are kept as individual entities (but the above simplifications will allow us to simulate very large economies, with $N_F = 100\,000$ firms or more).

3.1. Set-up of the model

The minimal version of the Mark 0 model is defined as follows. The salient features are the following:

- There are N_F firms in total and μN_F households, $\mu > 1$.[6] Each firm $i = 1 \ldots N_F$ pays a salary $W_i(t)$ and produces output $Y_i(t)$ by means of a one-to-one technology that uses only labor as an input. Productivity is chosen to be *constant in time* and fixed to $\alpha = 1$. We therefore neglect any productivity shocks in our model; interestingly, crises (when they occur) will be of endogenous origin. With $\alpha = 1$, $Y_i(t)$ is simply equal to the number of employees of firm i. Hence, the employment rate $\varepsilon(t)$ and the unemployment rate $u(t)$ are as follows:

$$\varepsilon(t) = \frac{1}{\mu N_F} \sum_i Y_i(t),$$
$$u(t) = 1 - \varepsilon(t). \tag{5}$$

- Households are described by their total accumulated savings $S(t)$ (which at this stage are always non-negative) and by their total wage income $\sum_i W_i(t) Y_i(t)$. At each time step, they set a total consumption budget[7]

$$C_B(t) = c \left[S(t) + \sum_i W_i(t) Y_i(t) \right] \tag{6}$$

which is distributed among firms using an intensity of choice model (see e.g. Anderson et al., 1992). The demand of goods for firm i is therefore

$$D_i(t) = \frac{C_B(t)}{p_i(t)} \frac{e^{-\beta p_i(t)/\overline{p}(t)}}{Z(t)}, \quad Z(t) = \sum_i e^{-\beta p_i(t)/\overline{p}(t)} \tag{7}$$

where β is the price sensitivity parameter determining an exponential dependence of households demand in the price offered by the firm; $\beta = 0$ corresponds to complete price insensitivity and $\beta \to \infty$ means that households select only the firm with the lowest price.[8] The normalization is such that $C_B(t) = \sum_i p_i(t) D_i(t)$, as it should be.

- Firms are described by their price $p_i(t)$, their salary $W_i(t)$, and their production $Y_i(t)$.
  - For simplicity, we fix the salary $W_i(t) \equiv 1$ – an extension that includes wage dynamics is discussed below.
  - For the price, we keep the Mark I price update rule (1), with the average production-weighted price:

$$\overline{p}(t) = \frac{\sum_i p_i(t) Y_i(t)}{\sum_i Y_i(t)}. \tag{8}$$

---

[6] Actually, households are treated as a unique aggregate variable, therefore μ is not a relevant parameter: one can see that its value is irrelevant and one can always set $\mu = 1$ for simplicity. Yet it is useful to think that the aggregate variables represents in an effective way a certain number of individual households, hence we keep the parameter μ explicit in the following.

[7] Of course, one could choose different c's for the fraction of savings and the fraction of wages devoted to spending, or any other non-linear spending schedule.

[8] In this sense, as long as $\beta > 0$ firms compete on prices. An averaged scatter plot of firms profits versus the price offered (not shown here) indeed displays a well-shaped concave profit function.

Note that this price update rule only makes sense if firms anticipate that households are price sensitive, i.e. if $\beta > 0$, which we will assume in the following. Still, the dynamics of the model as defined remains perfectly well-behaved when $\beta = 0$, even if in this limit, the rational behaviour of firms would be to increase their price indefinitely and produce very little.

– For production, we assume that firms are more careful with the way they deal with their workforce than posited in Mark I. Independent of their price level, firms try to adjust their production to the observed demand. When firms want to hire, they open positions on the job market; we assume that the total number of unemployed workers, which is $\mu N_F u(t)$, is distributed among firms according to an intensity of choice of model which depends on both the wage offered by the firm[9] and the same parameter β as it is for Eq. (7); therefore the maximum number of available workers to each firm is

$$\mu \tilde{u}_i(t) = \frac{e^{\beta W_i(t)/\overline{w}(t)}}{\sum_i e^{\beta W_i(t)/\overline{w}(t)}} \mu N_F u(t). \qquad (9)$$

where

$$\overline{w}(t) = \frac{\sum_i W_i(t) Y_i(t)}{\sum_i Y_i(t)}. \qquad (10)$$

In summary, we have

$$\begin{aligned}
\text{If } Y_i(t) < D_i(t) \Rightarrow & \begin{cases} Y_i(t+1) = Y_i(t) + \min\{\eta_+ (D_i(t) - Y_i(t)), \mu \tilde{u}_i(t)\} \\ \text{If } p_i(t) < \overline{p}(t) \Rightarrow p_i(t+1) = p_i(t)(1 + \gamma_p \xi(t)) \\ \text{If } p_i(t) \geq \overline{p}(t) \Rightarrow p_i(t+1) = p_i(t) \end{cases} \\
\text{If } Y_i(t) > D_i(t) \Rightarrow & \begin{cases} Y_i(t+1) = \max\{Y_i(t) - \eta_- [Y_i(t) - D_i(t)], 0\} \\ \text{If } p_i(t) > \overline{p}(t) \Rightarrow p_i(t+1) = p_i(t)(1 - \gamma_p \xi(t)) \\ \text{If } p_i(t) \leq \overline{p}(t) \Rightarrow p_i(t+1) = p_i(t) \end{cases}
\end{aligned} \qquad (11)$$

where $\eta_\pm \in [0, 1]$ are what we denote as the hiring/firing propensity of the firms. Note that this rule ensures that there is no overshoot in production; furthermore the max in the second rule is not necessary mathematically when $\eta_- \leq 1$, but we kept it for clarity. Each row of Eq. (11) specifies an adjustment mechanism for output and the individual price. According to this mechanism, there is an increase in output if excess demand is positive, and a decrease in output if excess demand is negative, i.e. if there is excess supply. The propensities to hire/fire $\eta_\pm$ can be seen as the sensitivity of the output change to excess demand/supply. These sensitivities are generally less than unity (i.e. the firm is not adjusting output one to one with excess demand/supply), because there are hiring (firing) costs of different kinds (real costs, time-to-hire, administrative constraints, inertia, etc.). Firing costs generate "labour hoarding", while hiring costs prevents the economy to adapt quickly to excess demand. Hiring and firing costs may not be identical. Therefore the sensitivity to excess demand (or hiring propensity η_+) may be different from the sensitivity to excess supply (or firing propensity η_-). Note that because of the $\min\{D_i(t) - Y_i(t), \mu u(t)\}$ term in Eq. (11), the total production of the model is bounded by $\sum_i Y_i(t) = \mu N_F \varepsilon(t) \leq \mu N_F$, as it should be because $\varepsilon(t) = 1$ corresponds to full employment and in that case $\sum_i Y_i(t) = \mu N_F$. However, in the following we will sometimes (when stated) remove this bound for a better mathematical tractability. This amounts to replace $\min\{D_i(t) - Y_i(t), \mu u(t)\} \to D_i(t) - Y_i(t)$ in Eq. (11) (it corresponds to choosing $\mu = \infty$). Removing the bound corresponds to a situation where labor resources can freely exceed the working population, such that the high employment phase translates into an exponential explosion of the economy output, which is of course unrealistic.

• *Accounting of firms and households*: Each firm $i = 1 \ldots N_F$ pays a total wage bill $W_i(t) Y_i(t)$ and has total sales $p_i(t) \min\{Y_i(t), D_i(t)\}$. Moreover, if the profit of the firm $\mathcal{P}_i(t) = p_i(t) \min\{Y_i(t), D_i(t)\} - W_i(t) Y_i(t)$ is positive, the firm pays a dividend $\delta \times \mathcal{P}_i(t)$ to the households.[10]

Note that if $D_i > Y_i$, the demand for goods of firm i cannot be immediately satisfied, and we assume that in this case households involuntarily save and delay their consumption until the next round (but still using Eq. (6) with the correctly updated savings). The *actual* consumption $C(t)$ (limited by production) is therefore given by

$$C(t) := \sum_{i=1}^{N_F} p_i(t) \min\{Y_i(t), D_i(t)\} \leq C_B(t) = \sum_{i=1}^{N_F} p_i(t) D_i(t). \qquad (12)$$

In summary, the accounting equations for total accumulated savings $S(t)$ and firms' net deposits $\mathcal{E}_i(t)$ – possibly negative

---

[9] Since at this stage wages are equal among firms the distribution is uniform. Below we allow firms to update their wage. A higher wage will then translate in the availability of a larger share of unemployed workers in the hiring process.

[10] We have also considered the case where firms distribute a fraction δ^+ of the profits *plus* the reserves. See below.

– are the following (here $\theta(x \geq 0) = 1$ and $\theta(x < 0) = 0$ is the Heaviside step function):

$$\mathcal{P}_i(t) = p_i(t) \min\{Y_i(t), D_i(t)\} - W_i(t)Y_i(t),$$
$$\mathcal{E}_i(t+1) = \mathcal{E}_i(t) + \mathcal{P}_i(t) - \delta\mathcal{P}_i(t)\theta(\mathcal{P}_i(t)),$$
$$S(t+1) = S(t) - \sum_i \mathcal{P}_i(t) + \delta\sum_i \mathcal{P}_i(t)\theta(\mathcal{P}_i(t)) \equiv S(t) + \mathcal{I}(t) - C(t), \quad (13)$$

where $\mathcal{I}(t)$ is the total income of the households (wages plus dividends) and $C(t)$ is the money actually spent by households, which is in general less than their consumption budget $C_B(t)$. This corresponds to *unvoluntary* households savings whenever production is below demand.

Note that total money $S(t) + \sum_i \mathcal{E}_i(t)$ is clearly conserved since $\Delta S(t) = -\Delta[\sum_i \mathcal{E}_i(t)]$ for cash-flow consistency in a closed economy.

- *Bank accounting*: As mentioned above, we allow the firms' net deposits to become negative, which we interpret as the firm being in need of an immediate extra line of credit. Depending on the financial fragility of the firm (defined below), the bank may or may not agree to restructure the debt and provide this extra credit. If it does, our accounting procedure can be rephrased in the following way. In case of negative net deposits $\mathcal{E}_i(t) < 0$ the bank provides the firm with the extra liquidity needed to pay the wages. The equity of the firm is therefore equal to its deposits when positive, and is close to zero, but still equal to what it needs, when the net deposits is negative. However, when the firm becomes too indebted, the bank will not provide the liquidity needed to pay the wages, leading to negative equity and bankruptcy. From an accounting perspective the matrix balance sheet of our model can be summarized as follows. We assume the bank's equity to be constant in time (let it be 0 for simplicity). The bank holds at the beginning of the simulation a quantity M of currency (issued by the central bank, not modelled here) and therefore has assets equal to M throughout the simulation. Households have deposits $S > 0$ while firms have either deposits (if $\mathcal{E}_i > 0$) or liabilities (if $\mathcal{E}_i < 0$). Defining $\mathcal{E}^+ = \sum_i \max(\mathcal{E}_i, 0)$ and $\mathcal{E}^- = -\sum_i \min(\mathcal{E}_i, 0)$, the balance sheet of the banking system is therefore

$$M + \mathcal{E}^-(t) = S(t) + \mathcal{E}^+(t) \equiv \mathcal{X}(t), \quad (14)$$

which means that the total amount of deposits $\mathcal{X}$ at any time is equal to initial deposits plus deposits created by the banking system when issuing loans. In this setting reserves M at the bank are kept unchanged when loans are issued (as outlined for example in Bank of England, 2014) and deposits increase (decrease) only when loans are issued (repaid). Correspondingly, the fundamental time-invariant macroeconomic accounting identity

$$S(t) + \mathcal{E}^+(t) - \mathcal{E}^-(t) = S(t) + \sum_i \mathcal{E}_i(t) = M \quad (15)$$

is obeyed and amounts to our money conserving equation.

- *Financial fragility and bankruptcy resolution:* We measure the indebtment level of a firm through the ratio of (negative) net deposits over payroll (equal, to a good approximation, to total sales):

$$\Phi_i = -\mathcal{E}_i/(W_iY_i) \quad (16)$$

which we interpret as a measure of financial fragility. (Implicitly, this assumes that the "real assets" of the firms – not modelled here – are proportional to total payroll.) If $\Phi_i(t) < \Theta$, i.e. when the level of debt is not too high compared to the size of the company, the firm is allowed to continue its activity. If on the other hand $\Phi_i(t) \geq \Theta$, the firm defaults.

When firms exceeds the bankruptcy threshold the default resolution we choose is the following. We first define the set $\mathcal{H}_i$ of financially "healthy" firms that are potential buyers for the defaulted firm i. The condition for this is that $\mathcal{E}_j(t) > \max(-\mathcal{E}_i(t), \Theta Y_j(t)W_j(t))$, meaning that the firm j has a strongly positive net deposits and can take on the debt of i without going under water:

- With probability $1 - f, f \in [0, 1]$ being a new parameter, a firm j is chosen at random in $\mathcal{H}_i$; j transfers to i the needed money to pay the debts, hence $\mathcal{E}_j \to \mathcal{E}_j + \mathcal{E}_i$ (remember that $\mathcal{E}_i$ is negative) and $\mathcal{E}_i \to 0$ after the transaction. Furthermore, we set $p_i = p_j$ and $W_i = W_j$, and the firm i keeps its employees and its current level of production.
- With probability f, or whenever $\mathcal{H}_i = \emptyset$, the firm i is not bailed out, goes bankrupt and its production is set to zero. In this case its debt $\mathcal{E}_i(t) < 0$ is transferred to the households' accumulated savings, $S(t) \to S(t) + \mathcal{E}_i(t)$ in order to keep total money fixed.[11]

Hence, when f is large, most of the bankruptcies load weighs on the households, reducing their savings, whereas when f is small, bankruptcies tend to fragilise the firm sector. (The Mark I+ model discussed above corresponds to $f = 1/2$.) It is important to stress that changing the details of the bankruptcy rules while maintaining proper money conservation does not modify the main qualitative message of our paper. The important point here is that the default costs are transferred to households and firms (to ensure money conservation) and have some repercussion on demand (through households accumulated savings) or on firms fragility (through firms net deposits). This can create default avalanches and crises.

- *Firm revival:* A defaulted firm has a finite probability φ per unit time to get revived; when it does so its price is fixed to $p_i(t) = \bar{p}(t)$, its workforce is the available workforce, $Y_i(t) = \mu u(t)$, and its net deposits is the amount needed to pay the

---

[11] One can interpret this by imagining the presence of a bank that collects the deposits of households and lend money to firms, at zero interest rate. If a firm goes bankrupt, the bank looses its loan, which means that its deposits (the households' savings) are reduced.

wage bill, $\mathcal{E}_i(t) = W_i(t)Y_i(t)$. This liquidity is provided by the households, therefore $S(t) \to S(t) - W_i(t)Y_i(t)$ when the firm is revived, again to ensure total money conservation. Note that during this bankrupt/revival phase, the households' savings $S(t)$ might become negative: if this happens, then we set $S(t) = 0$ and the debt of households is spread over the firms with positive liquidity,[12] proportional to their current value of $\mathcal{E}_i$, again in order to ensure total money conservation and $S(t) \geq 0$.

The above description contains all the details of the definition of the model, however for full clarification a pseudo-code of this minimal Mark 0 model is provided in Appendix B (together with the extensions discussed in Section 4). The total number of relevant parameter of Mark 0 is equal to 9: $c, \beta, \gamma_p, \eta_\pm, \delta, \Theta, \varphi, f$ plus the number of firms N_F. However, most of these parameters end up playing very little role in determining the *qualitative*, long-time aggregate behaviour of the model. Only two quantities play an important role, and turn out to be

1. The ratio $R = \eta_+/\eta_-$ between η_+ and η_-, which is meant to capture any asymmetry in the hiring/firing process. As noted above (see Fig. 2), a rising interest rate endogenously leads to such a hiring/firing asymmetry in Mark I and Mark I+. But other sources of asymmetry can also be envisaged: for example, overreaction of the firms to bad news and under-reaction to good news, leading to an over-prudent hiring schedule. Capital inertia can also cause a delay in hiring, whereas firing can be immediate.
2. The default threshold Θ, which controls the ratio between total debt and total circulating currency. In our minimal setting with no banks, it plays the role of a money multiplier. Monetary policy within Mark 0 boils down to setting of the maximum acceptable debt to payroll ratio Θ.

The other parameters change the phase diagram of the model quantitatively but not qualitatively. In decreasing order of importance, the most notable ones are f (the redistribution of debt over households or firms upon bankruptcies) and β (the sensitivity to price) – see below.

3.2. Numerical results and phase diagram

When running numerical simulations of Mark 0, we find (Fig. 3) that after a transient that can be surprisingly long,[13] the unemployment rate settles around a well defined average value, with some fluctuations (except in some cases where endogenous crises appear, see below).

We find the same qualitative phase diagram for all parameters $\beta, \gamma_p, \varphi, \delta, f$. For a given set of parameters, there is a critical value R_c of $R = \eta_+/\eta_-$ separating a full unemployment phase for $R < R_c$ from a phase where some of the labour force is employed for $R > R_c$. Here $R_c \leq 1$ is a value that depends on all other parameters. In Section 5 we explain how the phase transition at R_c and other features of the model can be understood by means of approximate analytical calculations.

The transition is particularly abrupt in the limit $\Theta \to \infty$ (no indebtment limit), in which the unemployment rate jumps all the way from 0 to 1 at R_c, see Fig. 3. The figure shows that indeed only the ratio $R = \eta_+/\eta_-$ is relevant, the actual values of $\eta_\pm$ only change the time scale over which the production fluctuates. Moreover, the phase diagram is almost independent of N_F, which confirms that we are effectively in a limit where the number of firms can be considered to be very large.[14]

Interestingly, for finite values of Θ, the phase diagram is more complex and shown in Fig. 4. Besides the full unemployment phase (region 1), which always prevails when $R < R_c$, we find three other different regions for $R > R_c$, that actually survive many extensions of Mark 0 that we have considered (see Section 4):

- At very large Θ (region 4, "FE"), the full employment phase persists, although a small value of the unemployment appears in a narrow region around $R \approx R_c$. The width of this region of small unemployment vanishes as Θ increases.
- At very low Θ (region 2, "RU"), one finds persistent "Residual Unemployment" in a large region of $R > R_c$. The unemployment rate decreases continuously with R and Θ but reaches values as large as 0.5 close to $R = R_c$ (see Fig. 4, bottom left).
- A very interesting *endogenous crises* phase appears for intermediate values of Θ (region 3, "EC"), where the unemployment rate is most of the time very close to zero, but endogenous crises occur, which manifest themselves as sharp spikes of the unemployment that can reach quite large values. These spikes appear almost periodically, and their frequency and amplitude depend on some of the other parameters of the model, in particular f and β, see Fig. 4, bottom right.

---

[12] Again, this can be interpreted by imagining that the firms with positive liquidity deposit their cash in the bank. When the bank needs to provide a loan to a revived firm, or loses money due to a bankrupt, it prefers to take this money from households' deposits, but if these are not available, then it takes the money from firms' deposits.

[13] Think of one time step as a quarter, which seems reasonable for the frequency of price and workforce updates. The equilibration time is then 20 years or so, or even much longer as for the convergence to the 'bad' state in Mark I, see Fig. 2. Albeit studying a very different ABM, similarly long time scales can be observed in the plots shown in Metzig and Gordon (2012). See also the discussion in the conclusion on this point

[14] Actually, in the specific case $\Theta = \infty$, even $N_F = 1$ provides a similar phase diagram, although the critical R_c slightly depends on N_F when this number is small.
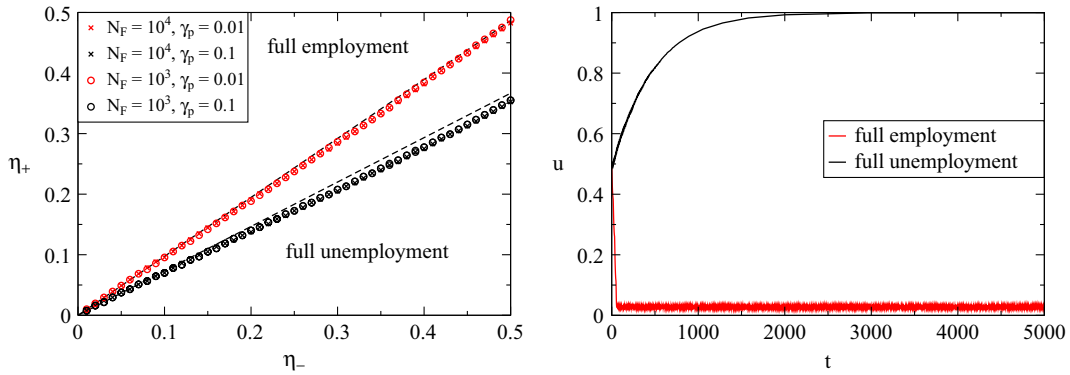
Fig. 3. (Left) Phase diagram of the basic Mark 0 model with $\beta = 2$ and $\Theta = \infty$. There are two distinct phases separated by a critical line which depends on the parameters γ_p and β. The dashed lines correspond to the analytical result in Eq. (29) which agrees well with numerical simulations for small values of η_+, η_- where the perturbative method of Section 3.3 and Appendix C.1 is justified. For simplicity we show here only results with $\beta = 2$ but Eq. (29) is in good agreement with numerical results up to $\beta \sim 4$. (Right) Two typical trajectories of $u(t)$ in the two phases ($R = 5/3$ in the full employment line and $R = 3/5$ in the full unemployment line). The other parameters are $N_F = 10\,000$, $\gamma_p = 0.1$, $\Theta = \infty$, $\varphi = 0.1$, $\beta = 2$, $c = 0.5$. $\delta = 0.02$ and $\varphi = 1$.
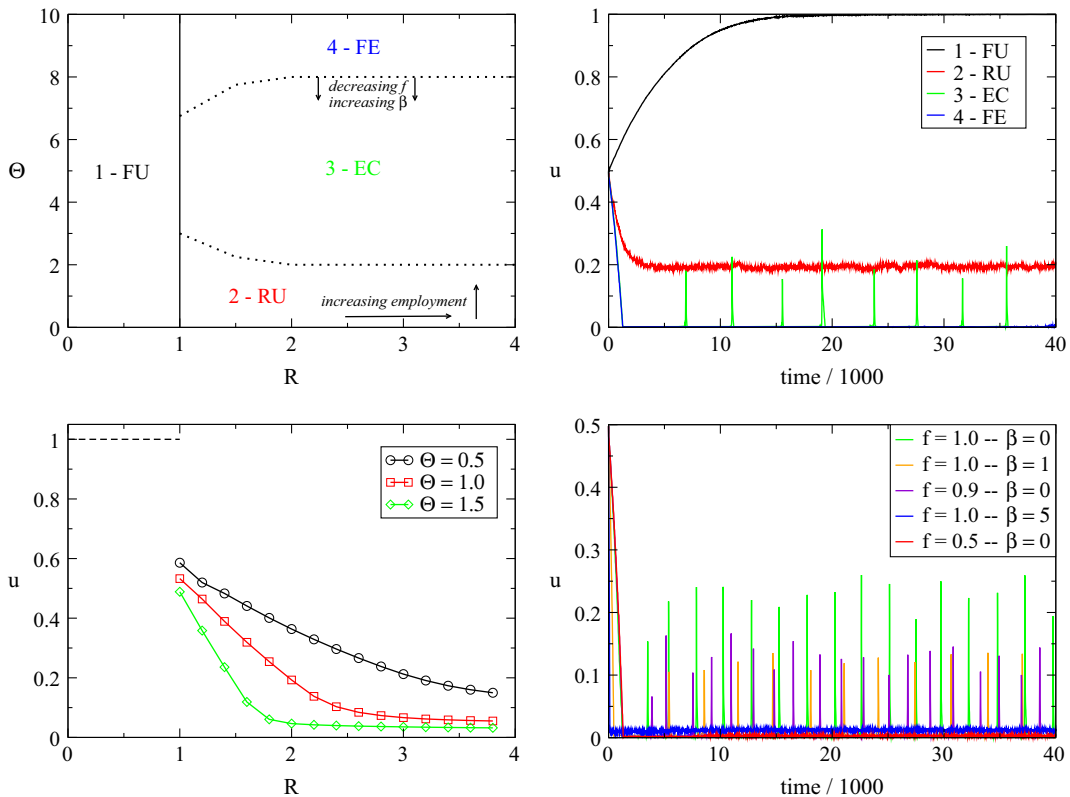

Fig. 4. (Top Left) Phase diagram of the basic Mark 0 model in the R–Θ plane, here with $N_F = 5000$, $c = 0.5$, $\gamma_p = 0.05$, $\delta = 0.02$, $\varphi = 0.1$. Here we keep the 'extreme' cases $\beta = 0$ and $f = 1$ since the dependence of phase boundaries on β and on f (here reported schematically) is shown in detail in Section 3.3.3. The other parameters are irrelevant. There are four distinct phases separated by critical lines. (Top Right) Typical time series of $u(t)$ for each of the phases. (Bottom Left) Stationary value of the unemployment rate as a function of R for different values of Θ in phase 2. (Bottom Right) Typical trajectories of $u(t)$ in the Endogenous Crisis phase (region 3) for different values of f and β, the other parameters are kept fixed. Note that increasing β or decreasing f lead to small crisis amplitudes.

The phase diagram in the plane R–Θ is presented in Fig. 4, together with typical time series of the unemployment rate, for each of the four phases.

We also show in Fig. 5 a trajectory of $u(t)$ in the good phase of the economy (FE, region 4) and zoom in on the small fluctuations of $u(t)$ around its average value. These fluctuations reveal a clear periodic pattern in the low unemployment phase; recall that we had already observed these oscillations within Mark I+. Oscillating patterns (perhaps related to the
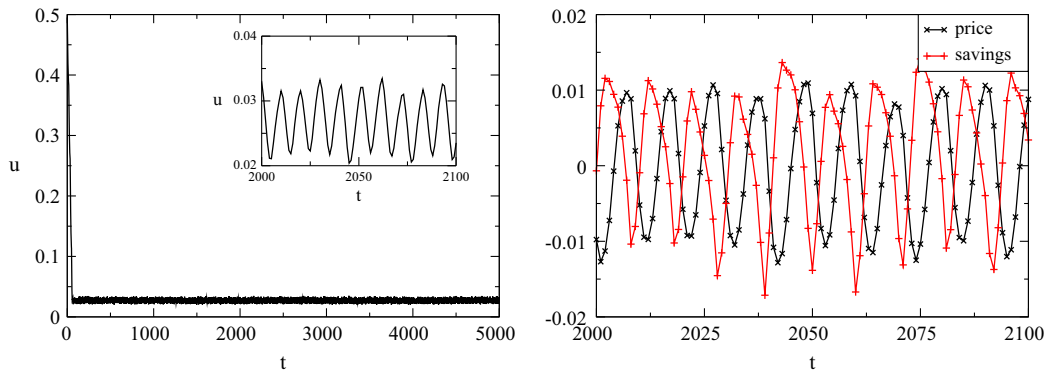
Fig. 5. Left: Typical time evolution of the unemployment u, starting from an initial condition $u=0.5$, for the basic Mark 0 model in the Full Employment phase (region 4). The trajectory leading to a "good" state of the economy is obtained for $\eta_+ = 0.5$ and $\eta_- = 0.3$. Note the clear endogenous "business cycles" that appear in that case. These runs are performed with $N_F = 10\,000$ firms, $c=0.5$, $\beta=2$, $\gamma_p=0.1$, $\varphi=0.1$, $\Theta=5$, $\delta=0.02$. Right: Oscillations for the average price and the average accumulated savings per household (each shifted by their average values for clarity) for the same run as in the Left figure. When prices are low, savings increase, while when prices are high, savings decrease. See Section 5 for a more detailed description.

so-called business cycle) often appear in simplified first order differential models of the macro-economy; one of the best known examples is provided by the Goodwin model (Goodwin, 1967; Flaschel, 2010), which is akin to a predator–prey model where these oscillations are well known. But these oscillations do also show up in other ABMs, see Dosi et al. (2010), Metzig and Gordon (2012). Note that these fluctuations/oscillations around equilibrium *do not regress* when the number of firms gets larger. We have simulated the model with $N_F = 10\,000$ firms and $N_F = 1\,000\,000$ firms with nearly identical amplitude and frequencies for these fluctuations. We will offer some insight on the origin of these oscillations in Section 5. Note that we do not expect such oscillations to remain so regular in the real economy, in particular because exogenous shocks are absent from our model and because we assume that all firms are characterized by the very same adaptation time scale γ_p.

3.3. Qualitative interpretation. Position of the phase boundaries

An important quantity that characterizes the behaviour of the model in the "good" phase of the economy (i.e. for $R > R_c$) is the global leverage (debt-to-equity) ratio k:

$$k(t) = \frac{\mathcal{E}^-(t)}{S(t) + \mathcal{E}^+(t)} = 1 - \frac{M}{S(t) + \mathcal{E}^+(t)} \tag{17}$$

where, due to money conservation, $k \leq 1$. The good state of the economy is characterized by a large average value of k reflecting the natural tendency of the economy towards indebtment, the level of which being controlled in Mark 0 by the parameter Θ (the average value of k increases with Θ).[15] Interestingly, in regions 2-RU and 4-FE $k(t)$ reaches a stationary state, whereas in region 3-EC its dynamics is characterized by an intermittent behaviour corresponding to the appearance of endogenous crises during which indebtment is released through bankruptcies.

3.3.1. The EC phase

This phenomenology can be qualitatively explained by the dynamics of the distribution of firms fragilities Φ_i. For $R > R_c$, firms overemploy and make on average negative profits, which means that the Φ_i's are on average drifting towards the bankruptcy threshold Θ:

- When Θ is small enough (i.e. in region 2-RU) the drift is continuously compensated by the reinitialization of bankrupted firms and the fragility distribution reaches a stationary state.
- For intermediate values of Θ, however, (i.e. in region 3-EC) the number of bankruptcies per unit of time becomes intermittent. Firms' fragilities now collectively drift towards the bankruptcy threshold; as soon as firms with higher fragilities reach Θ, bankruptcies start to occur. Since for f large enough, bankruptcies are mostly financed by households, demand starts falling which has the effect of increasing further the negative drift. This feedback mechanism gives rise to an avalanche of bankruptcies after which most of the firms are reinitialized with positive liquidities. This mechanism has the effect of *synchronizing* the fragilities of the firms, therefore leading to cyclical waves of bankruptcies, corresponding to the unemployment spikes showed in Fig. 4. The distribution of firms' fragilities does not reach a stationary state in this case.

---

[15] In a further version of the model, a central bank will be in charge of controlling k through a proper monetary policy.

- When $\Theta \gg 1$ (i.e. region 4-FE), households are wealthy enough to absorb the bankruptcy cost without pushing the demand of goods below the maximum level of production reached by the economy.[16] Hence, the economy settles down to a full employment phase with a constant (small) rate of bankruptcies.

The above interpretation is supported by a simple one-dimensional random walk model for the firms assets, with a drift that self-consistently depends on the number of firms that fail. This highly simplified model accurately reproduces the above phenomenology, and is amenable to a full analytical solution, which will be published separately.

The existence of the EC phase is a genuinely surprising outcome of the model, which was not put by hand from the outset (see our discussion in the first lines of the Introduction section). Crises occur there not as a result of sweeping parameters through a phase transition (as is the case, we argued, of Mark I with a time dependent interest rate) but purely as a result of the own dynamics of the system.

3.3.2. The role of β

The dependence of the phase boundaries on the different parameters is in general quite intuitive. The dependence of aggregate variables on β, for example, is interesting: everything else being kept equal, we find that increasing β (i.e. increasing the price selectivity of buyers) increases the level of unemployment (see inset of Fig. 6) and the amplitude of the fluctuations around the average value (a similar effect was noted in Dawid et al., 2011). Increasing γ_p increases the dispersion of prices around the average value and is thus similar to increasing β.

Increasing β has also, within the present setting of Mark 0, some counter-intuitive effects: it increases both the average price compared to wages and the profits of firms, hence stabilizing the FE phase and shifting its boundary with EC to lower values of Θ and R_c (see Fig. 6 for the amplitude of region 3-EC as a function of β). This effect can be understood by considering the demand-production gap $G_i = D_i - Y_i$ as a function of the price difference $\delta p_i = p_i - \bar{p}$. For a fixed value of G_i the rules for price and production updates are independent of β; however, the response of the demand to a price change is stronger for higher values of β: for small values of δp_i one finds approximately $G_i \approx -C(\beta)\delta p_i$, where $C(\beta) > 0$ is increasing with β. As a consequence, the absolute value of the gap, $|G_i|$ for a given value of $|\delta p_i|$ is on average increasing with β. Therefore, the total amount of unsold goods $\sum_i \max[Y_i - D_i, 0]$ and households accumulated savings S also increase with β: households involuntarily save more when they are more selective on prices. An increased amount of savings is in turn responsible for the average price increase while the households' higher wealth expands the FE region by shifting its boundary to lower values of Θ. Note that, however, the effect of β on the average price is numerically very small and depends sensitively on the precise consumption rule. For example, if we insist that households fully spend their consumption budget $C_B(t)$ by looking for available products at a higher price (as in Mark I), then the above effect disappears (see Section 5).

3.3.3. The role of f

A potentially more relevant discussion concerns the effect of the bankruptcies on the financial health of the firm sector. Decreasing f (i.e. the financial load taken up by households when bankruptcies occur) also stabilizes, as expected, the full employment phase. Such a stabilization can also be achieved by distributing a fraction δ^+ of the profit plus the total positive liquidity of the firms (instead of a fraction δ of the profits only), which has the obvious effect of supporting the demand. In fact, we find that as soon as $f \leq 0.81$ or $\delta^+ \geq 0.25$, the Endogenous Crisis phase disappears and is replaced by a continuous cross-over between the Residual Unemployment phase and the Full Employment phase (see Fig. 6).

3.4. Intermediate conclusion

The main message of the present section is that in spite of many simplifications, and across a broad range of parameters, the phase transition observed in Mark I as a function of the baseline interest rate is present in Mark 0 as well. We find that these macroeconomic ABMs generically display two very different phases – high demand/low unemployment vs. low demand/high unemployment, with a boundary between the two that is essentially controlled by the asymmetry between the hiring and the firing propensity of the firms (compare Figs. 2 and 3).

Moreover, in the Mark 0 model there is an additional splitting of the low unemployment phase in several regions characterized by a different dynamical behaviour of the unemployment rate, depending on the level of debt firms can accumulate before being forced into bankruptcy. We find in particular an intermediate debt region where endogenous crises appear, characterized by acute unemployment spikes. This Endogenous Crisis phase disappears when households are spared from the financial load of bankruptcies and/or when capital does not accumulate within firms, but is transferred to households through dividends. Clearly, these different phases will coexist if the model parameters themselves evolve with time, which one should expect in a more realistic version of the model.

---

[16] The amount of money circulating in the economy increases with Θ and for $R > R_c$ it is largely channelled to households savings since firms have on average negative liquidities.
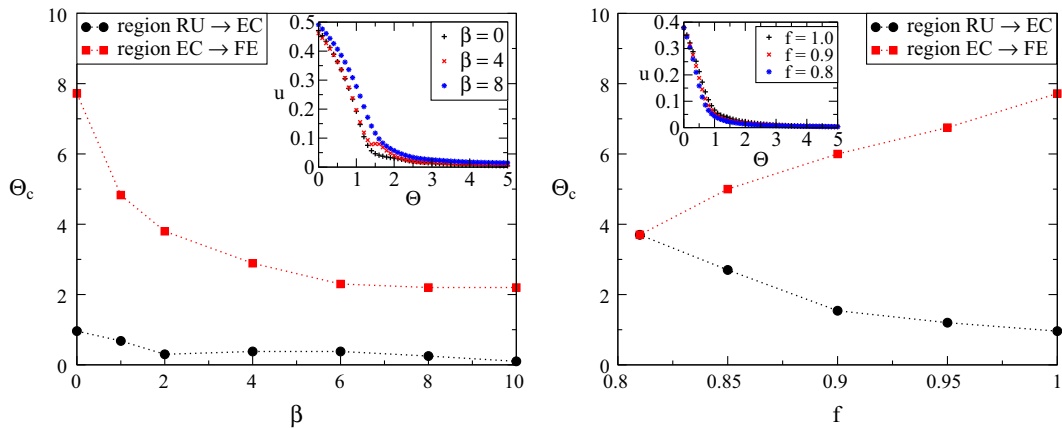
Fig. 6. Phase boundary location Θ_c as a function of the price sensitivity β (left) and of the debt share parameter f (right). Circles correspond to the location of the boundary between regions 2-RU and 3-EC while squares correspond to the boundary between regions 3-EC and 4-FE. In the inset of both figures we plot the average unemployment rate as a function of Θ for different values of β and f. For $f \leq 0.81$ region 4-EC disappears and the average value of the unemployment continuously goes to 0 without the appearance of endogenous crises. Increasing β shrinks the amplitude of region 3-EC which however remains finite for $\beta \geq 1$. As a criterion for being in region 3-EC we require that the amplitude of the crises $\max(u) - \min(u)$ stays above 5%. Parameters are $N = 5000$, $R = 3$ ($\eta_m = 0.1$), $\delta = 0.02$, $\varphi = 0.1$, $\gamma_p = 0.05$.

We now turn to the first extension of Mark 0 allowing for a wage dynamics that lead to long term inflation or deflation, absent in the above version of the model. Finally, several other extensions of the model will be briefly discussed and presented in full details in a separate publication.

4. Extension of Mark 0: wage update

As emphasized above, the Mark I+ and Mark 0 models investigated up to now both reveal a generic phase transition between a "good" and a "bad" state of the economy. However, many features are clearly missing to make these models convincing – setting up a full-blown, realistic macroeconomic Agent-Based Model is of course a long and thorny endeavour which is precisely what we want to avoid at this stage, focusing instead on simple mechanisms. Still, it is interesting to progressively enrich these simplified models not only to test for robustness of our phase diagram but also to investigate new effects that are of economic significance. We consider here wage dynamics, which is obviously an important ingredient in reality. Wage dynamics leads to some relevant effects, in particular the appearance of inflation. Wage dynamics is indeed an item missing from both the basic Mark 0 and Mark I models considered above, which assume fixed wages across time and across firms. Clearly, the ability to modulate the wages is complementary to deciding whether to hire or to fire, and should play a central role in the trajectory of the economy as well as in determining inflation rates.

Introducing wages in Mark 0 again involves a number of arbitrary assumptions and choices. Here, we follow (in spirit) the choices made in Mark I for price and production update, and propose that at each time step firm i updates its wage as

$$
\begin{aligned}
W_i^T(t+1) &= W_i(t)[1 + \gamma_w \varepsilon \xi_i'(t)] && \text{if } \begin{cases} Y_i(t) < D_i(t) \\ \mathcal{P}_i(t) > 0 \end{cases} \\
W_i(t+1) &= W_i(t)[1 - \gamma_w u \xi_i'(t)] && \text{if } \begin{cases} Y_i(t) > D_i(t) \\ \mathcal{P}_i(t) < 0 \end{cases}
\end{aligned}
\tag{18}
$$

where $u = 1 - \varepsilon$ is the unemployment rate and γ_w a certain parameter; $\mathcal{P}_i(t) = \min(D_i(t), Y_i(t))p_i(t) - W_i(t)Y_i(t)$ is the profit of the firm at time t and $\xi_i'(t)$ an independent $U[0, 1]$ random variable. If $W_i^T(t+1)$ is such that the profit of firm i at time t with this amount of wages would have been negative, $W_i(t+1)$ is chosen to be exactly at the equilibrium point where $\mathcal{P}_i = 0$, hence $W_i(t+1) = \min(D_i(t), Y_i(t))p_i(t)/Y_i(t)$; otherwise $W_i(t+1) = W_i^T(t+1)$.

The above rules are intuitive: if a firm makes a profit and it has a large demand for its good, it will increase the pay of its workers. The pay rise is expected to be larger if unemployment is low (i.e. if ε is large) because pressure on salaries is high. Conversely, if the firm makes a loss and has a low demand for its good, it will reduce the wages. This reduction is larger when unemployment is high because pressure on salaries is low. In all other cases, wages are not updated.

When a firm is revived from bankruptcy (with probability φ per unit time), its wage level is set to the production weighted average wage of all firms in activity.

The parameters $\gamma_{p,w}$ allow us to simulate different price/wage update timescales. In the following we set $\gamma_p = 0.05$ and $\gamma_w = z\gamma_p$ with $z \in [0, 1]$. The case $z = 0$ clearly corresponds to removing completely the wage update rule, such that the basic version of Mark 0 is recovered. The extended version of Mark 0 that we consider below is therefore characterized by *single additional parameter* $\gamma_w = z\gamma_p$, describing the frequency of wage updates.
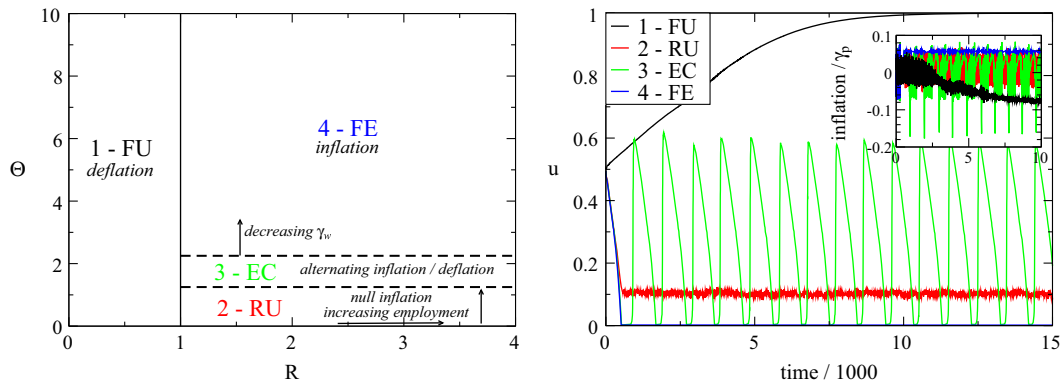
Fig. 7. (Left) Phase diagram in the R–Θ plane of the extended Mark 0 model with wage update with $\gamma_w = \gamma_p$. The parameters of the basic model are the same as in Fig. 4: $N_F = 5000$, $c = 0.5$, $\gamma_p = 0.05$, $\delta = 0.02$, $\varphi = 0.1$, $f = 1$, $\beta = 0$. As for the basic model, there are four distinct phases separated by critical lines. Wage update brings in inflation in the Full Employment phase, and deflation in the full unemployment phase. Endogenous crises are characterized by alternating cycles of inflation and deflation. Interestingly, we find that the location of phase boundaries is in this case almost untouched by changes in both β and f; the other parameters are irrelevant. Note that the wage update strongly stabilizes the full employment phase. (Right) A typical trajectory of $u(t)$ for each of the phases. In the inset, the price dynamics is shown, displaying inflation and deflation.

4.1. Results: variable wages and the appearance of inflation

In our *money conserving* toy economies,[17] a stationary inflation rate different from zero is possible as long as the ratio $\overline{p}(t)/\overline{W}(t)$ fluctuates around a steady value. In the absence of wage update, we have a fixed $W \equiv 1$ and inflation is therefore impossible. The main effect induced by wage dynamics is therefore the possibility of inflation.

Using the wage update rules defined below, we found that the average inflation rate depends on parameters such as the households propensity to consume c and the price/wage adjustment parameters $\gamma_{p,w}$. Most interestingly, we observe a strong dependence of the inflation rate upon the bankruptcy threshold Θ, with large Θ's triggering high inflation and low Θ's corresponding to zero inflation. For intermediate Θ's, periods of inflation and deflation may alternate and the model displays interesting instabilities.

We now analyze the influence of wage adjustments on the phase transitions discussed in the previous sections. The phase diagram for $z = 1$ ($\gamma_w = \gamma_p$) is reported in Fig. 7. The phenomenology that we find is again very similar to the simple Mark 0 without wage update, except for inflation. The most interesting effect is the appearance of inflation in the "good" phases of the economy, and deflation in the "bad" phases, as shown in Fig. 7. When $\Theta \gg 1$, we find again a first order critical boundary at $R = R_c$ that separates a high unemployment phase (with deflation) from a low unemployment phase (with inflation). For $R > R_c$ we see again two additional phases: "EC", with endogenous crises and, correspondingly, alternating periods of inflation and deflation but stable prices on the long run, and "RU", for small Θ, where there is no inflation but a substantial residual unemployment rate (see Fig. 7).

The appearance of endogenous crises is consistent with what discussed in Section 3.2 and is related to situations in which the debt-to-savings ratio $k(t)$ grows faster than prices. From this point of view, increasing γ_w allows firms to better adapt wages (and thus prices) and to absorb the indebtment through inflation; region EC indeed shrinks when increasing γ_w. Note that relating the parameters $\gamma_{p,w}$ to the flexibility of the labour and goods markets is not straightforward. In this sense, it would be instead useful to study the effects of *asymmetric* upward/downward wage and price flexibilities (for example by defining different $\gamma_{p,w}^{\pm}$) in order to understand whether improving the flexibility of the labour and goods markets destabilizes the economy (Eggertsson and Krugman, 2012; Greenwald and Stiglitz, 1993; Napoletano and et al., 2012). If, however, one considers the ratio $z = \gamma_w/\gamma_p$ as an indicator of wages flexibility (relative to prices flexibility), our results suggest that higher labour flexibility has a stabilizing effect.

We also observe that the oscillatory pattern found for the basic model persists as long as $\gamma_w \ll \gamma_p$, i.e. when wage updates are much less frequent that price updates. The power spectrum of the model for $R > R_c$ and $\Theta \gg 1$ is still characterized by the appearance of a peak (roughly corresponding to a period of 7 time steps). Interestingly, we find that upon increasing the ratio z, the peak in the frequency spectrum disappears but not in a monotonous fashion; intermediate values of z give rise to even more pronounced oscillations than for $z = 0$, before these oscillations disappear for $z > z_c \approx 0.25$.

In conclusion, the comparison between Figs. 7 and 4 demonstrates the robustness of our phase diagram against changes; introducing wages is a rather drastic modification since it allows inflation to set in, but still does not affect the phase transition at R_c, nor the overall topology of the phase diagram, which confirms its relevance. Interestingly, inflation is present in the good phases of the economy and deflation in the bad phases. Our analytical understanding of these effects is however still poor; we feel that it would be important to bolster the above numerical results by solving simpler "toy models" as we do for the basic Mark 0 (see Section 5 and Appendix C.1).

---

[17] Recall that the physical money is conserved in the model, but virtual money creation is still possible through firms' indebtment.
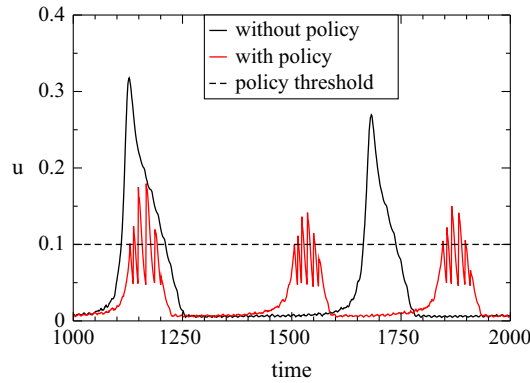
Fig. 8. Here we show an example of a "toy" policy experiment in the Mark 0 model without wage update (i.e $\gamma_w = 0$). We first run a simulation (the "without policy" line) with a constant value $\Theta = 2$ lying in the 3-EC region of Fig. 3. We then run the same simulation with a prototype central bank which increases Θ to 10 as long as u exceeds the threshold of 10%. Note that unemployment is partially contained, but the crisis frequency concomitantly increases. The other parameters values are $N = 10\,000$, $R = 2$, $\gamma_p = 0.05$, $\varphi = 0.1$, $\delta = 0.02$, $f = 1$.

4.2. Other extensions and policy experiments

One of the final goal of our study is to have a prototype framework where one is able to run meaningful policy experiments. In order to do that the model described so far is lacking a number of important ingredients, namely a central bank exogenously setting interest rates level and the amount of money in circulation. This is the project on which we are currently pursuing, with encouraging results (Gualdi et al., 2015a). Still, even within the simplistic framework of Mark 0, one can envisage an interesting prototype policy experiment which consists in allowing the "central bank" to temporarily increase the bankruptcy threshold Θ in times of high unemployment. Fig. 8 shows an example of this: the economy is, in its normal functioning mode, in the EC (region 3) of the phase diagram, with $R = 2$ and (say) $\Theta = 2$. This leads in general to a rather low unemployment rate u but, as repeatedly emphasized above, this is interrupted by acute endogenous crises. The central bank then decides that whenever u exceeds 10%, its monetary policy becomes accommodating, and amounts to raising Θ from its normal value 2 to – say – $\Theta = 10$. As shown in Fig. 8, this allows the bank to partially contain the unemployment bursts. However, quite interestingly, it also increases the crisis frequency, as if it did not allow the economy to fully release the accumulated stress. We expect that this phenomenology will survive in a more realistic framework.

We have also explored other potentially interesting extensions of Mark 0, for example adding trust or confidence, that may appear and disappear on time scales much shorter than the evolution time scale of any "true" economic factor, and can lead to market instabilities and crises (see e.g. Caccioli et al., 2009; Anand et al., 2010; Bouchaud, 2013). There are again many ways to model the potentially destabilizing feedback of confidence. One of the most important channels is the loss of confidence induced by raising unemployment, that increases the saving propensity of households and reduces the demand. The simplest way to encode this in Mark 0 is to let the "c" parameter, that determines the fraction of wages and accumulated savings that is devoted to consumption, be an increasing function of the employment rate $\varepsilon = 1 - u$. We indeed find that the confidence feedback loop can again induce purely endogenous swings of economic activity. Similarly, a strong dependence of c on the recent inflation rate can induce instabilities (Gualdi et al., 2015a).

5. Analytical description

We attempt here to describe analytically some aspects of the dynamics of Mark 0 in its simplest version, namely without bankruptcies ($\Theta = \infty$ for which φ and f become irrelevant), with $\beta \geq 0$, fixed wages $W = 1$, and no dividends ($\delta = 0$). For simplicity, we also fix $\mu = 1$ and $c = 1/2$. The only relevant parameters are therefore β, γ_p and $\eta_{\pm}$. The equations of motion of this very minimal model are

$$\text{If } Y_i(t) < D_i(t) \Rightarrow \begin{cases} Y_i(t+1) = Y_i(t) + \min\{\eta_+ [D_i(t) - Y_i(t)], 1 - \overline{Y}(t)\} \\ \text{If } p_i(t) < \overline{p}(t) \Rightarrow p_i(t+1) = p_i(t)(1 + \gamma_p \xi_i(t)) \\ \text{If } p_i(t) \geq \overline{p}(t) \Rightarrow p_i(t+1) = p_i(t) \end{cases}$$

$$\text{If } Y_i(t) > D_i(t) \Rightarrow \begin{cases} Y_i(t+1) = \max\{Y_i(t) - \eta_- [Y_i(t) - D_i(t)], 0\} \\ \text{If } p_i(t) > \overline{p}(t) \Rightarrow p_i(t+1) = p_i(t)(1 - \gamma_p \xi_i(t)) \\ \text{If } p_i(t) \leq \overline{p}(t) \Rightarrow p_i(t+1) = p_i(t) \end{cases}$$

$$D_i(t) = \frac{c}{p_i(t)} \frac{e^{-\beta(p_i(t) - \overline{p}(t))}}{Z(t)} \left[\max\{s(t), 0\} + \overline{Y}(t) \right], \quad Z(t) := \sum_i e^{-\beta(p_i(t) - \overline{p}(t))}$$

$$\mathcal{E}_i(t+1) = \mathcal{E}_i(t) - Y_i(t) + p_i(t)\min\{Y_i(t), D_i(t)\}$$

$$s(t) = M_0 - \overline{\mathcal{E}}(t). \tag{19}$$

Here M_0 is the total money in circulation, whose precise value is irrelevant for this discussion, and $s = S/N_F$ are the savings per agent. Overlines denote an average over firms, which is flat for $\overline{Y} = N_F^{-1}\sum_i Y_i(t)$ and $\overline{\mathcal{E}} = N_F^{-1}\sum_i \mathcal{E}_i(t)$ while it is weighted by production for $\overline{p}$, see Eq. (8). Note that the basic variables here are $\{p_i, Y_i, \mathcal{E}_i\}$, all the other quantities are deduced from these ones.

In the high employment phase, the model admits a stationary state with $\overline{Y}_{st} \sim 1$. To show this, let us focus on the case where γ_p is very small, but not exactly zero, otherwise of course the price dynamics is frozen. The stationary state is attained after a transient of duration $\sim 1/\gamma_p$, and in the stationary state fluctuations between firms are very small in such a way that $p_i \sim \overline{p}$ and $Y_i \sim \overline{Y}$. A stationary state of Eq. (19) has $Y_i(t+1) = Y_i(t)$ and therefore $Y_i = D_i$, hence $Y_i = D_i = \overline{Y}$. Furthermore, from $\mathcal{E}_i(t+1) = \mathcal{E}_i(t)$ we deduce that $p_i = \overline{p} = 1$. Finally, we have $D_i = \overline{Y} = (s+\overline{Y})/2$, which gives $s = \overline{Y}$ and $\mathcal{E} = M_0 - \overline{Y}$. One obtains therefore a continuum of stationary states, because the production $\overline{Y}$ or equivalently the employment is not determined by requiring stationarity. However, we will show in the following that these equilibria become unstable as soon as fluctuations are taken into account ($\gamma_p > 0$). We will see that fluctuations can induce either an exponentially fast decrease of $\overline{Y}$ towards zero (corresponding to the full unemployment phase), or an exponentially fast growth of $\overline{Y}$, which is therefore only cutoff by the requirement $\overline{Y} \le \mu = 1$, corresponding to full employment. It is therefore natural to choose the stationary state with $\overline{Y} = 1$ as a reference, and study the effect of fluctuations around this state. We therefore consider the stationary state with $Y_i = D_i = p_i = 1$, and $\mathcal{E}_i = M_0 - 1$ in such a way that $s = \overline{Y} = 1$.

For analytical purposes, we will focus below on the limit in which $\eta_{\pm}, \gamma_p \to 0$, in such a way that we can expand around the high employment stationary state and obtain results for the phase transition point. Numerically, the fact that γ_p and/or $\eta_{\pm}$ are small does not change the qualitative behaviour of the model.

In order to consider small fluctuations around the high employment stationary state, we define the following variables:

$$Y_i(t) = 1 - \gamma_p \zeta_i(t),$$
$$p_i(t) = e^{\gamma_p \lambda_i(t)},$$
$$\mathcal{E}_i(t) = M_0 - 1 - \gamma_p \alpha_i(t). \tag{20}$$

Note that in the following, overlines over ζ, α, λ always denote flat averages over firms. Using $p_i(t) = 1 + \gamma_p \lambda_i(t) + (\gamma_p \lambda_i(t))^2/2$ for $\gamma_p \to 0$, one finds that, to order γ_p^2,

$$\frac{e^{-\beta p_i(t)}}{Z(t)} \approx 1 - \beta\gamma_p\left(\lambda_i(t) - \overline{\lambda}(t)\right) - \frac{1}{2}\beta\gamma_p^2\left[(1-\beta)\left(\lambda_i^2(t) - \overline{\lambda^2}(t)\right) + 2\beta\overline{\lambda}(t)\left(\lambda_i(t) - \overline{\lambda}(t)\right)\right]. \tag{21}$$

Since we will find later that $\overline{\lambda}(t)$ is itself of order γ_p, we will drop the part of the last term in the above expression, which is $O(\gamma_p^3)$, i.e.

$$\frac{e^{-\beta p_i(t)}}{Z(t)} \approx 1 - \beta\gamma_p\left(\lambda_i(t) - \overline{\lambda}(t)\right) - \frac{1}{2}\beta(1-\beta)\gamma_p^2\lambda_i^2(t). \tag{22}$$

5.1. Stability of the high employment phase

To study the stability of the high employment phase, we make two further simplifications.

1. We neglect the fluctuations of the savings and fix $\alpha_i \equiv 0$. This is justified if the employment rate varies slowly over the time scale $\tau_c = -1/\ln(1-c)$ that characterizes the dynamics of the savings. But, as we shall show below, the dynamics of employment becomes very slow in the vicinity of the phase transition, hence this assumption seems to be justified.
2. In order to proceed analytically, we also consider the limit $\eta_+, \eta_- \to 0$ with $\eta_+ = R\eta_-$ and R of order one. In fact, we will even assume below that $\eta_{\pm} \ll \gamma_p$. Then, it is not difficult to see that ζ_i is a random variable of order η_+ and is therefore also very small. We define $\zeta_i = \eta_+ z_i$. We believe that these approximations are actually quite accurate, as is confirmed by the comparison of the theoretical transition line with numerical data (see Fig. 3).

With these further simplifications (i.e. setting $\alpha_i = 0$ and assuming that $\gamma_p, \eta_{\pm}$ are small and $z_i = O(1)$), we have, together with Eq. (20),

$$D_i(t) = 1 - \gamma_p\left(\lambda_i(t) + \beta\left(\lambda_i(t) - \overline{\lambda}(t)\right)\right) + \frac{1}{2}\gamma_p^2\lambda_i^2(t)\left(1 + \beta + \beta^2\right) + O\left(\gamma_p\eta_+\right) + O\left(\gamma_p^3\right),$$
$$\overline{p}(t) = 1 + \gamma_p\overline{\lambda}(t) + \frac{1}{2}\gamma_p^2\overline{\lambda^2}(t) + O\left(\gamma_p\eta_+\right) + O\left(\gamma_p^3\right), \tag{23}$$

If, as announced above, we further assume that $\eta_+ \ll \gamma_p$, it is justified to keep the terms of order γ_p^2 while neglecting the terms of order $\gamma_p\eta_+$ and γ_p^3.

We also note that if $|\lambda_i|$ is bounded by a quantity of order 1 (which we will find below), then the condition $\lambda_i(t) + \beta(\lambda_i(t) - \overline{\lambda}(t)) - \frac{1}{2}\gamma_p\lambda_i(t)^2 > 0$ is equivalent to $\lambda_i(t) + \beta(\lambda_i(t) - \overline{\lambda}(t)) > 0$. Then the first two lines of Eq. (19) become

$$
\text{If } \lambda_i(t) < \frac{\beta}{1+\beta}\overline{\lambda}(t) \Rightarrow \begin{cases} z_i(t+1) = z_i(t) - \min\{-\lambda_i(t) - \beta(\lambda_i(t) - \overline{\lambda}(t)) + \frac{1}{2}\gamma_p\hat{\beta}\lambda_i^2, \overline{z}(t)\} \\ \lambda_i(t+1) = \lambda_i(t) + \xi_i(t) - \frac{1}{2}\gamma_p\xi_i(t)^2 \end{cases}
$$

$$
\text{If } \lambda_i(t) > \frac{\beta}{1+\beta}\overline{\lambda}(t) \Rightarrow \begin{cases} z_i(t+1) = z_i(t) + \frac{1}{R}\left[\lambda_i(t) + \beta(\lambda_i(t) - \overline{\lambda}(t)) - \frac{1}{2}\gamma_p\hat{\beta}\lambda_i^2\right] \\ \lambda_i(t+1) = \lambda_i(t) - \xi_i(t) - \frac{1}{2}\gamma_p\xi_i(t)^2 \end{cases} \tag{24}
$$

where $\hat{\beta} = 1 + \beta + \beta^2$. Therefore, in this limit, the evolution of λ decouples from that of ζ (or z). Note that from Eq. (24) one has $-1 - (\beta/(1+\beta))\overline{\lambda}(t) - \gamma_p/2 < \lambda < 1 - (\beta/(1+\beta))\overline{\lambda}(t) - \gamma_p/2$ (which we rewrite in the form $\lambda_{\min} < \lambda < \lambda_{\max}$).

From the simplified evolution equation Eq. (24) we now obtain the evolution of the probability distribution $P_t(\lambda)$ and for $\overline{z}(t)$, which are by definition

$$
P_{t+1}(\lambda) = \int_0^1 d\xi \int_{(\beta/(1+\beta))\overline{\lambda}(t)}^{\lambda_{\max}} d\lambda' P_t(\lambda')\delta\left(\lambda - \lambda' + \xi + \frac{\gamma_p}{2}\xi^2\right) + \int_0^1 d\xi \int_{\lambda_{\min}}^{(\beta/(1+\beta))\overline{\lambda}(t)} d\lambda' P_t(\lambda')\delta\left(\lambda - \lambda' - \xi + \frac{\gamma_p}{2}\xi^2\right), \tag{25}
$$

and

$$
\overline{z}(t+1) = \overline{z}(t) + \frac{1}{R}\int_{(\beta/(1+\beta))\overline{\lambda}(t)}^{\lambda_{\max}} d\lambda\, P_t(\lambda)\left(\lambda + \beta(\lambda - \overline{\lambda}) - \frac{1}{2}\gamma_p\hat{\beta}\lambda^2\right)
$$
$$
- \int_{\lambda_{\min}}^{(\beta/(1+\beta))\overline{\lambda}(t)} d\lambda\, P_t(\lambda)\min\left\{-\lambda - \beta(\lambda - \overline{\lambda}) + \frac{1}{2}\gamma_p\hat{\beta}\lambda^2, \overline{z}(t)\right\}. \tag{26}
$$

Since the dynamics of λ is decoupled from the one of $\overline{z}$, we can assume that $P_t(\lambda)$ reaches a stationary state. In Appendix C.1 we show that, at first order in γ_p, we have for the stationary distribution

$$
P_{\text{st}}(\lambda) = \left(1 - \left|\lambda + \frac{\beta\gamma_p}{4}\right| - \frac{\gamma_p}{2}\text{sgn}(\lambda)\lambda^2\right). \tag{27}
$$

Hence, using the condition $P_{\text{st}}(\lambda) \geq 0$ we find that $|\lambda| \leq 1 + O(\gamma_p)$, consistently with the assumptions we made above. From this result, we also find that $\overline{\lambda} = -\gamma_p(1+\beta)/4$.

Note that the average price here decreases as β increases, confirming the discussion in Section 3.3.2. Indeed, we have neglected the fluctuations of S in the above calculation, while, as explained in Section 3.3.2, it is the increase of involuntary savings (and thus of S) with β that leads to the average price increase. As a further confirmation, one can easily simulate Eqs. (24) with $\alpha_i \neq 0$, and find $d\overline{\lambda}/d\beta > 0$ in this case.

Let us now discuss the dynamics of $\overline{z}(t)$. If at some time t we have $\overline{z}(t) = 0$, then it is clear from Eq. (26) that $\overline{z}(t)$ will grow with time, and it will continue to do so unless the last term in the equation becomes sufficiently large. Recalling that $|\lambda| \leq 1 + O(\gamma_p)$, it is clear that even if $\overline{z}(t)$ becomes very large, the last term in Eq. (26) can be at most given by $\int_{-\infty}^0 d\lambda\, P_{\text{st}}(\lambda)\left(-\lambda - \beta(\lambda - \overline{\lambda}) + \frac{1}{2}\gamma_p\lambda^2\right)$. Therefore, if

$$
\frac{1}{R}\int_{-\gamma_p\beta/4}^{\lambda_{\max}} d\lambda\, P_{\text{st}}(\lambda)\left(\lambda + \beta(\lambda - \overline{\lambda}) - \frac{1}{2}\gamma_p\hat{\beta}\lambda^2\right) > \int_{\lambda_{\min}}^{-\gamma_p\beta/4} d\lambda\, P_{\text{st}}(\lambda)\left(-\lambda - \beta(\lambda - \overline{\lambda}) + \frac{1}{2}\gamma_p\hat{\beta}\lambda^2\right), \tag{28}
$$

then $\overline{z}(t)$ will continue to grow with time, and $\overline{Y}(t) = 1 - \gamma_p\eta_+\overline{z}(t)$ will become very small and the economy will collapse. Conversely, if the condition in Eq. (28) is not satisfied, then Eq. (26) admits a stationary solution with $0 < \overline{z}_{\text{st}} < \infty$. Using Eq. (27), the critical boundary line finally reads[18]:

$$
R_c = \frac{\eta_+}{\eta_-} = \frac{\int_{-\gamma_p\beta/4}^{\lambda_{\max}} d\lambda P_{\text{st}}(\lambda)\left(\lambda + \beta(\lambda - \overline{\lambda}) - \frac{1}{2}\gamma_p\hat{\beta}\lambda^2\right)}{\int_{\lambda_{\min}}^{-\gamma_p\beta/4} d\lambda P_{\text{st}}(\lambda)\left(-\lambda - \beta(\lambda - \overline{\lambda}) + \frac{1}{2}\gamma_p\hat{\beta}\lambda^2\right)} \approx 1 - \gamma_p\frac{(2+\beta)^2}{2(1+\beta)} + \cdots \tag{29}
$$

where we used the condition $\gamma_p < 4/(\beta+2)$ which is always verified for $0 < \beta < +\infty$ and $\gamma_p \to 0$. As one can see in Fig. 3 this result is in good agreement with numerical result for $\beta = 2$. We therefore find, interestingly, that increasing β *decreases* the value of R_c, i.e. stabilizes the high employment phase, as indeed discussed above, see Fig. 4 and the discussion around it.

In other words, the ζ_i perform a *biased* random walk, in the presence of a noise whose average is given by the last two terms in Eq. (26) (of course, when $\overline{\zeta}$ is too small the minimum in the last term is important, because it is there to prevent $\overline{\zeta}$ from becoming negative). The system will evolve towards full employment, with $\overline{\zeta} = O(\eta_+)$ and $\overline{Y} = 1 - O(\gamma_p\eta_+)$, whenever the average noise is negative, and to full collapse, with $\overline{\zeta} \to \infty$ and $\overline{Y} = 0$, otherwise. The critical line is given by the equality condition, such that the average noise vanishes. Therefore, right at the critical point, the unemployment rate makes an

---

[18] It would be interesting to compute the first non trivial corrections in η to the transition line. We leave this for a later study.

unbiased random walk in time, meaning that its temporal fluctuations are large and slow. This justifies the "adiabatic" approximation[19] made above, that leads us to neglect the dynamics of the savings.

5.2. Oscillations in the high employment phase

As discussed above, in the FE phase macroeconomic variables display an oscillatory dynamics, see Fig. 5. Intuitively, the mechanism behind these oscillations is the following. When prices are low, demand is higher than production and firms increase the prices. But at the same time, households cannot consume what they demand, so they involuntarily save: savings increase when prices are low. These savings keep the demand high for a few rounds even while prices are increasing, therefore prices keep increasing above their equilibrium value. When prices are too high, households need to use their savings to consume, and therefore savings start to fall. Increase of prices and decrease of savings determine a contraction of the demand. At some point demand falls below production and prices start to decrease again, with savings decreasing at the same time. When prices are low enough, demand becomes again higher than production and the cycle is restarted. An example is shown in Fig. 5.

Based on this argument, it is clear that to study these oscillations, we need to take into account the dynamics of the savings so we cannot assume $\alpha_i = 0$ as in the previous section. However, here it is enough to consider the first order terms in γ_p. In terms of the basic variables in Eq. (20), the other variables that appear in Eq. (19) are easily written as follows:

$$\overline{Y}(t) = 1 - \gamma_p \overline{\zeta}(t),$$
$$s(t) = 1 + \gamma_p \overline{\alpha}(t),$$
$$\overline{p}(t) = 1 + \gamma_p \overline{\lambda}(t), \tag{30}$$

where for ζ, α, λ, overlines denote flat averages over firms, and

$$D_i(t) = 1 + \gamma_p \left(\tfrac{1}{2} \overline{\alpha}(t) - \tfrac{1}{2} \overline{\zeta}(t) - \lambda_i(t) \right). \tag{31}$$

Inserting this in Eq. (19), we arrive at the following equations that hold at the lowest order in γ_p:

$$\text{If } -\zeta_i(t) < \tfrac{1}{2}\overline{\alpha}(t) - \tfrac{1}{2}\overline{\zeta}(t) - \lambda_i(t) \Rightarrow \begin{cases} \zeta_i(t+1) = \zeta_i(t) - \min\{\eta_+ \left[\tfrac{1}{2}\overline{\alpha}(t) - \tfrac{1}{2}\overline{\zeta}(t) - \lambda_i(t) + \zeta_i(t) \right], \overline{\zeta}(t)\} \\ \text{If } \lambda_i(t) < \overline{\lambda}(t) \Rightarrow \lambda_i(t+1) = \lambda_i(t) + \xi_i(t) \\ \alpha_i(t+1) = \alpha_i(t) - \lambda_i(t) \end{cases}$$

$$\text{If } -\zeta_i(t) > \tfrac{1}{2}\overline{\alpha}(t) - \tfrac{1}{2}\overline{\zeta}(t) - \lambda_i(t) \Rightarrow \begin{cases} \zeta_i(t+1) = \zeta_i(t) - \eta_- \left[\tfrac{1}{2}\overline{\alpha}(t) - \tfrac{1}{2}\overline{\zeta}(t) - \lambda_i(t) + \zeta_i(t) \right] \\ \text{If } \lambda_i(t) > \overline{\lambda}(t) \Rightarrow \lambda_i(t+1) = \lambda_i(t) - \xi_i(t) \\ \alpha_i(t+1) = \alpha_i(t) - \zeta_i(t) - \tfrac{1}{2}\overline{\alpha}(t) + \tfrac{1}{2}\overline{\zeta}(t) \end{cases} \tag{32}$$

From an analytical point of view, the above model is still complex to make progress. We make therefore a further simplification, by assuming that

$$\zeta_i = C \lambda_i, \tag{33}$$

where C is a certain numerical constant. In terms of the original Mark 0 variables, this approximation is equivalent to, roughly speaking,

$$p_i - \overline{p} \propto (Y_i - D_i) - (\overline{Y} - \overline{D}). \tag{34}$$

i.e. fluctuations of the prices are proportional to supply–demand gaps. Although numerical simulations only show a weak correlation, the approximation (33) allows us to obtain a more tractable model that retains the basic phenomenology of the oscillatory cycles and reads:

$$\text{If } 2(1-C)\lambda_i(t) < \overline{\alpha}(t) - C\overline{\lambda}(t) \Rightarrow \begin{cases} \text{If } \lambda_i(t) < \overline{\lambda}(t) \Rightarrow \lambda_i(t+1) = \lambda_i(t) + \xi_i(t) \\ \alpha_i(t+1) = \alpha_i(t) - \lambda_i(t) \end{cases}$$

$$\text{If } 2(1-C)\lambda_i(t) > \overline{\alpha}(t) - C\overline{\lambda}(t) \Rightarrow \begin{cases} \text{If } \lambda_i(t) > \overline{\lambda}(t) \Rightarrow \lambda_i(t+1) = \lambda_i(t) - \xi_i(t) \\ \alpha_i(t+1) = \alpha_i(t) - C\lambda_i(t) - \tfrac{1}{2}\overline{\alpha}(t) + \tfrac{1}{2}C\overline{\lambda}(t) \end{cases} \tag{35}$$

This very minimal model, when simulated numerically, indeed gives persistent oscillations, independent of N, when $C > C^* \approx 0.45$, and can also be partially investigated analytically, see Appendix C.1.

---

[19] This refers, in physics, to a situation where a system is driven by an infinitely slow process. One can then consider the system to be always close to equilibrium during the process.

5.3. The "representative firm" approximation

To conclude this section, we observe that there is a further simplification that allows one to retain some of the phenomenology of Mark 0. It consists in describing the firm sector by a unique "representative firm", $N_F = 1$, with production $\overline{Y}(t)$, price $p(t)$ and demand $\overline{Y}(t)/p(t)$. The dynamics of the production and the price are given by the same rule as above, but now the dynamics of the price completely decouples:

$$
p(t) < 1 \Rightarrow \begin{cases} \overline{Y}(t+1) = \overline{Y}(t)\left(1+\eta_+\left(\frac{1}{p(t)}-1\right)\right) \\ p(t+1) = p(t)(1+\gamma\xi(t)) \end{cases}
$$

$$
p(t) > 1 \Rightarrow \begin{cases} \overline{Y}(t+1) = \overline{Y}(t)\left(1+\eta_-\left(\frac{1}{p(t)}-1\right)\right) \\ p(t+1) = p(t)(1-\gamma\xi(t)). \end{cases} \tag{36}
$$

Of course, this simple model misses several important effects: most notably those associated with Θ, hence the transition between the Full Employment, Endogenous Crises, and Residual Unemployment phases in Fig. 4. In particular, endogenous crises are never present in this case, because of the absence of a bankrupt/revival mechanism, and also the oscillatory pattern in the Full Employment phase disappears, because in this model savings are not considered. Still, this model is able to capture the transition between the Full Unemployment and Full Employment regions as a function of R (see Fig. 4), as confirmed by the analytical solution, which is in fact identical to one of the models with $N_F > 1$ when $\beta = 0$ and η is small. And since the model is so simple, one can hope that some of the extensions discussed in the previous section can be at least partly understood analytically within this "representative firm" framework (we will give a few explicit). This would be an important step to put the rich phenomenology that we observe on a firmer basis.

6. Summary, conclusion

The aim of our work (which is part of the CRISIS project and still ongoing) was to explore the possible types of *phenomena* that simple macroeconomic Agent-Based Models can reproduce, and to map out the corresponding *phase diagram* of these models, as Figs. 7 and 4 exemplify. The precise motivation for our study was to understand in detail the nature of the macro-economic fluctuations observed in the "Mark I" model devised by Delli Gatti and collaborators (Gaffeo et al., 2008; Delli Gatti et al., 2008, 2011). One of our central findings is the generic existence, in Mark I (and variations around that model) of a first order, discontinuous phase transition between a "good economy" where unemployment is low, and a "bad economy" where unemployment is high. By studying a simpler hybrid model (Mark 0), where the household sector is described by aggregate variables and not at the level of agents,[20] we have argued that this transition is induced by an *asymmetry between the rate of hiring and the rate of firing* of the firms. This asymmetry can have many causes at the micro-level, for example different hiring and firing costs. In Mark I, for example, it reflects the reluctance of firms to take loans when the interest rate is too high. As the interest rate increases, the unemployment level remains small until a tipping point beyond which the economy suddenly collapses. If the parameters are such that the system is close to this transition, any small fluctuation (for example in the level of interest rates) is amplified as the system jumps between the two equilibria. It is actually possible that the central bank policy (absent in our current model), when attempting to stabilize the economy, in fact brings the system close to this transition. Indeed, too low interest rate leads to overheating and inflation, and too high interest rate leads to large unemployment. The task of the central bank is therefore to control the system in the vicinity of an instability and could therefore be a natural realization of the enticing 'self-organized criticality' scenario recently proposed in Patzelt and Pawelzik (2011) (see also Bak, 1996).

Mark 0 is simple enough to be partly amenable to analytic treatments, that allow us to compute approximately the location of the transition line as a function of the hiring/firing propensity of firms, and characterize the oscillations and the crises that are observed. Mark 0 can furthermore be extended in several natural directions. One is to allow this hiring/firing propensity to depend on the financial fragility of firms – hiring more when firms are financially healthy and firing more when they are close to bankruptcy. We find that in this case, the above transition survives but becomes *second order*. As the transition is approached, unemployment fluctuations become larger and larger, and the corresponding correlation time becomes infinite, leading to very low frequency fluctuations. There again, we are able to give some analytical arguments to locate the transition line (Gualdi et al., 2015a). Other stabilizing mechanisms, such as the bankruptcy of indebted firms and their replacement by healthy firms (financed by the accumulated savings of households), lead to a similar phenomenology.

The role of the bankruptcy threshold Θ, which is the only "proto-monetary" effect in Mark 0, turns out to be crucial in the model, and leads to the phase diagram shown in Fig. 4. We generically find not one but three "good" phases of the economy: one is where full employment prevails (FE), the second one is where some residual unemployment is present (RU), and the third, intermediate one is prone to *acute endogenous crises* (EC), during which the unemployment rate shoots up before the economy recovers. Note again that there are no exogenous productivity shocks in Mark 0; interestingly, crises (when they occur) are of endogeneous origin and not as a result of sweeping parameters through a phase transition (as is

---

[20] For a recent study exploring the idea of hybrid models, see Assenza and Delli Gatti (2013).

the case, we argued, of Mark I with a time dependent interest rate) but purely as a result of the own dynamics of the system. The existence of endogenous crises driven by feedback loops in such simple settings is quite interesting from a general standpoint (see also Dosi et al., 2010; Dawid et al., 2011), and reinforces the idea that many economic and financial crises may *not* require large exogenous shocks (see Minsky, 2008; Challet et al., 2005; Delli Gatti et al., 2011; Sornette, 2005; Anand et al., 2011; Lorenz et al., 2009; Bouchaud, 2011, 2013; Bonart et al. (2014) for related discussions). We have shown that the endogenous crises can be defanged (in the model) if the household sector does not carry the full burden of firms bankruptcies and/or when the profits of firms are efficiently re-injected in the economy (see e.g. Fig. 6).

We have then allowed firms to vary the wages of their employees according to some plausible rules of thumb (wages in Mark I and Mark 0 are fixed); this leads to inflation or deflation but leaves the above picture essentially unchanged (see Fig. 7). Several other extensions are of obvious interest, and we plan to study them in the near future in the same stylized way as above. The most obvious one is to understand how a central bank that prints money and sets exogenously the interest rate can control the unemployment rate and the inflation rate in the vicinity of an unstable point, as we mentioned just above (Gualdi et al., 2015a). Other interesting topics are modeling research and innovation, allowing firms to produce different types of goods, and introducing a financial sector and a housing market (see e.g. Geanakoplos et al., 2012).

Beyond the generic phase diagram discussed in the whole paper, we found another notable, robust feature: the low unemployment phase of all the ABMs we considered is characterized by endogenous oscillations that do not vanish as the system size becomes large, with a period corresponding, in real time, to ~ 5–10 years (Dosi et al., 2010). It is tempting to interpret these oscillations as real and corresponding to the "business cycle", as they arise from a very plausible loop between prices, demand and savings. These oscillations actually also appear in highly simplified models, where both the household and the firm sectors are represented by aggregate variables (Goodwin, 1967; Gualdi et al., 2015b), as well as in network models (Bonart et al., 2014).

Building upon this last remark, a very important question, it seems to us, is how much can be understood of the phenomenology of ABMs using "mean-field" approaches (DiGuilmi et al., 2008; Aoki and Yoshikawa, 2011), i.e. dynamical equations for aggregate variables of the type considered, for example, in Goodwin (1967) and Flaschel (2010)? A preliminary analysis reveals that the dynamical equations corresponding to Mark 0 or Mark I already lead to an amazingly complex phase-diagram (Gualdi et al., 2015b). Are these mean-field descriptions quantitatively accurate? When do we really need agents and when is an aggregate description sufficient? The answer to this question is quite important, since it would allow one to devise faithful "hybrid" ABMs, where whole sectors of the economy would be effectively described in terms of these aggregate variables, only keeping agents where they are most needed.

Another nagging question concerns the calibration of macroeconomic ABMs. It seems to us that before attempting any kind of quantitative calibration, exploring and making a catalogue of the different possible qualitative "phases" of the model are mandatory. Is the model qualitatively plausible or is the dynamics clearly unrealistic? In what "phase" is the true economy likely to be? On this point, one of the surprise of the present study is the appearance of very long time scales. For example, even in the case of perfectly stable economy with wage update rule (18) and all γ parameters equal to 10% (a rather large value), the equilibrium state of the economy (starting from an arbitrary initial condition) is only reached after ≈ 200 time steps. If one thinks that the elementary time scale in these models is of the order of three months, this means that the physical equilibration time of the economy is 20–50 years, or even much longer, see e.g. Figs. 4 and 7. But there is no reason to believe that on these long periods all the micro-rules and their associated parameters are stable in time. Therefore, studying the *stationary state* of macroeconomic ABMs might be completely irrelevant to understand the real economy. The economy could be in a perpetual transient state (aka "non-ergodic"), unless one is able to endogenize the time evolution of all the relevant parameters governing its evolution (see the conclusion of Fagiolo and Roventini, 2012 for a related discussion).

If this is the case, is there any use in studying ABMs at all? We strongly believe that ABMs would still be genuinely helpful. ABMs allow us to experiment and scan the universe of possible outcomes – not missing important scenarios is already very good macroeconomics. Human imagination turns out to be very limited, and that is the reason we like models and equations, that help us guessing what *can* happen, especially in the presence of collective effects that are often very counterintuitive. In this respect, ABMs provide extremely valuable tools for generating scenarios, that can be used to test the effect of policy decisions (see e.g. the pleas by Buchanan, 2008, and Farmer and Foley, 2009). In order to become more quantitative, we think that ABMs will have to be calibrated at the micro-level, using detailed behavioural experiments and field studies to fine tune all the micro-rules needed to build the economy from bottom up (see e.g. for a review on behavioural experiments Hommes, 2013; Geanakoplos et al., 2012 for work in this direction.) Calibrating on historical data the emergent macro-dynamics of ABMs will most probably fail, because of the dimensionality curse and of the Lucas critique (i.e. the feedback between the trajectory of the economy and policy decisions that dynamically change the parameters of the model).

Appendix A. Pseudo-code for Mark I+

We describe here the pseudo-code of our version of Mark I, which we call Mark I+. To keep the length reasonable, a few irrelevant details will be omitted, but the information given here is enough to reproduce the results presented in the paper. In particular, we describe here only the part of the code that is needed to generate the dynamical evolution of the model, and we omit the part that is needed to generate the data output. The source code is available on the site of the CRISIS project (www.crisis-economics.eu).

A.1. Notations

We describe the algorithm in an object-oriented fashion, where the different agents are described as representatives of a few classes. We use an object oriented syntax. This syntax should be very intuitive and easy to follow. However, it is useful to clarify a few notational conventions:

- The declaration of a variable a (for example of integer type `int`) will be written in C syntax as `int`a.
- If a is an object of some class, then $a \cdot f(x)$ means that we are calling the method (function) $f(x)$ of object a with argument x.
- For simplicity, in the **for** loops we will use the C syntax where, for example, **for** $(t \leftarrow 1; t \leq T; t \leftarrow t+1)$ means that t is set to one before the loop starts, t is increased by 1 at the end of each iteration, and the loop continues if the condition $t \leq T$ is true.
- Instead of arrays we will use *vectors* of objects, and we will follow the notation of C++.[21] For example, `vector < int >` will denote an ordered set (array) of integers. Moreover,

  - A vector of size N will be declared as `vector < int >` $A(N)$, and by default the declaration `vector < int >` A means that the set A is initiated as empty.
  - $|A|$ will denote the size of the set A.
  - The notation $A[i]$ will denote the i-th element of the set (with $i = 0, \ldots, |A| - 1$).
  - The notation $A \leftarrow a$ will denote the operation of adding the element a to the set A, therefore increasing its size by one (this correspond to the "push_back" operation in C++).
  - The notation $a \leftarrow_R A$ will denote the extraction of a random element from A, which is set equal to a. Note that the element is not removed from A so that the size of A remains constant. The notation $A' \leftarrow_{R,M} A$ denotes extraction at random of M different elements from the set A, that constitute the set A'.
  - The notation $A \nrightarrow a$ will denote the removal of element a from set A, therefore reducing the size of A by one.
  - To denote an ordered iteration over an ordered set A, we will use a loop **for** $(a \in A)$.
  - We define a function average $(a \cdot f(), a \in A)$ that returns the average of the $|A|$ numbers $a[i] \cdot f(), i = 0, \ldots, |A| - 1$
- For logical operators, we use the C++ convention[22] which should be quite transparent.

A.2. Classes

The classes are the following:

- *firm;*
- *household;*
- *bank.*

The role of the bank is very limited at this level and this class is mainly included for the purpose of future extension. One household has a special role as it is the "owner" of the firms, which will pay dividends to this one household. The main loop

---

is described in Algorithm 1 below. The implementation of the firm class is in Algorithm 3, the household class in Algorithm 6 and the bank class in Algorithm 7.

Algorithm 1. Main loop of Mark I+.

Require: N_F Number of firms; N_H Number of households; ρ_0 baseline interest rate; T total evolution time;
    ▹ **Initialization**
`vector <household> H`(N_H)
`household O`   ▹ The owner of all firms
`vector <firm> F`(N_F)
`bank B`

$\overline{p} \leftarrow 1$ ▹ Average price
for $f \in F$ **do**
  f.set_owner(O)
end for
B.set_$\rho(\rho_0)$
    ▹ **Main loop**
for $(t \leftarrow 1; t \le T; t \leftarrow t+1)$ **do**

  `vector <firm> E,D`
  for $f \in F$ **do**   ▹ **Firm decide new strategy on prices and production**
    f.set_new_strategy$(\overline{p})$
    f.get_loans(B)
    f.compute_interests()
    f.define_labor_demand()
    if f.n_vacancies() > 0 **then** $D \leftarrow f$ ▹ Firms in D demand workforce
    else if f.n_vacancies() < 0 **then** $E \leftarrow f$ ▹ Firms in E have an excess of workforce
    end if
  end for

  for $f \in E$ **do**   ▹ **Job market and production**
    while f.n_vacancies() < 0 **do** f.fire_random_worker()   ▹ Firms with excess workforce fire random workers
    end while
  end for
  `vector <household> U`
  for $h \in H$ **do**
    if !h.working() **then** $U \leftarrow h$ ▹ U is the set of unemployed households
    end if
  end for
  while $|U| > 0$ && $|D| > 0$ **do**   ▹ Random match of unemployed households and demanding firms
    $h \leftarrow_R U$
    $f \leftarrow_R D$
    f.hire(h)
    $U \nrightarrow h$
    if f.n_vacancies()$==0$ **then** $D \nrightarrow f$
    end if
  end while

Algorithm 2. Main loop of Mark I+

  for $f \in F$ **do**   ▹ Firms produce and pay workers
    f.produce()
    f.pay_workers()
    if f.age()< 100 **then** f.markup_rule()   ▹ Young firms apply a markup rule to avoid bankrupt
    end if
  end for
    ▹ **Goods market**
  $H \leftarrow$ random_permutation(H)
  for $h \in H$ **do** h.consume(F) ▹ Consume in random order
end for
    ▹ **Accounting and bankrupts**
  bad_debts $\leftarrow 0$
  `vector <firm> L`
  for $f \in F$ **di**
    f.accounting(B)
    if f.liquidity() < 0 **then**   ▹ Firms with negative liquidity go bankrupt
      bad_debts $\leftarrow$ bad_debts $+ f$.liquidity()   ▹ Note: bad_debts is negative!
      $\overline{L} \leftarrow f$ ▹ $\overline{L}$ is the sent of bankrupt firms
    else $L \leftarrow f$ ▹ L is the set of healthy firms
    end if
  end for
  if $|L| == 0$ **then** break   ▹ If all firms are bankrupt, exit the program

end if

$\overline{p}_b \leftarrow$ average($f[i]$.price(), $f \in L$)

$\overline{Y}^T \leftarrow$ average($f[i]$.target_production(), $f \in L$)

$\overline{Y} \leftarrow$ average ($f[i]$.production(), $f \in L$)

for $f \in \overline{L}$ **do** f.reinit($\overline{p}_b, \overline{Y}^T, \overline{Y}$)   ▷ Bankrupt firms are reinitialized with the average parameters of healthy firms

end for

total_liquidity $\leftarrow \sum_{i=0}^{N_F} f[i]$.equity() $+ \sum_{i=0}^{N_H} h[i]$.wealth()

for $f \in F$ **do** f.get_money(bad_debts * $f[i]$.equity() / total_liquidity)

end for

for $h \in H$ **do** h.get_money(bad_debts * $h[i]$.wealth() /total_liquidity)

end for   ▷ Bad debt is spread over firms and households proportionally to their wealth

$\overline{p} \leftarrow \frac{\sum_{i=0}^{|F|-1} f[i].\text{price}()f[i].\text{sales}()}{\sum_{i=0}^{|F|-1} f[i].\text{sales}()}$ ▷ Update average price

end for

Algorithm 3. The class `firm`.

Parameters: $W=1$, $\alpha=1$, $\gamma_p = 0.1$, $\gamma_y = 0.1$, $\mu=0$, $\delta=0.2$, $\tau=0.05$

Dynamic variables: `vector <household>` E; `household` O; $p, Y, Y^T, D, \mathcal{L}, v, \mathcal{D}^T, t$

Dynamic variables: (**auxiliary**): $L_d, \rho, \mathcal{I}$

                    ▷ **Initialization methods**

function INIT

    $E \leftarrow$ **empty**   ▷ The set E is the list of employees and is initialized as empty

    $p \leftarrow 1$ ▷ Price

    $Y \leftarrow 1$ ▷ Production

    $Y^T \leftarrow 1$ ▷ Target production

    $D \leftarrow 1$ ▷ Demand

    $\mathcal{L} \leftarrow 50$ ▷ Liquidity

    $v \leftarrow 0$ ▷ Number of vacancies

    $\mathcal{D}^T \leftarrow 0$ ▷ Total debt

    $t \leftarrow 0$ ▷ Internal clock

end function

function SET_OWNER(`household` $\tilde{O}$)

    $O \leftarrow \tilde{O}$

end function

function REINIT ($\tilde{p}, \tilde{Y}^T, \tilde{Y}$)

    $p \leftarrow \tilde{p}$

    $Y \leftarrow \tilde{Y}$

    $Y^T \leftarrow \tilde{Y}^T$

    $D \leftarrow 0$

    $\mathcal{L} \leftarrow \min\{O.\text{wealth}(),\ Y/\alpha\}$ ▷ The owner injects money to restart the bankrupt firm

    O.get_money($-\mathcal{L}$)

    $v \leftarrow 0$

    $\mathcal{D}^T \leftarrow 0$

    $t \leftarrow 0$

    $v \leftarrow 0$

    for $h \in E$ **do**

        fire(h)

    end for

end function   ▷ **Output methods**

function PRICE

    return p

end function

function PRODUCTION

    return Y

end function

function STOCK

    return $Y - D$

end function

function SALES

    return D

end function

function TARGET_PRODUCTION

    return Y^T

end function

function EQUITY
   return $\mathcal{L} - \mathcal{D}^T$
end function

function LIQUIDITY
   return $\mathcal{L}$
end function

function N_VACANCIES
   return ν
end function

Algorithm 4. The class `firm`.

function AGE
   return t
end function
       ▷ **Accounting methods**

function GET_MONEY ($\tilde{m}$)
   $\mathcal{L} \leftarrow \mathcal{L} + \tilde{m}$
end function

function GET_LOANS(bank $\tilde{B}$)
   $\mathcal{L}_n \leftarrow \mathcal{W}\mathcal{L}_d - \mathcal{L}$ ▷ Financial need
   if $\mathcal{L}_n > 0$ **then**
     $\ell \leftarrow (\mathcal{D}^T + \mathcal{L}_n)/(\mathcal{L} + 0.001)$ ▷ This is the leverage
     $\rho_{\text{offer}} \leftarrow \tilde{B}.\text{compute\_offer\_rate}(\ell)$ ▷ New offered interest rate
     $\mathcal{D}^c \leftarrow \mathcal{L}_n \mathcal{F}(\rho_{\text{offer}})$ ▷ The function $F(\rho)$ can be whatever decreasing function of ρ, see Eq. (4)
     if $\mathcal{D}^c > 0$ **then**
       $\rho \leftarrow \rho_{\text{offer}}$ ▷ If new credit is contracted, the interest rate is updated
       $\mathcal{D}^T \leftarrow \mathcal{D}^T + \mathcal{D}^c$ ▷ Total debt is increased by current debt $\mathcal{D}^c$
       $\mathcal{L} \leftarrow \mathcal{L} + \mathcal{D}^c$
       $\tilde{B}.\text{get\_money}(-\mathcal{D}^c)$
     end if
   end if
end function

function COMPUTE_INTERESTS
   $\mathcal{I} \leftarrow \rho \mathcal{D}^T$ ▷ Interests to be paid in this round
end function

function PAY_WORKERS
   for $h \in E$ **do**
     $h.\text{get\_money}(W)$
   end for
   $\mathcal{L} \leftarrow \mathcal{L} - W|E|$
end function

function MARKUP_RULE
   if $Y > 0$ **then**
     $p_{\text{markup}} \leftarrow (1+\mu)(W|E| + \mathcal{I})/Y$
     $p \leftarrow \max\{p, p_{\text{markup}}\}$
   end if
end function

function ACCOUNTING (bank $\tilde{B}$)
   $\mathcal{L} \leftarrow \mathcal{L} - \mathcal{I} - \tau \mathcal{D}^T$ ▷ Firm pays interests and repays a fraction τ of its debt
   $\tilde{B}.\text{get\_money}(\mathcal{I} + \tau \mathcal{D}^T)$
   $\mathcal{D}^T \leftarrow (1-\tau)\mathcal{D}^T$
   $\mathcal{P} \leftarrow pD - W|E| - \mathcal{I}$ ▷ Profit
   if $\mathcal{P} > 0$ **then**
     $O.\text{get\_money}(\delta\mathcal{P})$ ▷ Firm pays dividends to the owner
     $\mathcal{L} \leftarrow \mathcal{L} - \delta\mathcal{P}$
   end if
end function

function SELL($\tilde{q}$)
   $D \leftarrow D + \tilde{q}$
   $\mathcal{L} \leftarrow \mathcal{L} + p\tilde{q}$
end function

Algorithm 5. The class `firm`.

        ▹ **Production and job market methods**

function SET_NEW_STRATEGY $(\tilde{p})$

  $t \leftarrow t+1$

  if $Y = D$ && $p \geq \tilde{p}$ **then** $Y^T \leftarrow Y(1+\gamma_y \text{ random})$ ▹ This is Eq. (1) in the main text

  else if $Y = D$ && $p < \tilde{p}$ **then** $p \leftarrow p\,(1+\gamma_p \text{random})$

  else if $Y > D$ && $p \geq \tilde{p}$ **then** $p \leftarrow p\,(1-\gamma_p \text{ random})$

  else if $Y > D$ && $p < \tilde{p}$ **then** $Y^T \leftarrow Y\,(1-\gamma_y \text{ random})$

  end if

  $Y^T \leftarrow \max\{Y^T, \alpha\}$

  $L_d \leftarrow \text{ceil}(Y^T/\alpha)$

end function

function DEFINE_LABOR_DEMAND

  $L_d \leftarrow \min\{L_d, \text{floor}(\mathcal{L}/\mathcal{W})\}$

  $L_d \leftarrow \max\{L_d, 0\}$

  $v \leftarrow L_d - |E|$

end function

function PRODUCE

  $Y = \min\{Y^T, \alpha|E|\}$

  $D = 0$ ▹ The demand is reset to zero at each production cycle

end function

function HIRE(household h)

  $E \leftarrow h$

  $h.\text{get\_job}(W)$

  $v \leftarrow v-1$

end function

function FIRE (household h)

  $h.\text{lose\_job}()$

  $Ev \rightarrow h$

  $v \leftarrow v+1$

end function

function FIRE_RANDOM_WORKER

  if $|E| > 0$ **then**

  $h \leftarrow_R E$

  fire(h)

  end if

end function

Algorithm 6. The class `household`.

Parameters: $M = 3$, $c = 0.8$

Dynamic variables: S, W

      ▹ **Initialization methods**

function INIT

  $S \leftarrow 0$ ▹ Savings

  $W \leftarrow 0$ ▹ Salary

end function

      ▹ **Accounting methods**

function GET_MONEY $(\tilde{m})$

  $S \leftarrow S + \tilde{m}$

end function

      ▹ **Output methods**

function WEALTH

  return S

end function

function WORKING

  if $W > 0$ **then return** True

  else return False

  end if

end function

▷ **Job and goods market methods**
function GET_JOB($\tilde{W}$)
  $W \leftarrow \tilde{W}$
end function

function LOSE_JOB
  $W \leftarrow 0$
end function

function CONSUME(vector < firm > $\tilde{F}$)
  budget $\leftarrow cS$
  if budget > 0 **then**
    $F_c \leftarrow_{R,M} \tilde{F}$ ▷ Extract M random firms from $\tilde{F}$ and put them in the set F_c
    $F_c \leftarrow$ sort $(f \in F_c, f.\text{price}())$ ▷ Order the set F_c according to firms' prices
    spent $\leftarrow 0$
    for $(i \leftarrow 0; i < |F_c| \ \&\& \ \text{spent} < \text{budget}; i \leftarrow i+1)$ **do**
      $s \leftarrow f[i].\text{stock}()$ ▷ s is the stock available from this firm
      if $s > 0$ **then**
        $q \leftarrow$ (budget - spent)$/f[i].\text{price}()$ ▷ Maximum possible consumption from this firm
        if $s > q$ **then**
          $f[i].\text{sell}(q)$
          spent $\leftarrow$ budget   ▷ The household has finished the budget, the loop ends
        else
          $f[i].\text{sell}(s)$
          spent $\leftarrow$ spent $+ s f[i].$ price()
        end if
      end if
    end for
  end if
  $S \leftarrow S -$ spent
end function

Algorithm 7. The class bank.

Dynamic variables: E, ρ_b
  function INIT
    $E \leftarrow 0$ ▷ Bank liquidity
    $\rho_b \leftarrow 0$ ▷ Baseline interest rate
  end function

  function SET_$\rho(\tilde{\rho})$
    $\rho_b \leftarrow \tilde{\rho}$
  end function

  function COMPUTE_OFFER_RATE($\tilde{\ell}$)
    return $\rho_b G(\tilde{\ell})$ ▷ We chose $G(\ell) = 1 + \log(1 + \ell)$
  end function

  function GET_MONEY($\tilde{m}$)
    $E \leftarrow E + \tilde{m}$
  end function

Appendix B. Pseudo-code of Mark 0

We present here the pseudo-code for the Mark 0 model discussed in Sections 3.1 and 4. The source code is available on the site of the CRISIS project (www.crisis-economics.eu).

Algorithm ~ 8. The basic Mark 0.

Require: N_F Number of firms; $\mu, c, \beta, \gamma_p, \eta_+^0, \eta_-^0, \delta, \Theta, \varphi, f$; T total evolution time;
    ▷ **Initialization**
  for $(i \leftarrow 0; i < N_F; i \leftarrow i+1)$ **do**
    $W[i] \leftarrow 1$ ▷ Salaries are always fixed to one
    $p[i] \leftarrow 1 + 0.2(\text{random} - 0.5)$
    $Y[i] \leftarrow \mu[1 + 0.2(\text{random} - 0.5)]/2$ ▷ Initial employment is 0.5
    $\mathcal{E}[i] \leftarrow W[i]Y[i] \ 2 \ \text{random}$
    $a[i] \leftarrow 1$
  end for

  $S \leftarrow N_F - \sum_i \mathcal{E}[i]$

▷**Main loop**
for $(t \leftarrow 1; t \leq T; t \leftarrow t+1)$ **do**
$\quad u \leftarrow 1 - \frac{1}{\mu N_F}\sum_i Y[i]$
$\quad \varepsilon \leftarrow 1 - u$
$\quad \overline{p} \leftarrow \frac{\sum_i p[i]Y[i]}{\sum_i Y[i]}$
$\quad \overline{w} \leftarrow \frac{\sum_i W[i]Y[i]}{\sum_i Y[i]}$
$\quad \tilde{u}[i] \leftarrow \frac{\exp(\beta W[i]/\overline{w})}{\sum_i a[i]\exp(\beta W[i]/\overline{w})}N_F u$

$\qquad$ ▷ **Firms update prices, productions and wages**
for $(i \leftarrow 0; i < N_F; i \leftarrow i+1)$**do**

$\quad$ **if** $a[i]==1$ **then**

$\qquad$ **if** $Y[i] < D[i]$ **then** $\qquad$ ▷ Wage update
$\qquad\quad$ **if** $\mathcal{P}[i] > 0$ **then**
$\qquad\qquad W[i] \leftarrow W[i][1 + \gamma_w \varepsilon \ \texttt{random}]$
$\qquad\qquad W[i] \leftarrow \min\{W[i], P[i]\min[D[i], Y[i]]/Y[i]\}$
$\qquad\quad$ **end if** $\qquad$ ▷ This is Eq. (11) in the main text
$\qquad\quad Y[i] \leftarrow Y[i] + \min\{\eta_+ (D[i] - Y[i]), \mu\tilde{u}[i]\}$
$\qquad\quad$ **if** $p[i] < \overline{p}$ **then** $p[i] \leftarrow p[i](1 + \gamma_p\texttt{random})$
$\qquad$ **end if**

$\qquad$ **else if** $Y[i] > D[i]$ **then**
$\qquad\quad$ **if** $\mathcal{P}[i] < 0$ **then**
$\qquad\qquad W[i] \leftarrow W[i][1 - \gamma_w u\texttt{random}]$
$\qquad\quad$ **end if**
$\qquad\quad Y[i] \leftarrow \max\{0, Y[i] - \eta_- (D[i] - Y[i])\}$
$\qquad\quad$ **if** $p[i] < \overline{p}$ **then** $p[i] \leftarrow p[i](1 - \gamma_p \ \texttt{random})$
$\qquad$ **end if**
$\quad$ **end if**
end if
end for

$\quad u \leftarrow 1 - \frac{1}{\mu N_F}\sum_i Y[i]$ $\qquad$ ▷ Update u and $\overline{p}$
$\quad \overline{p} \leftarrow \frac{\sum_i p[i]Y[i]}{\sum_i Y[i]}$

Algorithm~9. The basic Mark0.

$\qquad$ ▷ **Households decide the demand**
$C \leftarrow c(\max\{S, 0\} + \sum_i W[i]Y[i])$

for $(i \leftarrow 0; i < N_F; i \leftarrow i+1)$ **do**
$\quad D[i] \leftarrow \frac{Ca[i]\exp(-\beta p[i]/\overline{p})}{p[i]\sum_i a[i]\exp(-\beta p[i]/\overline{p})}$ $\qquad$ ▷ Inactive firms have no demand
end for

$\qquad$ ▷ **Accounting**
for $(i \leftarrow 0; i < N_F; i \leftarrow i+1)$ **do**
$\quad$ **if** $a[i]==1$ **then**
$\qquad \mathcal{P}[i] \leftarrow p[i]\min\{Y[i], D[i]\} - W[i]Y[i]$
$\qquad S \leftarrow S - \mathcal{P}[i]$
$\qquad \mathcal{E}[i] \leftarrow \mathcal{E}[i] + \mathcal{P}[i]$
$\qquad$ **if** $\mathcal{P}[i] > 0$ && $\mathcal{E}[i] > 0$ **then** $\qquad$ ▷ Pay dividends
$\qquad\quad S \leftarrow S + \delta\mathcal{P}[i]$
$\qquad\quad \mathcal{E}[i] \leftarrow \mathcal{E}[i] - \delta\mathcal{P}[i]$
$\qquad$ **end if**
$\qquad$ **if** $\mathcal{E}[i] > \Theta W[i]Y[i]$ **then** $\qquad$ ▷ Set of healthy firms
$\qquad\quad \mathcal{H} \leftarrow i$
$\qquad$ **end if**
$\quad$ **end if**
end for

$\qquad$ ▷ **Defaults**
deficit $= 0$
for $(i \leftarrow 0; i < N_F; i \leftarrow i+1)$ **do**
$\quad$ **if** $a[i]==1$ && $\mathcal{E}[i] < -\Theta Y[i]W[i]$ **then**
$\qquad j \leftarrow_R \mathcal{H}$
$\qquad$ **if** $\texttt{random} < 1 - f$ && $\mathcal{E}[j] > -\mathcal{E}[i]$ **then** $\qquad$ ▷ Bailed out
$\qquad\quad \mathcal{E}[j] \leftarrow \mathcal{E}[j] + \mathcal{E}[i]$
$\qquad\quad \mathcal{E}[i] \leftarrow 0$
$\qquad\quad p[i] \leftarrow p[j]$

```
            W[i]←W[j]
       else        ▷ Bankrupted
            deficit ← deficit − ℰ[i]
            a[i]←0
            Y[i]←0
       ℰ[i]←0
       end if
    end if
end for


            ▷ Revivals
ℰ⁺←0
for (i←0; i < N_F; i←i+1) do
    if a[i] == 0 && random < φ then        ▷ Reactivate firm
         a[i]←1
         p[i]←p̄
         Y[i]←μu random
         ℰ[i]←W[i]Y[i]
         deficit←deficit +ℰ[i]
    end if
    if a[i] == 1 && ℰ[i] > 0 then        ▷ Firms total savings
         ℰ⁺←ℰ⁺ +ℰ[i]
    end if
end for


            ▷ Debt
if deficit > S then        ▷ Households cannot be indebted
    deficit ← deficit −S
    S←0
    for (i←0; i < N_F; i←i+1) do
         if a[i] == 1 && ℰ[i] > 0 then
              ℰ[i]←ℰ[i]− ℰ[i]/ℰ⁺ deficit
         end if
    end for
else
    S←S− deficit
end if

end for
```

Appendix C. Perturbative solution of the schematic model

C.1. Stationary distribution at first order in γ_p

From Eq. (25), the stationary distribution $P_{st}(\lambda)$ satisfies

$$P_{st}(\lambda) = \int_0^1 d\xi \int_0^\infty d\lambda' \, P_{st}(\lambda')\delta\left(\lambda-\lambda'+\xi+\frac{\gamma_p}{2}\xi^2\right) + \int_0^1 d\xi \int_{-\infty}^0 d\lambda' \, P_{st}(\lambda')\delta\left(\lambda-\lambda'-\xi+\frac{\gamma_p}{2}\xi^2\right). \tag{C.1}$$

We recall that the dynamics of $\lambda_i(t)$ is

$$\begin{cases} \lambda_i(t)<0 & \lambda_i(t+1)=\lambda_i(t)+\xi_i(t)-\frac{1}{2}\gamma_p\xi_i(t)^2 \\ \lambda_i(t)>0 & \lambda_i(t+1)=\lambda_i(t)-\xi_i(t)-\frac{1}{2}\gamma_p\xi_i(t)^2 \end{cases} \tag{C.2}$$

with $\xi_i(t)$ a random variable in $[0,1]$. From these equations it is clear that positive and negative λ_i are pushed towards the origin, and after some transient time one necessarily has

$$-1-\tfrac{1}{2}\gamma_p \leq \lambda_i(t) \leq 1-\tfrac{1}{2}\gamma_p. \tag{C.3}$$

We first set $\gamma_p = 0$ and therefore restrict $\lambda \in [-1,1]$ in Eq. (C.1). We obtain

$$P_{st}(\lambda) = \int_{\max(0,-\lambda)}^{1-\max(\lambda,0)} d\xi \, P_{st}(\lambda+\xi) + \int_{\max(0,\lambda)}^{1-\max(-\lambda,0)} d\xi \, P_{st}(\lambda-\xi) = \mathcal{L}_0 P_{st}, \tag{C.4}$$

where we call $\mathcal{L}_0$ the linear operator that appears on the right-hand side. When one computes the action of $\mathcal{L}_0$ on the functions $g_n(\lambda) = \text{sign}(\lambda)\lambda^n$ and $h_n(\lambda) = \lambda^n$, one finds, for $\lambda > 0$

$$(n+1)\mathcal{L}_0 g_n(\lambda) = \begin{cases} 1-(-1)^n - \sum_{k=0}^{n-1} D_n^k(-1)^k\lambda^{n-k}, & \lambda > 0 \\ 1-(-1)^n + \sum_{k=0}^{n-1} D_n^k\lambda^{n-k}, & \lambda < 0 \end{cases} \tag{C.5}$$

with C_n^k the binomial coefficients and $D_n^k = \sum_{j=k}^n C_j^k$, and

$$(n+1)\mathcal{L}_0 h_n(\lambda) = \begin{cases} 1 - \lambda^{n+1} - (\lambda-1)^{n+1} = 1 + (-1)^n - 2\lambda^{n+1} - \sum_{k=1}^n C_{n+1}^k (-1)^k \lambda^{n+1-k}, & \lambda > 0, \\ (1+\lambda)^{n+1} + \lambda^{n+1} + (-1)^n = 1 + (-1)^n + 2\lambda^{n+1} + \sum_{k=1}^n C_{n+1}^k \lambda^{n+1-k}, & \lambda < 0. \end{cases} \tag{C.6}$$

Let us focus on small values of n, useful in the following:

$$\begin{aligned} n = 0 &\to \mathcal{L}_0 g_0 = 0, \quad \mathcal{L}_0 h_0 = 2(h_0 - g_1), \\ n = 1 &\to \mathcal{L}_0 g_1 = 1 - g_1, \quad \mathcal{L}_0 h_1 = -g_2 + h_1, \\ n = 2 &\to \mathcal{L}_0 g_2 = h_1 - g_2, \quad \mathcal{L}_0 h_2 = \tfrac{2}{3}(h_0 - g_3) + h_2 - g_1. \end{aligned} \tag{C.7}$$

In particular, one has

$$\mathcal{L}_0(h_0 - g_1) = 2(h_0 - g_1) - 1 + g_1 = (h_0 - g_1). \tag{C.8}$$

This shows that $P_0(\lambda) = h_0(\lambda) - g_1(\lambda) = 1 - |\lambda|$ (called the "tent") is an eigenvector of $\mathcal{L}_0$ with eigenvalue 1, i.e. this is the stationary state for $\gamma_p = 0$.

This basic solution allows us to obtain perturbatively the stationary solution for small γ_p. Recalling that the support of $P_{\text{st}}(\lambda)$ is given by Eq. (C.3), Eq. (C.1) can be written as

$$P_{\text{st}}(\lambda) = \int_{a_-}^{a_+} d\xi \, P_{\text{st}}\left(\lambda + \xi + \tfrac{1}{2}\gamma_p \xi^2\right) + \int_{b_-}^{b_+} d\xi \, P_{\text{st}}\left(\lambda - \xi + \tfrac{1}{2}\gamma_p \xi^2\right) = \mathcal{L}_{\gamma_p} P_{\text{st}}, \tag{C.9}$$

with

$$\begin{aligned} a_- &= \max\left\{0, -\lambda\left(1 + \tfrac{\gamma_p}{2}\lambda\right)\right\}, \quad a_+ = \min\left\{1, 1 - \lambda - \tfrac{\gamma_p}{2}\left(1 + (1-\lambda)^2\right)\right\}, \\ b_- &= \max\left\{0, \lambda\left(1 + \tfrac{\gamma_p}{2}\lambda\right)\right\}, \quad b_+ = \min\left\{1, 1 + \lambda + \tfrac{\gamma_p}{2}\left(1 + (1+\lambda)^2\right)\right\}, \end{aligned} \tag{C.10}$$

these integration bounds coming from the combination of the support in Eq. (C.3) and the integration bounds on λ' in Eq. (C.1). Writing $P_{\text{st}} = P_0 + \gamma_p P_1$ where P_0 is the tent solution, and $\mathcal{L}_{\gamma_p} = \mathcal{L}_0 + \gamma_p \mathcal{L}_1$, we obtain at first order in γ_p an equation of the form:

$$(1 - \mathcal{L}_0)P_1 = \mathcal{L}_1 P_0 = S. \tag{C.11}$$

The source term S can be most easily computed by computing $\mathcal{L}_{\gamma_p} P_0$ and expanding the result at first order in γ_p. We find $S = \lambda/2 - \text{sign}(\lambda)\lambda^2 = \tfrac{1}{2}h_1 - g_2$. This is nice because one can look for a solution involving only $n=1$ and $n=2$ that closes the equation:

$$P_1 = \alpha h_1 + \beta g_2 \tag{C.12}$$

Using $\mathcal{L}_0 g_2 = h_1 - g_2 = \mathcal{L}_0 h_1$, one finds

$$(1 - \mathcal{L}_0)P_1 = \alpha h_1 + \beta g_2 - (\alpha + \beta)(h_1 - g_2) = -\beta h_1 + (\alpha + 2\beta)g_2 = S = \tfrac{1}{2}h_1 - g_2, \tag{C.13}$$

so the result for the coefficients is

$$\beta = -\tfrac{1}{2}, \quad \alpha = 0. \tag{C.14}$$

The final solution for the stationary state to first order in γ is

$$P = 1 - |\lambda| - \tfrac{1}{2}\gamma_p \, \text{sign}(\lambda)\lambda^2. \tag{C.15}$$

Note that P is still normalized, as it should and goes to zero (to first order in γ_p) at the boundary of the support interval given in Eq. (C.3).

C.2. Perturbative analysis of the oscillations

In order to understand oscillations, we start from Eq. (35) and in the following we assume that $0 < C < 1$. We introduce in this appendix (not to be confused with the firms deposits in the main text) $\mathcal{E}(t) = (\overline{\alpha}(t) - C\overline{\lambda}(t))/2(1-C)$, $\Lambda(t) = \min\{\overline{\lambda}(t), \mathcal{E}(t)\}$ and $\Omega(t) = \max\{\overline{\lambda}(t), \mathcal{E}(t)\}$, and we write Eq. (35) equivalently as

$$\alpha_i(t+1) = \alpha_i(t) - \min\{\lambda_i(t), C\lambda_i(t) + \tfrac{1}{2}\left(\overline{\alpha}(t) - C\overline{\lambda}(t)\right)\}$$

$$\begin{cases} \text{If } \lambda_i(t) < \Lambda(t) \Rightarrow \lambda_i(t+1) = \lambda_i(t) + \xi_i(t) \\ \text{If } \Omega(t) < \lambda_i(t) < \Lambda(t) \Rightarrow \lambda_i(t+1) = \lambda_i(t) \\ \text{If } \lambda_i(t) > \Omega(t) \Rightarrow \lambda_i(t+1) = \lambda_i(t) - \xi_i(t) \end{cases} \tag{C.16}$$

The master equation for the distribution of λ reads

$$P_{t+1}(\lambda') = \int_{-\infty}^{\Lambda(t)} d\lambda \int_0^1 d\xi\, P_t(\lambda)\delta(\lambda'-\lambda-\xi) + \int_{\Lambda(t)}^{\Omega(t)} d\lambda\, P_t(\lambda)\delta(\lambda'-\lambda') + \int_{\Omega(t)}^{\infty} d\lambda \int_0^1 d\xi P_t(\lambda)\delta(\lambda'-\lambda+\xi)$$

$$= \int_{-\infty}^{\Lambda(t)} d\lambda \int_0^1 d\xi\, P_t(\lambda)\delta(\lambda'-\lambda-\xi) + \theta(\Lambda(t)\le \lambda' \le \Omega(t))P_t(\lambda') + \int_{\Omega(t)}^{\infty} d\lambda \int_0^1 d\xi\, P_t(\lambda)\delta(\lambda'-\lambda+\xi). \tag{C.17}$$

while the evolution equation for $\overline{\alpha}$ is

$$\overline{\alpha}(t+1) = \overline{\alpha}(t) - \int_{-\infty}^{\infty} d\lambda\, P_t(\lambda)\, \min\{\lambda, C\lambda + \tfrac{1}{2}\left(\overline{\alpha}(t) - C\overline{\lambda}(t)\right)\}. \tag{C.18}$$

C.2.1. Stationary state

Numerically, we observe that $\Lambda(t) \sim \Omega(t)$ and their variations are much smaller than the width of the distributions of λ and α. In the limit where

$$\Lambda(t) = \Omega(t) = \overline{\lambda}(t) = \mathcal{E}(t) = \frac{\overline{\alpha}(t) - C\overline{\lambda}(t)}{2(1-C)} \tag{C.19}$$

we obtain

$$P_{t+1}(\lambda') = \int_{-\infty}^{\Lambda(t)} d\lambda \int_0^1 d\xi\, P_t(\lambda)\delta(\lambda'-\lambda-\xi) + \int_{\Lambda(t)}^{\infty} d\lambda \int_0^1 d\xi\, P_t(\lambda)\delta(\lambda'-\lambda+\xi), \tag{C.20}$$

whose stationary solution is $P(\lambda) = P_0[\lambda - \overline{\lambda}^*]$ where $P_0(x) = (1-|x|)\theta(1-|x|)$ is the tent function, with $\overline{\lambda}^*$ undetermined at this stage. Plugging this result in Eq. (C.18), we have

$$\overline{\alpha}(t+1) = \overline{\alpha}(t) - \overline{\lambda}^* + \frac{1-C}{6}. \tag{C.21}$$

The fixed point is therefore $\overline{\lambda}^* = (1-C)/6$ and from the condition (C.19) we get $\overline{\alpha}^* = (2-C)\overline{\lambda}^* = (2-C)(1-C)/6$.

C.2.2. Oscillation around the stationary state – a simple approximation

In order to study the small oscillations we can make a very simple approximation, namely that at each time t we have $P_t(\lambda) = P_0[\lambda - \overline{\lambda}(t)]$. If we inject this approximation in Eq. (C.18) we get, provided $\overline{\lambda}(t) - \overline{\lambda}^*$ is not too large, that

$$\overline{\alpha}(t+1) = \overline{\alpha}(t) - \int_{-\infty}^{\infty} d\lambda\, \min\{\lambda, C\lambda + (1-C)\mathcal{E}(t)\}P_0\left[\lambda - \overline{\lambda}(t)\right]$$

$$= \overline{\alpha}(t) - \overline{\lambda}(t) - \frac{1-C}{6}[-1 + 3A - 3A^2 + A^3\,\mathrm{sgn}(A)]_{A=\mathcal{E}(t)-\overline{\lambda}(t)} \tag{C.22}$$

Next, injecting this approximation in Eq. (C.17) we get

$$\overline{\lambda}(t+1) = \overline{\lambda}(t) + \frac{1}{2}\int_{-\infty}^{\Lambda(t)} d\lambda\, P_0\left[\lambda - \overline{\lambda}(t)\right] - \frac{1}{2}\int_{\Omega(t)}^{\infty} d\lambda\, P_0\left[\lambda - \overline{\lambda}(t)\right]$$

$$\int_{\Omega(t)-\overline{\lambda}(t)}^{\infty} dx\, T[x]$$
$$= \frac{1}{2}\left[\Lambda(t) + \Omega(t) - \frac{1}{2}[\Lambda(t) - \overline{\lambda}(t)]^2\,\mathrm{sgn}\left[\Lambda(t) - \overline{\lambda}(t)\right] - \frac{1}{2}[\Omega(t) - \overline{\lambda}(t)]^2\,\mathrm{sgn}\left[\Omega(t) - \overline{\lambda}(t)\right]\right] \tag{C.23}$$

We therefore get a closed system of two equations for $\overline{\lambda}(t)$ and $\overline{\alpha}(t)$. For small C, this system of equations converges quickly to the fixed point. However, for C sufficiently close to 1 ($C \sim 0.94$) it has a limit cycle of period 2, followed by an exponential divergence.

C.2.3. Oscillation around the stationary state – perturbative computation

Another strategy to characterize the oscillations is to look again for a perturbed solution around the stationary state, with $\overline{\lambda}(t) = (1-C)/6 + \delta\overline{\lambda}_1(t)$ and $\overline{\alpha}(t) = (2-C)\overline{\lambda} + 2(1-C)\delta\overline{\lambda}_2(t) + C\delta\overline{\lambda}_1(t)$. The distribution of λ_i at time t is now $P_t(\lambda) = P_0(\lambda - (1-C)/6) + \delta P_t(\lambda - (1-C)/6)$. The shifted variable $\lambda - (1-C)/6$ will be denoted as x. The aim is to write an evolution for δP_{t+1} after one time step. The $\mathcal{L}_0$ operator is the same as before, as well as the set of functions g_n and h_n. We also introduce $\Delta(x)$ as a δ-function slightly spread out, but of unit area (its precise width is irrelevant if it is small enough).

To lowest order, one has

$$\mathcal{L}_0\Delta \approx \frac{1}{2}h_0. \tag{C.24}$$

We will also need, as found above,

$$\mathcal{L}_0 h_0 = 2(h_0 - g_1), \quad \mathcal{L}_0 g_0 = 0, \quad \mathcal{L}_0 g_1 = h_0 - g_1. \tag{C.25}$$

We now introduce $d = |\overline{\delta\lambda}_2 - \overline{\delta\lambda}_1|$ and $s = \overline{\delta\lambda}_2 + \overline{\delta\lambda}_1$. Assuming $d, s \ll 1$, one finds, to first order,

$$\delta P_{t+1} = \mathcal{L}_0 \delta P_t - \tfrac{1}{2} d_t h_0 + \tfrac{1}{2} s_t g_0 + d_t \Delta \tag{C.26}$$

Note that the integral $[\delta P] = \int_{-1}^{1} dx \, \delta P(x)$ is zero and conserved in time, as it should be (one has $[h_0] = 2$, $[g_0] = 0$ and $[\Delta] = 1$).

The idea now is to expand δP in terms of the h, g functions. It turns out that only three of them are needed, plus the Δ contribution. Indeed, assume

$$\delta P_t = A_t (h_0 - 2\Delta) + B_t g_0 + C_t \left(g_1 - \tfrac{1}{2} h_0\right) \tag{C.27}$$

Then, using the dynamical equation and the algebra above, one finds

$$\mathcal{L}_0 \delta P_t = A_t \left(2(h_0 - g_1) - h_0\right) = -2A_t \left(g_1 - \tfrac{1}{2} h_0\right) \tag{C.28}$$

Hence, the evolution is closed on itself, with the following evolution rules:

$$A_{t+1} = -\tfrac{1}{2} d_t, \quad B_{t+1} = \tfrac{1}{2} s_t, \quad C_{t+1} = -2A_t. \tag{C.29}$$

Note that $A_t - B_t = -\tfrac{1}{2}(d_{t-1} + s_{t-1}) = -\max\left(\overline{\delta\lambda}_{1t-1}, \overline{\delta\lambda}_{2t-1}\right) = -M_{t-1}$.

Now, by definition $\overline{\lambda}_t = (1 - C)/6 + \int_{-1}^{1} dx \, x\delta P_t$, but since h_0 and g_1 are even, the only contribution comes from the g_0 component. Therefore

$$\overline{\delta\lambda}_{1t} = B_t \int_{-1}^{1} dx \, x g_0(x) = B_t. \tag{C.30}$$

This leads to a first evolution equation:

$$\overline{\delta\lambda}_{1t+1} = \tfrac{1}{2} s_t = \tfrac{1}{2} \left(\overline{\delta\lambda}_{1t} + \overline{\delta\lambda}_{2t}\right). \tag{C.31}$$

The second equation comes from the evolution of the α_i's. From

$$\overline{\alpha}(t+1) = \overline{\alpha}(t) - C\overline{\delta\lambda}_1(t) - \frac{1-C}{2}\overline{\delta\lambda}_2(t) - (1-C)\int_{-1}^{0} dx \, x\delta P_t(x), \tag{C.32}$$

we finally find

$$\overline{\delta\lambda}_{2t+1} = \frac{3}{4}\overline{\delta\lambda}_{2t} - \frac{1}{4}M_{t-1} + \frac{7}{24}d_{t-2} - \frac{C}{4(1-C)}\left(\overline{\delta\lambda}_{1t} + \overline{\delta\lambda}_{2t}\right). \tag{C.33}$$

The coupled set of iterations for $\overline{\delta\lambda}_{1t}$ and $\overline{\delta\lambda}_{2t}$, Eqs. (C.31) and (C.33), leads to damped oscillations for $C < C^{**}$ and sustained oscillations for $C > C^{**}$, with $C^{**} \approx 0.91$. The numerical value of C^{**} however does not coincide with that of $C^* \approx 0.45$, obtained from the direct simulation of Eq. (35). This might be due to the neglecting higher order corrections, which appear to be numerically large: the oscillations generated by Eq. (35) are of amplitude ~ 0.1 at the onset, suggesting a sub-critical bifurcation. For small C, on the other hand, the oscillation amplitude is small and the above equations appear to be quantitatively correct, validating the above calculations.

References

Anand, K., Gai, P., Marsili, M., 2010. The rise and fall of trust networks. In: Progress in Artificial Economics. Lecture Notes in Economics and Mathematical Systems, vol. 645, p. 77.

Anand, K., Kirman, A., Marsili, M., 2011. Epidemics of rules, rational negligence and market crashes. Eur. J. Financ 19.5, 438–447 (2013).

Anderson, S.P., De Palma, A., Thisse, J.F., 1992. Discrete Choice Theory of Product Differentiation. MIT Press, New York and Bouchaud (2013) for a recent discussion and further references.

Aoki, M., Yoshikawa, H., 2011. Reconstructing Macroeconomics: A Perspective from Statistical Physics and Combinatorial Stochastic Processes. Cambridge University Press, Cambridge (UK).

Ashraf, Q., Gershman, B., Howitt, P., 2011. Banks, Market Organization, and Macroeconomic Performance: an Agent-Based Computational Analysis. no. w17102. National Bureau of Economic Research.

Assenza, T., Delli Gatti, D., 2013. E Pluribus Unum: Macroeconomic Modelling for Multi-agent Economies. UCSC Working Paper.

Axtell, R.L., 2000. Why agents? On the varied motivations for agent computing in the social sciences. In: Macal, C.M., Sallach, D. (Eds.), Proceedings of the Workshop on Agent Simulation.

Bak, P., 1996. How Nature Works: The Science of Self-Organized Criticality. Copernicus, Springer, New York.

Ball, P., 2012. Why Society is a Complex Matter. Springer, Berlin (for an introductory review).

Balibar, S., 2007. The discovery of superfluidity. J. Low Temp. Phys. 146, 441–470.

Bank of England, 2014. Quarterly Bulletin, Q1 vol. 54, No. 1.

Bonart, J., Bouchaud, J.-P., Landier, A., Thesmar, D., 2014. Instabilities in Large Economies: Aggregate Volatility Without Idiosyncratic Shocks arxiv:1406.5022.

Bouchaud, J.-P., 2011. The endogenous dynamics of markets: price impact, feedback loops and instabilities. In: Berd, A. (Ed.), Lessons from the 2008 Crisis, Risk Books, Incisive Media, London.

Bouchaud, J.-P., 2013. Crises and collective socio-economic phenomena: simple models and challenges. J. Stat. Phys. 151, 567–606. and refs. therein.

Brock, W., Durlauf, S., 2001. Discrete choice with social interactions. Rev. Econ. Stud. 68, 235.

Buchanan, M., 2008. This Economy does not Compute, New York Times, October 1, 2008.

Buiter, W., 2009. The Unfortunate Uselessness of Most "State of the Art" Academic Monetary Economics. ⟨http://blogs.ft.com/maverecon/2009/03/⟩.

Caballero, R.J., 2010. Macroeconomics after the crisis: time to deal with the pretense-of-knowledge syndrome. J. Econ. Perspect. 24, 85–102.

Caccioli, F., Vivo, P., Marsili, M., 2009. Eroding market stability by proliferation of financial instruments. Eur. Phys. J. B 71, 467–479.

Challet, D., Marsili, M., Zhang, Y.C., 2005. Minority Games. Oxford University Press, Oxford (UK).

Cristelli, M., Pietronero, L., Zaccaria, A. 2011. CriticalOverview of Agent-based Models for Economics arxiv:1101.1847.

Dawid, H., Neugart, M., 2011. Agent-based models for economic policy design. East. Econ. J. 37, 44–50.

Dawid, H., Gemkow, S., Harting, P., van der Hoog, S., Neugarty, M., 2011. The Eurace@Unibi Model: An Agent-Based Macroeconomic Model for Economic Policy Analysis. Working Paper, Bielefeld University.

Deissenberg, C., van der Hoog, S., Dawid, H., 2008. Eurace: a massively parallel agent-based model of the european economy. Appl. Math. Comput. 204, 541–552.

Delli Gatti, D., Desiderio, S., Gaffeo, E., Cirillo, P., Gallegati, M., 2011. Macroeconomics from the Bottom-Up. Springer, Berlin.

Delli Gatti, D., Gaffeo, E., Gallegati, M., Giulioni, G., Palestrini, A., 2008. Emergent Macroeconomics: An Agent-Based Approach to Business Fluctuations. Springer, Berlin.

DiGuilmi, C., Gallegati, M., Landini, S., 2008. Economic dynamics with financial fragility and mean-field interaction: a model. Physica A: Stat. Mech. Appl. 387 (15), 3852–3861.

Dosi, G., et al., 2013. Income distribution, credit and fiscal policies in an agent-based Keynesian model. J. Econ. Dyn. Control 37 (8), 1598–1625.

Dosi, G., Fagiolo, G., Roventini, A., 2005. An evolutionary model of endogenous business cycles. Comput. Econ. 27, 3–24.

Dosi, G., Fagiolo, G., Roventini, A., 2010. Schumpeter meeting Keynes: a policyfriendly model of endogenous growth and business cycles. J. Econ. Dyn. Control 34, 1748–1767. and Fagiolo and Roventini (2012).

Eggertsson, G.B., Krugman, P., 2012. Debt, deleveraging, and the liquidity trap: a Fisher–Minsky–Koo approach. Q. J. Econ. 127, 1469–1513.

Fagiolo, G., Roventini, A., 2012. Macroeconomic policy in DGSE and agent-based models. Rev. l'OFCE 5, 67–116.

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

Farmer, D., Foley, D., 2009. The economy needs agent-based modelling. Nature 460, 685–686.

Flaschel, P., 2010. The Macrodynamics of Capitalism. Springer, Berlin.

Gaffeo, E., Delli Gatti, D., Desiderio, S., Gallegati, M., 2008. Adaptive microfoundations for emergent macroeconomics. East. Econ. J. 34, 441–463.

Gali, J., 2008. Monetary Policy, Inflation and the Business Cycle. Princeton University Press, Princeton (USA).

Geanakoplos, J., Axtell, R., Farmer, D.J., Howitt, P., Conlee, B., Goldstein, J., Hendrey, M., Palmer, N.M., Yang, C.-Y., 2012. Getting at Systemic Risk via an Agent-Based Model of the Housing Market. Cowles Foundation Discussion Papers 1852, Cowles Foundation for Research in Economics, Yale University.

Giardina, I., Bouchaud, J.-P., 2003. Bubbles crashes and intermittency in agent based market models. Eur. Phys. J. B 31, 421–437.

Goldenfeld, N., 1992. Lectures on Phase Transitions and the Renormalization Group. Addison-Wesley, Reading.

Goldstone, R., Janssen, M., 2005. Computational models of collective behavior. Trends Cogn. Sci. 9, 424–430.

Goodwin, R.M., 1967. A growth cycle. In: Feinstein, C.H. (Ed.), Socialism, Capitalism and Economic Growth, Cambridge University Press, Cambridge (UK).

Grazzini, J., Delli Gatti, D., 2012. In: the Crisis Workshop in Milan, November 2012. Its transcription is given by the pseudo-code in Appendix A.

Greenwald, B.C., Stiglitz, J.E., 1993. Financial market imperfections and business cycles. Q. J. Econ. 108, 77–114.

Gualdi, S., Tarzia, M., Zamponi, F., Bouchaud, J.-P., 2015a. Policy Experiments in Simple Macro-economic ABM. In preparation.

Gualdi, S., Tarzia, M., Zamponi, F., Bouchaud, J.-P., 2015b. On Simple Aggregate Descriptions of ABMs (and the real economy). In preparation.

Hommes, C.H., 2006. Heterogeneous Agent Models in Economics and Finance, Handbooks in Economics Series, vol. 23. Elsevier, North-Holland1109–1186.

Hommes, C., 2013. Behavioral Rationality and Heterogeneous Expectations in Complex Economic Systems. Cambridge University Press.

Kirman, A., 1992. Whom or what does the representative individual represent? J. Econ. Perspect. 6, 117–136.

Kirman, A., 2010. The economic crisis is a crisis for economic theory. CESifo Econ. Stud. 56, 498–535.

Kondor, I., Csabai, I., Papp, G., Mones, E., Czimbalmos, G., Sandor, M.Cs., 2014. Strong Random Correlations in Networks of Heterogeneous Agents. http://dx.doi.org/10.1007/s11403-014-0125-5, Springer.

LeBaron, B., Tesfatsion, L., 2008. Modeling macroeconomies as open-ended systems of interacting agents. Am. Econ. Rev.: Pap. Proc. 98, 246–250.

Leijonhufvud, A., 2006. Agent-based macro. In: Tesfatsion, L., Judd, K. (Eds.), Handbook of Computational Economics, pp. 1625–1636.

Lorenz, J., Battiston, S., Schweitzer, F., 2009. Systemic risk in a unifying framework for cascading processes on networks. Eur. Phys. J. B 71, 441.

Lye, R., Peng, J., Tana, L., Cheong, S.A., 2012. Understanding Agent-Based Models of Financial Markets: a Bottom-Up Approach Based on Order Parameters and Phase Diagrams. Physica A 391, 22.

Mandel, A., Furst, S., Lass, W., Meissner, F., Jaeger, C., 2009. Lagom generiC: an Agent-Based Model of Growing Economies. ECF Working Paper 1 /2009.

Metzig, C., Gordon, M. 2012. Heterogeneous Enterprises in a Macroeconomic Agent-Based Model arxiv:1211.5575.

Minsky, H., 2008. Stabilizing an Unstable Economy. McGraw-Hill, New York.

Napoletano, M., et al., 2012. Wage formation investment behavior and growth regimes: an agent-based analysis. Rev. l'OFCE 5, 235–261.

Patzelt, F., Pawelzik, K., 2011. Criticality of adaptive control dynamics. Phys. Rev. Lett. 107, 238103.

Pyka, A., Fagiolo, G., 2007. Agent-based modelling: a methodology for neo-Schumpeterian economics. In: Elgar Companion to Neo-Schumpeterian Economics, p. 467.

Raberto, M., Teglio, A., Cincotti, S., 2012. Debt, deleveraging and business cycles: an agent-based perspective. Economics: The Open-Access Open-Assessment E-Journal 6 (2012-27), 1–49.

Russo, L., 2004. The Forgotten Revolution: How Science Was Born in 300 BC and Why It Had to Be Reborn. Springer, Berlin (Italian-speaking readers might also consult the wonderful first chapter of L. Russo, Flussi e Riflussi: indagine sull'origine di una teoria scientifica, Feltrinelli (2003), unfortunately out of print but available in many public libraries).

Sims, C., 1980. Macroecon. Real. Econom. 48, 1–48.

Sornette, D., 2005. Endogenous versus exogenous origins of crises. In: Albeverio, S., Jentsch, V., Kantz, H. (Eds.), Extreme Events in Nature and Society, Springer, Heidelberg.

Tesfatsion, L., 2002. Agent-based computational economics: growing economies from botton up. Artif. Life 8, 55–82.