DREAM
Danish Research Institute for
Economic Analysis and Modelling

Copy or Deviate?
The Market Economy as a Self-Organizing System

Peter Stephensen

Working paper 2022:2
21 November 2022

www.dreamgruppen.dk

Copy or Deviate? The Market Economy as a Self-Organizing System*

Peter Stephensen†

November 21, 2022

Preliminary version. Do not quote.

Abstract

An agent-based model for a decentralized market economy with multiple sectors is set up, and it is demonstrated that the economy is self-organizing if four empirically relevant conditions apply: S-shaped production functions, search behavior on the demand side, competition-based price setting and free entry of new firms.

The model is inspired by Gintis (2007) but differs on several points, for example by *not* using evolutionary theory. Instead of the evolutionary mechanism of Gintis (2007), we argue for a social mechanism which we call *Copy-Deviate behavior*. We define Copy-Deviate behavior as a situation where satisfied agents *copy* each other (which ensures equilibrium), whereas dissatisfied agents *deviate* from each other and thereby set the system in motion (possibly towards a new equilibrium). It is argued that competition-based pricing known from marketing can be seen as an example of Copy-Deviate behavior.

It is demonstrated that the system is self-organizing by numeric simulation. It is shown that the system follows a balanced growth path where real wages and production grow with the exogenous growth rate. The system is exposed to three supply shocks, and it is demonstrated that the model's effects follows conventional economic theory. It is demonstrated that competition-based pricing makes it relatively easy to understand the system, both in and out of equilibrium.

*The research project is supported by Economic Policy Research Network (EPRN), Denmark.

†Danish Research Institute for Economic Analysis and Modeling, DREAM, Landgreven 4, 7th floor DK-1301 Copenhagen K, E-mail: psp@dreamgruppen.dk [v0.13].

1 Introduction

Since Walras (1874) there has been a clear dichotomy between the concepts of 'equilibrium' and 'stability of equilibrium' in economic theory (De Vroey, 2016). The question of the existence of equilibrium is typically treated independently of the question of which processes could possibly produce this equilibrium. There is a considerable literature on the existence of general equilibria (e.g. Debreu, 1952; Arrow and Debreu, 1954). Partly independent of this, there is also a considerable literature on stability of equilibria (Arrow et al. 1959; Hahn and Negishi 1962; Uzawa 1962; Smale 1976). The purpose of this article is to break this dichotomy. For this purpose, an agent-based model for a decentralized market economy is set up, and it is demonstrated that the economy is self-organizing[1] if four empirically relevant conditions apply: S-shaped production functions, search behavior on the demand side, competition-based price setting and free entry for new firms. The model is inspired by Gintis (2007) but differs on several points, for example by *not* using evolutionary theory.

The first element (S-shaped production functions) concerns the supply side. Firms typically have an *optimal size* measured in terms of employment or production. There is an extensive literature on heterogeneous firms that theoretically and empirically shows that firms within the same sector have significant differences in productivity and size (for an overview, see Syverson, 2011). This literature should basically be able to explain why firms with different productivity can exist at the same time. The typical explanation is a combination of increasing returns (often due to fixed costs) and a concave revenue function (Syverson 2011: 333). The revenue function is concave either because of imperfect competition or - if perfect competition - diminishing returns to scale in the production function. Well-known examples from this literature are Melitz (2003), Hopenheyn (1992) and Eaton & Kortum (2002).

We follow Hopenheyn (1992) and assume that the individual firm has an S-shaped production function due to increasing returns to scale at small production levels and decreasing returns to scale at large production levels. Even under perfect competition, such a firm will have a well-defined optimal size regardless of what the system's prices and wages are. This is important - and in clear contrast to a firm with constant returns to scale. A firm with constant returns to scale in a perfectly competitive market is a strange creature - both in equilibrium and out of equilibrium. In equilibrium, a firm with constant returns to scale will take the cost-determined price that precisely ensures zero profit. In this situation, the firm will be indifferent to the quantity produced[2]. Out

[1] By self-organizing it is meant that equilibrium occurs exclusively via interaction between autonomous economic agents.

[2] In a walrasian equilibrium the quantity (supply) is determined by the demand. How this information is passed from the demand side to the supply side is typically not explained.

of equilibrium, a firm with constant returns to scale will typically have negative profit (and therefore will not produce) or positive profit (and therefore will want to produce infinitely much). The assumption of an S-shaped production function is, in addition to being more empirically relevant, also instrumental in achieving a theoretical description of the firm that is robust both within and outside equilibrium.

The second element (search behavior) deals with the demand side. It is often assumed in modern macro theory that there is search behavior in the labor market (Pissarides, 2000; Mortensen and Pissarides, 1994). Recently, increasing attention has also been paid to *consumer good search* (Qiu and Ríos-Rull, 2019; Kaplan and Menzio, 2016; Storesletten, Bai and Rios-Rull, 2011). We assume here that there is search behavior in both the labor market and the goods markets. The search friction on the good markets means that even in a homogeneous market a firm can set a price that is different from the market's average price. If the firm sets a high price, it will lose customers. If it sets a low price, it can attract customers. In other words, the firm can use the price to manage its demand. This is an important dynamic mechanism that must not be confused with a static demand curve.

The third element deals with the firm as price-setter. If the two conditions described above apply, the firm has a well-defined optimal supply and an opportunity to influence demand. If the firm's perceived demand at the prevailing market price (the average price on the homogeneous market) is exactly equal to the optimal supply, the firm has good reason to set its price equal to the market price. In doing so, it retains current customers. If the firm has too few customers, it has good reason to set the price lower than the market price for a period until it has attracted the desired number of customers. Conversely, if the firm has too many customers, it can for a period set its price higher than the market price. This price strategy is actually widespread in reality and is called *competition-based pricing* within marketing (see e.g. Kotler et al., 2020). According to Hinterhuber (2008), this price strategy is the most widespread compared to the two alternative strategies customer-value-based pricing and cost-based pricing.

The fourth element deals with the supply side as a system. We follow the literature on heterogeneous firms mentioned above (Hopenheyn, 1992; Melitz, 2003; Eaton & Kortum, 2002; Syverson, 2011) and assume free entry of new firms. We assume that an investor starts new firms and closes them if they do not give rise to a positive profit. This mechanism gives rise to Darwinian selection (Syverson, 2011: 348). Only the best firms survive, and in the event of a shock to the economy, you may find that the low-productivity firms close.

Why do the above four characteristics lead to a self-organizing system? The first three conditions make the firm a well-functioning social entity that *i)* has a well-defined size (supply), *ii)* has an objective (demand = supply), and *iii)* has a mechanism available by which it, via interaction with the outside world, can achieve its goal. We will call this *Copy-Deviate behavior*. As mentioned above, a firm that is satisfied with its demand will set its price at the average market price. It will *copy* the other firms. If it is dissatisfied with its demand, it will either set the price above or below the market price. It will *deviate* from the other firms. This gives rise to a powerful social mechanism. If everyone is satisfied, everyone will copy each other and we will, by definition, be in an equilibrium. If many are dissatisfied, many will deviate from the average and the system will move - possibly towards a new equilibrium.

The model developed is an agent-based model (Testfatsion 2006; LaBaron and Tesfatsion 2008; Dawid and Della Gatti 2018). The basic idea in agent-based models is to model only agents and their interactions. This type of model is therefore particularly well-suited for analyzing self-organizing systems[3]. The economy consists of approximately 50,000 heterogeneous households that demand 10 goods and supply labor. There are about 2,500 heterogeneous firms spread over 10 sectors. Within each sector, the same homogeneous product is offered. Finally, there is an investor who starts new firms. The productivity of new firms is heterogeneous in two dimensions: the productivity is the product of an individual productivity drawn from a Pareto distribution and a macro-productivity that grows over time. In this way, there will potentially be growth in the overall system and all firms will in the long term be outcompeted by new firms, regardless of productivity.

It is demonstrated that the system is self-organizing in four ways. First, it is shown that the system follows a balanced growth path where real wages and production grow with the exogenous growth rate. Business cycles are observed which are mainly due to the behavior of the investor, but the use of competition-based pricing on both the goods and labor markets implies that these are quite weak, and they would probably be absent in a system with millions of agents. Second, a productivity shock is performed. The productivity of all firms are increased, and it is demonstrated that the economy relatively quickly moves to a new equilibrium where real wages and production have grown while the rest of the variables have returned to their original values. It is possible to explain in high detail how the market price and the market wage react and ensure that the real wage adjusts as standard economic theory would predict. Third, a random number of firms are destroyed. It is demonstrated that firms are recreated relatively quickly and that the economy moves back to its original state in the long run. As an interesting corollary, it takes many years (about 30 years) before average productivity and total production return to their original levels. This is because it takes many years to recreate the firms' original age structure. Fourth and finally, the productivity of

[3] Axel Leijonhufvud: "Agent-based methods provide the *only way* in which we can explore the self-regulatory cababilities of complex dynamic models and thus advance our understanding of the adaptive dynamics of actual economies." (Leijonhufvud, 2006: 1636-7).

a single sector is shocked. It is demonstrated that in the long term this means that the sector's price and employment fall while production grows. This is consistent with the literature on unbalanced growth (e.g. Baumol, 1967; Acemoglu and Guerrieri, 2008).

In relation to existing literature, this article is most reminiscent of Gintis (2007). In Gintis (2007) an agent-based model is set up for a Walrasian economy. There are certain points of similarity but also several fundamental differences. Similarities include that the firms' production functions are non-linear (an inverted U-shape), that households display search behavior on the goods market, that firms are price setters and that there is entry of new firms. The four central elements mentioned above are therefore in place. The most significant difference is that Gintis (2007) uses an evolutionary approach where 5 per cent of the firms in each period copy the behavioral parameters of the best firms. This means that the firms in the same sector set different prices ("private prices"), which contributes to making the system self-organizing. We achieve price heterogeneity in another way - namely via competition-based pricing. We therefore use a social mechanism (Copy-Deviate-behavior) in place of an evolutionary mechanism. In addition, it should be mentioned that our production technology and the achievement of firm heterogeneity are based on Melitz (2003) and Hopenheyn (1995).

The article is structured as follows: Section 2 describes the model and its agents. In section 3, it is demonstrated that the system is self-organizing via a base run and 3 shocks. Finally, section 4 is the conclusion.

2 The model

We use an *agent-based model* (Testfatsion 2006; LaBaron & Tesfatsion 2008; Dawid & Delli Gatti 2018) to describe the system. The basic idea in agent-based models is to model only agents and their interaction. You only model conditions that actually occur in the system you are considering. It is not allowed/possible to add abstract elements such as equilibrium conditions, a Walrasian auctioneer or rational expectations. This type of model is therefore particularly well-suited for analyzing self-organizing systems. One can choose to see it as an advantage that agent-based models force the modeler to take a concrete, non-abstract approach. An Occams-razor-like principle could be: It is beautiful to abstract, but one should only do so when all obvious concrete solutions have been tried. Instead of rational expectations or perfect foresight, in an agent-based model you would typically assume that the agents' expectations are based on historical data (as in reality). Search behavior can be modeled explicitly by the individual agent searching among other agents, i.e. without using an ad hoc matching function as is typically done in macroeconomic models (Mortensen and Pissarides 1994; Pissarides 2000). One can thereby provide a true network-based micro-foundation for search behavior.

In practical terms, this model uses a so-called Event-Pump-Architecture (Stephensen, 2015), which has been used for several years in the Danish microsimulation model SMILE (Hansen et al., 2021). The model is programmed in C# and can be downloaded from github[4].

There are three types of agents in the model: households, firms and the investor. There are many heterogeneous households and firms and one investor[5]. In the following, we will describe the agents one at a time.

2.1 Households

Households demand goods and supply labor. They exhibit search behavior in both the goods market and the labor market. In the goods market, the household looks for the firm with the lowest price and in the labor market, it looks for the firm with the highest wage. The search behavior in the two markets plays a central role in making the system self-organizing.

The households are heterogeneous in two dimensions: age and income. A household consists of 1 person. New households enter the economy at age 18. The household dies according to the Gompertz-Makeham law (Gavrilov and Gavrilova, 1992), as the age-dependent probability of death μ_a is log-linear:

$$log(\mu_a) = \gamma_\mu (a - 100), \gamma_\mu > 0, \tag{2.1}$$

[4]https://github.com/PeterStephensen/dream.agentbased.SOE2.git

[5]The investor is the models only representative agent.

where a is age. The parameterization has been chosen so that the maximum lifetime is 100 years. Each period, a fixed number of 18-year-old households are born into the economy, and the system adjusts to a stationary population with a fixed overlapping generational structure.

The household has two sources of income: wage and profit. Employed households receive wages. All households receive an equal share of the profit generated in the last period[6]. Assume the j'th household is employed in firm i. The household will then receive the wage in period t:

$$v_{jt} = w_{it} \cdot \rho_{jt},$$

where w_{it} is the wage per productivity unit in firm i and ρ_{jt} is the household's individual productivity. The purpose of the parameter ρ_{jt} is to create heterogeneity in the income, as known for example from Carroll (1997). It is assumed that $\rho_{jt} = 1$ when the household enters the economy at age 18. After this, it is assumed that:

$$log\left(\rho_{jt}\right) = log\left(\rho_{j,t-1}\right) + \varepsilon_{jt},$$

where ε_{jt} is normally distributed with mean 0. This dynamic results in a right-skewed log-normally distributed wage distribution in the overall system. Income inequality increases over the lifetime.

The total current income for the jth household is

$$y_{jt} = v_{jt} + \pi_t$$

where π_t is the profit paid per household (the same for all households) and $v_{jt} = 0$ if the household is unemployed.

It is assumed that households consume their entire current income. They are Hand-to-Mouth consumers as it is called in New-Keynsian terminology. This assumption is made in order to have as simple a credit market as possible. We concentrate on the market mechanisms of the goods and labor markets and the creation of new firms. It will be an area for future research to examine how the model works if households have access to a financial market and thus can smooth consumption.

2.1.1 Buying goods

A household has potential contact with two types of firms in a given period: the firm where it works and the firms where it buys goods.

In a given period (month), the standard behavior of the household is to buy from the same firms

[6] It has no effect on the model's results how the profit is distributed due to the very simple consumer behavior described below.

as last month. With a low probability, the household conducts a market survey for each sector i. It randomly draws n firms from sector i. If the household among these n firms can find a product that is cheaper than the current one, it switches provider.

It is assumed that household j has a CES utility function:

$$u_{jt} = \left[\sum_{i=1}^{S} \mu_{ji}^{\frac{1}{E}} c_{jit}^{\frac{E-1}{E}} \right]^{\frac{E}{E-1}} \tag{2.2}$$

where S is the number of sectors, μ_{ji} are distribution parameters and E is the elasticity of substitution between the S goods. We assume that $E < 1$ is the same for all households and that μ_{ji} is drawn from a rectangular distribution when the household is born. The households therefore have heterogeneous preferences.

If the household maximizes the utility function (2.2) given the income y_{jt} we have:

$$c_{jit} = \mu_{ji} \left(\frac{p_{it}}{P_{jt}^C} \right)^{-E} \frac{y_{jt}}{P_{jt}^C} \tag{2.3}$$

where the CES price index P_j^C is given by:

$$P_{jt}^C = \left[\sum_{i=1}^{S} \mu_{ji} p_{it}^{1-E} \right]^{\frac{1}{1-E}}.$$

In each period, the household starts by making a budget for its purchases of goods. It appears from (2.3) that the budget for the individual item $b_{jit} \equiv p_{it} \cdot c_{jit}$ is given by:

$$b_{jit} = \mu_{ji} \left(\frac{p_{it}}{P_{jt}^C} \right)^{1-E} y_{jt} \tag{2.4}$$

The household collects prices from current suppliers $(p_{1t}, .., p_{St})$ and calculate P_{jt}^C. Based on (2.4), it can then distribute the income y_{jt} to individual parts b_{jit}. The household purchases goods from sector i by contacting the relevant firm in sector i and asking to buy for b_{jit}. If the firm can deliver the desired quantity, it answers Yes. If the firm answers No, the household carries out a market survey to find a firm that can deliver. It is an important detail here that this communication means that the firm is able to count how many potential customers it has. It can therefore recognize if there is a excess demand for its product.

2.1.2 Labor market

The households are active on the labor market between the age of 18 and 66. At age 67 it retires and lives only on profits. As mentioned a few times, we are making sure that the income cycle is as simple as possible in this version of the model. It will be relatively easy to introduce a private or public pension system into the model. In the same way, it will be relatively easy to introduce an income tax system that can finance unemployment benefits and public pensions.

The standard behavior of employed households is to work in the same firm as in the last period (month) and receive the wage that the firm now pays. With a low probability, the household carries out an on-the-job search. It does this by drawing m random firms. If one or more of the firms pay a higher wage than the current firm, it sends a CanIHaveAJob message to the firm. The firm immediately answers either Yes or No. The firm's considerations in this regard are described below in the section on the firms. If the new firm answers Yes, the household sends an IQuit message to the current employer. This tells the employer that the household is not employed from the start of the next period.

If the household is unemployed, it carries out a market survey very similar to the one above. The only difference is that it does not compare the potential firms' wages with an existing wage (the reserve wage is 0). Among the firms offering a job, the household chooses the one with the highest wage.

Note that here we have modeled a search labor market without using a matching function. As mentioned above, matching functions are often used in modern macro theory. One of the strengths of agent-based modeling is that search behavior can be modeled explicitly quite easily. This is also seen in many other agent-based models (Dawid & Della Gatti 2018; Gintis, 2007). If there are many positions offered, it will be easier for the individual household to find a job. On the other hand, if there are many households looking for a job, it will be more difficult for the individual household to find a job. This is exactly what the matching function describes in a macro model.

2.2 Firms

It is not obvious who determines the price in a competitive market. If all are price takers, who sets the price? We solve the problem by assuming that the firms are price setters and follow a principle called *competition-based pricing*. This is inspired by the marketing literature and the search-theoretical concept Competitive Search Equilibrium (Wright et al. 2021).

In marketing, a distinction is made between three basic price strategies: customer value-based pricing, cost-based pricing and competition-based pricing (see any marketing text book, e.g. Kotler et al., 2021). In the three strategies, emphasis is placed on customers, own costs or competitors.

Markup pricing is an example of cost-based pricing. In agent-based macro models and New Keynesian models, markup pricing is the typical assumption (Dawid & Della Gatti 2018). Reference is typically made to Dixit & Stiglitz (1977) who explain markup pricing as the result of monopolistic competition.

As mentioned, there is an alternative pricing strategy in marketing called competition-based pricing. Here, the firm sets its price *relative* to its competitors. A mature firm in a homogeneous market that sells according to its capacity will typically set a price corresponding to the market price, as an expression of the fact that it is satisfied with the status quo. The market price must here be understood as the average price on the market, as one can easily imagine that not all firms charge the same price. If the firm has excess capacity (corresponding to lower demand than desired), it may decide to set the price lower than the market price for a period of time (Kotler et al., 2021). It does this to attract customers. Similarly, excess demand may cause the firm to set the price higher than the market price for a period of time.

New firms will use special competition-based pricing strategies: skimming or penetration (Kotler et al., 2021). Skimming is relevant if the new product has novelty interest (think of a new Apple product). You will then typically set the price high relative to comparable products, and then let it fall over time. This skims profit. If we are dealing with a homogeneous product (as in our case) you will typically initially set the price low to penetrate the market.

In Hinterhuber (2008), managers are interviewed about their choice of pricing strategy[7]. It turns out that 44 per cent use competition-based pricing, 37 per cent use cost-based pricing and 17 per cent use customer value-based pricing. According to Hinterhuber, customer value-based pricing within marketing is considered to be the best method, but also the method for which it is most difficult to obtain data. Cost-based pricing is considered to be the worst method because attention to external conditions is too limited.

It is interesting to note that cost-based pricing does not entail interaction with other agents. This makes the strategy easier to work with - both for firms and model builders. This does not apply to the other two strategies. Competition-based pricing gives rise to an interesting social mechanism we will call *Copy-Deviate behaviour*. As mentioned above, a mature firm with full capacity will be satisfied with the state it is in and will therefore set its price at the market average price. It will *copy* the other firms. If it is dissatisfied with the state it is in (experiencing excess or under-demand),

[7]Hinterhuber describes his quantitative method as follows: "A sample of 126 marketing managers, business unit managers, key account managers, pricing managers, and general managers were initially recruited for this study. These managers participated in in-house pricing workshops which the author conducted in the period 2006-2007. Companies represented included automotive, chemicals, information technology (IT), chemicals, industrial services and fast moving consumer goods. We held nine workshops at nine different companies in Germany, Austria, China, and the USA. The study design is thus cross-sectional, multi-country, and multi-industry". For additional information see Hinterhuber (2008).

it chooses a price that is above or below the market's average price. It will *deviate* from the other firms. If everyone is satisfied, everyone copies each other and we are in an equilibrium. If many are dissatisfied, many will deviate from the average and the system will therefore move. Possibly towards a new equilibrium.

If a firm sets a price lower than the sector's market price in order to increase its market share, it is quite important what the other firms do. If few other firms undercut the market price, the firm under consideration will achieve its goal and gain market share. If many firms want to be larger (e.g. after a positive shock to the productivity of all firms), then many firms will lower the price and this will cause the market price to fall. We thereby obtain a (social) theory of the price dynamics at the macro level. We will see several examples of this below.

It is taken for granted in the marketing literature that the firm can attract or repel customers by deviating from the market price. This allows the firm to use the price to control the amount of demand. This is a different point of view than in standard economic theory where the firm under perfect competition takes the price as given, but can sell any quantity. There is a search-theoretical term *Competitive Search Equilibrium* (Wright et al., 2021) which can rationalize the point of view of the marketing literature. In the competitive models in Wright et al. (2021) there are a large number of buyers and sellers. The sellers sell a homogeneous product and each seller has a given capacity. Buyers look for the seller who charges the lowest price. A seller can therefore attract customers by setting a low price relative to his competitors. The individual seller therefore faces a downward sloping demand curve. This relationship between price and quantity is *not* due to buyers' preferences (as we usually assume in economic theory, e.g. monopolistic competition). The correlation is due to search friction. In the models described in Wright et al. (2021) a matching function is used to describe the search activity. In an agent-based model, as mentioned, it is possible to model search behavior explicitly. As described in section 2.1.1 it is assumed that the household searches for the firm with the lowest price. It is demonstrated in section 3 that the combination of this search friction and competition-based pricing helps to make the system self-organizing.

If all agents are the same, it is difficult to imagine a well-functioning social system based on Copy-Deviate behavior. Either everyone copies each other or everyone deviates from each other - in the same way. Heterogeneity is therefore a necessary building block in a social theory based on Copy-Deviate behavior. In our model, the firms are heterogeneous in three dimensions. First, there are new and mature firms. As described above, the new firms will tend to undercut the market price to enter the market (penetration strategy), while the mature firms will tend to follow the market price. Second, firms are assumed to have heterogeneous productivity. High- and low-productivity firms will have different incentives and it will especially be low-productivity firms that are at risk of default (i.e. exit). It is empirically well-founded that firms in the same industry can have widely varying productivity (Syverson, 2011). It is also empirically well-founded that firms in the same industry can have very varying sizes measured in terms of production or employment (Syverson, 2011). Finally, we assume that there are several sectors in the economy.

We make two key assumptions about the firm in our model. Firstly, that the firms have heterogeneous productivity and secondly, that the production function is *S-shaped*, in that there are increasing returns to scale at small production levels and decreasing returns to scale at high production levels. These two conditions are related in that the second condition explains the reasonableness of the first. It is difficult to explain the coexistence of high- and low-productivity firms in an industry with constant returns to scale and perfect competition. A firm with constant returns to scale can be of any size and there is therefore nothing to prevent the most productive firm from taking over the entire market. In recent theories of firm heterogeneity, the coexistence of high- and low-productivity firms is typically explained based on the combination of increasing returns and either monopolistic competition or diminishing returns to scale (Melitz 2003; Hopenheyn 1992). Under these circumstances, the individual firm will have an *optimal size* given the market's prices and wages. Highly productive firms will be large and low-productive firms will be small - but they will be able to exist at the same time without the low-productive firms being out-competed from the market. In Melitz (2003), the firm's optimal size is explained by the combination of increasing returns to scale and monopolistic competition. Increasing returns to scale means that for a given output price the firm will want to produce as much as possible. However, monopolistic competition means that the firm must lower the price to get rid of more goods. This means that the firm has an optimal firm size which i.a. depends on the size of the market. In Hopenheyn (1992), increasing returns to scale are combined with decreasing returns to scale, so that the firm has an S-shaped production function. For a given market price (i.e. under perfect competition) such a firm will have an optimal size. In Melitz, the firm's optimal size is *externally* explained (market conditions), while at Hopenheyn the explanation is *internal* (technology). In the real world, it is probably a combination of these two explanations that is relevant. We choose Hopenheyn's internal explanation. This makes it possible to model a market that in a Walrasian model would have perfect competition. Such a market can be considered the most generic and basic and a good starting point for an agent-based theory of the market. The next step is obviously to model the interaction between the agents that leads to imperfect competition, but it is by no means obvious how this should be done in an agent-based theoretical framework.

In the marketing literature and in the theory of Competitive Search (Wright et al., 2021), it is assumed that the individual firm has a given capacity or optimal size. In the marketing literature, this is due to fixed costs and diminishing returns to scale (Kotler et al., 2021), and in Wright et al.

(2021) it is an explicit assumption. Our assumption of an S-shaped production function ensures that the firm has an optimal size corresponding to a given capacity.

2.2.1 Expectations

The firm typically has to make decisions at the beginning of a month based on expectations for what will happen in the coming and subsequent months. We assume that the firm only projects one period forward and that it uses the simplest possible version of exponential smoothing (Brown, 1956; Hold, 1957; Hyndman et al., 2008). If a variable x_t is observed in period t, the expected value in the next period x^e_{t+1} is calculated by:

$$x^e_{t+1} = \lambda x^e_t + (1-\lambda)x_t. \tag{2.5}$$

The parameter λ is called the smoothing factor. We will use this method in many contexts as we thereby ensure that the agents do not overreact on new information.

2.2.2 Management

The management calculates at the beginning of each month what the firm's optimal size is given expectations for prices and wages. Each period, a central statistical unit in the economy calculates the average price level $\overline{P}_t$ in the sector and the average wage $\overline{W}_t$ in the economy. On the basis of these, the expected market price and market wage in the next period are calculated:

$$P^e_{t+1} = \lambda_P P^e_t + (1-\lambda_P)\overline{P}_t$$

$$W^e_{t+1} = \lambda_W W^e_t + (1-\lambda_W)\overline{W}_t$$

As in Melitz (2003) and Hopenheyn (1992) the production function is given by:

$$y = \varphi \cdot max\{l^\alpha - \phi, 0\} \tag{2.6}$$

where y is production, l is employment, φ is productivity, $\phi > 0$ is a parameter for increasing returns to scale and $0 < \alpha < 1$ is a parameter for decreasing returns to scale. In Melitz (2003) we have $\alpha = 1$ and in Hopenheyn (1992) we have $\alpha < 1$. We follow Hopenheyn and assume that $\alpha < 1$.

Given expected sector price P^e and market wage W^e, the firm wants to maximize profit:

$$\pi = P^e y - W^e l \tag{2.7}$$

If (2.6) is inserted into (2.7), it can be shown that the profit is maximized if:

$$l^* = \left(\alpha \frac{\varphi}{W^e/P^e}\right)^{\frac{1}{1-\alpha}} \tag{2.8}$$

It can be seen that the optimal employment is decreasing in the expected product real wage W^e/P^e and increasing in the productivity φ. High-productivity firms are larger than low-productivity firms.

The optimal production is given by:

$$y^* = \varphi\left((l^*)^{\alpha} - \phi\right) \tag{2.9}$$

If (2.8) and (2.9) are inserted into (2.7), it can be seen that the optimal profit is given by:

$$\pi^* = P^e\left((1-\alpha)\left(\alpha \frac{\varphi}{W^e/P^e}\right)^{\frac{\alpha}{1-\alpha}} - \phi\right)\varphi \tag{2.10}$$

The optimal profit is therefore increasing in the productivity φ. From (2.10) it can then be calculated that $\pi^* \geq 0$ if and only if:

$$\varphi \geq \hat{\varphi} \equiv a_{\varphi} \frac{W^e}{P^e} \tag{2.11}$$

where the constant a_{φ} is given by:

$$a_{\varphi} = \frac{1}{\alpha}\left(\frac{\phi}{1-\alpha}\right)^{\frac{1-\alpha}{\alpha}}.$$

The individual firm only has positive profit if it has a productivity φ that is higher than a cut-off value $\hat{\varphi}$. From (2.11) it can be seen that this cut-off value depends on the expected real wage. If the real wage grows, the least productive firms will lose their profitability. This is a central mechanism in the model (just as it probably is in reality).

2.2.3 Production

In a given period, production is determined by the production function (2.6) and actual employment l at the start of the period (measured in productivity units). Production is therefore known at the beginning of the period.

The firm maintains a list of employees. The j'th employee has an individual productivity ρ_j.

Employment is calculated as

$$l = \sum_j \rho_j.$$

Note that this means that high-productivity and low-productivity employees are fully substitutable. This means that it makes sense to talk about a wage per unit of productivity w. The total labor costs are given by wl.

The firm's production and actual sales will often not coincide. There can be both over- and under-demand. In agent-based models, this is often handled by assuming some form of storage (Dawid & Della Gatti 2018). This is seen in this model as an unnecessary complication. The important thing is that the firm can sense whether there is over- or under-demand on the product market. This is ensured in a different way in this model. You can think of the firm as a service firm. The output y is the maximum potential sales determined by the number of employees. If d is the actual demand experienced by the firm during the current period, the realized sales s will be given by:

$$s = min\left\{y, d\right\}.$$

If $d > y$, the firm experiences excess demand. If $d < y$, the firm experiences under-demand and has excess capacity.

2.2.4 Marketing department

The marketing department's task is to choose a selling price that retains and/or attracts customers, so that the demand corresponds to the optimal production defined by the management. In the section on the household, it was described that households react relatively slowly if a firm's price deviates from the sector's market average. The marketing department can therefore use the price as a steering device. If it wants more customers, it can for a period of time set the price lower than the market price. If, on the other hand, the firm experiences significant excess demand, it can afford to raise the price above the market average for a period. As mentioned above, this pricing strategy is called *competition based pricing* (Kotler et al., 2020) in the marketing literature.

The firm's expected demand d^e and sales s^e are given by:

$$d_{t+1}^e = \lambda_d d_t^e + (1 - \lambda_d) d_t$$

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

If $d^e > y^*$ the firm expects excess demand (y^*, the optimal production, is given by (2.9)). If $s^e < y^*$ the firm expects under-demand.

The firms's price rule can be described as follows:

$$p = \begin{cases} \left(1+\eta_p^U\right) P^e & \text{if } d^e > y^* \\ \left(1-\eta_p^D\right) P^e & \text{if } s^e < y^* \end{cases} \tag{2.12}$$

where p is the firm's price and P^e is the expected average sector price. The parameters η_p^U and η_p^D are respectively the markup and 'markdown'. If the firm expects excess demand, it sets the price higher than the market price to limit demand (and raise earnings). If the firm expects insufficient demand, it lowers the price to attract customers.

We assume that the size of the parameters η_p^U and η_p^D depends on the degree of disequilibrium. Excess demand is measured by:

$$z_p^d \equiv \frac{d^e - y^*}{y^*}$$

and under-demand is measured by:

$$z_p^s \equiv \frac{y^* - s^e}{y^*}.$$

We assume that the markup η_p^U is given by:

$$\eta_p^U = \begin{cases} \overline{\eta}_p^U \frac{z_p^d}{\overline{z}_p^d} & \text{for } z_p^d < \overline{z}_p^d \\ \overline{\eta}_p^U & \text{for } z_p^d \geq \overline{z}_p^d \end{cases}$$

so that a small deviation from equilibrium results in a small increase in price. If the imbalance is above a certain limit ($z_p^d \geq \overline{z}_p^d$) the fixed maximum markup $\overline{\eta}_p^U$ is used.

Defined in the same way:

$$\eta_p^D = \begin{cases} \overline{\eta}_p^D \frac{z_p^s}{\overline{z}_p^s} & \text{for } z_p^s < \overline{z}_p^s \\ \overline{\eta}_p^D & \text{for } z_p^s \geq \overline{z}_p^s \end{cases}$$

It is assumed that the firm only adjusts its price with a certain probability ω_p. With the probability $1-\omega_p$ the same price as last period is maintained.

In agent-based models, price rules reminiscent of (2.12) often appear. However, these rules are typically based on the last period's price and not the market price (Dawit & Delli Gatti, 2018; Gintis, 2007). In these models, the question is typically asked: should I raise or lower my price? In our model, another question is asked: should I do like the others or should I deviate? As mentioned, this gives rise to the social mechanism that we call Copy-Deviate behavior.

2.2.5 Human resources department

The HR department's task is to advertise positions and choose a wage that retains employees and, if necessary, attracts new employees. It is assumed that all employees in the firm receive the same wage per unit of productivity.

If the actual employment is higher than the management's desired employment (given by 2.8), the excess number of employees is fired according to the last-in-first-out principle. If, on the other hand, the actual employment is less than the management's desired employment, the HR department advertises vacancies. It is assumed that a proportion γ_l of the target vacancies are advertised:

$$v = max\left\{\gamma_l\left(l^* - l\right), 0\right\} \tag{2.13}$$

where v is newly advertised vacancies, l^* is optimal employment and l is actual employment. Note that v is measured in productivity units. The firm hires employees until employment measured in productivity units has grown v units.

In each period there is a number q who quit and a number a who apply for a job in the firm. We assume that the firms's expectations for these variables are given by

$$q^e_{t+1} = \lambda_q q^e_t + \left(1 - \lambda_q\right) q_t$$

$$a^e_{t+1} = \lambda_a a^e_t + \left(1 - \lambda_a\right) a_t$$

If $a^e - q^e < v$ is true in a given period, the HR department will expect that it will be difficult to attract a suitable number of new employees. It will then set the wage higher than the market wage in order to attract labor. If, on the other hand, $v = 0 < a^e - q^e$, the firm will have an incentive to lower the wage.

The firm's wage rule is given by

$$w = \begin{cases} \left(1 + \eta^U_w\right) W^e & \text{if } a^e - q^e < v \text{ and } v > 0 \\ \left(1 - \eta^D_w\right) W^e & \text{if } a^e - q^e > 0 \text{ and } v = 0 \end{cases}$$

Excess demand (from the firm's point of view) is measured by:

$$z^d_w \equiv \frac{v + q^e - a^e}{l^*}$$

and under-demand is measured by:

$$z^s_w \equiv \frac{a^e - q^e}{l^*}.$$

We assume that the markup η_w^U is given by:

$$\eta_w^U = \begin{cases} \overline{\eta}_w^U \frac{z_w^d}{\overline{z}_w^d} & \text{for } z_w^d < \overline{z}_w^d \\ \overline{\eta}_w^U & \text{for } z_w^d \geq \overline{z}_w^d \end{cases}$$

Defined in the same way:

$$\eta_w^D = \begin{cases} \overline{\eta}_w^D \frac{z_w^s}{\overline{z}_w^s} & \text{for } z_w^s < \overline{z}_w^s \\ \overline{\eta}_w^D & \text{for } z_w^s \geq \overline{z}_w^s \end{cases}$$

As above, it is assumed that the firm only adjusts the wage with a certain probability ω_w. With the probability $1-\omega_w$ the same wage as last period is maintained.

2.2.6 New and mature firms

In each period, a number of new firms are started. We describe in the next section how the investor decides how many firms to start. The j'th new firm born in period t is assumed to have the productivity:

$$\varphi_{jt} = \Gamma_t \vartheta_j,$$

where Γ_t is the productivity level for firms born in period t and ϑ_j is the firm-specific productivity. As in Melitz (2003) and Hopenheyn (1992), the firm-specific productivity is drawn from a Pareto distribution[8]. Annual productivity is assumed to grow exponentially as a reflection of ongoing technological progress:

$$\Gamma_t = (1+g)\Gamma_{t-1}$$

This specification provides a kind of "exogenous creative destruction": any firm, regardless of its productivity, will be outcompeted by newer firms in the long run.

Productivity is only recognized after the start of the firm. This introduces an important element of uncertainty in the decision to start a new business, cf. the next section on the investor.

New firms have *entry costs*. It is assumed that the firm must hire a number of employees $\overline{n}$ for $\overline{T}$ periods. This is called the start-up period. Entry cost is the wage cost for these employees. The parameters $\overline{n}$ and $\overline{T}$ are the same for all firms and no production can be done in the $\overline{T}$ periods.

After the start-up period, the firm starts hiring new employees as described above. The firm will typically set a price lower than the market price in order to enter the market (penetration strategy).

[8]The Pareto distribution has been chosen to ensure comparability with the literature. We could easily use another distribution, e.g. a log-normal distribution.

If there are not enough applications, the firm may also set a wage higher than the market wage in order to attract employees.

After a while, the firm will have attracted the optimal number of customers and employees. If in this situation it earns a positive profit, the investor will let it live. The firm is then called a mature firm. If the firm does not make a profit at the current market price and market wage, it is closed.

2.3 The investor

The investor makes a living by starting firms, and you can think of her as a venture capitalist or 'angel investor'. The investor finances the firms's losses in the future and receives all profits. Investor can close the firm at any time. If this happens, all employees are fired. It is assumed that if the firm cannot earn a positive profit at the optimum at the current expected market price and wage, then the firm will be closed with a given probability. This specification implies that the firm will be closed after some months if it starts to generate losses.

The financial transactions between investor and firms can be seen as a simple credit market. This is the only credit given in the economy, since households consume all of their income.

2.3.1 Measurement of profitability

Because there is uncertainty about the new firms' productivity, the investor will buy a portfolio of firms (as in Melitz (2003) and Hopenheyn (1992)). The firms in the portfolio will both be from different sectors and have different productivity. This section deals with how the investor measures profitability. In the next section, this is used to decide how big the portfolio should be and how it should be distributed among sectors.

The investor makes decisions under fundamental uncertainty. We assume that the investor does not have access to knowledge about what will happen in the future. It is therefore necessary to use current and historical data.

Ideally, the investor should do a full cohort analysis for the given sector. Let $\left(\pi_{jt}, a_{jt}\right)$ be the profit and age of the j'th firm at time t. Assume the investor at time t has started n_t new firms and that the new firms are the first in the list of all firms ($j \in \{1,..,n_t\}$). The investor's expected, discounted real income is then given by:

$$E_t\left[\pi\right] = \sum_{j=1}^{n_t} \sum_{s=t}^{\infty} \frac{\pi_{js}}{P_s^C} \left(\frac{1}{1+r}\right)^{s-t}$$

where r is the investor's discount rate and P_t^C is the consumer price index in period s. It will typically apply that π_{jt} is negative at the start (entry cost), but for the successful firms profits increase in the long run. Note that if $E_t[\pi] = 0$, the return of the portfolio is r, whereas if $E_t[\pi] > 0$, the portfolio gives a better return than r. We will say that the portfolio earns a *normal return* if $E_t[\pi] = 0$.

The investor cannot calculate $E_t[\pi]$ because the future is unknown. An estimate of the sector's profitability can be obtained by looking at the *current* cross-section:

$$\hat{E}_t[\pi] = \sum_{j=1}^{N_t} \pi_{jt} \left(\frac{1}{1+r}\right)^{a_{jt}},$$

where $N_t > n_t$ are all existing firms at time t. Here, the following question is asked: if the current correlation between profitability and age is maintained in the future, how is the overall profitability of the portfolio? This can be seen as a measure of the sector's current profitability, as it is recognized that it may take many years before positive earnings from the mature firms are achieved. In a cross-sectional analysis, it is okay to disregard inflation if the analysis is carried out under the assumption that the market price and wage (which is included in the profit) grow corresponding to the consumer price index. This is a reasonable assumption in the closed system that we consider.

The size $\hat{E}_t[\pi]$ is an absolute measure and therefore difficult to relate to anything. We choose to normalize the expected, discounted profit with the standard deviation of the expected, discounted profit:

$$\hat{s}_t = \frac{\hat{E}_t[\pi]}{\sigma_t}$$

where σ_t is given by

$$\sigma_t^2 = Var\left[\pi_{jt}\left(\frac{1}{1+r}\right)^{a_{jt}}\right]$$

The value $\hat{s}_t$ is a so-called Sharpe Ratio (Sharpe, 1994) that measures a portfolio's return compared to a risk-free return, adjusted for risk[9]. We assume that the investor bases her actions on an expected Sharpe ratio:

$$s_{t+1}^e = \gamma_s s_t^e + (1-\gamma_s)\hat{s}_t$$

In what follows, we will call s_t^e the *expected profitability* in a given sector.

2.3.2 Portfolio size and distribution

The investor must choose how many new firms to invest in. As mentioned above, it is impossible to know in advance what productivity the individual firm has. First the firms are started and then the productivity is percieved (drawn from a Pareto distribution). We divide the problem into two parts: the size of the total portfolio and its division into sectors

[9] In our case, it is not a risk-free return, but the investor's discount rate.

We assume that the investor gradually adapts to an appropriate total portfolio size based on expected average profitability for all sectors s^e_{t+1}. If it is expected that the total portfolio will earn a normal return ($s^e_{t+1} = 0$), the same portfolio size as last period will be maintained. If the portfolio is expected to give an above-normal return ($s^e_{t+1} > 0$), the size of the portfolio is increased, etc. This behavior is achieved by assuming that

$$\frac{\Delta n_{t+1}}{n_t} = \xi \cdot s^e_{t+1}, \xi > 0, \tag{2.14}$$

where Δn_{t+1} is the change in portfolio size.

There are two things to comment on regarding the behavioral relationship (2.14). Firstly, the profit measure s^e_{t+1} is a relatively advanced, risk-corrected quantity based on potentially large amounts of data. In agent-based models, it is possible to let the agents perform complex analyzes on the model-generated data[10]. This is an example of this. Corresponding analyzes would be hard to perform in a classical macro-model. In such a model it would be necessary to make simplifying assumptions that ensure that aggregation is possible. Secondly, it should be emphasized that the interaction between investor behavior (2.14) and the remaining economy gives rise to an extremely robust overall system. In an economy with S-shaped production functions, it is to a large extent the number of firms that is responsible for scaling the total macro-production up and down (the external margin). The rule (2.14) does not in itself contain any information about the optimal size of the portfolio n_t, but it ensures, in interaction with the remaining system, that firms are started when it is worthwhile. This can be seen as a very central element in the system's market mechanism.

Given the size of the total portfolio n_{t+1}, the investor must choose the allocation to sectors $n_{j,t+1}$. We assume the following logit-like specification:

$$n_{j,t+1} = \frac{n_{jt} e^{\kappa \cdot s^e_{j,t+1}}}{\sum_i n_{it} e^{\kappa \cdot s^e_{i,t+1}}} n_{t+1}, \kappa > 0. \tag{2.15}$$

An essential property of this specification is that if $s^e_{j,t+1} = s^e_t$ for all j, then it applies that $n_{j,t+1}/n_{t+1} = n_{jt}/n_t$. If all sectors have the same expected profitability, then they all maintain their share of the portfolio. If a sector has a higher expected profitability than the other sectors, then its share of the portfolio will grow. A description of this portfolio specification can be found in Stephensen (2021).

[10]On the other hand, it is almost by definition impossible to assume perfect foresight or rational expectations.

2.4 Equilibrium and self-organization

Given the sub-elements of the system, we can seek to characterize an equilibrium of the system. This will later make it easier to interpret what is going on in the model.

We consider a situation where all firms in a sector charge the same market price and pay the same market wage. The investor opens a constant number of new firms in all sectors each period and a corresponding number of firms close each period because they have been outcompeted. The firms have therefore converted to an equilibrium state where the total number of firms is constant and where there is a fixed age structure. This is a relatively complicated steady-state structure.

In the labor market, we assume that there is constant structural employment, as search behavior creates a constant structural unemployment rate (this is supported by the results in section 3).

We saw in section 2.2.2 about firms' choice of optimal production and employment, that the individual firm has an optimal size given the real wage. Higher real wages result in lower optimal employment. If the investor opens many new firms (so that the total number of firms becomes high), many firms will have to share the structural employment that exists. Therefore, the real wage will be high in equilibrium in order to reduce the individual firm's employment.

A higher real wage means a lower profit. Total discounted profits (incl. entry cost) will therefore be low, so that the expected profitability will be low. *There is therefore a negative correlation between the number of new firms in each period and the expected profitability*. If there is an inflow of new firms that generates an expected profitability of zero, we are in equilibrium. We will see in the next section that this is typically the case.

The negative correlation between the inflow of new firms and expected profitability is a central source of self-organization. If the investors increase the inflow of firms in order to chase short-term profit, this will reduce the profit in the long term. If the investor has an inflow of firms above the equilibrium level for a longer period, she will experience that investments do not earn a normal return (negative expected profitability). This will according to the behavior described in section 2.3.2 cause the investor to reduce the inflow of new firms and therefore move back to the equilibrium level. We will see several examples of this in the following sections.

3 Results

The model is run in a base run and three supply shocks. The goal is to demonstate that the system is self-organizing. We will see that the overall system tends to oscillate around a steady state. When you shock the system, it returns to a state that is relatively close to the original one.

Two different permanent shocks to productivity and a temporary shock to the number of firms are performed. If productivity in all sectors increases permanently, one would typically expect total production to rise and real wages to grow. It is demonstrated that this is the case. In the temporary shock, a random proportion of the firms are closed. It is demonstrated that the system is self-organizing in the sense that roughly the same number of firms are opened in the long term. Finally, we shock productivity in just one sector and compare with the effects in Acemoglu and Guerrieri (2008).

The model is run for a population of approx. 50,000 households and approx. 2,500 firms. There are 10 sectors. The model runs for 150 years divided into months. Base run and shocks are run 150 times so that mean values and confidence intervals can be calculated.

The base run begin with three so-called burn-in periods. The first burn-in period lasts 20 years and it is assumed that there is a constant inflow of new firms and that no firms close. The second burn-in period lasts 15 years. Here it is assumed that there is a constant inflow of new firms, but that the firms are closed if they are not profitable. Finally, the third burn-in period lasts 50 years. Here the model runs freely, and it is thus assumed that the investor decides the number of new firms. The purpose of the third burn-in period is to allow the system to converge to a steady state. After the three burn-in periods, the model runs for 150 years.

Due to search activity and income dynamics, there is stochasticity in the model. As mentioned above, the model is run 150 times to assess the significance of this stochasticity. Stochastic is based on the random generator in the computer language C#. Such a generator forms from an integer (the 'seed') a uniquely infinite series of pseudo-random numbers. The counterfactuals are constructed by starting from the same seed as in the base run.

The most important of the model's parameters are shown in appendix A. The model's parameters are not calibrated to a specific data base. As mentioned, the model is intended as a principled example. A number of robustness tests have been carried out and the model is relatively robust. It is important for the convergence of the model to a steady-state that both prices and wages comply with the competition-based pricing strategy. The model will therefore be sensitive to alternative assumptions about the wage formation. This is an obvious topic for future research.

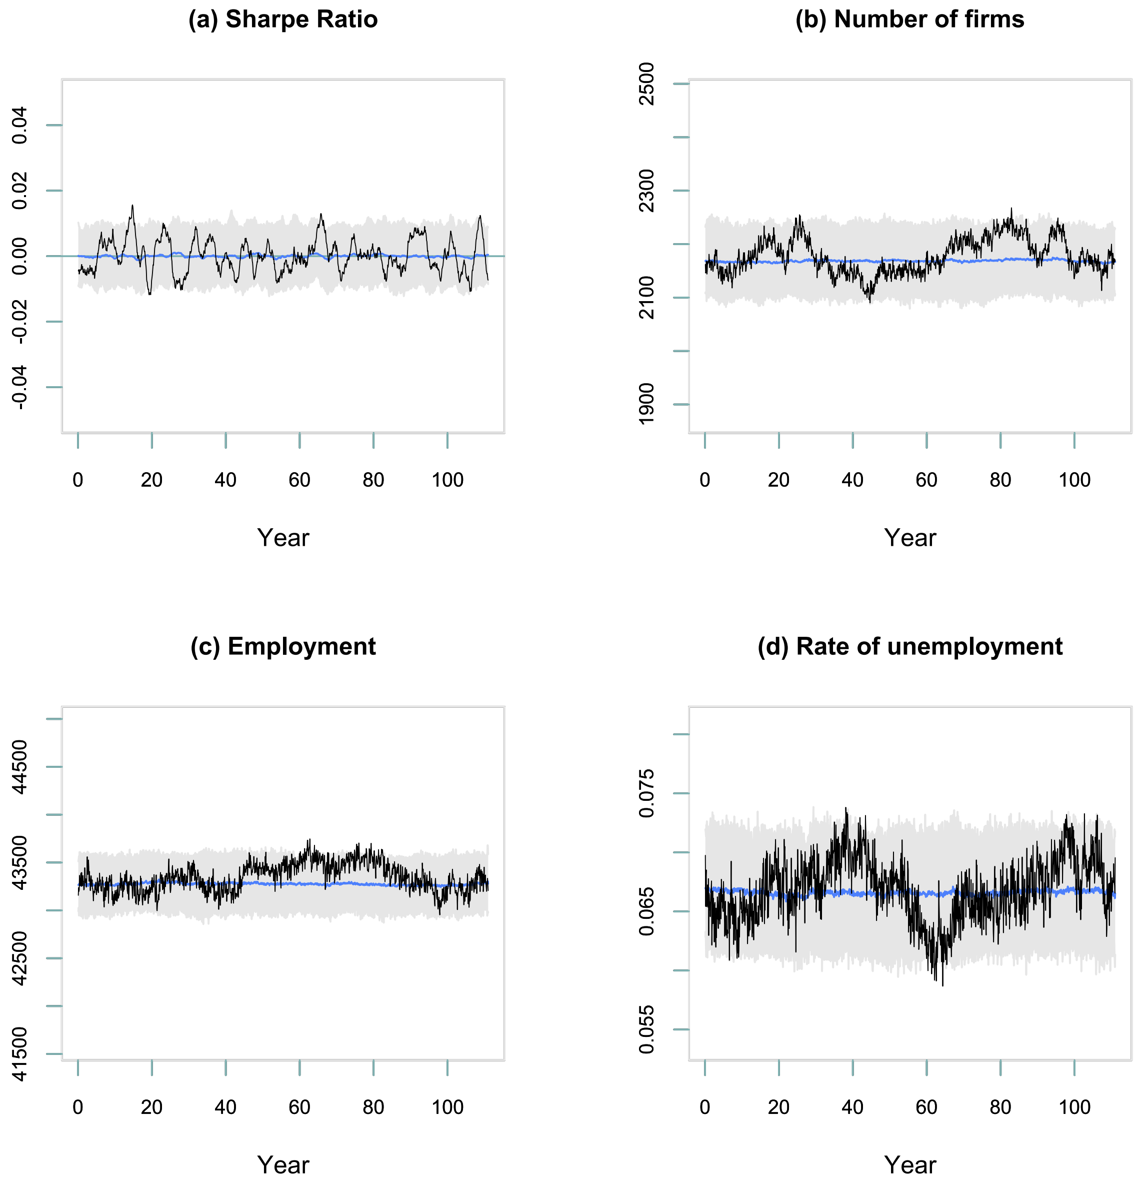

Figure 3.1: Base run. 110 years of monthly data.

The blue curve averages 150 runs. The black curve is an example of a specific run (the first). The shaded areas marks 95 percent confidence intervals.

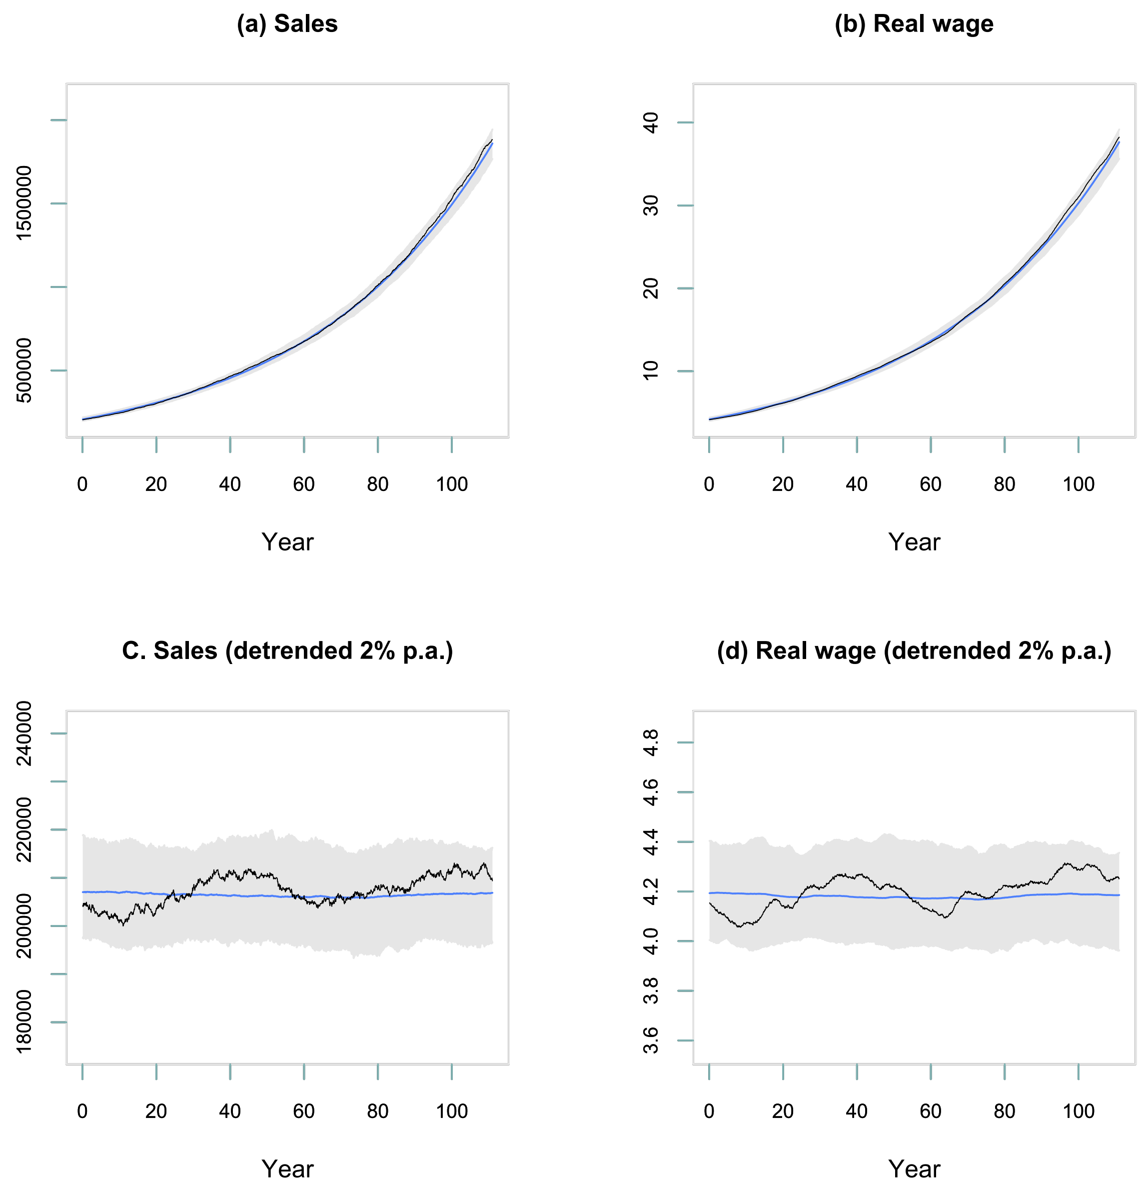

Figure 3.2: Base run. 110 years of monthly data.

The blue curve averages 150 runs. The black curve is an example of a specific run (the first). The shaded areas marks 95 percent confidence intervals.

3.1 Base run

The last approx. 100 years of the basic runs are shown in Figure 3.1 and 3.2. In each of the 8 sub-figures, the blue curve shows the mean of the 150 runs and the gray area shows the 95-percent confidence intervals. The black curve is an example of a random run (the first one).

The average over 150 runs (the blue curve) and the confidence intervals (the gray areas) can be seen as a description of the underlying structural properties of the system. It can be seen from the blue curves that the expected profitability and employment show significant stability. It can be seen for all variables that the widths of the confidence intervals are relatively stable. The confidence intervals are remarkably symmetrical.

In Figure 3.2 it can be seen that sales and real wages grow over time. This demonstates that the system's market mechanism works and that the system works as expected. Productivity of new firms grows by 2 per cent. per year. New firms will therefore slowly but surely outcompete mature firms. The growing real wage drives low-productivity firms out of the market so that average productivity increases. It can be seen from the Figure 3.1 that employment is stable. The higher productivity therefore means higher production. Higher production can be demanded precisely because the real wage grows for more or less constant employment.

In Figure 3.2 sales and real wages are shown detrended by 2 per cent. per year. It can be seen that the mean values (the blue curves) are stable, corresponding to sales and real wages growing by approx. 2 percent per year.

If the actual developments in Figure 3.1 and 3.2 are considered (the black curves), it is seen that the variables fluctuate around a level or a trend. The expected profitability (Figure 3.1(a)) fluctuates between approx. -1 per cent and 1 per cent and is thus relatively close to 0. This is a fundamental source of the system's self-organizing properties. The unemployment rate (Figure 3.1(d)) fluctuates between 6.0 per cent and 7.25 per cent with an average of 6.7 per cent. In general, cyclical fluctuations are seen in the system, but the deviation from the average is quite limited. If there were millions of households and 100,000s of firms, the cyclical fluctuations would be almost invisible. The conclusion is that the system is capable of generating business cycles, but that these are very limited in a model with competition-based pricing.

In Figure 3.3, cross-sectional output is shown in a random year in the first base run. The top left figure shows optimal and actual employment. As described in section 2.2.6, a firm must hire $\overline{n}$ people for $\overline{T}$ months before the firm starts (this is the entry cost). It is assumed that $\overline{n} = 10$ in the runs carried out. This explains that some firms have an optimal employment of 10. Most of the remaining more than 2 thousand firms lie close to the 45-degree line, corresponding to the firms' actual employment being close to the optimal. This shows that the market mechanism on the labor market is functioning well, as the firms are able to attract the necessary number of employees. In

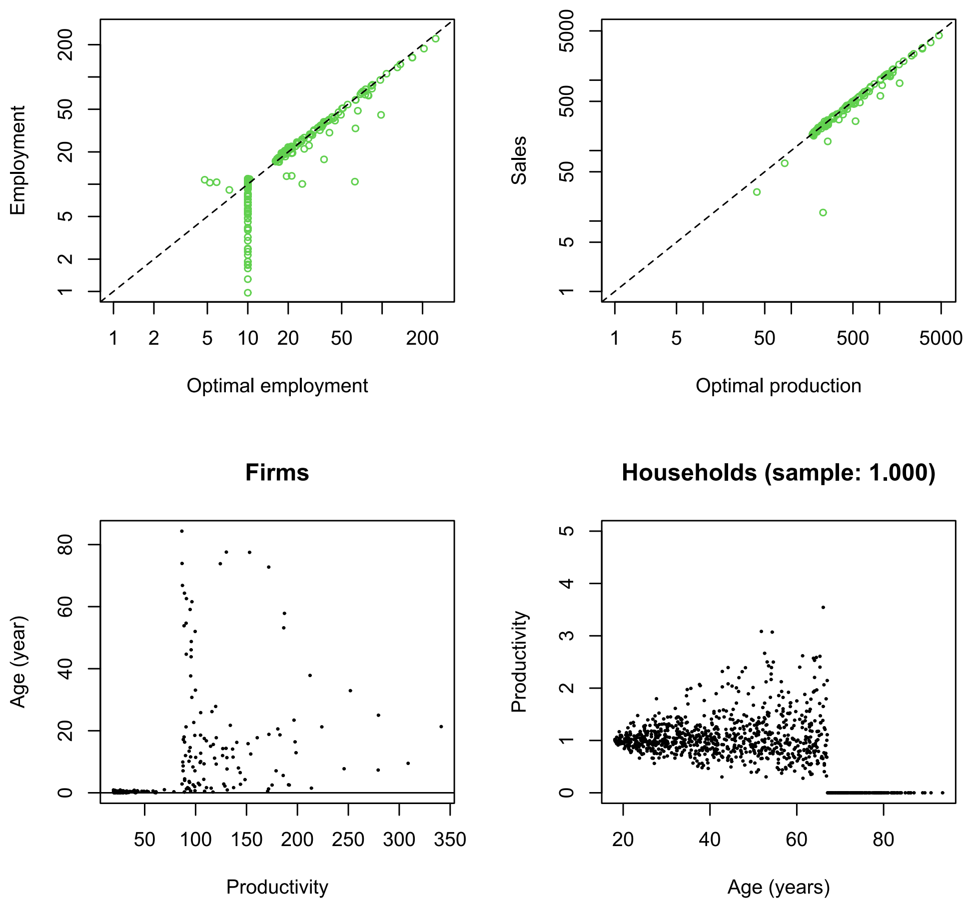

Figure 3.3: Cross section output. Random baseline year. Samples of 250 firms and 1.000 households.

the same way, it can be seen in the upper right figure that the actual and the optimal production are close to each other for most firms. This shows that the market mechanism is functioning well on the goods market, as the firms are able to attract the necessary number of customers.

In the lower left figure in Figure 3.3, the productivity and age of all firms can be seen plotted against each other (sample of 250 firms corresponding to approx. 10 per cent). It appears that the oldest firm is over 80 years old and that there is a cut-off value for productivity as demonstrated in section 2.2.2. A firm is only allowed to survive if its productivity is higher than this value. The cut-off value grows over time due to technological progress. The lower right figure in Figure 3.3 shows a sample of 1,000 households (out of approx. 50,000). The relationship between age and individual productivity is shown. As described in section 2.1, the individual household's log productivity is a random walk over life. Individual productivity will therefore be log-normally distributed with a variance that increases over life. At age 65, the household retires and productivity falls to 0.

3.2 Counterfactuals

Three shocks are performed to the system. A permanent shock to all sectors' productivity, a shock to one sector's productivity and a shock to the number of firms. The shocks are carried out on the 150 base runs described above. The average effects and confidence intervals are shown for the first 10 years in the Figure 3.4 and 3.6. In Walrasian equilibrium models, there will often be an element of black-box due to simultaneity. This is to a lesser extent the case in well-constructed agent-based models, where in principle everything can be be explainable. The assumption of competition-based pricing makes it relatively easy to explain the price and wage dynamics of the model. In order to demonstrate this, some of the explanations below are relatively detailed.

Productivity shock

The effect of a 10 per cent increase in productivity can be seen in Figure 3.4. The increase in productivity corresponds to an innovation that breaks through in all firms at the same time. Overall, the economy reacts relatively quickly as a new equilibrium is reached after 4-5 years. In the long term, the number of firms and employment is largely unchanged, while real wages and sales have grown approx. 10 per cent. This is completely as expected and emphasizes, just like the base run, that the system is self-organizing.

In the short term, the increase in productivity leads to an upswing of a few years, followed by a minor setback. On impact, firms can suddenly produce 10 per cent more. The supply immediately grows by 10 per cent whereas the demand (sales) is unchanged (Figure 3.4 (d)). There is therefore a significant over-supply on the goods market. This causes the individual firm to lower the price

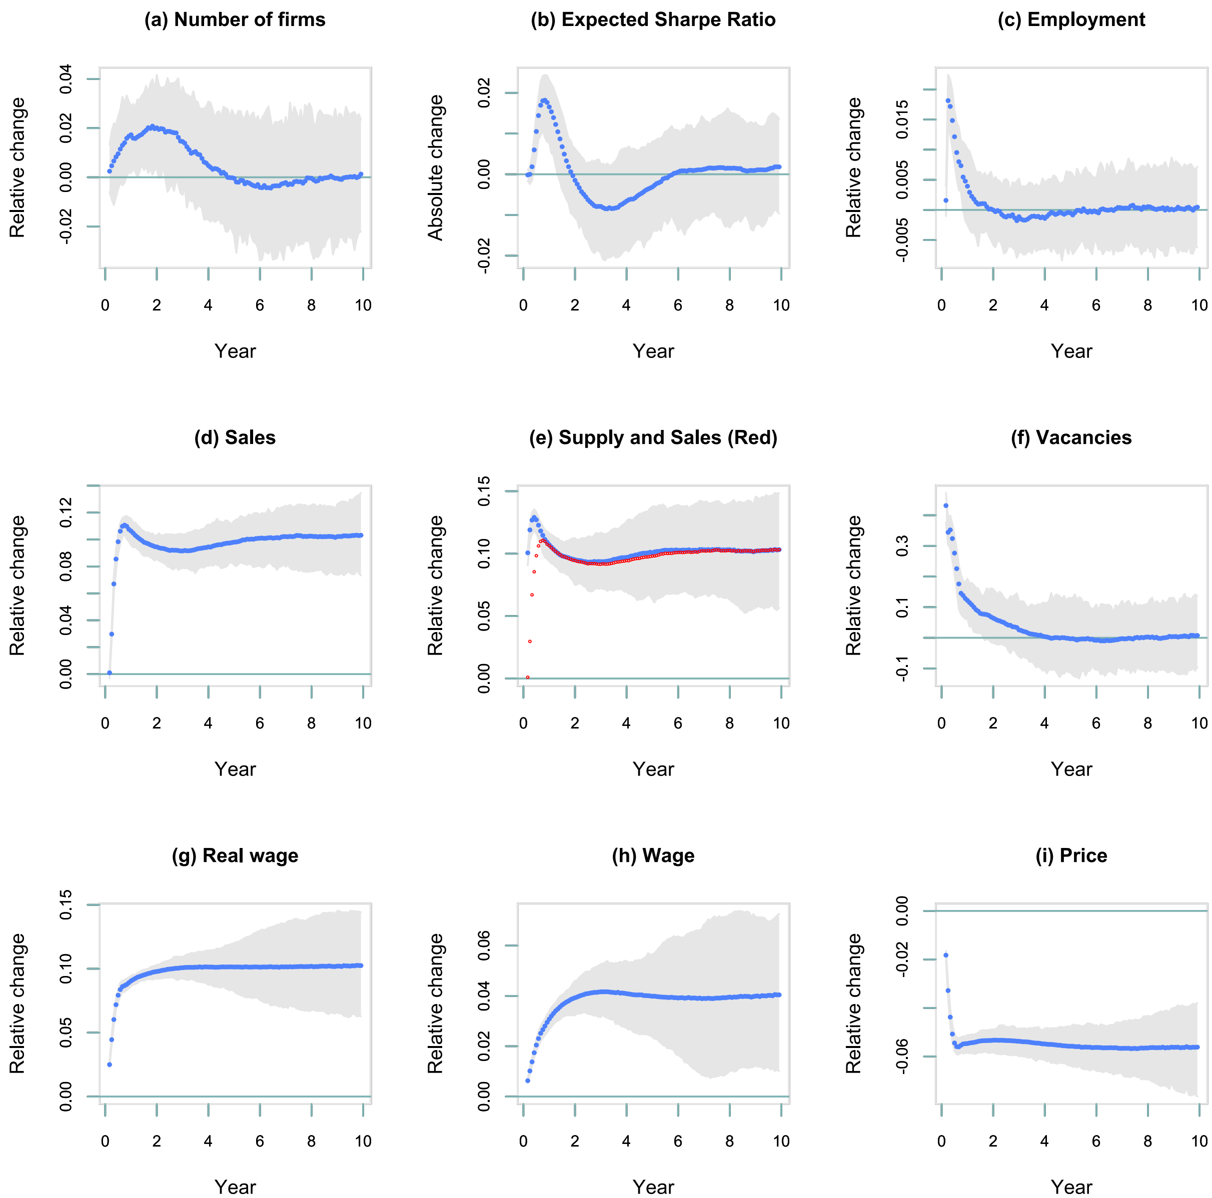

Figure 3.4: 10 pct. shock to productivity. 10 years of monthly data.

The blue points averages 150 runs. The shaded areas marks 95 percent confidence intervals.

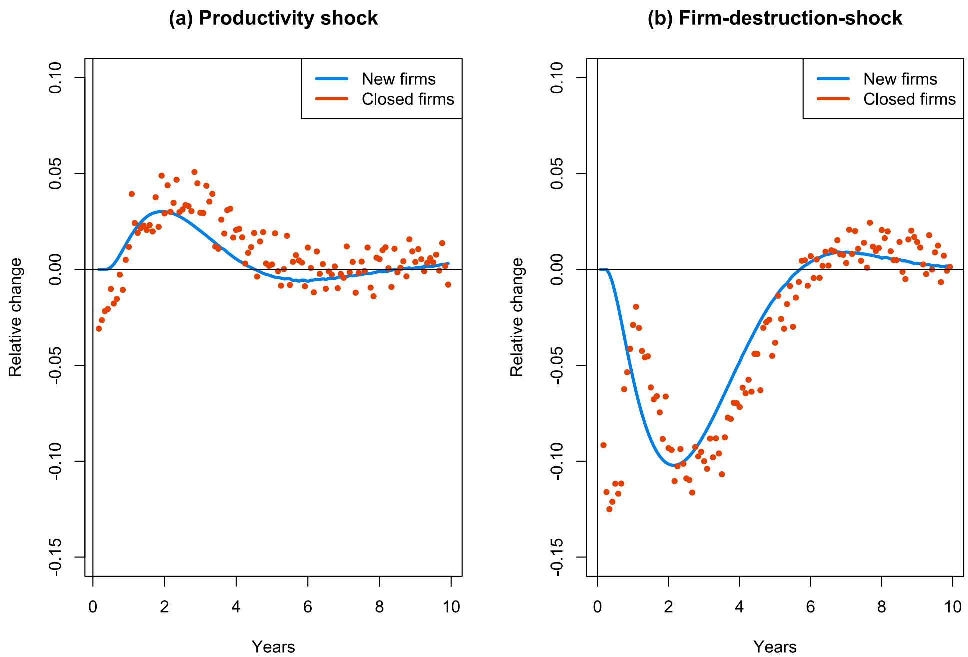

Figure 3.5: Effect on number of new and closed firms. Averages of 150 runs.

to attract customers. This drives the market price down. It can be seen from Figure 3.4 (i) that the market price falls rapidly in the first 4 months by approx. 5.5 per cent.

A higher productivity will, for a given real wage, result in the individual firm having a higher optimal employment and production (this is evident from the equations (2.8) and (2.9)). This can be seen as a short-term effect on the intensive margin. It can be seen in Figure 3.4 (f) that the number of vacancies grows sharply immediately. After one month, a similar increase in employment is seen. Some firms start raising wages to attract employees. This drives up the market wage (Figure 3.4 (h)). The increase in the market wage is more moderate than the fall in prices, with the combined changes in price and wages resulting in real wages growing by almost 9 per cent during half a year (Figure 3.4 (g)). This increase in the real wage largely neutralizes the effect on the intensive margin, and within a year employment is back at its original level and supply is close to an increase of 10 per cent that can be explained by the productivity increase.

As described above, the permanent increase in production leads to an upswing followed by a minor slowdown. The reason for this is the sluggish adjustment of prices and wages - just as it is known from New Keynesian models. If prices and wages did adjust immediately so that the real wage grew by 10 per cent, optimal output would grow by 10 per cent, and the optimal employment would be unchanged. As a result, the expected profitability would be unchanged. Therefore, the investor would not increase the number of firms and we would not see an effect on the external margin. Inertia in price and wage formation means that productivity increases profits in the short term. This means that expected profitability will be positive in the short term. This causes the investor to open new firms and we get an effect on the external margin. It can be seen in Figure 3.4 (a) that the expected profitability is positive in almost 2 years. After two years, however, the real wage adjusts to its long-term level. This puts pressure on the profitability of the firms so that the expected profitability will be negative in the following almost three years. This results in a decreasing number of firms and employment that is (slightly) below its structural level.

Figure 3.5 (a) shows the background for the development in the number of firms. The number of new firms (the blue curve) moves smoothly because it depends on the expected profitability, which is basically a macro variable. It can be seen that the number of closed firms (the red dots) initially decreases. This is because the higher productivity for an unchanged real wage makes some of the low-productivity firms survive. In the first 2 years, therefore, more firms are opened than are closed. The number of firms is therefore increasing. This puts pressure on expected profits and causes the number of new firms to fall again. After the first 2 years, more firms are closed than are opened. Several firms close because "too many" were opened at the beginning. In the long run, the number of new and closed firms converges. The system moves into equilibrium.

Firm destruction

In this shock, 10 per cent of the firms are closed. The closed firms are randomly selected and at the time of the shock all employees are fired and the investor loses her investment. The shock can, for example, be seen as a natural disaster without loss of human life. The interesting question is whether the system is sufficiently self-regulating for the firms to be restarted. The answer is yes. After 6 years, the number of firms has been re-established. On the other hand, it takes over 30 years before the original productivity and real wages are restored. This is because it takes many years to recreate the original firm structure (age distribution of the firms).

The effect of the shock in the first 10 years is shown in the Figure 3.6. The first 30 years are also shown in appendix B. It can be seen that the number of firms, employment, sales, supply and vacancies all fall by 10 per cent in the first period (a, c, d, e, f in Figure 3.6). The fall in employment means that the number of unemployed is growing considerably. After just one month, the firms see a large increase in the number of job applications. This explains why the firms are cutting wages. It appears from Figure 3.6 (h) that after 4-5 months the market wage has fallen by 2 per cent and after a year the wage has fallen by almost 4 per cent. The effect on the goods market is less. In Figure 3.6 (e) supply and sales are shown. The fact that demand (sales) falls by 10 per cent and remains low for 4-5 months is due to our assumption that households consume their entire income and that

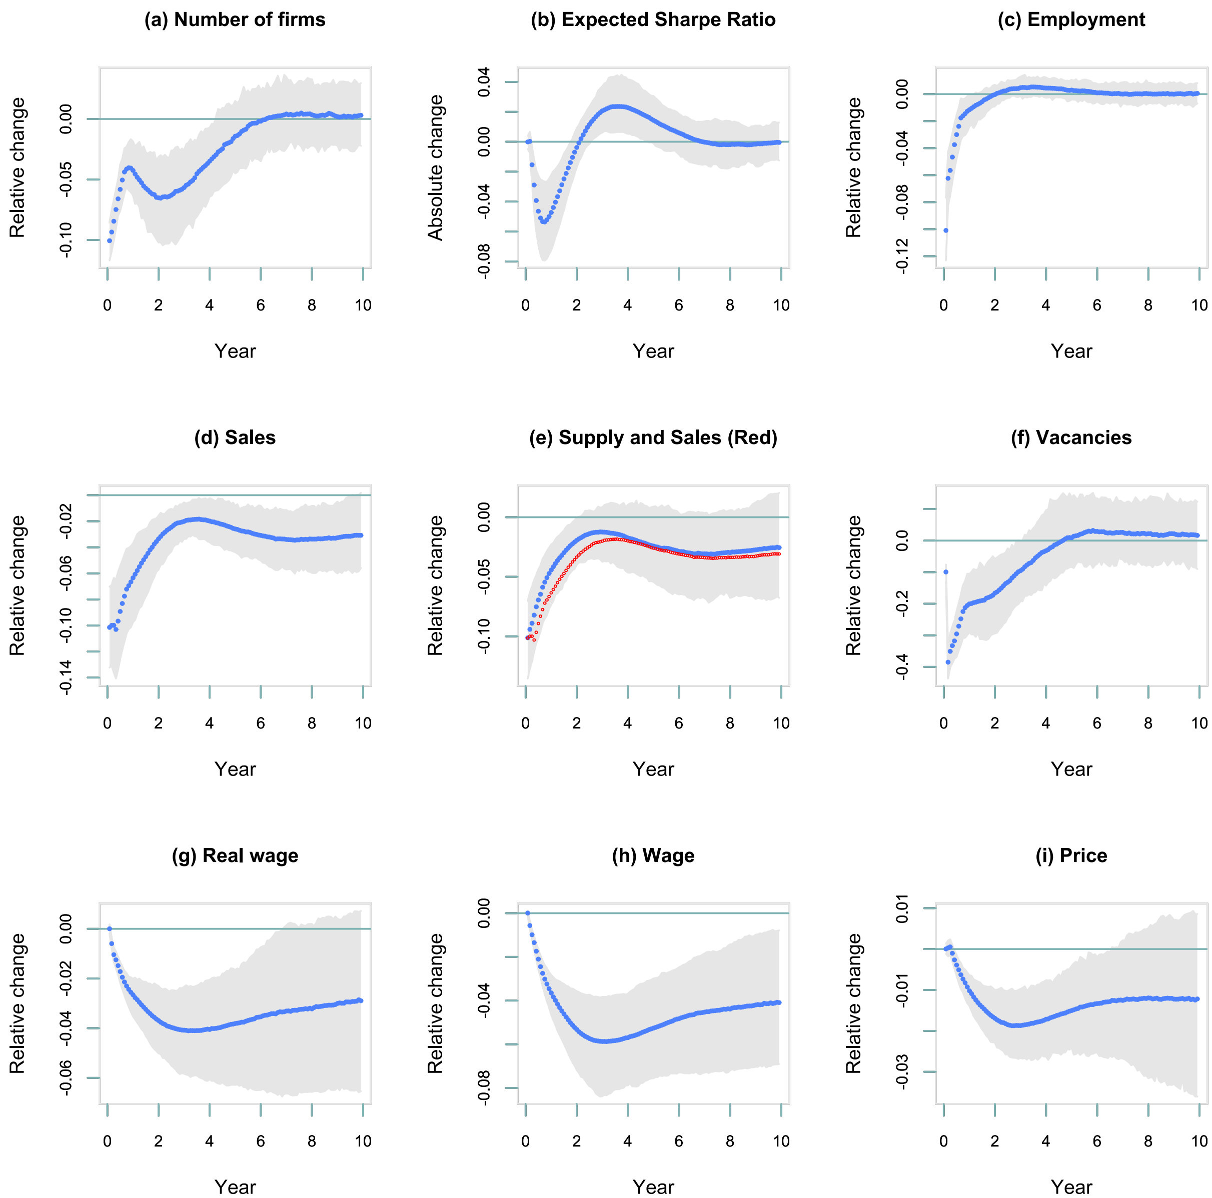

Figure 3.6: Destruction of 10 pct. of firms. 10 years of monthly data.

The blue points averages 150 runs. The shaded areas marks 95 percent confidence intervals.

there is no unemployment benefit. The model does not have any automatic stabilizers[11]. It can be seen from Figure 3.6 (e) that there is over-supply on the goods market throughout the process, which explains that the firms lower prices to attract customers. The market wage falls more than the market price, which explains that the real wage in Figure 3.6 (g) falls.

The many layoffs in a system without a safety net send many into poverty and demand is hit hard. The system is going through a crisis with unemployment, falling profits and low capacity utilization. However, the crisis is very short. Already after six months, employment has largely been restored (Figure 3.6 (c)). This is because around half of the closed firms have been replaced after six months (Figure 3.6 (a)) and that real wages are falling. Both the external and the internal margin react relatively quickly. The reason why the number of firms grows significantly in the first six months is shown in the Figure 3.5 (b). In the short term, the number of closed firms falls by approx. 10 per cent. This is because the number of firms has fallen by 10 per cent. However, the investor only changes the number of new firms slowly. Therefore, in the first half of the year, more new firms are started than are closed. Therefore, the number of firms is growing.

Figure 3.6 (b) shows that the expected profitability is negative for the first 2 years. The negative profitability is due to the fact that the real wage reacts sluggishly and that the firms experience an over-supply and therefore only make money on part of their capacity. As long as the expected profitability is negative, the number of new firms falls (the blue curve in Figure3.5 (b)). This is due to investor's behavior according to equation (2.14).

Productivity shock in one sector

In the last shock, we analyze the effect of changing productivity in one sector. We have balanced growth in the model, but this shock hits the sectors asymmetrically and it is interesting to see if we get results that are consistent with the literature on unbalanced growth (Baumol, 1967; Acemoglu and Guerrieri, 2008). Acemoglu and Guerrieri (2008: 469) describes a two-sector model where the sectors' productivity growth is different. They show that if the demand-side elasticity of substitution between the two goods is less than 1, then the price and employment in the high-productivity sector will grow more slowly than in the low-productivity sector. Employment will move into the low-productivity sector. This is a technological explanation of the historic pattern that employment in agriculture moved into manufacturing, and - later - that employment in manufacturing moved into the service sectors. It is a central assumption that the elasticity of substitution between the two goods is less than 1, as this implies that the two goods are difficult to substitute for each other for the demanders.

[11] As previously mentioned, it will be straight forward to create an unemployment benefit scheme, but the idea with this model is to model the markets as cleanly as possible.

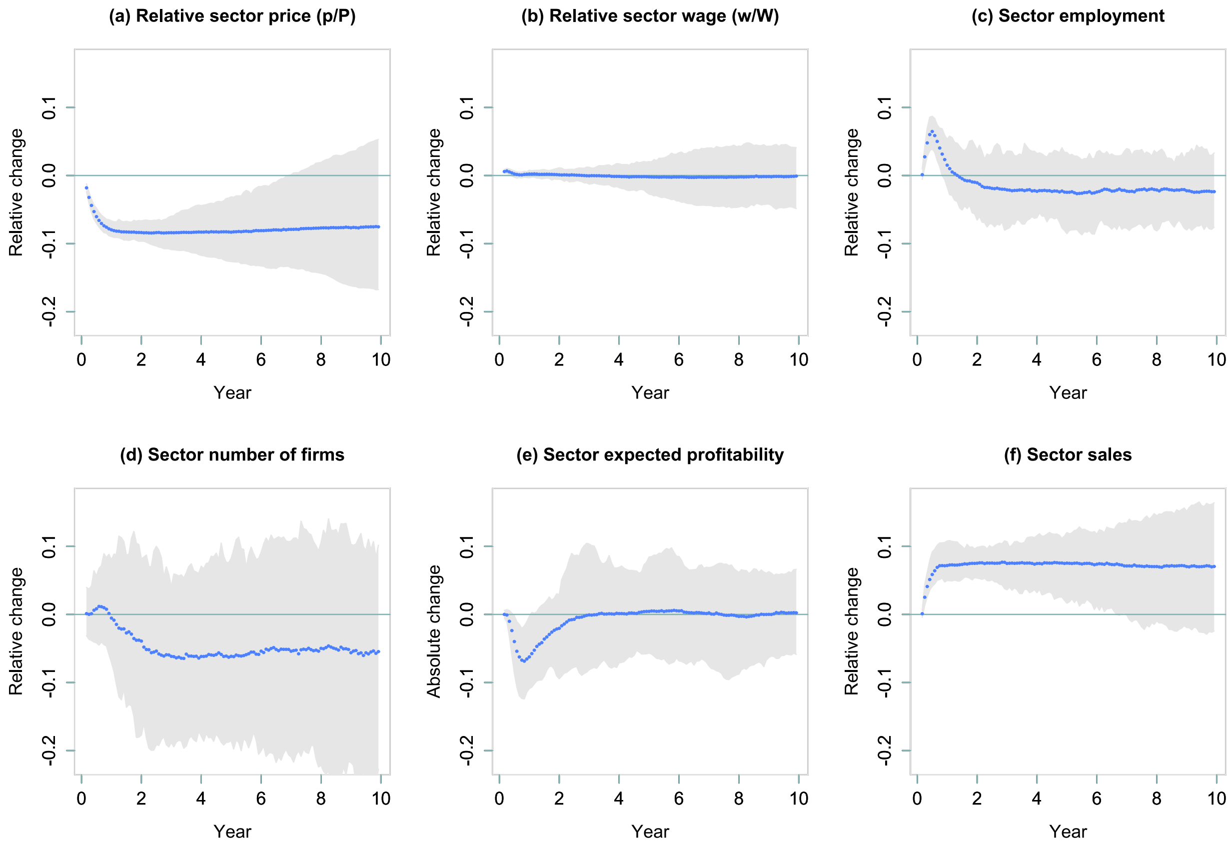

Figure 3.7: 10 pct. shock to productivity in one sector. Results for the affected sector. 10 years of monthly data.

The blue points averages 150 runs. The shaded areas marks 95 percent confidence intervals.

We permanently increase the productivity of new firms by 10 per cent in one of the sectors of the economy. The sector therefore has higher productivity growth than the other sectors in a single period. As mentioned earlier, we assume that the households' elasticity of substitution between all the sectors' goods is less than 1. We therefore also get long-term results that are consistent with Acemoglu and Guerrieri (2008). Figure 3.7(a) shows the considered sector's price relative to the economy's average price p/P. In the same way, the sector's relative wage w/W is shown in Figure 3.7(b). The macroeconomic effects of the shock are shown in appendix C. It can be seen that the sector's price falls relative to the other sectors' prices and that the relative wage is largely unchanged. In the long term, it can be seen that employment and the number of firms fall (Figure 3.7(c)-(d)). Since the economy's total employment is not affected in the long term (see Appendix C), this means that labor moves into the other sectors. This fully corresponds with the intuition of Baumol (1967) and Acemoglu and Guerrieri (2008).

4 Conclusion and further research

By building an agent-based model, it is demonstrated that a decentralized market economy becomes a self-organizing system if four empirically relevant conditions apply: S-shaped production functions, search behavior on the demand side, competition-based price setting and free entry of new firms. The system is robust and the business cycles generated are relatively small. This is due to competition-based pricing on both the goods and labor markets.

In the current model, it is assumed that households consume their entire income and that the investor's returns are determined by an exogenously given discount factor. The obvious next step is to assume that households have one or more savings motives and in this connection introduce a bank that stores household wealth and lends money to investors. In this way, an endogenous interest rate can be achieved and a correlation is generated between the firms' return and the return on the household's wealth. Another obvious area of development is to look at alternative wage structures. As mentioned, it is assumed in the current model that the wage is set based on competition-based pricing. Finally, adding a public sector will make it possible to add axation, unemployment benefits and to carry out policy shocks to the system.

Saving households, inertia in wage formation and a public sector must be expected to generate larger business cycles. The model described here can be seen as a prototype of an ideal market economy. The next step is to add realistic phenomena and institutions, and in that connection we will probably observe far more Keynesian effects.

References

- Acemoglu, D. and V. Guerrieri (2008) *Capital Deepening and Non-Balanced Economic Growth*. Journal of Political Economy, Vol. 116, No. 3, 2008, pp.

- Baumol, W. (1967) *Macroeconomics of Unbalanced Growth The Anatomy of Urban Crisis*. American Economic Review, 57, 415-426.

- Brown, Robert G. (1956). *Exponential Smoothing for Predicting Demand*. Cambridge, Massachusetts: Arthur D. Little Inc. p. 15.

- Carroll, Christopher D. (1997), *Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis*, Quarterly Journal of Economics 112, 1-56.

- Dawid, Herbert and Delli Gatti, Domenico (2018), *Agent-Based Macroeconomics*, Handbook of Computational Economics Volume 4, 2018, Pages 63-156.

- Debreu, Gérard (1959). *The theory of value: an axiomatic analysis of economic equilibrium* New York: Wiley.

- De Vroey, Miche, (2016) *A History of Macroeconomics from Keynes to Lucas and Beyond*, Cambridge University Press

- Dixit, Avinash K. and Stiglitz, Joseph E. (1977) *Monopolistic Competition and Optimum Product Diversity*. The American Economic Review Vol. 67, No. 3 (Jun., 1977), pp. 297-308.

- Eaton, Jonathan and Kortum, Samuel (2002) . *Technology, Geography, and Trade*. Econometrica, Vol. 70, No. 5. (Sep., 2002), pp. 1741-1779.

- Gavrilov, L. A and Gavrilova, N. S. (1992) *The Biology of Life Span: a Quantitative Approach*. Journal of the Institute of Actuaries , Volume 119 , Issue 2 , pp. 379 - 380.

- Gintis, Herberty (2007) *The Dynamics of General Equilibrium.* The Economic Journal. Volume117, Issue 523 (October 2007) pp. 1280-1309.

- Hansen, M., Zangenberg, J., Markeprand, T., Kargo, A. and Sønderby, M. (2021) *Demografi, socioøkonomi og boligstruktur i danske kommuner* (in Danish). DREAM Report. Link

- Hinterhuber, A. (2008) *Customer value-based pricing strategies: why companies resist*. Journal of Business Strategy, Vol. 29 No. 4, pp. 41-50.

- Holt, Charles C. (1957). *Forecasting Trends and Seasonal by Exponentially Weighted Averages*. Office of Naval Research Memorandum. 52.

- Hopenhayn, Hugo A. (1992) *Entry, exit, and firm dynamics in long run equilibrium*. Econometrica, Vol. 60, pp. 11127-1150.

- Hyndman, R.J., Koehler, A.B., Ord, J.K., and Snyder, R.D. (2008) *Forecasting with exponential smoothing: the state space approach*, Springer-Verlag.

- Kaplan, Greg and Guido Menzio (2016) *Shopping externalities and self-fullling unemployment fluctuations*. Journal of Political Economy 124.3, pp. 771–825.

- Kotler, Philip, Armstrong, G., Harris, L.C., He, H. (2020) *Principles of Marketing. Eighth European Edition*. Pearson.

- LaBaron, Blake, Tesfatsion, Leigh (2008) *Modeling Macroeconomies as Open-Ended Dynamic Systems of Interacting Agents*. The American Economic Review Vol. 98, No. 2, Papers and Proceedings of the One Hundred Twentieth Annual Meeting of the American Economic Association (May, 2008), pp. 246-250

- Leijonhufvud, Axel (1968) *On Keynesian economics and the economics of Keynes: a study in monetary theory*. New York: Oxford University Press.

- Leijonhufvud, Axel (2006) *Agent-Based Macro*. In L. Tesfatsion and K.L. Judd (eds.) handbook of Computational Economics. Vol II. Amsterdam, Elsevier: 1625-1638.

- Melitz, Marc J. (2003) *The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity.* Econometrica, Vol. 71, No. 6. (Nov., 2003), pp. 1695-1725.

- Mortensen, Dale; Pissarides, Christopher (1994). Job creation and job destruction in the theory of unemployment. Review of Economic Studies. 61 (3): 397–415.

- Pissarides, Christopher (2000) *Equilibrium Unemployment Theory* (2nd ed.). MIT Press.

- Qiu, Zhesheng and José-Víctor Ríos-Rull (2019) *Directed search, nominal rigidities and markup cyclicality*. Working paper.

- Sharpe, William F. (1994). *The Sharpe Ratio*. The Journal of Portfolio Management. 21 (1): 49–58.

- Stephensen, Peter (2015) *The Event-Pump: An Agent Based approach to Microsimulation*. Paper presented at the 5th World Congress of the International Microsimulation Association (IMA) Luxembourg, 2-4 September 2015. *Link*

- Stephensen, Peter (2022) *A short note on Perceived Risk*. DREAM Working paper. *Link*

- Storesletten, Kjetil & Jose-Victor Rios Rull & Yan Bai (2011) *Demand Shocks that Look Like Productivity Shocks* 2011 Meeting Papers 99, Society for Economic Dynamics.

- Syverson, Chad (2011) *What Determines Productivity?* Journal of Economic Literature 49 (2): 326–65.

- Testfatsion, L. (2006) *Agent-Based Computational Modeling and Macroeconomics*. In Colander, D. (ed.). Post-Walrasian Macroeconomics: Beyond the Dynamic Stocastic General Equilibrium Model. Cambridge, Combridge University Press: 175-201.

- Wright, R., Kircher, P., Julien, B., and Guerrieri, V. (2021) *Directed Search and Competitive Search Equilibrium: A Guided Tour*. Journal of Economic Literature, 59 (1): 90-148.

Appendix A: Selected model parameters

Parameter	Text (Equation)	Value
γ_μ	Scale parameter in Gompertz law (2.1)	0.1
α	Decreasing returns in production function (2.6)	0.5
ϕ	Increasing returns in production function (2.6)	2.0
φ_0	Minimum productivity in firms' Pareto-distribution	0.5
k	Shape parameter in firms' Pareto-distribution	2.5
$\lambda_P, \lambda_W, \lambda_d, \lambda_s, \lambda_q, \lambda_a$	Smoothing factor in firms' expectations	0.4
n	Number of firms in search for lower price	75
m	Number of firms contacted when searching for job	4
$\overline{\eta}_p^U$	Maximal markup in price-setting	0.05
$\overline{\eta}_p^D$	Maximal 'markdown' in price-setting	0.05
$\overline{z}_p^d$	Parameter in price-setting	0.15
ω_p	Probability of recalculating price	0.5
γ_l	Proportion of optimal vacancies advertised	0.1
$\overline{\eta}_w^U$	Maximal markup in wage-setting	0.05
$\overline{\eta}_w^D$	Maximal 'markdown' in wage-setting	0.05
$\overline{z}_p^d$	Parameter in wage-setting	0.15
ω_w	Probability of recalculating wage	0.5
$\overline{n}$	Employment during startup period	10
$\overline{T}$	Duration of startup period (months)	6
g	Productivity growth p.a. in new firms	0.02
r	Risk free rate of return p.a.	0.05
ξ	Investor sensitivity (2.14)	0.15
γ_s	Smoothing factor in investors expectations	0.7

Appendix B: Firm destruction shock. 30 years.

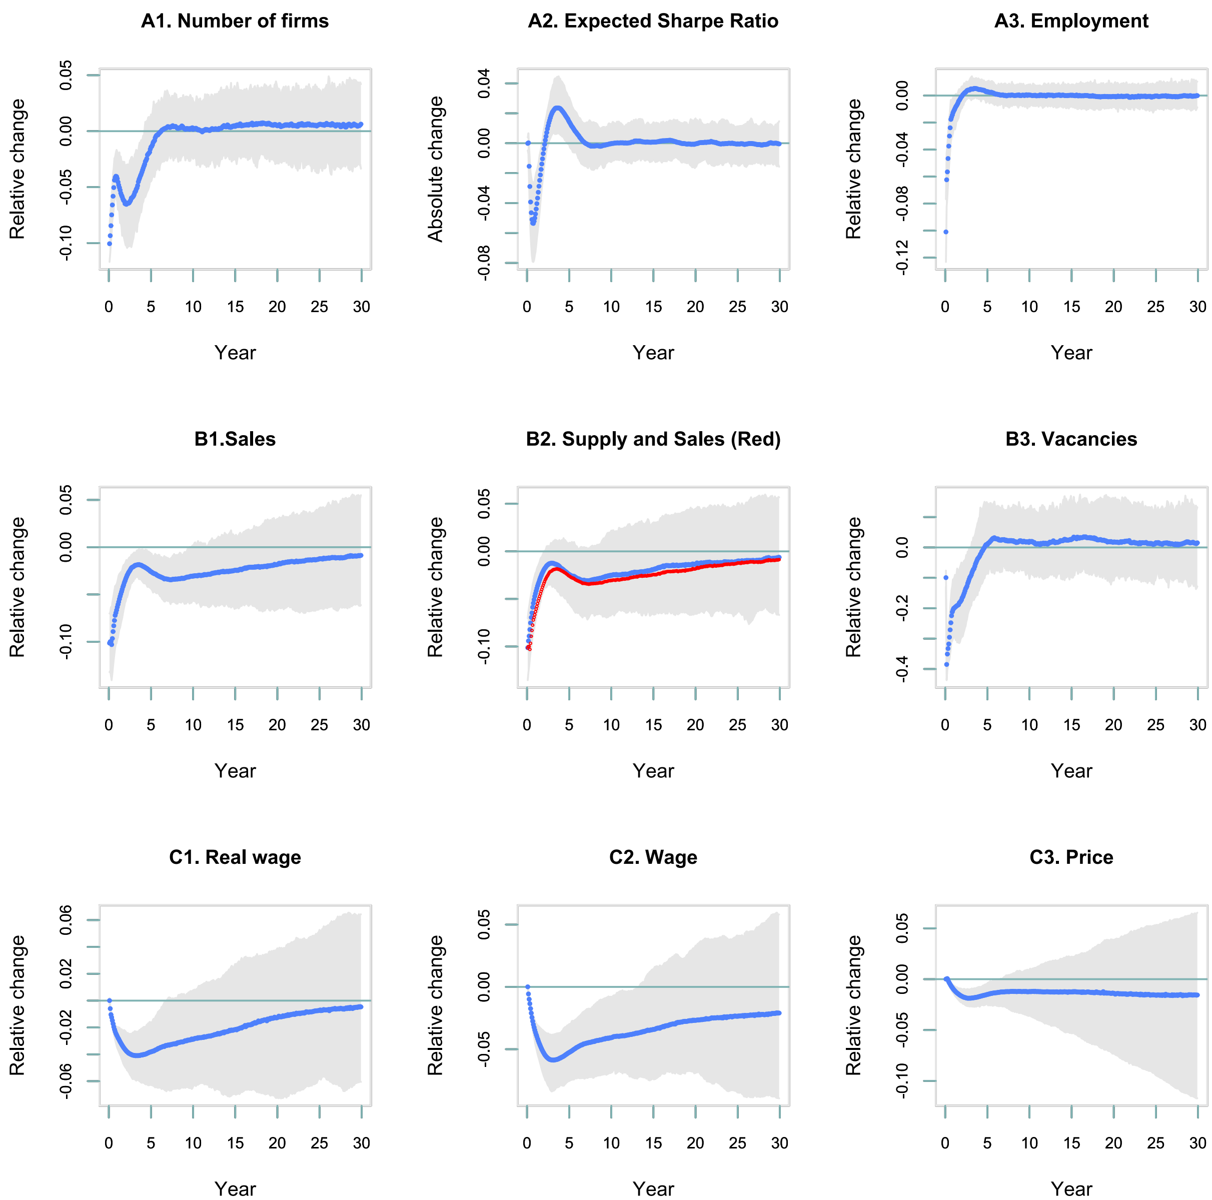

The blue points averages 150 runs. The shaded area marks 95 percent confidence intervals.

Appendix C: Productivity shock in one sector. Macro effects.

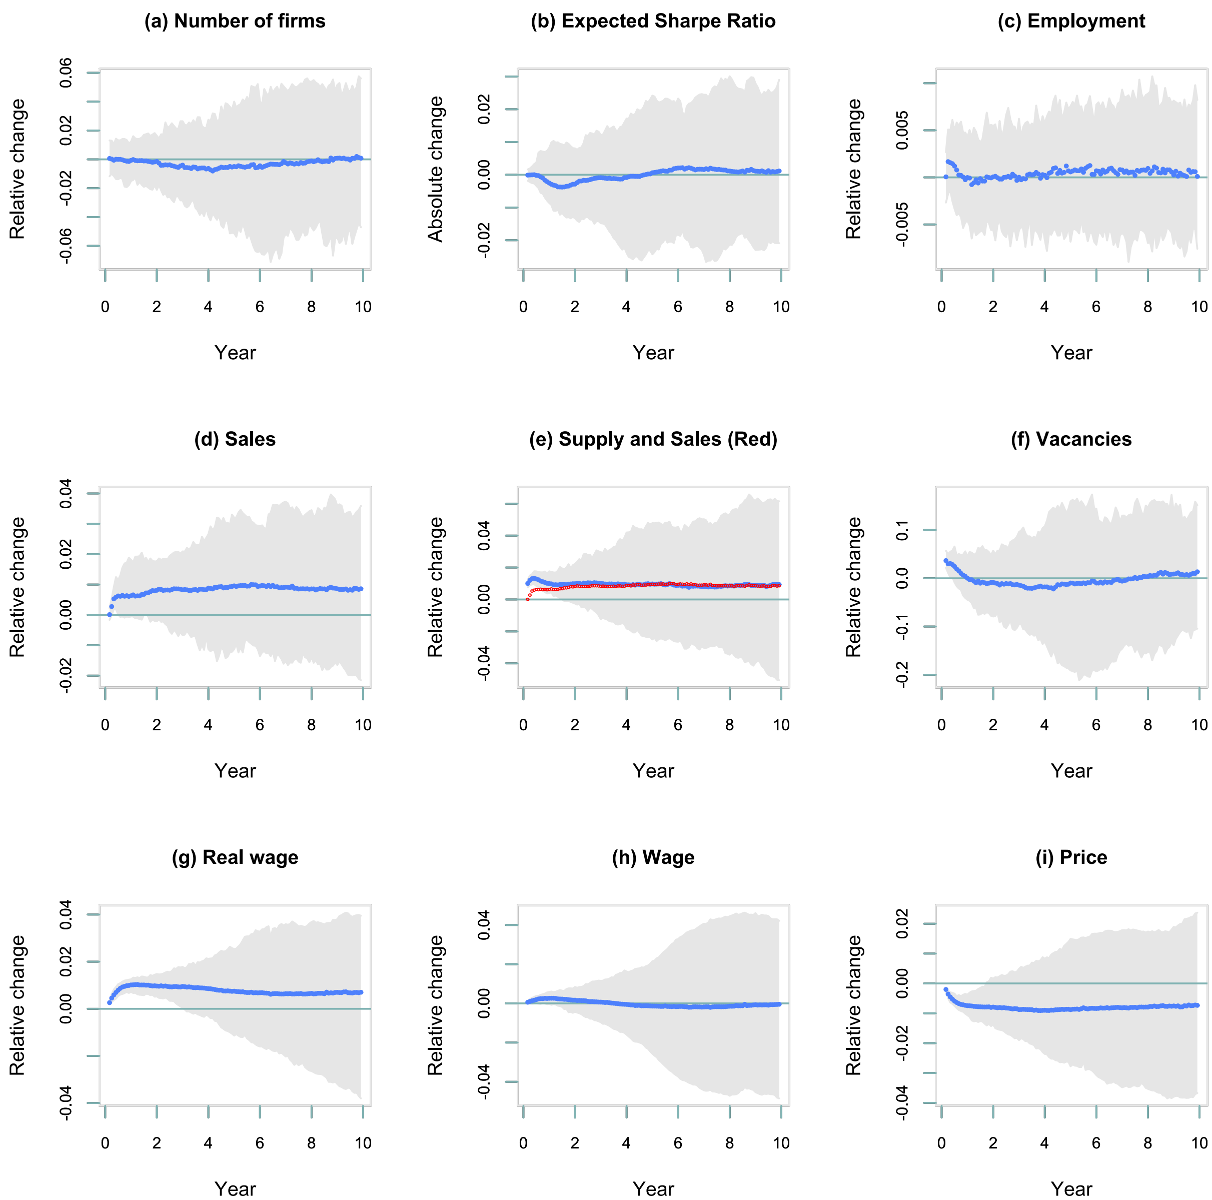

The blue points averages 150 runs. The shaded area marks 95 percent confidence intervals.