

DREAM

Danish Research Institute for
Economic Analysis and Modelling

DREAM


The Market Economy as a Self-Organizing System. An Agent Based Theory

Peter Stephensen

Working paper 2026:02

18 June 2026

www.dreamgruppen.dk



The Market Economy as a Self-Organizing System. An Agent Based Theory^{*}

Peter Stephensen^{*}

Danish Research Institute for Economic Analysis and Modeling, DREAM and Department of Economics, University of Copenhagen.

Abstract

A central difficulty in economic theory is to connect existence claims about competitive equilibrium with an account of the actual decentralized processes through which quantities and allocations become mutually consistent in real time. General equilibrium analysis clarifies what joint consistency requires, but it often leaves the coordinating mechanism implicit. This paper takes a constructive approach: it builds a large-scale agent-based economy in which all trading interactions are explicit, expectations are disciplined by observable histories rather than by rational-expectations shortcuts, and long-run regularities emerge from micro behavior rather than from auxiliary equilibrium operators. The agent-based system incorporates several novel features: *competition-based pricing* from the marketing literature, *network-based search* behavior that generates endogenous gravitational effects, household *behavioral buffer-stock* budgeting, a *communication protocol* between agents, and a competitive banking sector in which an investor governs firm entry through a Sharpe-type profitability signal. The paper argues that competition-based pricing can be viewed as an instance of a social mechanism termed *copy–deviate behavior*: satisfied agents imitate one another—supporting coordination toward an equilibrium—whereas dissatisfied agents deviate from prevailing practice—thereby generating potential movement toward a new equilibrium.

The economy comprises heterogeneous overlapping-generations households and heterogeneous firms with S-shaped technologies. Monte Carlo simulations show the economy tracking a balanced growth path. Two supply-side experiments confirm self-organizing dynamics: a permanent productivity increase raises output and real wages proportionally, while random firm destruction triggers re-entry that restores firm counts within years.

Keywords: Agent-based model, Self-organization, Competition-based pricing, Copy–deviate behavior, Decentralized market economy

JEL codes: C63, D50, E10, L11

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

^{*}DREAM, Landgreven 4, 7th floor DK-1301 Copenhagen K, E-mail: psp@dreamgruppen.dk

1. Introduction

A central difficulty in economic theory is to connect existence claims about competitive equilibrium with an account of the actual decentralized processes—prices, wages, search, entry, and exit—through which quantities and allocations become mutually consistent in real time (Leijonhufvud, 1968; De Vroey, 2016). General equilibrium analysis clarifies what joint consistency requires, but it often leaves the coordinating mechanism implicit (Walras, 1854; Debreu, 1952; Arrow and Debreu, 1954). Largely separate from this, there is also a substantial literature on the stability of equilibrium (Arrow et al., 1959; Hahn and Negishi, 1962; Uzawa, 1962; Smale, 1976).

This paper takes a constructive approach: it builds a large-scale agent-based economy in which all trading interactions are explicit, expectations are disciplined by observable histories rather than by rational-expectations shortcuts, and long-run regularities emerge from micro behavior rather than from auxiliary equilibrium operators. The agent-based system incorporates several novel features: *competition-based pricing* from the marketing literature, *network-based search* behavior that generates endogenous gravitational effects, household *behavioral buffer-stock budgeting*, a *communication protocol* between agents, and a competitive banking sector in which an investor governs firm entry through a risk adjusted Sharpe-type profitability signal.

The economy contains heterogeneous households, heterogeneous firms, a banking sector that links household wealth to firm creation, and an investor who finances entry. Households consume goods, supply labor while young, save for retirement, and die stochastically according to an overlapping-generations demographic law. On the goods and labor sides, households engage in *network-based search*: they consult other households to learn where to shop and work, which endogenously directs demand and job applications toward larger sellers and employers in proportion to their customer and employment bases. Entrants overcome initial anonymity through a simple advertising channel.

Firms produce with an S-shaped technology in employment, in the spirit of workhorse models of firm dynamics with fixed costs and decreasing returns at the intensive margin (Melitz, 2003; Hopenhayn, 1992; Eaton and Kortum, 2002). Management chooses a target scale under expected prices and wages; marketing and human-resource departments then set prices and wages using *competition-based pricing rules* adapted from the marketing literature (see e.g. Kotler et al., 2021). Under competition-based pricing, mature firms that are content with demand and staffing track sector averages, whereas dissatisfied firms temporarily mark down or mark up relative to the market. This pattern generates a social mechanism the paper terms *copy-deviate behavior*: copying supports coordination when many agents are near their targets, while deviation generates motion when

many are not. Firm-level heterogeneity in productivity is coupled with creative destruction from technological progress among entrants and with exit when profitability fails.

The financial side is deliberately stylized but internally coherent. Household wealth is intermediated; banks compete until profits are passed through as deposit income, which delivers a short nominal interest rate tied to aggregate bank earnings relative to outstanding wealth. The representative bank acts as an investor that finances startups, bears losses, receives profits, and adjusts the flow of new firms in response to a risk-adjusted measure of expected profitability based on a Sharpe-type index (Sharpe, 1994) computed from the cross section of firms. Together with goods- and labor-market rules, this entry margin provides the external adjustment mechanism that scales the number of active producers.

Expectations are formed without perfect foresight or rational expectations. Agents update views using *exponential smoothing* (Brown, 1956; Holt, 1957; Hyndman et al., 2008), with smoothing parameters chosen to separate high-frequency forecasting from slow recognition of structural level shifts in key aggregates.

Goods trade and labor reallocation are implemented through an explicit *message protocol*: households send purchase and job inquiries, firms accept or reject depending on inventories and vacancies, and separations and deaths update employer rosters. This design makes excess demand and hiring frictions observable objects inside the simulation rather than reduced-form residuals. Inventories depreciate quickly by assumption, which forces a tight connection between expected sales, production, and customer service in each month.

Household finances combine life-cycle motives with a buffer-stock consumption rule linked to perceived permanent income, expected real interest rates, and a slowly moving wealth target. This can be seen as a behavioral version of Carroll (1997). Unemployment interrupts accumulation; retirees decumulate on a fixed horizon rule. Inheritance transfers wealth across generations. These ingredients are chosen to keep the demand side transparent while still generating realistic dispersion in assets and incomes.

The contribution is twofold. Substantively, the paper shows that a decentralized market economy with empirically recognizable pricing and search protocols can track a stable growth path with moderate endogenous fluctuations and can undergo transparent adjustment to permanent productivity shifts and to large disruptions that destroy a random subset of firms. Methodologically, it demonstrates the conditions that must be in place to explain why a market economy is self-regulating. The firm plays a central role here. The firm can be viewed as a social unit with a *desired size* (explained here by S-shaped production functions) and a well-functioning social mechanism for achieving that size (*competition-based pricing*). We call this *copy-deviate behavior*. In addition, it is important to have a well-functioning financial sector that finances new firms and ensures

that older, low-productivity firms are closed down. The importance of a well-functioning financial sector is not new; what is new is the interaction between the financial sector and copy-deviate behavior. This—together with *network search*—creates a highly robust self-regulating system.

Relative to the existing literature, the article is closest to Gintis (2007), which presents an agent-based Walrasian economy. There are overlaps but also major differences. Similarities include nonlinear (inverted-U) firm production functions, household search on goods markets, firms as price-setters, and entry of new firms. The main difference is that Gintis (2007) adopts an evolutionary rule under which 5 per cent of firms each period copy the behavioral parameters of the best performers. That generates heterogeneity in “private prices” within a sector and helps make the system self-organizing. We obtain price dispersion through a different channel—competition-based pricing—and rely on a social mechanism (copy–deviate behavior) rather than an evolutionary one. In addition, our production technology and the source of firm heterogeneity draw on Melitz (2003); Hopenhayn (1992).

The quantitative implementation targets an economy large enough for granular heterogeneity to matter, yet small enough to replicate extensively. The baseline experiments reported later use on the order of tens of thousands of households and hundreds of firms, evolve at monthly frequency, and are repeated many times with fixed pseudo-random seeds across scenarios so that differences between baseline and counterfactual paths reflect economic mechanisms rather than sampling noise. A multi-phase initialization routine first builds a plausible firm-age structure and stationary demographics before endogenizing the short nominal interest rate and recording long post-equilibrium histories. The total simulated horizon is long by macro standards, reflecting the slow dynamics of productivity distributions after organizational destruction.

The remainder of the paper is organized as follows. Section 2 presents agents, expectation formation, firm behavior, finance, and the equilibrium logic linking investor entry to expected profitability. Section 3 describes the computational design, reports the baseline simulation, and studies counterfactual supply shocks. Section 4 concludes.

2. The Economy

The system is implemented as an *agent-based model* (Tesfatsion, 2006; LeBaron and Tesfatsion, 2008; Dawid and Delli Gatti, 2018). The core idea is to represent only agents and their interactions, imposing only mechanisms and constraints that are explicitly present in the economy under study. By construction, one does not append free-standing abstract devices—such as equilibrium conditions, a Walrasian auctioneer, or rational expectations—that are not themselves generated by agent behaviour. For this reason, agent-based modelling is especially well suited to studying self-organising systems¹. A modeller may also regard it as a virtue that the approach forces a concrete specification rather than relying on reduced-form abstraction. An Occam’s-razor-style guideline might read: abstraction is valuable, but it should be reserved until straightforward concrete representations have been exhausted. Instead of rational expectations or perfect foresight, agent-based models typically specify expectations from observable histories, as in applied work on learning and adaptation. Likewise, search can be implemented explicitly, with individual agents searching over other agents, rather than through a reduced-form matching function of the kind common in macro labour models (Mortensen and Pissarides, 1994; Pissarides, 2000). This permits a network-based microfoundation for search and matching.

Operationally, the code uses an event-pump architecture (Stephensen, 2015), which has for several years underpinned the Danish microsimulation model SMILE (Hansen et al., 2021). The model is written in C# and is available at github².

The economy is closed and contains three agent types: households, firms, and a banks. There are many heterogeneous households and firms and a single representative bank³. The following subsections describe each agent type in turn.

2.1. Expectations

Before describing the agents, we provide a brief overview of expectation formation in the model. As noted above, agents do not know future outcomes: there is neither perfect foresight nor rational expectations. Instead, they form expectations about a range of variables—such as interest rates, inflation, prices, and wages—using information contained in the model’s own history.

We distinguish short-run from long-run expectations. Short-run expectations are represented by conventional one-period-ahead forecasts (defined below). Long-run expectations require agents

¹Axel Leijonhufvud: “Agent-based methods provide the only way in which we can explore the self-regulatory capabilities of complex dynamic models and thus advance our understanding of the adaptive dynamics of actual economies.” (Leijonhufvud, 2006, p. 1636).

²<https://github.com/PeterStephensen/dream.agentbased.SOE2.git>

³The bank is the only representative agent in the model.




Figure 1: Real interest rate. Data and filtered. Generated by the model.

to form views about the economy’s “natural” or structural levels of key variables. Agents observe monthly data spanning a century or more and are therefore better informed, in this respect, than most real-world decision makers. They must both extract the underlying structural level from noisy historical series—i.e., filter the data—and detect shifts in those structural levels when they occur. In short, agents must learn from history without being lost in the past.

We assume that the agents use the statistical method *Exponential Smoothing*⁴ (Brown, 1956; Holt, 1957; Hyndman et al., 2008). With Simple Exponential Smoothing (SES), expectations are updated each period with the equation:

$$y_t^e = \lambda y_{t-1}^e + (1 - \lambda) y_t \quad (1)$$

where y_t^e is the expected value and y_t is the actual value (Hyndman and Athanasopoulos, 2024). We call λ the smoothing parameter.

⁴This method is at least as widespread and effective as ARIMA (Hyndman and Athanasopoulos, 2024).

It can be shown that the above difference equation (1) has the solution :

$$y_t^e = (1 - \lambda) \sum_{s=0}^{\infty} \lambda^s y_{t-s}$$

The expected value y_t^e is thus a weighted average of the entire historical record, with weights that decline exponentially as observations recede into the past. Larger values of the parameter λ imply that more weight is placed on distant observations—i.e., agents effectively look further back in time.

What value should λ take? For short-run expectations, we adopt the standard statistical procedure used for simple exponential smoothing (SES) (Hyndman and Athanasopoulos, 2024). Estimated values of λ are typically modest (little smoothing), because the objective is optimal one-step-ahead prediction.

For long-run expectations, we follow a different principle: we seek as much smoothing as possible (high λ , strong filtering) while still allowing structural breaks to be detected. As a working rule—admittedly somewhat arbitrary—we require that 80% of a permanent level shift be reflected in expectations within three years (36 months). Straightforward calculation shows that this is satisfied for $\lambda = 0.96$ (Appendix A). The relatively high numerical value reflects the monthly frequency: with quarterly data, the analogous choice would be $\lambda = 0.87$; with annual data, $\lambda = 0.58$.

Figure 2.1 illustrates a representative simulation. The black series plots the realised real interest rate (defined below) over 50 years. A productivity shock hits in year 10 (described in detail later). When λ is estimated as in Hyndman and Athanasopoulos (2024), we obtain $\lambda = 0.83$; the orange series plots the corresponding expectations. That specification filters out short-run noise and delivers accurate short-horizon forecasts. The blue series shows long-run expectations under our calibration $\lambda = 0.96$. Relative to the orange series, the blue series is much smoother, and expectations track the level shift in the real interest rate induced by the productivity shock.

2.2. Households

Households demand goods and supply labor. They exhibit search behavior in both the goods market and the labor market. Modern macroeconomic theory often assumes search in the labor market (Pissarides, 2000; Mortensen and Pissarides, 1994). Consumer search in product markets has received increasing attention more recently (Qiu and Rios-Rull, 2019; Kaplan and Menzio, 2016; Storesletten et al., 2011). 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. We assume that the household conducts *network search* as described in section (2.2.1). This means that the individual

household consults other households for advice when choosing firms for work and purchasing goods.

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:

$$\log(\mu_a) = \gamma_\mu (a - \bar{a}), \gamma_\mu > 0, a \leq \bar{a} \quad (2)$$

where a is age and \bar{a} is the maximal 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 three sources of income: wage, capital income and inheritance. Employed households receive wages. Assume the j 'th household is employed in firm i . The household will then receive the wage in period t :

$$v_{jt} = w_{it} \cdot \rho_{jt}, \quad (3)$$

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 from Carroll (1997) and many others. It is assumed that $\rho_{jt} = 1$ when the household enters the economy at age 18. After this, it is assumed that:

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

where ε_{jt} is normally distributed⁵. This dynamics 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 j th household is

$$y_{jt} = v_{jt} + i_t A_{j,t-1} + h_t$$

where i_t is the nominal interest rate, $A_{j,t-1}$ is household wealth ultimo last month and h_t is inheritance from dead parents. $v_{jt} = 0$ if the household is unemployed.

⁵The mean is given by $\mu = -\sigma^2/2$ where σ^2 is the variance. This implies that $E[\rho_{jt}] = E[\rho_{j,t-1}]$ at any age.

2.2.1. Network search and advertising

A household has potential contact with two types of firms in a given period: the firm where it works and the firm where it buys goods. We assume that the household find these to firms using *network search*.

It is common in agent-based models to use search. Typically, this is done by randomly selecting a number of firms from the list of firms (Dawid and Delli Gatti, 2018; Gintis, 2007). This means that all firms have the same probability of being chosen - regardless of their size. This has the unfortunate effect that there will be a tendency for small firms to be overwhelmed by demand, while large firms are underwhelmed by customers. In reality, large firms will automatically experience more attention than small firms. This is partly due to advertising, and partly because households communicate with each other regarding their choice of firms. We model both aspects.

We achieve that large firms experience greater demand than small firms by using *network search*⁶. In practical terms, this is done by extracting a number of random households instead of a number of random firms. The household obtains a list of firms by asking the list of households which firms they use for work or shopping. This makes the probability of a firm being selected directly proportional to its number of employees or customers. In a model with heterogeneous firms, a *gravitational* mechanism like this is important for the system's self-regulating properties.

As mentioned, we also model advertising. This is a necessary consequence of network search. Under network search, it would be difficult for new firms to enter the market, since no households know them. We therefore assume that new firms send spam mails with their price to a number of random households. After the firm has entered the market, it stops advertising and relies solely on households' network search. More on this in the next section on firms.

Note that here we have modeled search 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 and Delli Gatti, 2018; Gintis, 2007). If there are many positions offered on the labor market, 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.

⁶An alternative method is to assume that the individual firm's probability of selection is a logit function. This could in principle depend on the firm's size - although this is typically not the case, see Dawid & Delli Gatti (2018).

2.2.2. *Buying goods*

In a given period (month), the standard behavior of the household is to buy from the same firm as last month. With a low probability, the household conducts a market survey. It randomly draws n firms using network search. If the household among these n firms can find a product that is cheaper than the current one, it switches provider.

The household purchases goods from a firm by sending a *CanIBuy* message to the firm. The demanded amount b is part of the message. If the firm's price is p and if the firm can deliver the desired quantity $q = b/p$, 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 implies that the firm is able to count how many potential customers it has.* It can therefore recognize if there is an excess demand for its product.

2.2.3. *Labor market*

The households are active on the labor market between the age of 18 and 66. At age 67 it retires and lives from its wealth. 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 network-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 *JobApplication* 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 compares the potential firms' wages with a reservation wage. The reservation wage is the wage from last employment period, discounted with a fixed rate every unemployment period. Among the firms offering a job, the household chooses the one with the highest wage (if any).

2.2.4. *Saving behaviour*

Since there are no public pensions, households must save for old age. We assume that households are credit-constrained and build up their wealth through a *behavioral buffer-stock strategy*. The method can be seen as a behavioral variant of Carroll (1997). Carroll shows that a rational

household under uncertainty (and certain assumptions) will have a savings behavior that resembles a buffer-stock rule of thumb: the household builds up a buffer of a certain size as a safeguard against shocks. It is thus shown that rational behavior can resemble a rule of thumb. We simply assume that households follow this rule of thumb.

The household calculates its expected labor market income v_t^e as described in section 2.1. One can think of v_t^e as an estimate of the long-term income or the 'permanent income', as it is often referred to.

For employed households that are still active in the labor market, we assume that the consumption budget is given by the nominal consumption function:

$$b_t = r_t^e A_{t-1} + v_t^e + \xi_v (v_t - v_t^e) + \xi_r (r_t - r_t^e) A_{t-1} + \xi_A (A_{t-1} - A_t^*) \quad (4)$$

where A_t^* is the desired long-term wealth and r_t^e is the expected real interest rate (which we assume is the same for all agents). One can think of r_t^e as an estimate of the long-term real interest rate. The expected real interest rate is based on calculations of both the expected nominal interest rate and the expected inflation. There is thus an implicit assumption that the expected inflation is valid going forward. In this sense, consumer behavior is influenced by inflation expectations.

The first two terms on the right in (4) is the expected real capital income and expected wage income ("permanent income"). If we only had these two terms ($b_t = r_t^e A_{t-1} + v_t^e$), the consumption budget would be determined by the permanent income level and real capital income. This would be a cautious strategy, where consumption is determined by the perceived income level, and where only that part of capital income is consumed which ensures that wealth (in steady state) grows with expected inflation (and thus preserves purchasing power).

The next two terms define the short-run marginal propensity to consume out of labor market income and capital income, respectively. The last term defines the buffer-stock effect that ensures that the long-term wealth is equal to the desired level.⁷

The desired long-term wealth is assumed to be determined by the permanent income:

$$A_t^* = \gamma_A \cdot v_t^e \quad (5)$$

This means that the long-term wealth slowly adjusts to the income level of the individual household. One can relatively easily show that under (4) and (5), the wealth will converge to the desired level if the underlying variables are constant in the future. In practice, almost all households will experience growing incomes in the long run, as there is approximately 2% yearly growth due to

⁷It is assumed that $\xi_v = 0.95$, $\xi_r = 0.1$ and $\xi_A = 0.005$.

technological progress. By setting γ_A relatively high and ξ_A low, we obtain reasonable builds of wealth over the households' working lives. In the current simulations, we have set $\gamma_A = 30$ and $\xi_A = 0.005$.

Unemployed households that have not yet retired are assumed to cease accumulating wealth and simply follow the rule⁸:

$$b_t = \gamma_u \cdot v_t^e, \gamma_u < 1.$$

This will typically reduce the wealth (a genuine buffer-stock effect). If the wealth potentially becomes negative, the consumption budget that brings the wealth to 0 is chosen. The household has then essentially gone bankrupt. We do not model this explicitly, but assume that consumption is 0 until a new job is found. In the current simulations, this happens very rarely.

Retirees also follow a rule of thumb. In each period, the household calculates what constant consumption can be maintained for H months given the current wealth. This calculation is made given the expected interest rate. It applies that

$$b_t = \omega(r_t^e, H) \cdot A_{t-1}$$

where

$$\omega(r, H) \equiv \frac{r}{1+r} \frac{1}{1 - \left(\frac{1}{1+r}\right)^H}$$

In the simulations shown, it is assumed that the time horizon is 10 years ($H = 12 \cdot 10$). With this specification, consumption is decreasing (in real terms), and there is a gradual depreciation of wealth.

At some point the household dies. In this case, the wealth is given to heirs. In the current implementation, it is assumed there are 2 heirs randomly chosen from households aged between 50 and 65.

2.3. 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).

⁸It is assumed that $\gamma_u = 0.7$.

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 textbook, 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 and Delli Gatti, 2018). Reference is typically made to Dixit and 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 strategy⁹. 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

⁹Hinterhuber 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).

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), 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 2.2.2 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 two 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).

We assume that the production function is *S-shaped*, because of increasing returns to scale at small production levels and decreasing returns to scale at high production levels. S-shaped production functions can be seen as an explanation of heterogeneity. At first glance 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; Hopenhayn, 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 Hopenhayn (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 the Melitz model, the firm's optimal size is *externally* explained (market conditions), while in the Hopenhayn model the explanation is *internal* (technology). In the real world, it is probably a combination of these two explanations that is relevant. We follow the *internal* technology explanation emphasized by Hopenhayn. 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.

In the following, the organization and behavior of firms are described. This is done by describing management, production, communication protocol, the marketing department, the human resources department, and entry and exit of firms.

2.3.1. 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 \bar{P}_t and the average wage \bar{W}_t in the economy. On the basis of these, the expected market price P_{t+1}^e and market wage in the next period W_{t+1}^e are calculated using short-run Exponential Smoothing (see section 2.1):

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

$$y = \varphi \cdot \max \{l^\alpha - \phi, 0\} \quad (6)$$

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 Hopenhayn (1992) we have $\alpha < 1$. We follow Hopenhayn (1992) and assume that $\alpha < 1$.

We interpret y as the maximum production given the employment l . The actual sales s will be stochastic. If the firm wants to be able to deliver goods to its customers, it must have a maximum production that is higher than the expected sales s^e . We assume that the firm, as a precautionary principle, follows the rule:

$$y = \frac{s^e}{\gamma_y} \quad (7)$$

where $\gamma_y < 1$.

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

$$\pi = P^e s^e - W^e l$$

or from (7):

$$\pi = P^e \gamma_y y - W^e l \quad (8)$$

If (6) is inserted into (8), it can be shown that the profit is maximized if:

$$l^* = \left(\alpha \frac{\gamma_y \varphi}{W^e / P^e} \right)^{\frac{1}{1-\alpha}} \quad (9)$$

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:

$$y^* = \varphi \left((l^*)^\alpha - \phi \right) \quad (10)$$

If (9) and (10) are inserted into (8), it can be seen that the optimal profit is given by:

$$\pi^* = P^e \left((1-\alpha) \left(\alpha \frac{\gamma_y \varphi}{W^e / P^e} \right)^{\frac{\alpha}{1-\alpha}} - \phi \right) \gamma_y \varphi \quad (11)$$

The optimal profit is therefore increasing in the productivity φ . From (11) it can then be calculated that $\pi^* \geq 0$ if and only if:

$$\varphi \geq \hat{\varphi} \equiv a_\varphi \frac{W^e}{P^e} \quad (12)$$

where the constant a_φ is given by:

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

The individual firm only has positive profit if it has a productivity φ that is higher than a cut-off value $\hat{\varphi}$. From (12) 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).

2.3.2. Production and the communication protocol

In a given period, production is determined by the production function (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 production is at the start of the month added to the stock via the accumulation equation:

$$S_t = (1 - \lambda_S)S_{t-1} + y_t$$

where S_t is the stock, λ_S is the rate of depreciation, and y_t is production. In a manufacturing firm, S_t can be interpreted as inventory, and in a service firm, S_t can be interpreted as contracts. We will assume that λ_S is relatively high (25 per cent per month), so that it is necessary for the firm to ensure that supply is relatively close to demand.

A protocol has been defined for communication between households and firms. Four messages have been defined: *JobApplication*, *IQuit*, *Death*, and *CanIBuy*. The first two concern the labor market (see section 2.3.4). The *Death* message is sent to the firm when the household dies, so that the firm can remove the household from its list of employees.

Finally, the *CanIBuy* message is sent from the household to the firm when the household wants to buy goods. Along with the message, the budget b is included. The quantity is then given by $x = b/p$ where p is the firm's price. If $S_t \geq x$, the firm responds with *Yes*, delivers the desired quantity, and reduces S_t by the quantity x . If $0 \leq S_t < x$, the firm responds with *No*, delivers the quantity S_t , and reduces S_t to 0. After this, the firm has sold out and cannot deliver to more customers in this period.

At the end of the period, the firm can calculate the total potential demand d_t during the period and the actual sales $s_t \leq d_t$.

2.3.3. The 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., 2021) in the marketing literature.

Implementing competition-based pricing may sound straightforward. In practice, however, building a system that is robust in a volatile stochastic environment requires attention to a number of details. We assume – first – that the firm sets a target price (p_t^*), which it adjusts to gradually

according to:

$$p_t = \gamma_p p_{t-1} + (1 - \gamma_p) p_t^*$$

This is done to avoid unrealistic swings in the price (we assume that $(\gamma_p = 0.8)$).

The next step is to determine the target price (p_t^*). Doing so requires knowledge of the firm's expected market situation in the next period. The firm's expected demand d_{t+1}^e in the next period is given by:

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

If $d_{t+1}^e > \gamma_y y_t^*$, the firm expects excess demand at the current optimal production level (the optimal production y^* is given by (10)). As noted above, the parameter $\gamma_y < 1$ reflects a precautionary principle. The firm wants to produce more than expected demand so that it can absorb fluctuations in demand (we assume that $\gamma_y = 0.9$).

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

$$p_t^* = \begin{cases} (1 + \eta_{p,t}^U) P_t^e & \text{if } d_{t+1}^e > \gamma_y y_t^* \\ (1 - \eta_{p,t}^D) P_t^e & \text{if } d_{t+1}^e \leq \gamma_y y_t^* \end{cases} \quad (13)$$

where P_t^e is the expected market price in the economy. The parameters $\eta_{p,t}^U$ and $\eta_{p,t}^D$ 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 $\eta_{p,t}^U$ and $\eta_{p,t}^D$ depends on the degree of disequilibrium. Excess demand is measured by:

$$z_t^p \equiv \frac{d_{t+1}^e - \gamma_y y_t^*}{\gamma_y y_t^*}$$

We assume that the markup $\eta_{p,t}^U$ is given by:

$$\eta_{p,t}^U = \bar{\eta}_p^U \cdot \min \{ \psi_p^U z_t^p, 1 \}, z_t^p > 0,$$

so that a small deviation from equilibrium results in a small change in price. ψ_p^U is a sensitivity parameter and $\bar{\eta}_p^U$ is the maximum markup.

In the same way we define::

$$\eta_{p,t}^D = \bar{\eta}_p^D \cdot \min \{ -\psi_p^D z_t^p, 1 \}, z_t^p < 0.$$

In agent-based models, price rules reminiscent of (13) often appear. However, these rules are

typically based on the last period's price and not the market price (Dawid and 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.

2.3.4. The Human Resources department

The HR department's task is to post vacancies 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.

Workers are laid off if:

$$l_t > \gamma l_t^*$$

where l_t is the actual employment (measured in productivity units), $\gamma > 1$ is a markup and l_t^* is optimal employment, defined by (9). The markup γ reflect a precautionary stance and reluctance to dismiss employees. The excess number of employees is fired according to the last-in-first-out principle. The firm deletes the laid-off workers from the list of employees and dispatches a *YouAreFired* message to the relevant households.

The number of vacancies is given by:

$$v_t = \max \{ \gamma l_t^* - l_t + q_{t+1}^e, 0 \}$$

where q_{t+1}^e is the expected number of quitters.

During each period, the firm may receive three job-related messages from households: *JobApplication*, *IQuit*, and *Death*. Messages may be accompanied by relevant data. For these three message types, the relevant data is the household's productivity, ρ .

At the beginning of a period, the firm initializes a variable $l^{\text{nextPeriod}} = l_t$, where l_t denotes employment at the start of the period (measured in productivity units). Upon processing *IQuit* or *Death*, the household is removed from the firm's list of employees, and productivity units are subtracted from the employment counter: $l^{\text{nextPeriod}} \rightarrow l^{\text{nextPeriod}} - \rho$.

If the firm receives a *JobApplication* message, it responds *Yes* whenever remaining vacancies satisfy $v_t \geq \rho$. In that case, it updates $l^{\text{nextPeriod}} \rightarrow l^{\text{nextPeriod}} + \rho$ and $v_t \rightarrow v_t - \rho$, and adds the household to the employee list.

At the start of the subsequent period, employment is set to $l_{t+1} = l^{\text{nextPeriod}}$. If enough applications arrive, the firm can fill all vacancies. Otherwise, realised employment l_{t+1} falls short of the planned level—and hence will typically lie below desired employment, γl_{t+1}^* .

We model competition-based wage-setting in much the same way as when marketing sets the

price (see 2.3.3). We assume that the firm sets a target wage (w_t^*), which it adjusts to gradually according to:

$$w_t = \gamma_w w_{t-1} + (1 - \gamma_w) w_t^*$$

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

$$w_t^* = \begin{cases} (1 + \eta_{w,t}^U) W_t^e & \text{if } l_t < \gamma l^* \\ (1 - \eta_{w,t}^D) W_t^e & \text{if } l_t \geq \gamma l^* \end{cases}$$

The deviation of actual employment from planned employment is measured by:

$$z_t^w \equiv \frac{\gamma l^* - l_t}{\gamma l^*}$$

The markup:

$$\eta_{w,t}^U = \bar{\eta}_w^U \cdot \min \{ \psi_w^U z_t^w, 1 \}, z_t^w > 0,$$

and the 'markdown':

$$\eta_{w,t}^D = \bar{\eta}_w^D \cdot \min \{ -\psi_w^D z_t^w, 1 \}, z_t^w < 0,$$

2.3.5. *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); Hopenhayn (1992), the firm-specific productivity is drawn from a Pareto distribution¹⁰. 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. We assume that Γ_t grows 2

¹⁰The Pareto distribution has been chosen to ensure comparability with the literature. We could easily use another distribution, e.g. a log-normal distribution.

percent per year.

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 \bar{n} for \bar{T} periods to “build” the firm. This is called the start-up period. Entry cost is the wage cost for these employees. The parameters \bar{n} and \bar{T} are the same for all firms and no production can be done in the \bar{T} periods¹¹.

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). If there are not enough job 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 with a given probability.

2.4. *The financial sector*

In the model, the financial sector plays the same role as in reality: to safeguard households’ wealth and earn a return on it, and to finance firms’ investment. We represent this financial circuit between households and firms by a representative bank that practises *fractional-reserve banking*. Motivated by the desire to keep this bank as simple and generic as possible, we set up a stylised model with many identical banks and perfect competition. The representative bank is defined as the symmetric equilibrium of this system.

The rationale for the banking model draws on the historical development in London in the seventeenth century. There, the city’s goldsmiths (re)invented fractional-reserve banking and laid the foundations for the system we have today (Quinn, 1997). At the beginning of the seventeenth century, well-to-do citizens increasingly began to store their gold with goldsmiths. They did so solely as a precaution. Goldsmiths offered secure storage in their vaults. Customers paid a custody fee— analogous to paying for a safe-deposit box today. In return, they received a receipt (goldsmith’s note) acknowledging the deposit.

Goldsmiths observed that only a very small fraction of the gold was withdrawn at any given time. They therefore began to lend out a share of the gold at interest. Fractional-reserve banking had emerged. Goldsmiths were now in a position where the more deposits they attracted, the more they earned. They therefore had an incentive to attract additional deposits, first by eliminating the

¹¹We assume $\bar{n} = 10$ and $\bar{T} = 6$.

fee and later by paying interest on deposits. A circuit thereby arose in which part of the interest earned by goldsmiths on their lending was passed through to depositors.

2.4.1. *The Bank model*

We assume there are several identical banks that compete. Households can store wealth in a bank's safe. That arrangement need not entail interest paid from bank to household—on the contrary. The bank supplies a security service to the household. Household deposits leave each bank holding substantial funds. That allows fractional-reserve banking. In this economy there is only one way to invest: starting new businesses. Each bank has an investment arm that supplies capital to new firms (see 2.4.2).

The return on banks' investments is pure profit. If the bank pays no interest to households, its only outlays are for premises and staff—relatively small next to the return on deployed wealth—so we ignore them below.

When a bank earns pure profit, it has a strong incentive to attract as much household wealth as possible, for example by paying interest on deposits. Even if banks were not required to pay interest initially, they still have good reason to do so.

That yields competition among banks. We assume they bid up the interest rate on deposits until pure profit is eliminated. We model the banking system as a representative agent with zero profit and no costs. It invests in new firms and passes the resulting profit to households as interest on their wealth. That yields a simple expression for the economy's short nominal interest rate i_t :

$$i_t = \frac{\Pi_t}{A_{t-1}} \quad (14)$$

where Π_t is the firms total profit earnings in period t and A_{t-1} is the households' total wealth ultimo period $t - 1$.

In practice there are many distinct assets (including government bonds), and banks' margins are not competed down to zero. The financial sector splits pure profit with households but, as is well known, retains a sizeable share. None of this undermines the point that the simple interest-rate story in (14) gives an intuitive, attractive bridge between the real and financial sides of the economy.

2.4.2. *The Investor*

As mentioned above, we assume that the representative bank has an investment department (hereafter the investor). The investor makes a living by starting firms, and you can think of her as a venture capitalist or 'angel investor'. To simplify matters, we assume that investor finances the firms's losses in the future and receives all profits. Investor can close the firm at any time. If that 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 (assumed to be 0.5). 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 are assumed to be credit constrained.

Measurement of profitability. Because there is uncertainty about the new firms' productivity, the investor will buy a portfolio of firms (as in Melitz (2003); Hopenhayn (1992)). This section deals with how the investor measures profitability. In the next section, this is used to decide how big the portfolio should be.

The investor makes decisions under fundamental uncertainty as 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. Let (π_{jt}, a_{jt}) 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 ($j \in \{1, \dots, n_t\}$). The investor's expected, discounted real income is then given by:

$$E_t[\pi] = \sum_{j=1}^{n_t} \sum_{s=t}^{\infty} \frac{\pi_{js}}{P_s^C} \frac{R_s}{R_{t-1}}$$

where the discount factor R_s is given by:

$$R_s = \prod_{v=0}^s \frac{1}{1 + i_s}$$

where i_s is the nominal interest 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 portfolio's return is the same as what can be achieved through the future nominal interest rate i_s , whereas if $E_t[\pi] > 0$, the portfolio gives a better return than the future nominal interest rate. 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 current profitability can be obtained by looking at the *current* cross-section:

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

where $N_t > n_t$ are all existing firms at time t and r_t is the current real interest rate. 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 current profitability, as it is recognized that it may take many years before positive earnings from the mature firms are achieved.

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

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

The value \hat{s}_t is a so-called Sharpe Ratio (Sharpe, 1994) that measures a portfolio's return compared to a normal return, adjusted for risk. 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*.

Portfolio size. 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 perceived (drawn from a Pareto distribution).

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

$$\frac{\Delta n_{t+1}}{n_t} = \xi \cdot s_{t+1}^e, \quad \xi > 0, \quad (15)$$

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

There are two things to comment on regarding the behavioral relationship (15). Firstly, the profit measure s_{t+1}^e 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 data¹². 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

¹²On the other hand, it is almost by definition impossible to assume perfect foresight or rational expectations.

interaction between investor behavior (15) 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 (15) 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.

2.5. *Equilibrium and self-organization*

Given the sub-elements, we can characterize an equilibrium of the system. Doing so will facilitate the interpretation of the model's mechanisms below.

Consider a configuration in which all firms charge the same market price and pay the same market wage. Each period, the investor opens a constant number of new firms, and an equal number of firms exit because they are outcompeted. The system is then in a steady state with a constant total number of firms and a time-invariant age distribution. This steady state is relatively elaborate.

In the labour market, we observe constant structural employment and unemployment: search frictions imply a stable non-zero unemployment rate (as documented when we run the model in the next section 3).

As shown in Section 2.3.1, when firms choose optimal output and employment, each firm has an optimal scale conditional on the real wage: higher real wages imply lower optimal employment. If investor open many new firms (raising the total number of active firms), the fixed pool of structurally employed workers must be shared across more producers. Equilibrium therefore requires a higher real wage, which compresses each firm's desired employment.

Higher real wages reduce profits. Hence discounted profits net of entry costs are lower, and expected profitability falls. *There is thus a negative relationship between the per-period inflow of new firms and expected profitability.* An equilibrium obtains when the inflow is such that expected profitability is zero.

This negative relationship between firm entry and expected profitability is a central mechanism of self-organization. If investor raise entry to pursue short-run gains, long-run profitability is eroded. If entry remains above its equilibrium level for an extended period, investments fail to earn a normal return (expected profitability becomes negative). As specified in Section 2.4.2, this induces investors to scale back entry, restoring the equilibrium inflow. Subsequent sections provide several illustrations.

3. Results

The model is simulated in a baseline run and under two supply-side shocks. The aim is to demonstrate that the economy exhibits self-organising behaviour.

We implement (i) a permanent Harrod-neutral (labour-augmenting) technology shock and (ii) a temporary shock to the number of active firms. When labour productivity rises permanently in a closed economy, standard intuition suggests higher aggregate output and real wages that rise in proportion to the shock; the simulations confirm that this pattern emerges in the model. In the second experiment, a random fraction of firms is forced to exit. The results show that the system is self-organising in the sense that firm entry eventually restores the long-run number of firms and key macroeconomic aggregates return to their pre-shock levels.

The economy comprises approximately 17,000 households and 750 firms. Each simulation spans 100 years at monthly frequency. The baseline and both shock scenarios are replicated 200 times, enabling computation of means and confidence intervals. The Monte Carlo exercise is executed in parallel on the Danish supercomputer Computerome. The model is programmed in C#; open-source code is publicly available¹³.

Owing to search frictions, household income dynamics and stochastic mortality, the model exhibits stochastic behavior. As noted above, we execute 200 independent replications to quantify the role of this randomness. Random draws are produced by the pseudo-random number generator in C#. Given an integer seed, the generator yields a deterministic infinite sequence of pseudo-random numbers. Each shock scenario is paired with its baseline counterpart by initializing both simulations from the same seed, ensuring a controlled comparison of outcomes.

The parameters of the model are not calibrated to a particular database or economy; the model is intended as a stylized illustration. Several robustness checks have been conducted; overall, the qualitative results appear robust. Stable convergence to a steady state requires that both prices and wages satisfy the competition-based pricing rules. The framework will therefore be sensitive to alternative assumptions about wage setting—an important direction for future research and a natural channel through which Keynesian mechanisms could enter.

All simulations begin with three burn-in phases. Phase 1 lasts 25 years: new firms enter at a constant rate, no exits occur, and the interest rate is fixed (exogenous—not chosen by the investor). Phase 2 lasts 15 years: loss-making firms may exit. Phase 3 lasts 60 years: firm creation is determined by the investor, while the interest rate remains fixed. At the end of Phase 3 (100 years cumulatively), the interest rate is made endogenous and the economy runs without these restrictions. After an additional 100 years—to allow convergence to equilibrium—shocks are imposed

¹³Github-henvisning.

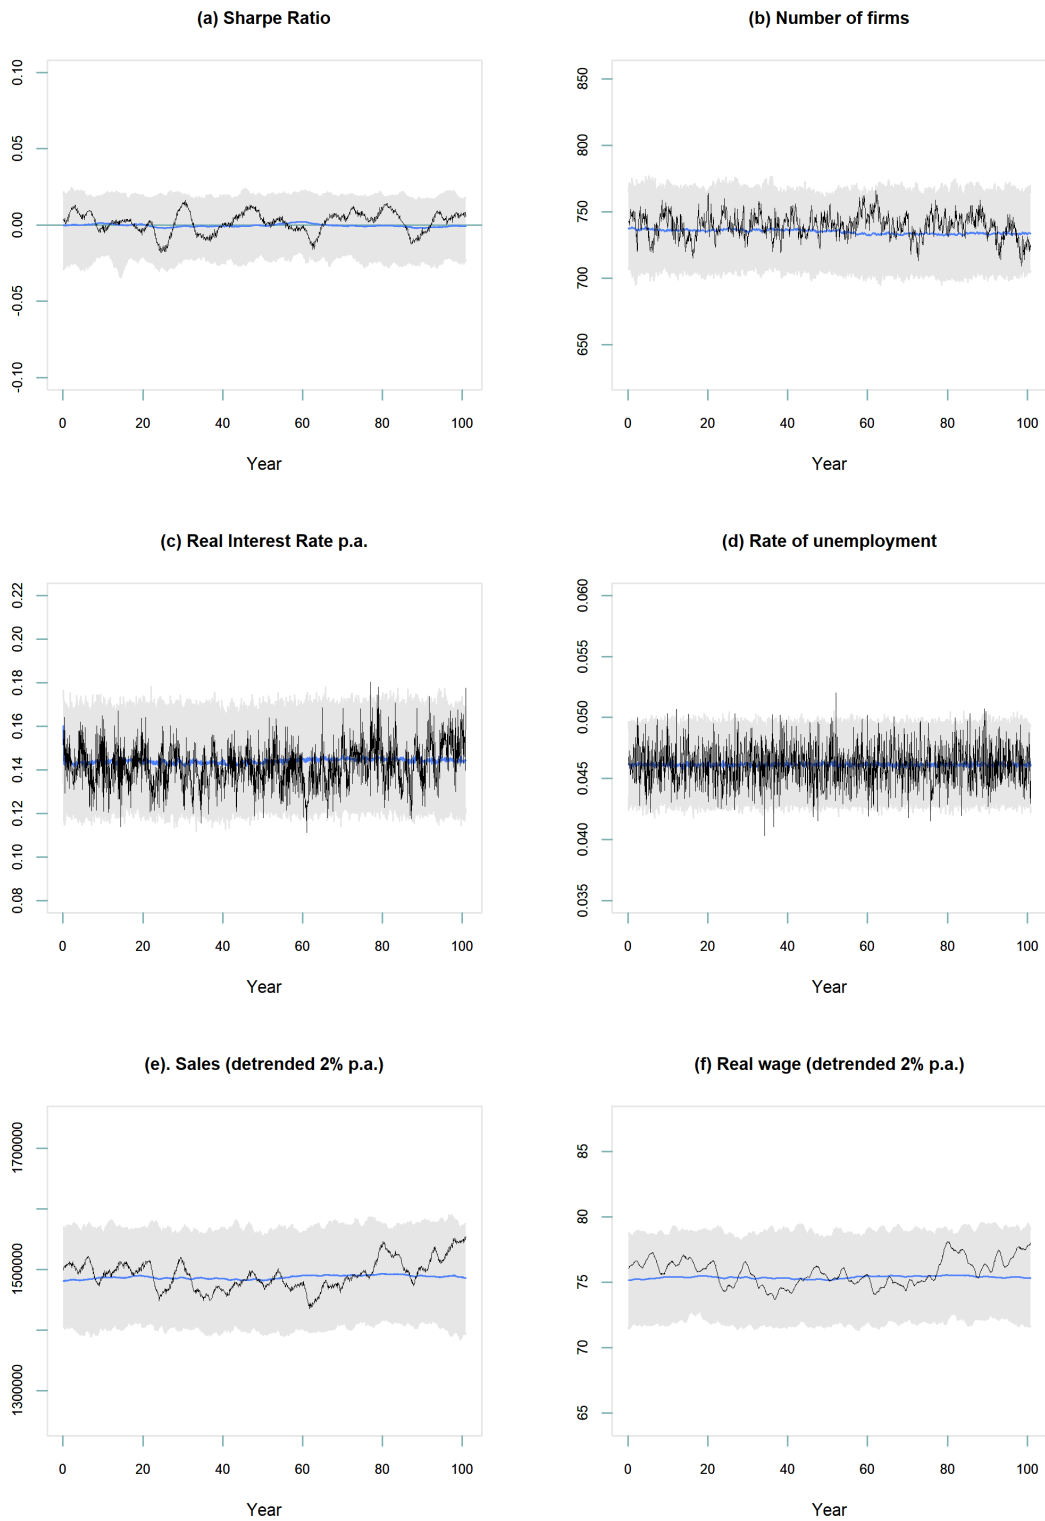


Figure 2: Base run. 100 years of monthly data. The blue curve averages 200 runs. The black curve is an example of a specific run. The shaded areas marks 95 percent confidence intervals.

and outcomes over the subsequent 100 years are recorded. In total, each run spans 300 years.

3.1. Base run

The baseline simulations are summarised in Figure 2. In each of the six panels, the blue line plots the mean across 200 replications, the grey band denotes the 95% confidence interval, and the black line shows one representative draw (the first replication). The mean paths are remarkably stable. For all variables, confidence bands remain fairly constant in width and are close to symmetric around the mean.

Sales and the real wage in Figure 2 are displayed after removing a deterministic 2 per cent annual trend. The near-constancy of the detrended series therefore indicates that sales and the real wage grow at roughly 2 per cent per year in the model. This suggests that market clearing operates as intended and that the qualitative dynamics match prior expectations. Entrants' productivity grows at 2 per cent per year, so new firms gradually displace older, less productive incumbents. Rising real wages induce exit among low-productivity producers, raising average productivity. Figure 2 also shows stable employment. Higher productivity therefore maps into higher output, which can be absorbed by aggregate demand because real wages rise while employment remains approximately unchanged. Self-organising forces keep the economy very close to steady state.

Inspecting individual paths (the black lines) reveals fluctuations around a level or trend. Expected profitability (Figure 2a) oscillates between about -2 per cent and 2 per cent, remaining close to zero—a central mechanism underlying self-organisation. The unemployment rate (Figure 2d) ranges from 4.25 per cent to 5.0 per cent, with a mean of 4.7 per cent. Thus, the economy exhibits cyclical variation, yet deviations from central tendency are modest. With millions of households and hundreds of thousands of firms, such cycles would appear negligible at aggregate frequency. In sum, the framework can generate business-cycle dynamics, but their amplitude is small under competition-based pricing. Equivalently, the economy is strongly self-organising.

3.2. Counterfactuals

Two experiments are conducted. The first is a permanent productivity shock; the second is a shock to the number of active firms. Both are superimposed on each of the 200 baseline replications described above. Mean responses and 95% confidence intervals for the first six years are reported in Figures 3 and 4. In appendix the first 50 years is shown.

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 explainable. The assumption of competition-based pricing makes it relatively easy to explain the price and wage dynamics of the model. Hopefully, this will be clear from what follows.

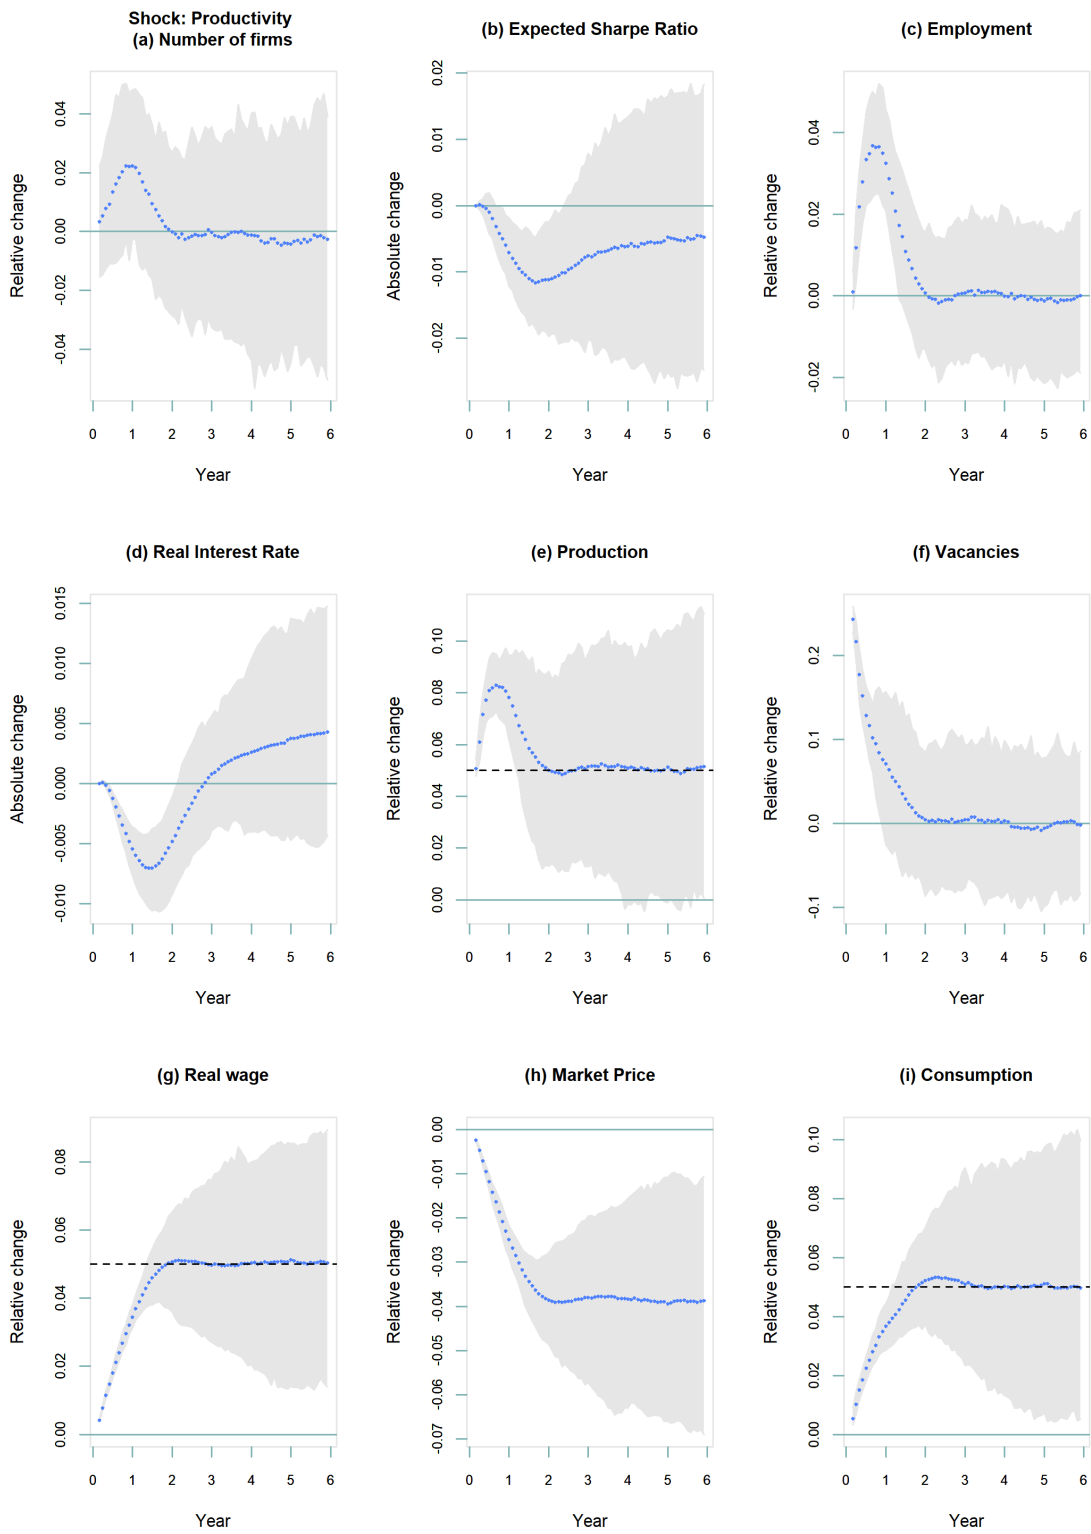


Figure 3: 5 pct. shock to productivity. 6 years of monthly data.
The blue points averages 200 runs. The shaded areas marks 95 percent confidence intervals.

Permanent productivity shock

The effect of a 5 per cent permanent increase in labor-augmenting productivity is shown in Figure 3. The shock can be interpreted as an innovation that diffuses simultaneously to all firms. Overall, the economy adjusts fairly quickly: a new equilibrium is reached within about two to four years. In the long run, the number of firms and employment are largely unchanged, while real wages and output rise by roughly 5 per cent. Although the model is not neoclassical, it behaves in this respect as a neoclassical model would. This underscores that the system is self-organizing.

In the short run, the productivity gain produces a boom lasting a few years. On impact, firms can suddenly produce about 5 per cent more. Goods supply therefore jumps by roughly 5 per cent, whereas consumer demand (sales) is almost unchanged (Figure 3, panels (e) and (i)). Consumption is unchanged because, for given prices, wages, and wealth, households do not change their behaviour. The goods market is thus left with substantial excess supply. Each firm cuts its price to attract customers, which pulls the market price down—the 'deviate' part of copy–deviate behaviour. Figure 3 (h) shows the market price falling by about 4 per cent over the first two years.

Higher productivity implies, for a given real wage, higher optimal employment and output at the firm level (see (9) and (10))—a short-run change in the intensive margin. Figure 3 (f) shows a sharp immediate rise in vacancies. Some firms bid up wages to hire workers, which lifts the market wage. The wage increase is more muted than the price decline; taken together, the movements in prices and wages raise real wages by about 5 per cent over the first two years. That real-wage response largely offsets the intensive-margin effect: within two years, employment returns to its initial level and supply is close to the 5 per cent increase implied by the productivity increase.

Figure 3 (a) and (b) show that the number of firms rises in the short run while expected profitability (the expected Sharpe ratio) falls. That may seem counterintuitive at first, since expected profitability governs how many new firms are started. The reason is that the increase in the number of firms is not due to more new firms opening, but to a decline in the number of firms that exit. Because aggregate productivity is higher, firms that would otherwise have been shut down on account of low idiosyncratic productivity are allowed to survive longer.

Appendix B shows the long-run dynamics (50 years). It reports that the expected Sharpe ratio, employment, and vacancies return to their initial levels in the long run. Output and consumption are, in the long run, 5 per cent higher—corresponding exactly to the size of the shock. The real interest rate rises by about 0.2 percentage points. In a textbook neoclassical Ramsey model, a permanent Harrod-neutral productivity shock would not change the real interest rate in the long run. However, it is well known from the literature (Diamond, 1965; Galor, 1988; Galor and Ryder, 1989; de la Croix and Michel, 2002; Auerbach and Kotlikoff, 1987) that the real interest rate can be endogenous in models with overlapping generations. A permanent increase in Harrod-neutral

productivity implies in this model a (small) positive long-run effect on the real interest rate. This explains why the real wage rises somewhat less than 5 per cent and why the number of firms declines slightly in the long run. The Sharpe ratio measures profitability relative to the real interest rate; a higher real interest rate therefore raises the required profitability.

Firm destruction

In this shock, 5 per cent of firms are shut down. The closed firms are drawn at random; at the time of the shock, all employees are laid off and the investor loses her investment. The shock may be interpreted, for example, as a natural disaster with no loss of human life. The interesting question is whether the system is sufficiently self-regulating for firms to be restarted. The answer is yes. After less than two years, the number of firms has been restored. By contrast, it takes more than 30 years before original productivity and real wages are recovered. This reflects the long time needed to rebuild the original firm structure (the age and productivity distribution of firms).

The effect of the shock over the first six years is shown in Figure 4¹⁴. It can be seen that the number of firms, employment, output, and vacancies all fall by 5 per cent upon impact (a, c, e, and f in Figure 4). The decline in employment implies a substantial increase in unemployment. After just one month, firms therefore observe a large rise in job applications. This explains why surviving firms reduce vacancies: with many unemployed workers, firms need to post fewer vacancies for a given number of hires (typical search behaviour). As can be seen from Figure 4 (f), vacancies fall by 10 per cent in period 2. Easier access to labour leads each firm to wish to set its wage below the market wage—but because all firms do so, the market wage falls. It appears from Figure 4 (g) that after 12 months the real wage has fallen by 2 per cent.

On the goods market, output falls immediately by 5 per cent (e in Figure 4). Private consumption falls by slightly less than 5 per cent upon impact (i in Figure 4) and by only about 1 per cent in period 2. The limited decline in period 2 reflects partly that many unemployed households find jobs (a competition-based labour market is very effective), and partly that unemployed households run down savings so that wealth declines. In the early years, output falls by more than consumption, so there is excess demand. This leads firms to wish to set their price above the market price, and the market price therefore rises (h in Figure 4). The higher prices imply a positive expected Sharpe ratio, and the number of new firms increases (a and b in Figure 4).

Appendix C shows that, over a 50-year horizon, all variables return close to their initial levels. The shock is non-permanent, and the system returns to its original state. This demonstrates a substantial degree of self-organization. For output, private consumption, the real wage, and the market price, it takes roughly 30 years before equilibrium is restored. The reason is the overlapping

¹⁴The first 50 years are also shown in appendix C

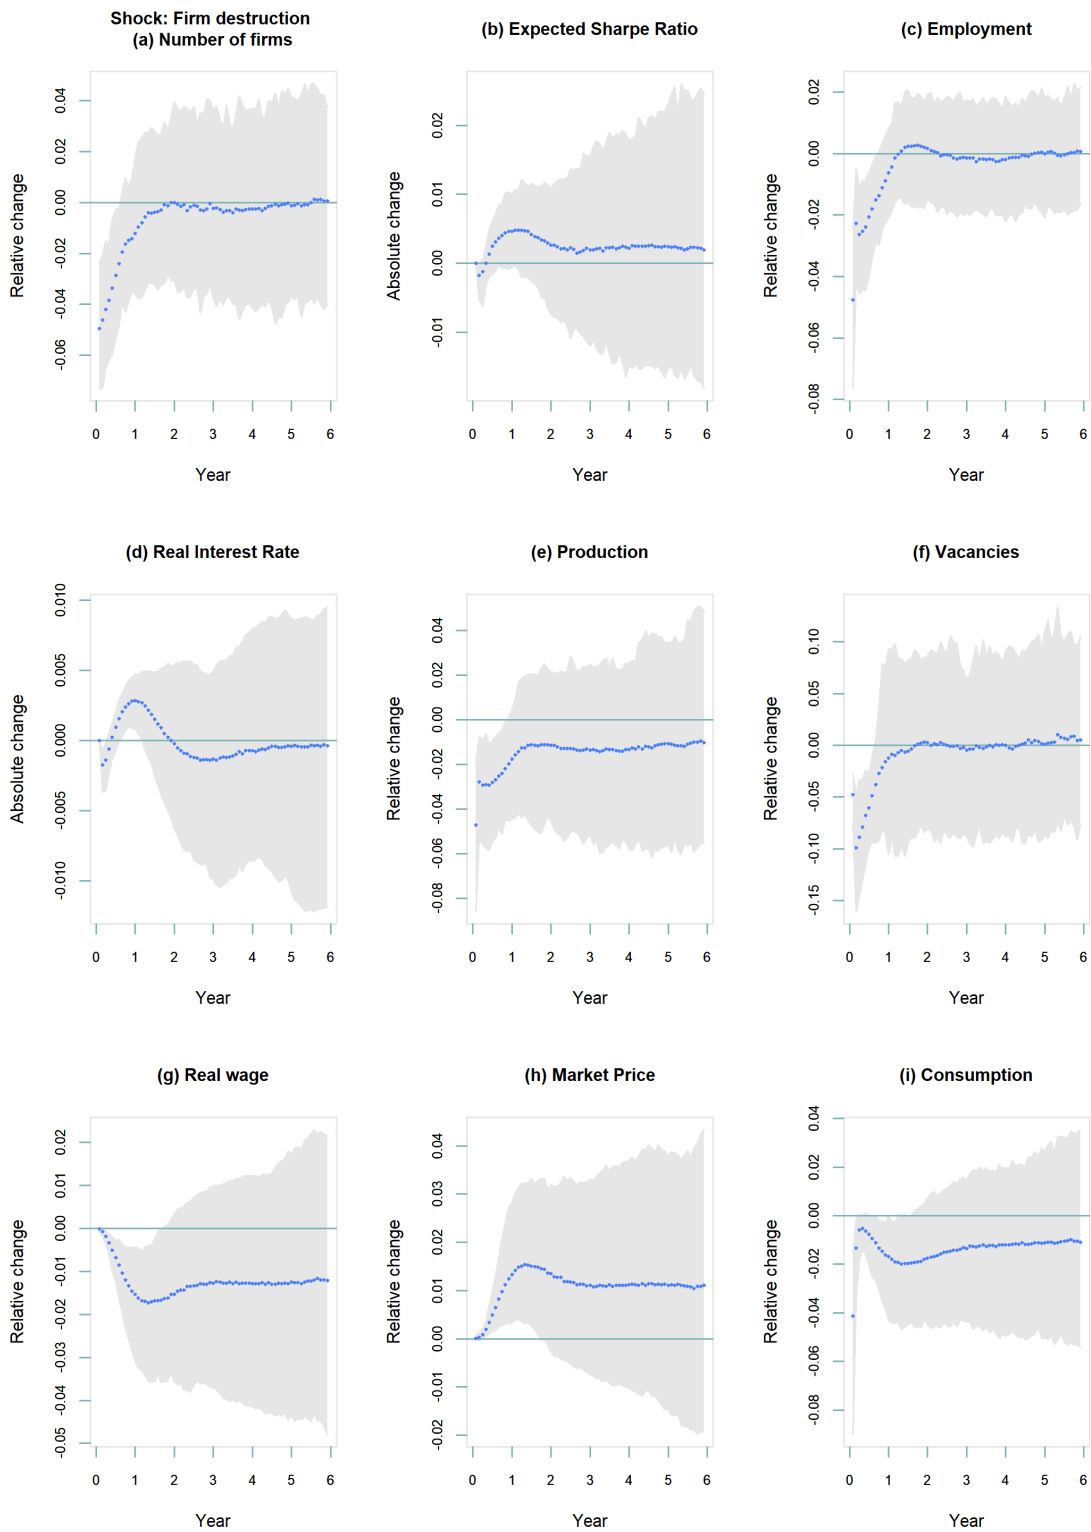


Figure 4: Destruction of 5 pct. of firms. 6 years of monthly data.
The blue points averages 200 runs. The shaded areas marks 95 percent confidence intervals.

generations of both households and firms. It takes many years to re-establish the distributions that prevails in equilibrium. This entails a decline in average productivity, which explains lower output (for given employment) and a lower real wage. The large destructive shock leaves persistent scars that fade over the long run.

4. Conclusion

This paper develops an agent-based macroeconomic model in which decentralized interaction—rather than an external equilibrium construct—coordinates a large economy with heterogeneous households and firms. The analysis highlights four interacting mechanisms: S-shaped firm-level technologies that pin down target scale, explicit network search (with advertising for entrants) in product and labor markets, competition-based pricing and wage setting that formalize copy–deviate social dynamics, and investor-driven entry governed by a Sharpe-style profitability signal together with a banking sector that transfers firm earnings to depositors through a short nominal interest rate.

The numerical results support two main claims. First, in long baseline simulations with exponential-smoothing expectations and competition-based rules, the economy remains close to a balanced growth path: detrended series are stable, expected profitability hovers near zero, structural unemployment is nearly constant, and realized cycles are modest—consistent with strong self-stabilizing forces in prices and wages. Second, supply-side experiments behave in economically interpretable ways. A uniform permanent productivity increase raises real wages and output in proportion while other stationary ratios revert after a short transition driven by temporary excess supply, price cuts, and vacancy dynamics; the short-run firm count response can reflect lower exit rather than higher entry. A random destruction of a fraction of firms is followed by rapid re-entry that restores the long-run number of firms within several years, but the recovery of aggregate productivity and real wages can take multiple decades because replicating the pre-shock firm age composition is inherently slow.

The framework is a principled prototype, not a tight empirical match to a particular country or sample. Its purpose is to show that self-organization can be demonstrated with explicit protocols and replicated simulations, and to make price and wage dynamics traceable in shock responses. The most important sensitivity noted in the analysis is that long-run coherence relies on competition-based pricing on both sides of the market; alternative wage-setting assumptions are a natural direction for extending the model, together with richer fiscal and social insurance institutions already suggested by the household side of the environment.

Open-source code accompanies the project, reflecting the view that large decentralized models should be inspectable. Future work can build on the same event architecture to add modules—savings motives, credit constraints, nominal rigidities, or public budgets—while preserving the paper’s central question: which market processes, made explicit, suffice to organize aggregate outcomes?

Finally, the exercise is meant to complement—not replace—more aggregated equilibrium analyses. Where DSGE models summarize matching and pricing in compact equations, this framework

exposes the same economic forces as operational rules that can be stressed, audited, and extended module by module. The payoff is a clearer map from assumptions about behavior to time-series outcomes, at the cost of heavier computation and a heavier burden of documentation. That trade-off is increasingly central as dynamic economics continues to move toward heterogeneous-agent and computational environments.

References

- Arrow, K.J., Block, H.D., Hurwicz, L., 1959. On the stability of the competitive equilibrium, II. *Econometrica* 27, 82–109.
- Arrow, K.J., Debreu, G., 1954. Existence of an equilibrium for a competitive economy. *Econometrica* 22, 265–290.
- Auerbach, A.J., Kotlikoff, L.J., 1987. *Dynamic Fiscal Policy*. Cambridge University Press.
- Brown, R.G., 1956. *Exponential Smoothing for Predicting Demand*. Arthur D. Little Inc., Cambridge, Massachusetts.
- Carroll, C.D., 1997. Buffer-stock saving and the life cycle/permanent income hypothesis. *Quarterly Journal of Economics* 112, 1–56.
- de la Croix, D., Michel, P., 2002. *A Theory of Economic Growth: Dynamics and Policy in Overlapping Generations*. Cambridge University Press.
- Dawid, H., Delli Gatti, D., 2018. Agent-based macroeconomics, in: *Handbook of Computational Economics*. volume 4, pp. 63–156.
- De Vroey, M., 2016. *A History of Macroeconomics from Keynes to Lucas and Beyond*. Cambridge University Press.
- Debreu, G., 1952. A social equilibrium existence theorem. *Proceedings of the National Academy of Sciences of the United States of America* 38, 886–893.
- Diamond, P.A., 1965. National debt in a neoclassical growth model. *American Economic Review* .
- Dixit, A.K., Stiglitz, J.E., 1977. Monopolistic competition and optimum product diversity. *American Economic Review* 67, 297–308.
- Eaton, J., Kortum, S., 2002. Technology, geography, and trade. *Econometrica* 70, 1741–1779.
- Galor, O., 1988. The long-run implications of a hicks-neutral technical progress. *International Economic Review* 29, 177–183.
- Galor, O., Ryder, H.E., 1989. Existence, uniqueness, and stability of equilibrium in an overlapping generations model with productive capital. *Journal of Economic Theory* 49, 360–375.
- Gavrilov, L.A., Gavrilova, N.S., 1992. The biology of life span: A quantitative approach. *Journal of the Institute of Actuaries* 119, 379–380.

- Gintis, H., 2007. The dynamics of general equilibrium. *Economic Journal* 117, 1280–1309.
- Hahn, F.H., Negishi, T., 1962. A theorem on non-tatonnement stability. *Econometrica* 30, 463–469.
- Hansen, M., Zangenberg, J., Markeprand, T., Kargo, A., Sonderby, M., 2021. Demografi, socioekonomi og boligstruktur i danske kommuner. Technical Report. DREAM. URL: <https://dreamgruppen.dk/publikationer/2021/maj/demografi-socioekonomi-og-boligstruktur-i-danske-kommuner/>. in Danish.
- Hinterhuber, A., 2008. Customer value-based pricing strategies: Why companies resist. *Journal of Business Strategy* 29, 41–50.
- Holt, C.C., 1957. Forecasting Trends and Seasonal by Exponentially Weighted Averages. Memorandum 52. Office of Naval Research.
- Hopenhayn, H.A., 1992. Entry, exit, and firm dynamics in long run equilibrium. *Econometrica* 60, 1127–1150.
- Hyndman, R.J., Athanasopoulos, G., 2024. Forecasting: Principles and Practice. 3 ed., OTexts, Melbourne, Australia.
- Hyndman, R.J., Koehler, A.B., Ord, J.K., Snyder, R.D., 2008. Forecasting with Exponential Smoothing: The State Space Approach. Springer-Verlag.
- Kaplan, G., Menzio, G., 2016. Shopping externalities and self-fulfilling unemployment fluctuations. *Journal of Political Economy* 124, 771–825.
- Kotler, P., Armstrong, G., Harris, L.C., He, H., 2021. Principles of Marketing. Eighth european edition ed., Pearson.
- LeBaron, B., Tesfatsion, L., 2008. Modeling macroeconomies as open-ended dynamic systems of interacting agents. *American Economic Review* 98, 246–250. Papers and Proceedings.
- Leijonhufvud, A., 1968. On Keynesian Economics and the Economics of Keynes: A Study in Monetary Theory. Oxford University Press, New York.
- Leijonhufvud, A., 2006. Agent-based macro, in: Tesfatsion, L., Judd, K.L. (Eds.), *Handbook of Computational Economics*. Elsevier, Amsterdam. volume 2, pp. 1625–1638.
- Melitz, M.J., 2003. The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica* 71, 1695–1725.

- Mortensen, D., Pissarides, C., 1994. Job creation and job destruction in the theory of unemployment. *Review of Economic Studies* 61, 397–415.
- Pissarides, C., 2000. *Equilibrium Unemployment Theory*. 2 ed., MIT Press.
- Qiu, Z., Rios-Rull, J.V., 2019. Directed search, nominal rigidities and markup cyclicity. Working paper.
- Quinn, S., 1997. Goldsmith-banking: Mutual acceptance and interbanker clearing in restoration london. *Explorations in Economic History* 34, 411–432.
- Sharpe, W.F., 1994. The sharpe ratio. *The Journal of Portfolio Management* 21, 49–58.
- Smale, S., 1976. Exchange processes with price adjustment. *Journal of Mathematical Economics* 3, 211–226.
- Stephensen, P., 2015. The event-pump: An agent based approach to microsimulation. URL: <https://dreamgroup.dk/publications/2015/september/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.
- Storesletten, K., Rios-Rull, J.V., Bai, Y., 2011. Demand shocks that look like productivity shocks, in: 2011 Meeting Papers, Society for Economic Dynamics.
- Syverson, C., 2011. What determines productivity? *Journal of Economic Literature* 49, 326–365.
- Tesfatsion, L., 2006. Agent-based computational modeling and macroeconomics, in: Colander, D. (Ed.), *Post-Walrasian Macroeconomics: Beyond the Dynamic Stochastic General Equilibrium Model*. Cambridge University Press, Cambridge, pp. 175–201.
- Uzawa, H., 1962. Production functions with constant elasticities of substitution. *Review of Economic Studies* 29, 291–299.
- Walras, L., 1854. *Éléments d'économie politique pure, ou théorie de la richesse sociale*. Guillaumin, Paris.
- Wright, R., Kircher, P., Julien, B., Guerrieri, V., 2021. Directed search and competitive search equilibrium: A guided tour. *Journal of Economic Literature* 59, 90–148.

Appendix A: Determination of smoothing parameter for long-run forecast

Expectations are defined by Simple Exponential Smoothing:

$$y_t^e = \lambda y_{t-1}^e + (1 - \lambda) y_t$$

Assume we are in a situation where $y_0^e = y_0 = 1$. For $t \geq 1$, there is a permanent increase in y_t such that $y_t = 1 + \phi$, $t \geq 1$. We want to know what λ should be for $y_t = 1 + \eta\phi$ for a given $t = T$. Or, in other words: what should λ be for a fraction η of the structural shift ϕ to be "recognized" at time T.

It can be shown using standard calculus that

$$y_t^e = 1 + \phi - \phi\lambda^t$$

for $t \geq 1$.

If we want $y_t^e = 1 + \eta\phi$ for $t = T$, it must hold that

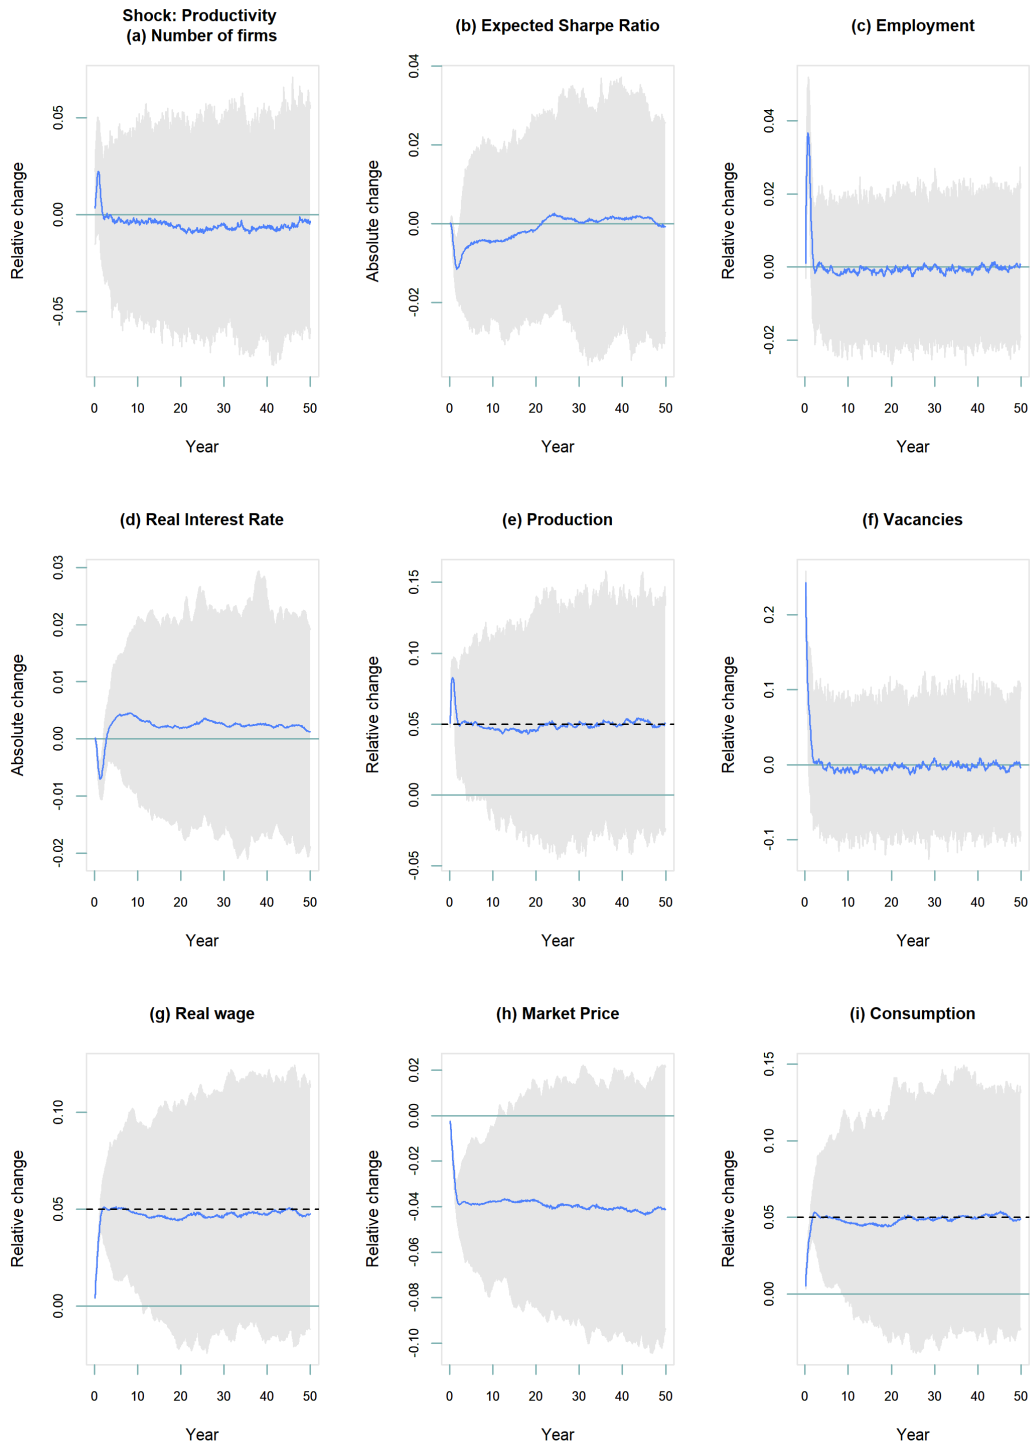
$$1 + \eta\phi = 1 + \phi - \phi\lambda^T$$

such that

$$\lambda = (1 - \eta)^{\frac{1}{T}}$$

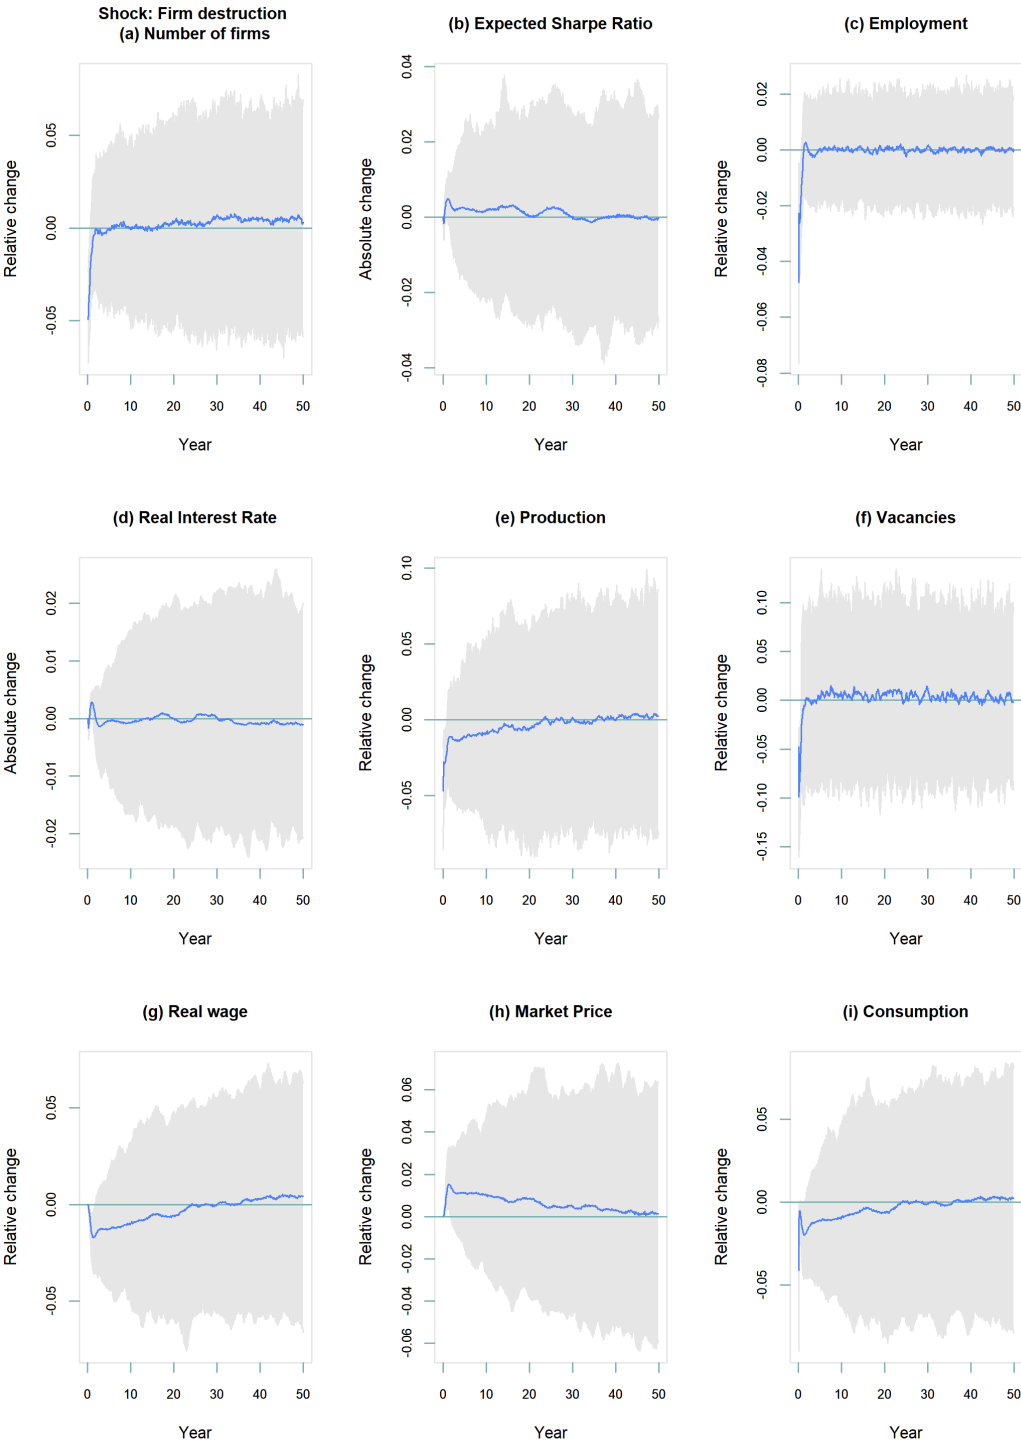
If $\eta = 0.8$ and $T = 3 \cdot 12 = 36$ we have that $\lambda = 0.96$. Note that this result holds regardless of the value of the structural change ϕ .

Appendix B: Productivity shock. 50 years.



The blue points averages 200 runs. The shaded area marks 95 percent confidence intervals.

Appendix C: Firm destruction shock. 50 years.



The blue points averages 200 runs. The shaded area marks 95 percent confidence intervals.