Equal chances, unequal outcomes? Network-based evolutionary learning and the industrial dynamics of superstar firms

Jan Schulz¹ · Daniel M. Mayerhoffer¹

Accepted: 12 May 2021 / Published online: 29 May 2021
© The Author(s) 2021

Abstract

With the advent of platform economies and the increasing availability of online price comparisons, many empirical markets now select on relative rather than absolute performance. This feature might give rise to the ‘winner takes all/most’ phenomenon, where tiny initial productivity differences amount to large differences in market shares. We study the effect of heterogeneous initial productivities arising from locally segregated markets on aggregate outcomes, e.g., regarding revenue distributions. Several of those firm-level characteristics follow distributional regularities or ‘scaling laws’ (Brock in *Ind Corp Change* 8(3):409–446, 1999). Among the most prominent are Zipf’s law describing the largest firms’ extremely concentrated size distribution and the robustly fat-tailed nature of firm size growth rates, indicating a high frequency of extreme growth events. Dosi et al. (*Ind Corp Change* 26(2):187–210, 2017b) recently proposed a model of evolutionary learning that can simultaneously explain many of these regularities. We propose a parsimonious extension to their model to examine the effect for deviations in market structure from global competition, implicitly assumed in Dosi et al. (2017b). This extension makes it possible to disentangle the effects of two modes of competition: the global competition for sales and the localised competition for market power, giving rise to industry-specific entry productivity. We find that the empirically well-established combination of ‘superstar firms’ and Zipf tail is consistent only with a knife-edge scenario in the neighbourhood of most intensive local competition. Our model also contests the conventional wisdom derived from a general equilibrium setting that maximum competition leads to minimum concentration of revenue (Silvestre in *J Econ Lit* 31(1):105–141, 1993). We find that most intensive local competition leads to the highest concentration, whilst the lowest concentration appears for a mild degree of (local) oligopoly. Paradoxically, a level playing field in initial conditions might induce extreme concentration in market outcomes.

Keywords Agent-based modelling · Replicator dynamics · Pareto distribution · Fat tails · Evolutionary learning · Competition

JEL Classification C63 · D21 · D43 · L11 · L13 · L14

1 Introduction

Within an increasing number of markets, an individual firm's fate is no longer determined by absolute performance but by its performance relative to its competitors. Put differently, individual success in the market is a function of the now widespread availability of price comparisons on the internet (Akerman et al. 2021) and platform competition (Autor et al. 2020). In those 'winner takes all/most' markets, tiny differences in initial productivity can manifest themselves into large differences in market shares, typically leading to a high emergent concentration of market power. Somewhat surprisingly, the determinants of initial productivity at market entry have received little scholarly attention. We systematically explore the aggregate effects of heterogeneous initial conditions by exploiting a plausible notion of industry-specific productivity within locally segregated markets. Our approach builds on the intuition that comparisons of relative performance are seldom global and typically localised. For example, Uber and Alphabet not competing within the same submarket, even though their business strategies building on network effects and intangibles appear to be very similar. Our results suggest that market segmentation exhibits sizeable and counter-intuitive effects on the distribution of market shares, on firm growth and firm survival.

Empirically, these aforementioned firm-level characteristics follow distributional regularities or 'scaling laws' (Brock 1999), whose underlying mechanisms require explanation. Dosi et al. (2017b) recently proposed a model of industrial dynamics that features evolutionary learning: Individual firms innovate and increase their productivity but compete for market shares according to a global selection mechanism based on productivity. This *evolutionary learning* mechanism combines cumulative learning with a 'winner takes all/most' market structure. Despite its bare-bones, partial equilibrium nature, the model is able to explain a surprising number of stylised facts in industrial dynamics, such as strongly heterogeneous size distributions, scaling between size growth rates and their variance as well as (persistent) heterogeneity in productivity. The model has been applied successfully and essentially unchanged in a macroeconomic setting both for explanatory purposes and policy experiments in the 'Keynes meets Schumpeter' (K+S) modelling approach (Dosi et al. 2010, 2013, 2015, 2017a). Apart from the K+S approach, the distributional regularities, which even the partial model produces, have been identified to be of great macroeconomic relevance (Gabaix 2011; di Giovanni et al. 2011).

We propose a parsimonious extension to this model to examine the effect for deviations in market structure from global competition. Namely, we introduce a network structure of localised competition and innovation. This extension makes it possible to disentangle effects from two modes of competition: global competition for sales and localised competition for market power, giving rise to industry-specific productivity differences. Our contribution is thus twofold: Firstly, we test the benchmark model results' robustness for different market structures,

as defined by local competition for market shares and localised entry, where the entry process' precise nature has recently been identified as the most important driver of aggregate outcomes within the model (Dosi et al. 2018). By this, we are able to constrain the range of possible competitive mechanisms in light of the empirical evidence in more detail. Secondly, we take a complementary approach to the macroeconomic implementations and examine even further the micro processes that different competition structures induce. These microeconomic considerations have consequences for regulatory policy and implications even at the managerial level, in particular, for market entry timings.

Besides the cumulative idiosyncratic learning mechanism from our benchmark, for a non-complete network, our model also features a second process that distinguishes between firms' different productivity levels: When a new firm enters the model market, it acquires an industry-specific productivity level, modelled as the weighted average productivity of its direct competitors. This mechanism applies to a firm only once in its lifetime, namely at foundation/market entry. Nonetheless, it can crucially shape the entire life of a firm, since initial productivity determines whether an entrant can stabilise its position in the market or is quickly forced out of it again: We show that successful entrants typically join highly productive product markets. Hence, from a management perspective, our findings underline the importance of timing market entrance and thorough search prior to entry. Namely, our model can explain, why Schlichte et al. (2019) find the most successful firms to be fast followers in innovative markets rather than the original innovators themselves.

We find that the empirically well-established combination of 'superstar firms' and Zipf tails is consistent only with a knife-edge scenario in the neighbourhood of most intensive local competition. Moreover, our contests the conventional wisdom derived from a general equilibrium setting that maximum competition leads to minimum concentration of revenue (Silvestre 1993). Instead, we find that most intensive local competition leads to the highest concentration and the lowest concentration appears for a mild degree of (local) oligopoly. Relating to a different notion of competition, this finding might also be interpreted as evidence that 'winner takes all/most' markets are far from the ordoliberal ideal which considers "competition [to be] the most ingenious disempowerment instrument in history" (Böhm 1960, p. 22, author's translation). By contrast, it is precisely the ordoliberal demands for a 'level playing field' in combination with 'performance based competition' which lead to the highest concentration of revenue and hence power asymmetries within such markets that ordoliberals hope to avoid (Dold and Krieger 2019).

The remainder of this paper is organised as follows: We firstly discuss the stylised empirical facts that we intend to study as well as concepts and models on which we build (Sect. 2), followed by a detailed introduction of our model, based on Dosi et al. (2017b) (Sect. 3). Thereafter, we present the core simulation outcomes and explain their generating mechanisms (Sect. 4). Finally, we situate our findings in economic, policy and business discussions, drawing practical as well as normative implications, before closing with proposals for further research (Sect. 5).

2 Models of evolutionary learning as a representation of empirical findings

2.1 Stylised facts in industrial dynamics

As a selection criterion for the parameter range that our proposed model spans, we use a set of stylised facts from traditional microeconomic literature and the industrial dynamics literature on distributional regularities in various firm-specific variables. The most prominent of these regularities is the finding of Zipf's law, originally based in linguistics (Zipf 1949), for the upper tail of firm size distributions. This implies that the size distributions of the largest firms are extremely concentrated, where the second-largest firm has only approximately one half of the size of the largest, the third-largest only a third and so on. Zipf's law in firm sizes appears to be a genuine and universal characteristic of market economies. A non-exhaustive list of studies on Zipf's findings include Axtell (2001) for the US, di Giovanni et al. (2011) for France, Pascoal et al. (2016) for Portugal, Okuyama et al. (1999) for Japan, Kang et al. (2011) for the Republic of Korea, Zhang et al. (2009) along with Heinrich and Dai (2016) for China, and Fujiwara et al. (2004) for several European countries. This empirical regularity not only constrains the set of possible generating mechanisms, but Zipf's law has also been linked to several important economic phenomena, such as the surge in CEO payments in recent decades (Gabaix and Landier 2008), the explanation of aggregate fluctuations from the micro-level together with increases in aggregate volatility (Gabaix 2011) and the welfare effects of barriers to entry and trade liberalisation (Di Giovanni and Levchenko 2013). This whole strand of macroeconomic literature takes Zipf's law as their starting point but does not examine possible conditions for which it emerges. Our findings on the determinants for Zipf's law might also provide insights on how to influence this wide range of phenomena, from CEO payments to aggregate fluctuations and international trade, which from a structural perspective this literature takes as given.

Another phenomenon of similar attributed economic relevance is the recent emergence of 'superstar firms', that operate in 'winner takes most/all' markets. They have experienced substantial and sustained increases in revenue over relatively short amounts of time (Autor et al. 2020). Anecdotal examples for this behavior are Alphabet and Uber. The rise of these firms has been proposed as an explanation for the recent decline in the labour share of national income (Autor et al. 2020), otherwise famously staying constant throughout the most part of recorded history of capitalist economies (Kaldor 1961) and the rise of wage inequality (Gabaix and Landier 2008).

Growth in various measures of size such as gross sales, total assets or number of employees for the whole range of firms has also been shown to be fat-tailed with relatively frequent extreme events, where empirical growth rate densities display a characteristic 'tent shape' on a semi-logarithmic scale, implying an exponential power functional form (Amaral et al. 1997; Bottazzi et al. 2001, 2002; Bottazzi and Secchi 2005, 2006; Alfarano and Milaković 2008; Bottazzi

et al. 2011; Erlingsson et al. 2013; Mundt et al. 2015). These distributions have frequently been identified as Laplacian (Kotz et al. 2012), which though has recently been challenged theoretically and empirically (Mundt et al. 2015). We stick to the exponential power or Subbotin (1923) shape and focus on the fat-tailed nature of size growth rates.

Concerning the firm age distribution, empirical findings are scarcer. However, the (limited) consensus appears to be that firm age is approximated well by an exponential distribution, as shown by a number of studies: Coad and Tamvada (2008) for several developing countries; Kinsella (2009) for Irish firms; Coad (2010) for the plant level in the US and Daepf et al. (2015) for publicly listed firms in the US. This exponential age distribution also has the crucial theoretical implication that death rates are independent of firm age, as Daepf et al. (2015) confirm empirically.

The main stylised facts which our model aims to replicate are thus: a Zipf law in the upper tail of firm sizes;¹ fat-tailed growth behaviour; and a high frequency of ‘superstar’ high growth events coupled with an exponential age distribution with a common insolvency probability for all firms, irrespective of age.

2.2 Agent-based models

Competition within industries, which produces the stylised facts discussed here, constitutes a socio-economic system. To study the dynamics in this system and find candidates for mechanisms that lead to empirically observed facts, an agent-based model (ABM) forms an adequate approach (Klein et al. 2018): Due to competitive interactions between individual firms, one cannot properly describe the system by additive aggregation of the model, but observed phenomena are emergent (Coleman 1990). Agent-based models can highlight and explain emergent phenomena and open the black box of competitive interactions on a system level in order to uncover interactions in and structures of subsystems (Hedström and Ylikoski 2010). Given the high level of idealisation, we do not intend to make quantitative predictions, but are nevertheless confident that the model can reveal central qualitative features of its economic target system (Grüne-Yanoff 2009).

2.3 An agent-based model of learning and selection

We introduce a layer of locality to the model of learning and selection of industrial dynamics by Dosi et al. (2017b).² This model understands learning as an increase in productivity by a random factor, as detailed in Sect. 2.4 for details. Gabaix (2009) identifies this stochastic process featuring a multiplicative component as an adequate generating mechanism for the empirically observed power-law distribution of firm size.

¹ The focus on the upper tail is partially motivated by the fact that about a third of variations in US GDP growth can be explained by the idiosyncratic destinies of the 100 largest firms (Gabaix 2011).

² Here, we only describe the basic notion of the benchmark model and our extension. For details and all equations, see Sect. 3.

Given the resulting heterogeneous and dynamic levels of productivity, market shares are allocated accordingly. Our proposed allocation mechanism makes use of a biological metaphor, the Darwinian ‘survival of the fittest’ principle, now in the form of a ‘replicator dynamics’ approach (Fisher 1930). The fittest or most productive firms grow to dominate the market, while less productive firms fall victim to competition and are driven out (Cantner 2017). In our formal description, we remain agnostic about the precise nature of the mechanism translating productivity increases into growing market shares to allow for a reasonably general application. These translation mechanisms by which higher market shares might accrue due to enhanced productivity include: increased product quality for given unit costs; decreasing unit costs for products of equal quality; freed up funds for increased marketing spending; or any other mechanism.

More specifically, in our representation the market share of a firm grows or shrinks according to how its productivity compares with the weighted average productivity of all firms in the model; thereafter, firms whose shares have fallen below a threshold leave the market and are replaced by new entrants. This constitutes a selective replicator dynamics process for which Cantner and Krüger (2008) as well as Cantner et al. (2012) present empirical evidence.

Dosi et al. (2017b) use these replicator dynamics in their model: Initially, all firms have equal market shares and productivity levels. At each time step, firms increase their productivity by a random factor, following which they gain or lose market share depending on how their own productivity compares with global average productivity. Firms whose share falls below a threshold value are replaced by new entrants. These new entrants start with the market share with which firms were initialised (and shares of all incumbents are adjusted so that the aggregate market size remains constant), but have their productivity level set to the current weighted global average productivity in the model. We carefully extend this model by adding a network layer to capture actual competitive interactions between agents. There are various ABM that emphasise the role of locality in economic interaction (Tesfatsion 2017). Our careful extension allows to gain more specific insights on the impact of localised competition without losing track of the core mechanisms. Hence, we validate the model by showing that the specific case of a complete network, which resembles the model by Dosi et al. (2017b), also displays a similar behaviour and yields the same stylised facts.

2.4 Productivity gain through stochastic learning

Our model features two channels of learning. Firstly, incumbent firms increase their productivity periodically. Secondly, the initial productivity of entrant firms depends on the localised market that they enter and hence they learn from their link neighbours.

The periodical learning describes the efforts of each firm to improve its productivity continually. In economics, the general concept of learning as a belief update justified by self-collected or socially acquired evidence (Zollman 2010) is often understood as a rational endeavour that agents explicitly control, as Evans

and Honkapohja (2013) point out in their overview. Moreover, approaches such as Bray and Savin (1986) or Milani (2007) reveal a close connection between learning and rational expectations. While such detailed understanding of learning is appropriate when investigating a learning process itself, a macroscopic approach seems sufficient for a study of industrial dynamics, where the outcome of learning contributes to one of many mechanisms. This macroscopic approach focuses on the productivity gain that any learning activities of firms yield. Thus, the model abstracts from details of the learning process and does not distinguish where (e.g. product improvements, production efficiency, marketing) or how (e.g. new inventions, imitation of others, deliberate management choices) the productivity gain takes place. When abstracting from subject-specific features of learning, one can treat success as being randomly distributed among individual learners and hence understand learning as an increase of productivity by a random factor. This stochastic learning seems to be an appropriate way to capture actual outcomes, as empirical findings are approximately represented (Luttmer 2007). Moreover, replicator dynamics are also consistent with an understanding of learning as imitating more successful others' behaviour. Schlag (1998) demonstrates this analytically by showing that the aggregate population behaviour follows a replicator dynamics whenever agents choose the individually most successful learning rule.

Accordingly, we follow Dosi et al. (2017b) and deliberately keep the learning process purely stochastic—do not explicitly include rational expectations or adaption to other firms—in order to focus on the network structure effects. With a purely stochastic process, we circumvent the problem that the precise form and effect of an innovation is per definition unpredictable and thus resort to much more modest statistical assumptions about the average rate of technological progress (Arrow 1991).

Besides the periodical stochastic learning of incumbent firms, the network layer and namely the localised market entry that it implies constitutes a second implicit mechanism of learning, which depends on asymmetric innovation. ABM studies that employ multiple or asymmetric learning processes in other contexts reveal unexpected system behaviour and have high explanatory power. For example, Klein and Marx (2018) and Klein et al. (2019) show in a model how asymmetric learning and information cascades shape individual estimates of how likely political revolution is. Asymmetric learning also plays a role in iterated games, as Macy and Flache (2002) show. Mayerhoffer (2018) runs a variant of the Hegselmann-Krause bounded confidence model (Hegselmann and Krause 2002) parallel to a network-based opinion update procedure and finds that the coexistence of both learning mechanisms can explain group-specific attitudes towards queerness among adolescents. In their model of Humean moral theory, Will and Hegselmann (2014) also employ explicit and implicit asymmetric learning in parallel. Models of learning and knowledge diffusion in networks also find application in business science, where they can provide explanations for competitive advantage, as Greve (2009) shows for shipbuilders and shipping companies. In these structures, Skilton and Bernardes (2015) find that successful market entry empirically depends on the network layout (Table 1).

Table 1 Parameters, initial values and calculation

Parameter	Variable	(Initial) values	Equation
Population size	N	150	<i>Constant</i>
Linking probability	p	0.01–1 (in steps of 0.01)	<i>Constant</i>
Local market competitors	K	Dependent on p	–
Firm Productivity	a	1	(1), (6) (at entry)
Global market share	s	$1/N$	(2)
Localised market power	l	$1/K$	(4)

The interpretation of variables is given in Sect. 3.1

3 Model description

This section provides a content-oriented presentation, see Table 1 for an overview of the chosen parameters; for technical details see the commented model, which is appended electronically³ and the description following the accompanying ODD protocol.⁴

3.1 Model properties and initialisation

The model observes a population of 150 firms that constitute an economy. We adopt this intuitively low number from Dosi et al. (2017b), but sensitivity analyses showed that our results also hold for larger populations. In this economy, firms try to maximise their sales revenue by improving their productivity through learning. However, whilst a firm does act rationally, this is only in a bounded manner due to its possession of imperfect information and environmental complexity. Hence, it does not adapt to the behaviour of other firms or form expectations. Undirected links connect some firms, but the firms themselves have no perception of their links.

Links between firms do not mean that they cooperate in research and development or in production; on the contrary, each link represents a direct competitive relationship between two firms in their selling of products that are (almost) perfect substitutes for each other. Pellegrino (2019) recently used the same methodology in a general equilibrium setting to identify aggregate trends and welfare costs of market power in the US. Thus, the model adds a level of locality to competition by linking firms. Clusters of densely linked firms represent an industry with aggravated competition. With this modelling approach, we combine two concepts of market structure that have enjoyed great success in the macroeconomic literature: Chamberlinian monopolistic competition (Chamberlin 1949; Robinson 1969) and the concept of a product space first introduced by Hidalgo and Hausman (Hidalgo and Hausmann 2009). From monopolistic competition, we take the notion that market power

³ The computer simulation was implemented in Netlogo (Wilensky 1999).

⁴ We describe the model according to the protocol guidelines by Grimm et al. (2010) and the extension proposed by Müller et al. (2013).

in differentiated, localised product markets is consistent with strong global competition, as indicated in the Chamberlinian concept by zero long-run profits. From Hidalgo and Hausmann (2009), we take the idea that similarity of products can be formalised by a network structure in a product space, where a network linkage indicates similarity.

In this context, to formalise local monopolistic competition within the product space, a Erdős–Rényi (ER) random network structure (Erdős and Rényi 1960) seems most appropriate because it does not call for an assumption that firms deliberately form competitive links. To generate the random network at initialisation, a link appears between any pair of firms with probability p . This probability is global, exogenously set and constant. $p = 1$ represents the model by Dosi et al. (2017b). One major advantage of ER networks is the myriad of analytical results pertaining to the network structure for different p . We make use of two results in particular. Firstly, the degree distribution of the network is asymptotically Poisson with thin tails (Newman 2005). This result provides assurance that our findings on the fat-tailed nature of growth rates and the size distribution are no artefacts of the network structure we impose, but rather a genuine emergent feature of interaction within the model.⁵ Secondly, there almost surely exists a single giant component for the range of network probabilities between 1 and 100% as the parameter range we consider. All other components have, almost surely, size of order $O(\log(N))$, where N is the number of firms. Hence, we can also examine the relevance of highly heterogeneous competitive environments, where the intensity of competition depends on whether or not a firm is connected to the single giant component or not (Erdős and Rényi 1960).

Besides its position in the network, each firm possesses three attributes that may vary over time. Firstly, there is the global market share s , which one can understand as sales revenue generated by each firm; it measures a firm's level of success and ultimately determines its survival. Initial shares of firms are equal: $\forall i(s_i(t_0)) = 1/N$. Secondly, local competition represented by the network structure means that firms also possess a localised market power l , which measures how productive a firm is in comparison with its immediate competitors that are the link-neighbours in our model. There is no immediate relationship between localised market power and sales revenue; high localised market power does not necessarily imply a high global share; for example, firms with great power in a small or unproductive industry may be small at the level of the whole market. Initial localised market powers are calculated following the same logic as global shares $\forall i(l_i(t_0)) = 1/|K_i|$, where K_i is the set containing all link-neighbours of i and i itself. Thirdly, productivity or level of competitiveness a of a firm indicates how well this firm is equipped for selling its products. It includes a firm's technological and business knowledge along with its skill base, but it could also be shaped by a specific demand for products that the

⁵ Note that even though empirical degree distributions are often fat-tailed warranting an explanation on their own (Johnson et al. 2014), there is no reason to impose a specific network structure a priori, as to the best of our knowledge, there exists no evidence on topologies of competition networks. We thus opt for as minimal assumptions as possible, leaving us with ER-type networks.

firm offers. This attribute improves through learning and in turn, the productivity of a firm impacts the global share and localised market power of the firm itself as well as of other firms. However, for the purposes of our model, firms start with an equal level of competitiveness: $\forall i(a_i(t_0)) = 1$.

3.2 Events during the simulation

The simulation proceeds in discrete time steps, within each of which the following processes take place in sequential order:

1. Learning: Firms (potentially) increase their productivity.
2. Assessment of global shares and localised market powers.
3. Entry and exit: Firms with low global market shares leave the market and new entrants replace them

3.2.1 Learning

In the model, learning incorporates all processes that improve a firm's level of competitiveness. This includes an intentional quest for innovations in product design, efficiency of production, and supply-chain management as well as marketing. At the same time, according to this concept, a firm can also 'learn' if customers grow more interested in its products independent of deliberate actions by the firm (e.g., products become popular due to some trends set by third parties). We subsume this variety of aspects under a firm-specific and idiosyncratic learning mechanism that we take from Dosi et al. (2017b), while the general concept of this learning dates back to earlier work by Dosi et al. (1995) who propose this as a baseline condition in their model. For this multiplicative stochastic process, each firm i determines its productivity a_i as follows:

$$a_i(t) = a_i(t-1)(1 + \pi_i(t)), \quad (1)$$

where $\pi_i(t)$ describes a firm's learning parameter and is drawn from a rescaled beta distribution with $\alpha = 1.0$, $\beta = 5.0$, $\beta_{\min} = 0.0$, $\beta_{\max} = 0.3$ and an upper notional limit $\mu_{\max} = 0.2$. This notional limit ensures that the maximum productivity growth rate is indeed 0.2, as we fix all drawn values higher than that for the notional limit. Firms do not experience negative learning because their productivity is measured in absolute terms rather than being compared with that of other firms at this stage. Learning in this sense does not entail failures which would imply negative productivity gains. This deliberate modelling choice excludes planning mistakes on the part of individual firm management, and isolates effects generated by the interplay between stochastic learning and market selection. Furthermore, absolute productivity growth depends on the firm's previous productivity level, meaning that the expected productivity gain grows proportionally to its past productivity. However, learning is

independent of firm size, and hence, there is no (direct) amplifier that would reward larger firms with a potential for higher rates of productivity gain.

3.2.2 Assessment of global shares

A replicator dynamics formulation reproducing the one used by Dosi et al. (2017b) determines the global market share:

$$s_i(t) = s_i(t-1) \left(1 + \frac{a_i(t) - \bar{a}(t)}{\bar{a}(t)} \right). \quad (2)$$

For any market containing at least two firms, it holds that $0 < s_i(t) < 1$. The global parameter $\bar{a}$ is the weighted average productivity of all firms N in the global market:

$$\bar{a}(t) = \sum_j a_j(t) s_j(t-1). \quad (3)$$

A firm's global share depends not only on its own productivity level, but also the weighted productivity levels of all other firms in the market. The weighting ensures that larger rather than smaller firms shape the market more distinctively. The sales revenue of a firm grows (or shrinks) according to how much the productivity level of this firm exceeds (or undercuts) the weighted average productivity level.

3.2.3 Assessment of localised market power

The calculation of localised market power l_i is similar to that of global market share:

$$l_i(t) = l_i(t-1) \left(1 + \frac{a_i(t) - \bar{a}_i^L(t)}{\bar{a}_i^L(t)} \right). \quad (4)$$

However, agents now compare their productivity level only to those of their link neighbours and the firm itself, the set K :

$$\bar{a}_i^L(t) = \sum_{j \in K_i} a_j(t) l_j(t-1). \quad (5)$$

Consequently, a completely unconnected firm (i.e., one with a local monopoly on all its products) has a localised market power of 1, while the value for each other firm may be above or below its global share.

3.2.4 Entry and exit

When the global share of a firm drops below the threshold of 0.001, it leaves the market. Using the global share rather than the localised market power here makes sense because a firm becomes unprofitable or goes bankrupt if its sales revenue

is too low. This may happen even to firms that possess high power in an industry which ultimately proves to be too small and unsustainable, while conversely firms that are small players in a large industry may create a high revenue. Firms leaving the market disappear from the model and also destroy all their links, which improves the local position of their old link-neighbours (i.e., former direct competitors).

Each departing firm is immediately replaced by a new entrant. This entrant links with incumbents and other new entrants with the same probability p used for initial network generation. The global share of entrants is $1/N$, and their localised market power is $1/|K_i|$ (with K_i again being the set of all link-neighbours and the entrant i itself). However, entrants do not start with a low productivity value of 1, but instead acquire the specific productivity level of the industry they enter (or of the whole market as a fallback should the entrant have no links) altered by the common learning parameter:

$$a_i(t_{\text{entry}}) = \begin{cases} \overline{a}_i^L(t_{\text{entry}}) + \pi_i(t_{\text{entry}}) & \text{if } K_i \setminus i \neq \emptyset \\ \overline{a}(t_{\text{entry}}) + \pi_i(t_{\text{entry}}) & \text{if } K_i \setminus i = \emptyset \end{cases} \quad (6)$$

Generally, entrants benefit from past technological and management innovations as well as local market conditions that shape competitiveness in the industries that they enter. Locally stronger incumbents in the industry play a more central role here. The productivity level of link-neighbours weighted by their localised market power reflects this notion. Here, localised market power is used as the weighting instead of global share, because the local importance of an incumbent firm matters for entrants and not absolute sales revenue. This again reflects the boundedly rational nature of our model, where entrants learn indirectly from the experience of local incumbents, particularly in regard to their tacit product domain knowledge (Glauber et al. 2015). Some entrants are uniquely innovative, meaning that they increase their own productivity levels beyond those of their environments.

Because entrants' global shares are higher than those of market leavers, the sum of all global shares now exceeds 1. To correct for this, the share of each firm is reduced proportionally (divided by the sum of all shares). Likewise, the entry/exit process altered the network structure and thus each firm divides its localised market power by the sum of all localised market powers in K . These adjustments normalise the corresponding values and ensure comparability over time.

4 Simulation results

To ensure comparability, the results presented in this section generally follow the parametrisation in Dosi et al. (2017b): in particular, we opt for 50 Monte Carlo iterations of simulation runs with time-series length $T = 200$ and for $N = 150$ firms each. Given that the trajectories of all aggregate statistics for all linkage probabilities converge to a stationary state very quickly, we present our estimations for

Table 2 Summary of test statistics and estimated p values for a fitted power-law distribution to the simulated sales shares with p in increments of 0.1 for the whole parameter space

p	Test statistic KS	p value KS	Test statistic CvM	p value CvM
0	0.0496	0.327	0.265	0.171
0.1	0.0165	0.785	0.062	0.802
0.2	0.0225	0.347	0.140	0.421
0.3	0.0218	0.356	0.161	0.357
0.4	0.0124	0.930	0.0428	0.918
0.5	0.0210	0.569	0.0939	0.615
0.6	0.0178	0.935	0.0501	0.876
0.7	0.0196	0.452	0.194	0.279
0.8	0.0149	0.860	0.066	0.776
0.9	0.0211	0.515	0.152	0.383
1	0.0565	0.143	0.320	0.119

“CvM” refers to the Cramér–von-Mises test, and “KS” to the Kolmogorov–Smirnov test. For both tests and all reported linkage probabilities p , the hypothesis of a power-law cannot be rejected at the usual significance level of 5%

$t = 200$ only, but for the pooled Monte Carlo simulation runs. For $p = 1$, which in its model assumptions fully corresponds to the model set-up in Dosi et al. (2017b), we derive results that are also in full qualitative agreement with the results obtained by Dosi et al. (2017b), indicating that we are indeed generalising their case.⁶ For all p , we find that the productivity distribution is very heterogeneous with fat tails and hence consistent with Dosi et al. (2017b) and the empirical evidence cited therein. The negative relationship between size and variance in growth rates, though, holds only for $p = 1$ and its neighbourhood, strengthening the case we are building in this section to confirm that the fully connected network is indeed the empirically relevant benchmark.⁷ In contrast to non-parametric descriptive statistics, we analyse our three main results by fitting parametric distributions to the data that allow us to examine the generating mechanisms pertaining to the variable in question.

4.1 Size distribution

We find that the upper tail of the firm size distribution (measured by market shares) is for all p characterised by a power-law distribution. The distribution cannot be statistically distinguished by standard non-parametric tests such as the Cramér–von-Mises test and the Kolmogorov–Smirnov test from a continuous power-law

⁶ We are strongly indebted to Giovanni Dosi and Marcelo Pereira for their helpful comments and sharing their code with us to establish this comparability in all subtle details.

⁷ The results on the variance-size scaling, as well as the productivity distribution, are available upon request.

Fig. 1 Empirical complementary cumulative distribution function of firm sizes for $p = 0.05, 0.9$ and 1

distribution (Anderson and Darling 1952; Anderson 1962; Smirnov 1948). We report the corresponding test statistics and p values in Table 2 below.

Furthermore, for all p , the power-law regime spans approximately two orders of magnitude and hence meets the common minimum requirement for a power-law to be present (Stumpf and Porter 2012). This is equivalent to saying that the discrete distribution of shares can be approximated by this probability density function (PDF) for the continuous analogue of this power-law distribution as

$$p(s) = Cs^{-\alpha}, \quad \text{for } s \geq s_{\min}, \quad (7)$$

where: $s_{\min}$ denotes the minimum share from which on the power-law applies; α denotes the characteristic exponent of the power-law distribution, and C is a normalising constant letting the probability density integrate to 1. Notice that $\alpha \geq 1$ is an inverse measure of concentration, where a lower α indicates a higher degree of inequality. Figure 1 shows the complementary cumulative distribution function (CCDF) of the upper tail firm size distribution for three different linkage probabilities. The minimum was determined by the standard procedure in this field first outlined by Clauset et al. (2009), essentially by: fitting a reverse-order-statistic to a power-law with increasing sample sizes; obtaining the Kolmogorov–Smirnov test statistic; and choosing the $s_{\min}$ that minimises it. This method has been shown to outperform other

Fig. 2 Estimated $\hat{\alpha}$ for all p in increments of 1%. Lines between estimates are visual aids only

methods robustly, such as minimising the Bayesian information criterion (Clauset et al. 2009). Indeed, as the non-parametric tests also suggest, all three distributions display an approximately linear behaviour on a double-logarithmic scale (Newman 2005). The slope for each p is approximately α , where a lower α indicates that the CCDF does decay more slowly, indicating a higher concentration with a higher frequency of large shares.

Given the set-up of our model, this is not surprising as it essentially comprises a stochastically multiplicative process with an entry-exit mechanism that has been shown to be the most promising candidate for generating power-laws (Gabaix 2009). Hence, for our set-up the path-dependent stochastically multiplicative process seems to remain the most critical feature of the model, irrespective of the underlying network structure. Moreover, regardless of the underlying mode of local competition, we want to highlight that this extremely heterogeneous power-law distribution implies a situation that is far from the perfect competition usually assumed as a benchmark for general equilibrium models.

The functional form for the upper tail of empirical firm size distributions is thus seemingly broadly consistent with all connectivity patterns for the underlying localised network. However, the empirical consensus that the upper tail of firm sizes is characterised by Zipf's law with an estimated $\hat{\alpha}$ not statistically different from 1 constrains the permissible p to a much more narrow range. In Fig. 2, we show the

Fig. 3 Empirical density of firm size growth rates g on a semi-logarithmic scale for $p = 0.05, 0.9$ and 1

behaviour of estimated $\hat{\alpha}$ for the whole range of p in our model using 1% increments and with error bands corresponding to two sample standard deviations upwards and downwards, implying that the plotted intervals span the true α with 95% confidence.⁸

Two features are striking in the plot: firstly, Zipf's law is consistent only with the two knife-edge scenarios of an extremely sparse network in the (narrow) neighbourhood of $p = 0$ and the other extreme of a very dense network in the (narrow) neighbourhood of $p = 1$. This, in turn, implies that the empirical evidence constrains us to these two extremes. Secondly, contrary to economic intuition built within general equilibrium models, measured concentration is maximal—Zipf—for the highest degree of local competition and lowest for a mild (local) oligopoly around $p = 0.9$.

4.2 Growth rates of market shares

Another focal point of the industrial dynamics literature is the presence of fat-tailed growth rate distributions in sales. In more colloquial terms, this implies that jumps in firms' market shares are relatively more frequent than one would expect from a Gaussian distribution. Note that the presence of non-Gaussian growth rate distributions alone indicates that the growth process is not

⁸ The estimation of $\hat{\alpha}$ was carried out by using the associated maximum likelihood estimator (MLE) or Hill estimator that has been shown to be less biased compared to OLS methods or fitting a linear function onto the power-law on a double-logarithmic scale. Cf. also Goldstein et al. (2004) for a more rigorous analysis of different graphical methods and their respective shortcomings compared to an MLE. The standard errors were obtained exploiting the asymptotic Gaussianity of the Hill estimator (De Haan and Resnick 1998).

independent in time. According to the central limit theorem, this would induce Gaussian growth rates. Of course, stochastically multiplicative growth processes like ours responsible for emergence of the power-law in levels are actually path-dependent and thus violate independence. Indeed, Dosi et al. (2018) produce robust findings which support fat-tailed growth rates for the baseline specification we use together with a vast range of different specifications and parameter constellations. Their baseline model corresponds to our $p = 1$ parametrisation. We also find fat-tailed growth rate distributions for p different than 1, as can be seen for $p = 0.05$ and 0.9 in Fig. 3. However, at least for the not fully connected network, the fat-tailed nature of growth rates is primarily due to extreme losses, rather than frequent extreme growth events which are at odds with the presence of superstar firms. We want to highlight also that this fat-tailed nature is a different concept from mere ‘dispersion’. While dispersion does indeed seem to decline with p , being fat-tailed refers to the frequency of extreme events *relative* to the frequency of events closer to the expected growth rates, for which inference by visual inspection is a much harder task. Within Fig. 3, frequent extreme growth events are present only for the fully connected network. This finding, though, might merely be an artefact of the three network connectivities under consideration and thus will not hold for the whole parameter space. We need to explore the full parameter space to identify possible switching behaviour concerning the source of fat tails in the simulated growth rate distributions.

A standard procedure for identifying fat tails and quantifying the degree of ‘fat-tailedness’ in growth rate distributions is to fit a Subbotin distribution (Subbotin 1923) to the data and take its shape parameter b as a measure of heavy tails’ strength. The Subbotin density includes the Gaussian for $b = 2$, the Laplacian for $b = 1$, the Dirac-Delta for $b \rightarrow 0$ (from above), and the uniform distribution for $b \rightarrow \infty$ as special cases. Consequently, we define fat-tailed behaviour for all $b \geq 0$ significantly smaller than 2 for the Gaussian case. As the contemporary relevance of superstar firms is central to our concerns within this study, we are primarily interested in extreme growth events as opposed to extreme losses. We opt for an asymmetric variant of the Subbotin distribution, introduced by Bottazzi (2014), to distinguish extreme growth events from extreme losses. The PDF is given by:

$$p(g; a_r, a_l, b_l, b_r, m) = C^{-1} \text{Exp} \left(-\frac{1}{b_l} \left| \frac{g-m}{a_l} \right|^{b_l} \Theta(m-g) + \frac{1}{b_r} \left| \frac{g-m}{a_r} \right|^{b_r} \Theta(g-m) \right), \quad (8)$$

where $\Theta(\cdot)$ denotes the Heaviside theta function; m is a centrality parameter; a_l and a_r are the scale parameters of the left and right tails, respectively; whilst b_l and b_r are shape parameters for both tails with the analogous interpretation as in the symmetric case. In the language of this distributional analysis, ‘superstar-like’ behaviour is obtained for relatively frequent extreme growth events, that is, a fat right tail of the growth rate distribution with b_r significantly lower than 2. We estimate both

Fig. 4 Estimated $\hat{b}_l$ and $\hat{b}_r$ for all p in increments of 1%. Black estimates for b_l , gray estimates for b_r . Lines between estimates are visual aids only

parameters for the growth rate distributions of market shares by MLE for each p in 1% increments. The corresponding standard errors are obtained by utilising the Fisher information (Ruppert 2014).⁹ Figure 4 shows the values of b_l and b_r as a function of p with p increasing in 1% increments.

The figure highlights two distinct regimes with respect to the growth rate distributions. Taken as a whole, all growth rate distributions seem to be fat-tailed in agreement with empirical studies. However, the source of this fat-tailed behaviour differs between regimes. While for the broadest range of p between 0 and about 0.93, thus between a completely sparse and a very dense network, relatively frequent *extreme losses* are responsible for the fat tails, the situation changes dramatically in the neighbourhood of a fully connected network, where relatively frequent *extreme growth* dominates. Superstar-like behaviour is thus consistent only with extremely dense networks implied by a p close to 1.

4.3 Age

Finally, for age, our model can mimic the empirically observed exponential distribution in age levels for all p . This can be seen in Fig. 5, where the three age distributions being considered display approximately linear behaviour on a semi-logarithmic scale, consistent with an exponential functional form.¹⁰ This emergent

⁹ In particular, we employ the freeware *Subbotools 1.3.0* specifically designed for the estimation of different flavours of the Subbotin distribution (Bottazzi 2014), which delivered by far the most efficient parameter estimates for different samples of data we simulated.

¹⁰ The exponential might appear to not fit the age distribution too well for $p = 0.9$. However, this impression is mainly an artefact of the semi-logarithmic scale and pertains only to the largest 0.1% of values. For the remaining 99.9%, the fit is extremely good, leaving us confident that the exponential is a reasonable choice here.

Fig. 5 Empirical density of firm age on a semi-logarithmic scale for $p = 0.05, 0.9$ and 1 with exponential fits obtained by MLE

exponential stationary age distribution coupled with stable population levels has, in itself, an important implication: The exit probability or insolvency rate is common and constant for all firms (Artzrouni 1985). Thus, every firm irrespective of its age has the same probability of becoming insolvent in any period and, consequently, has the same expected age. The relatively stable distributions of size and growth rates within time at the meso-level are consistent with a very dynamical economic system underlying these regularities and high rates of ‘churning’ in the composition of firms, where even local and global market leaders face the same certain prospect of insolvency at some point.

While the functional form of the emergent age distribution is constant for different p , its estimated parameter $\hat{\lambda}$ as the insolvency rate changes with p . As can be seen, the firms’ life expectancies vary widely with p . For $p = 0.05$, the firms cluster around a very young age, while exhibiting much wider dispersion and higher expected age for $p = 1$. Life expectancies thus seem to increase in the network connectivity, but also grow more heterogeneous.

When examining the whole parameter space of p , the insolvency rate falls monotonically with p as we show in Fig. 6.¹¹

Thus, while all firms irrespective of their age face the same estimated insolvency probability $\hat{\lambda}$ per regime or per p , this insolvency probability differs widely between

¹¹ The estimated $\hat{\lambda}$ was estimated through MLE. The standard errors were obtained by utilising the fact that $\hat{\lambda}$ is just the inverse of the sample mean and that the associated sample standard deviation is therefore $(\hat{\lambda}\sqrt{N})^{-1}$ (Lehmann and Casella 2006).

Fig. 6 Estimated $\hat{\lambda}$ for all p in increments of 1%. Lines between estimates are visual aids only

the regimes. Given that expected age is just the inverse of $\hat{\lambda}$, this implies that with a higher p , firms tend to stay in the market for a much longer time, and there is less ‘churning’ between periods.

4.4 Generating mechanisms

Dosi et al. (2017b) explain their model outcomes based on the idiosyncratic learning process and replicator dynamics; this explanation also straightforwardly applies for the complete network in our model. However, the tails of the global share distribution, the growth distribution and market exit probability of firms react in a highly elastic way towards changes in the network topology. Thus, these results suggest the presence of a second model mechanism that depends on network density and the implied distribution of localised market power. Because the learning of incumbents and assessment of their global shares work irrespective of network layout, the success of entrants remains the sole candidate for such a driving mechanism.

Since all entrants have identical initial market shares, the individual success of each entrant depends largely on its initial level of competitiveness. To gain a high level of initial competitiveness, an entrant must connect to as many highly productive incumbents with high localised market power as possible (i.e., join a thriving industry). Such connections become less likely for smaller link probabilities p . Thus, in sparse networks, most entrants start with low productivity. Furthermore, since the assessment of global market share compares the productivity of the firm in question with the weighted

global average productivity, these relatively unproductive entrants quickly lose market share in the first periods of their lives. That explains the fat left tail of the growth rate distribution for small p .

At the same time, those entrants that connect to highly productive and powerful incumbents have a comparative advantage, increase their sales quickly, and manage to catch up with even the most successful firms in the market. Hence, firms within the power-law tail exhibit more homogeneous sizes, the sparser the network is. The low maximum firm age and high probability of exit for sparse networks is a corollary of these two aspects: even successful firms are challenged, find themselves outperformed by productive younger competitors and finally leave the market, while most entrants do so after only a few simulation periods. Without explicit targeting, we are able to replicate ‘imprinting’ behaviour or the empirically well established phenomenon that founding conditions exhibit lasting effects on the entrants’ survival probabilities (Gerowski et al. 2010).

For higher linking probabilities, the rate of entrants with a high initial productivity level grows, making the left tails of the growth distribution thinner. However, the most productive entrants are also hindered by the higher average productivity level and consequently have a harder time becoming superstars; thus, the inequality within the power-law tail decreases even more.

Furthermore, the importance of the birth productivity mechanism, which favours few entrants and lets many suffer, becomes weaker the denser the network is and hence the more similar localised market power as well as global market share become. If the entrants’ fate is no longer determined at birth, learning becomes more important. Thereby, the replicator dynamics of global share assessment means a fat right tail of the growth distribution and tails of the firm size distribution in accordance with Zipf’s law.

To summarise, two distinct mechanisms govern the productivity of firms and consequently their commercial success. The first is a process of learning that occurs within each period and is equally strong for all network layouts, but its effects depend on attendant productivity levels. The second mechanism is the allocation of initial productivity based on link-neighbours, which applies only once to each entrant at birth. The mode of operation and the strength of this second mechanism depends to a great extent on network density. For least dense networks, it dooms most of the entrants to a fast market exit while it is at the same time also subsidising a few of them in an extreme way, prolonging their accumulation of market shares. For denser networks, more firms share this subsidy and hence the most successful firms become more equal in terms of their size. Furthermore, it is noteworthy that the second mechanism takes precedence over the first for all but the densest networks, according to simulation outcomes. That is the case because birth productivity also implies a path-dependency: For an unequal birth productivity distribution of entrants, learning stabilises, and amplifies this inequality due to a higher productivity, also meaning a potential for a higher absolute gain through learning.

5 Discussion

We introduce a network-structure to the bare-bones model of a ‘winner takes most/all’ market proposed by Dosi et al. (2017b). This extension generates surprisingly rich dynamics and intriguing implications when deviating from the benchmark of a fully competitive localised market. In particular, we have been looking to highlight both the positive and normative implications we draw from our modelling exercise and their practical relevance for economic regulation as well as management decisions in the case of a single firm.

Empirically, we find that the stylised facts of industrial dynamics, namely Zipf’s law in the firm-size distribution and fat-tailed, ‘superstar-like’ firm growth rates are consistent only with a situation very close to the benchmark of a fully connected network, meaning most intensive localised competition. All other network connectivities lead to significant deviations from the stylised facts in at least one regard. Hence, if we can accurately identify parts of the empirical mechanism—and there exists evidence that replicator dynamics play an essential role in empirical markets (Cantner and Krüger 2008; Cantner et al. 2012), our results will point to product markets that are relatively undifferentiated. Thus, market power comes from global rather than localised dynamics. These results are in stark contrast to our initial expectations of low concentration and high rates of ‘churning’ for relatively high degrees of localised competition. This indicates that anecdotal insights gained from analysing static frameworks of competition do not necessarily transfer well to situations where strong non-linearities and feedback mechanisms are present.

In our model, the coexistence and partial interaction of two learning mechanisms and replicator dynamics explain the results: (1) stochastic productivity improvements in each period for each incumbent firm constitute the first way of learning; (2) works indirectly at market entry because an entrant’s initial productivity depends on the weighted average productivity of the incumbents that it links to, meaning within its specific industry. The less densely connected a network, the fewer entrants form connections to highly productive incumbents; hence, their initial productivity is low, and consequently, their market shares decrease, which explains the fat left tail of growth rate distributions and lower average firm age. However, those entrants connected to highly productive incumbents thrive because their initial productivity is high in comparison to most incumbents. Thus, they can catch up with even the most successful incumbents and market concentration decreases. Methodologically, our model of networked competitive interaction can thus be thought as a complement to the foundational theoretical study by Cantner et al. (2019) who study collaboration in networks to explain especially the instability of early-lifecycle firms by lock-in effects within suboptimal value chains. Namely, our model suggests a mechanism that may be present in addition to ‘failures of selection’ (Cantner et al. 2019) and cause the high rates of churning and volatility in market shares of young firms already observed by Mazzucato (1998). We demonstrate that such instability can also emerge for functioning selection and industry-specific initial productivity, as long as markets are locally

segregated or, equivalently, the competitive network exhibits rather low density. We aim to investigate the interplay of both the collaborative and competitive network channels in further research.

Since inequalities of initial productivity shrink with increasing network density and do not exist for the complete network, a lower level of competitiveness implies lower market concentration. Furthermore, for the complete network, other sources of inequality in initial productivity could replace the market entry learning mechanism: The absence of inequality in starting conditions leads to the most successful firms acquiring a greater market share. Consequently, one must accept the success of these superstars as an outcome if the aim is to create full equality of opportunity; otherwise, to avoid high market concentration by the most successful firms, one must deliberately create inequality of opportunity. In less abstract terms, our model suggests the common fear that active industrial policy creating unequally favourable starting conditions for specific firms and thus being anticompetitive to be at least partially misguided (cf. Sokol (2014) for a vocal proponent of this view): The relevant metric for consumers is perhaps ex post concentration in market shares, indicating that one can accept or even foster ex ante inequality in starting conditions to decrease such concentration after the fact. Active industrial policy enhancing the productivity of incumbents can even lead to positive productivity spill-overs, since entrants benefit from the average productivity of the market they enter. The more relevant trade-off within such markets appears to be between decreasing concentration in market power (lowering p) or decreasing the amount of turnover in the market (increasing p), with 'turnover' typically also implying (transient) increases in unemployment and the destruction of firm-specific capital and knowledge. In this way, our model can help to identify the relevant trade-offs for regulatory policy and contribute to a richer view apart from standard static efficiency considerations.

Besides these global findings, the model also suggests that there are localised cycles of productivity and firm size: If incumbents have acquired high localised market power and a high productivity level, new entrants joining the market segment and engaging in competitive interaction (i.e., linking to the productive incumbents) also start with a high productivity rate. Consequently, the industry in question becomes even more productive until it overheats, and incumbents are repressed from the market while the high productivity shifts to another (possibly new), related market segment. This effect is entirely in line with the empirical study, in which Schlichte et al. (2019) show that the timing of entry to highly specified submarkets between two technology waves is crucial for the success of new firms. Moreover, our model supports their finding that there is a first-follower advantage (in our model represented by successful entrants) as opposed to a first-mover advantage (moderately successful incumbents that are nonetheless outperformed by entrants) because of growing consumer acceptance of new technology (Davis 1989). These findings might be of particular interest to practitioners in Venture Capital and are consistent with their empirical emphasis on 'deal selection' compared to other phases of the investment process (Gompers et al. 2020). However, there exists no consensus on the correct selection strategy, with some trend-following venture capitalists selecting 'hot sectors' and other contrarian ones avoiding them (Gompers et al. 2020). In principle, our model points to the trend-following strategy to benefit from the high initial

local productivity of the relevant submarket. This is still no guarantee for success, though, as the submarket in question might be on the brink of overheating, also providing a rationale for the contrarian view. Venture capitalists, in our model, should thus pick sectors with high expected *growth* in contrast to present size in *levels* to avoid entering markets near the end of a technology wave.

The validity of our model depends to a great extent on the validity of the baseline model by Dosi et al. (2017b), which we assume to be given. However, since we re-implement the mechanisms from scratch and include the baseline model as a special case reproducing its findings, we can affirm the internal validity of the baseline model and our extension. With regard to external validity, we hope that including a network structure of localised competition can facilitate resemblance (Mäki 2009) between model and real-world economies. Our explanans can actually be true and the cause for the observed empirical fact. Hence, our proposed mechanism fulfils the minimum conditions for a good epistemically possible how-possibly explanation formulated by Grüne-Yanoff and Verreault-Julien (2021). Nevertheless, with the inclusion of localised market power, the nature of our model and thus, the implied mode of analysis remains highly stylised. Hence, the validity of the model is based on its “qualitative agreement with empirical macrostructures” (Fagiolo et al. 2019, p. 771), namely the replication of the stylised empirical facts that our model successfully attempts. Put differently, we develop a specific parallel reality (Sugden 2009) that features generating mechanisms for empirical findings in our reality and hence our results present a candidate explanation for the stylised empirical facts (Epstein 1999). Consequently, there may be different, more adequate, parallel realities featuring either these or even better mechanisms, despite to the best of our knowledge there being no existing models that fulfil these characteristics.

Alternative mechanisms firstly concern the network that we use. While we test for any network density, we limit ourselves to random link formation as we are not aware of empirical evidence for any specific network topology in our context. However, a non-random (e.g., preferential attachment or spatial-dependent) link formation may impact simulation results, especially for low network densities. Moreover, we distinctly interpret links as indicators for localised competition that only matters for a firm’s initial productivity level. One could further explicate such localised competition and track it over time. Alternative or additional layers of links could also represent cooperation between firms or their products being complements. Our model’s most apparent limitation concerns the baseline replicator dynamics equation, though, which implies that the emergent concentration is ‘good concentration’ (Covarrubias et al. 2020) and fully justifiable by productivity differences. Empirically, it is questionable if concentration indeed only reflects productivity (Covarrubias et al. 2020), with firms erecting artificial barriers to entry or acquiring competitors and discontinuing their innovative product lines in so-called ‘killer acquisitions’ (Cunningham et al. 2021) leading to ‘bad concentration’. Since it is at least conceivable that a high concentration of the good type is preferable to lower bad concentration, the inclusion of strategic anticompetitive behaviour might alter the policy conclusions of our model and tilt them more towards antitrust measures, which might like in our baseline model induce high ex-post concentration purely based on productivity differences.

Furthermore, the stylised nature of our findings points to obvious extensions and avenues for further research. To analyse not only the consistency with stylised facts, but also for quantitative predictions and even policy experiments, the ABM community has recently developed certain new methods. These are aimed at bringing modelling closer to the data and calibrating parameters (Hassan et al. 2010), particularly by utilising the method of simulated moments (Gourieroux et al. 1996) and related approaches (Bargigli et al. 2018). Given the partial equilibrium nature of our model, this would probably also necessitate allowing for a variable total number of firms over time by including mergers and acquisitions as well as consumer demand, a state sector and even financial markets for meaningful policy experiments. With the benchmark model by Dosi et al. (2017b), this was attempted by the K+S ABM (Dosi et al. 2010). An extension of this sort would enable us in future research to conduct policy experiments and quantify welfare effects for different market structures.

We believe our model to be a valuable contribution to the discussion on market structures and a key step towards a unifying explanation for both the microeconomic evidence on ‘superstar’ firms and the distributional findings in industrial dynamics. To the best of our knowledge, these two strands of literature have not as yet been linked. We find that for the replicator dynamics approach from industrial dynamics to be consistent with the existence of superstar firms, shown by microeconomic studies, there needs to be an almost perfect level of competition. These outcomes emerge because in each simulation period, the firms improve their productivity by idiosyncratic stochastic learning, while new entrants adopt the specific productivity level of their sub-market (i.e., the productivity of their immediate competitors weighted by localised market power). Thus, the model suggests that new firms are most successful if they join existing, highly productive submarkets with high growth potential. Accordingly, while accepting bounds on rationality, we can single out the strategy which a perfectly rational market entrant would pick when faced with a certain market structure. This not only fosters an understanding of market dynamics, but can also be applied to highlight the importance of market intelligence for the management in new firms and new technology markets.

Acknowledgements The authors thank Giovanni Dosi and Marcello Pereira for sharing their code and explaining their rationale; for helpful comments on style and content (in alphabetical order) Anna Gebhard, Bettina Gregg, Torsten Heinrich, Dominik Klein, Johannes Marx, Alessandro Lucini Paioni, Roland Parr, Simon Scheller, Jürgen Sirsch, Moritz Schulz and Jan Weber as well as the participants of the WICK 2020 and EAEPE 2020 conferences. We are especially indebted to two anonymous reviewers and the editor Tim Weitzel for multiple insightful comments and suggestions to improve our paper. All remaining errors are our own, of course. This paper first circulated under the title “Innovation of Incumbents and Entrants in Localised Competition”. We are immensely grateful to Tim Weitzel for pointing us to a more appropriate title.

Funding Open Access funding enabled and organized by Projekt DEAL.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is

not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Akerman A, Leuven E, Mogstad M (2021) Information frictions, broadband internet and the relationship between distance and trade. *Appl Econ Am Econ J* (forthcoming)
- Alfarano S, Milaković M (2008) Does classical competition explain the statistical features of firm growth? *Econ Lett* 101(3):272–274
- Amaral LAN, Buldyrev SV, Havlin S, Maass P, Salinger MA, Stanley HE, Stanley MH (1997) Scaling behavior in economics: the problem of quantifying company growth. *Phys A* 244(1):1–24
- Anderson TW (1962) On the distribution of the two-sample Cramer-von Misés criterion. *Ann Math Stat* 33(3):1148–1159
- Anderson TW, Darling DA (1952) Asymptotic theory of certain “goodness of fit” criteria based on stochastic processes. *Ann Math Stat* 23(2):193–212
- Arrow K (1991) The dynamics of technological change. In: OECD (ed) *Technology and productivity: the challenge for economic policy*. OECD, Paris, pp 473–476
- Artzrouni M (1985) Generalized stable population theory. *J Math Biol* 21(3):363–381
- Autor D, Dorn D, Katz LF, Patterson C, Van Reenen J (2020) The fall of the labor share and the rise of superstar firms. *Q J Econ* 135(2):645–709
- Axtell RL (2001) Zipf distribution of US firm sizes. *Science* 293(5536):1818–1820
- Bargigli L, Riccetti L, Russo A, Gallegati M (2018) Network calibration and metamodeling of a financial accelerator agent based model. *J Econ Interact Coord* 15(2):413–440
- Böhm F (1960) Demokratie und unternehmerische Macht. In: Institut für ausländisches und internationales Wirtschaftsrecht (ed) *Kartelle und Monopole im modernen Recht*. C. F. Müller, Karlsruhe, pp 1–24
- Bottazzi G (2014) Subbotools user’s manual. Laboratory of Economics and Management Working Paper Series, Sant’Anna School of Advanced Studies, Pisa 2004/01:1–23
- Bottazzi G, Secchi A (2005) Growth and diversification patterns of the worldwide pharmaceutical industry. *Rev Ind Organ* 26(2):195–216
- Bottazzi G, Secchi A (2006) Explaining the distribution of firm growth rates. *Rand J Econ* 37(2):235–256
- Bottazzi G, Dosi G, Lippi M, Pammolli F, Riccaboni M (2001) Innovation and corporate growth in the evolution of the drug industry. *Int J Ind Organ* 19(7):1161–1187
- Bottazzi G, Cefis E, Dosi G (2002) Corporate growth and industrial structures: some evidence from the Italian manufacturing industry. *Ind Corp Change* 11(4):705–723
- Bottazzi G, Coad A, Jacoby N, Secchi A (2011) Corporate growth and industrial dynamics: evidence from French manufacturing. *Appl Econ* 43(1):103–116
- Bray MM, Savin NE (1986) Rational expectations equilibria, learning, and model specification. *Econometrica* 54(5):1129–1160
- Brock W (1999) Scaling in economics: a reader’s guide. *Ind Corp Change* 8(3):409–446
- Cantner U (2017) Foundations of economic change: an extended Schumpeterian approach. In: Cantner U, Pyka A (eds) *Foundations of economic change*. Springer, pp 9–49
- Cantner U, Krüger JJ (2008) Micro-heterogeneity and aggregate productivity development in the German manufacturing sector. *J Evolut Econ* 18(2):119–133
- Cantner U, Krüger JJ, Söllner R (2012) Product quality, product price, and share dynamics in the German compact car market. *Ind Corp Change* 21(5):1085–1115
- Cantner U, Savin I, Vannuccini S (2019) Replicator dynamics in value chains: explaining some puzzles of market selection. *Ind Corp Change* 28(3):589–611
- Chamberlin EH (1949) *Theory of monopolistic competition: a re-orientation of the theory of value*. Oxford University Press, London
- Clauset A, Shalizi CR, Newman ME (2009) Power-law distributions in empirical data. *SIAM Rev* 51(4):661–703

- Coad A (2010) Investigating the exponential age distribution of firms. *Economics* 4:1–30
- Coad A, Tamvada JP (2008) The growth and decline of small firms in developing countries. *Papers on Economics and Evolution*, Max-Planck-Institut für Ökonomik 0808:1–32
- Coleman JS (1990) *Foundations in social theory*. Harvard University Press, Massachusetts
- Covarrubias M, Gutiérrez G, Philippon T (2020) From good to bad concentration? US industries over the past 30 years. *NBER Macroecon Annu* 34(1):1–46
- Cunningham C, Ederer F, Ma S (2021) Killer acquisitions. *J Polit Econ* 129(3):649–702
- Daepf MI, Hamilton MJ, West GB, Bettencourt LM (2015) The mortality of companies. *J R Soc Interface* 12(106):20150120
- Davis FD (1989) Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS Q* 13(3):319–340
- De Haan L, Resnick S (1998) On asymptotic normality of the Hill estimator. *Stoch Models* 14(4):849–866
- Di Giovanni J, Levchenko AA (2013) Firm entry, trade, and welfare in Zipf's world. *J Int Econ* 89(2):283–296
- di Giovanni J, Levchenko AA, Rancière R (2011) Power laws in firm size and openness to trade: Measurement and implications. *J Int Econ* 85(1):42–52
- Dold M, Krieger T (2019) *Ordoliberalism and European economic policy: between realpolitik and economic Utopia*. Routledge, London
- Dosi G, Marsili O, Orsenigo L, Salvatore R (1995) Learning, market selection and the evolution of industrial structures. *Small Bus Econ* 7(6):411–436
- Dosi G, Fagiolo G, Roventini A (2010) Schumpeter meeting Keynes: a policy-friendly model of endogenous growth and business cycles. *J Econ Dyn Control* 34(9):1748–1767
- Dosi G, Fagiolo G, Napoletano M, Roventini A (2013) Income distribution, credit and fiscal policies in an agent-based Keynesian model. *J Econ Dyn Control* 37(8):1598–1625
- Dosi G, Fagiolo G, Napoletano M, Roventini A, Treibich T (2015) Fiscal and monetary policies in complex evolving economies. *J Econ Dyn Control* 52:166–189
- Dosi G, Napoletano M, Roventini A, Treibich T (2017a) Micro and macro policies in the Keynes + Schumpeter evolutionary models. *J Evolut Econ* 27(1):63–90
- Dosi G, Pereira MC, Virgillito ME (2017b) The footprint of evolutionary processes of learning and selection upon the statistical properties of industrial dynamics. *Ind Corp Change* 26(2):187–210
- Dosi G, Pereira MC, Virgillito ME (2018) On the robustness of the fat-tailed distribution of firm growth rates: a global sensitivity analysis. *J Econ Interact Coord* 13(1):173–193
- Epstein JM (1999) Agent-based computational models and generative social science. *Complexity* 4(5):41–60
- Erdős P, Rényi A (1960) On the evolution of random graphs. *Publ Math Inst Hung Acad Sci* 5(1):17–60
- Erlingsson EJ, Alfarano S, Raberto M, Stefánsson H (2013) On the distributional properties of size, profit and growth of Icelandic firms. *J Econ Interact Coord* 8(1):57–74
- Evans GW, Honkapohja S (2013) *Learning as a rational foundation for macroeconomics and finance*, vol 68. Princeton University Press, Princeton
- Fagiolo G, Guerini M, Lamperti F, Moneta A, Roventini A (2019) Validation of agent-based models in economics and finance. In: Beisbart C, Saam NJ (eds) *Computer Simulation validation: fundamental concepts, methodological frameworks, and philosophical perspectives*. Springer International Publishing, Cham, pp 763–787
- Fisher R (1930) *The genetical theory of natural selection*. Clarendon Press, Oxford
- Fujiwara Y, Di Guilmi C, Aoyama H, Gallegati M, Souma W (2004) Do Pareto-Zipf and Gibrat laws hold true? An analysis with European firms. *Phys A* 335(1–2):197–216
- Gabaix X (2009) Power laws in economics and finance. *Annu Rev Econ* 1(1):255–294
- Gabaix X (2011) The granular origins of aggregate fluctuations. *Econometrica* 79(3):733–772
- Gabaix X, Landier A (2008) Why has CEO pay increased so much? *Q J Econ* 123(1):49–100
- Geroski PA, Mata J, Portugal P (2010) Founding conditions and the survival of new firms. *Strateg Manag J* 31(5):510–529
- Glauber J, Wollersheim J, Sandner P, Welpel IM (2015) The patenting activity of German universities. *J Bus Econ* 85(7):719–757
- Goldstein ML, Morris SA, Yen GG (2004) Problems with fitting to the power-law distribution. *Eur Phys J B Condens Matter Complex Syst* 41(2):255–258

- Gompers PA, Gornall W, Kaplan SN, Strebulaev IA (2020) How do venture capitalists make decisions? *J Financ Econ* 135(1):169–190
- Gourieroux M, Gourieroux C, Monfort A, Monfort DA (1996) Simulation-based econometric methods. Oxford University Press, Oxford
- Greve HR (2009) Bigger and safer: the diffusion of competitive advantage. *Strateg Manag J* 30(1):1–23
- Grimm V, Berger U, DeAngelis DL, Polhill JG, Giske J, Railsback SF (2010) The ODD protocol: a review and first update. *Ecol Model* 221(23):2760–2768
- Grüne-Yanoff T (2009) Learning from minimal economic models. *Erkenntnis* 70(1):81–99
- Grüne-Yanoff T, Verreault-Julien P (2021) How-possibly explanations in economics: anything goes? *J Econ Methodol* 28(1):114–123
- Hassan S, Pavón J, Antunes L, Gilbert N (2010) Injecting data into agent-based simulation. In: Takadama K, Cioffi-Revilla C, Deffuant G (eds) *Simulating interacting agents and social phenomena*. Springer Japan, Tokyo, pp 177–191
- Hedström P, Ylikoski P (2010) Causal mechanisms in the social sciences. *Annu Rev Sociol* 36(1):49–67
- Hegselmann R, Krause U (2002) Opinion dynamics and bounded confidence models, analysis, and simulation. *J Artif Soc Soc Simul* 5(3):1–33
- Heinrich T, Dai S (2016) Diversity of firm sizes, complexity, and industry structure in the Chinese economy. *Struct Change Econ Dyn* 37:90–106
- Hidalgo CA, Hausmann R (2009) The building blocks of economic complexity. *Proc Natl Acad Sci* 106(26):10570–10575
- Johnson SL, Faraj S, Kudaravalli S (2014) Emergence of power laws in online communities: the role of social mechanisms and preferential attachment. *MIS Q* 38(3):795–808
- Kaldor N (1961) Capital accumulation and economic growth. In: Hague DC (ed) *The theory of capital*. Springer, pp 177–222
- Kang SH, Jiang Z, Cheong C, Yoon S-M (2011) Changes of firm size distribution: the case of Korea. *Phys A* 390(2):319–327
- Kinsella S (2009) The age distribution of firms in Ireland, 1961–2009. Technical report
- Klein D, Marx J (2018) Wenn du gehst, geh ich auch! Die Rolle von Informationskaskaden bei der Entstehung von Massenbewegungen. *PVS Politische Vierteljahresschrift* 58(4):560–592
- Klein D, Marx J, Fischbach K (2018) Agent-based modeling in social science, history, and philosophy. An introduction. *Hist Soc Res* 43(163):7–27
- Klein D, Marx J, Scheller S (2019) Rational choice and asymmetric learning in iterated social interactions—some lessons from agent-based modeling. In: Markler K, Schmitt A, Sirsch J (eds) *Demokratie und Entscheidung*. Springer, pp 277–294
- Kotz S, Kozubowski T, Podgorski K (2012) *The Laplace distribution and generalizations: a revisit with applications to communications, economics, engineering, and finance*. Springer Science & Business Media, Berlin
- Lehmann EL, Casella G (2006) *Theory of point estimation*. Springer Science & Business Media, Berlin
- Luttmer EG (2007) Selection, growth, and the size distribution of firms. *Q J Econ* 122(3):1103–1144
- Macy MW, Flache A (2002) Learning dynamics in social dilemmas. *Proc Natl Acad Sci* 99(suppl 3):7229–7236
- Mäki U (2009) Missing the world models as isolations and credible surrogate systems. *Erkenntnis* 70:29–43
- Mayerhoffer DM (2018) Raising children to be (in-)tolerant: influence of church, education, and society on adolescents' stance towards queer people in germany. *Hist Soc Res* 43(1):144–167
- Mazzucato M (1998) A computational model of economies of scale and market share instability. *Struct Change Econ Dyn* 9(1):55–83
- Milani F (2007) Expectations, learning and macroeconomic persistence. *J Monet Econ* 54(7):2065–2082
- Müller B, Bohn F, Dreßler G, Groeneveld J, Klassert C, Martin R, Schlüter M, Schulze J, Weise H, Schwarz N (2013) Describing human decisions in agent-based models-ODD+ D, an extension of the odd protocol. *Environ Model Softw* 48:37–48
- Mundt P, Alfaro S, Milaković M (2015) Gibrat's law redux: think profitability instead of growth. *Ind Corp Change* 25(4):549–571
- Newman MJE (2005) Power laws, Pareto distributions and Zipf's law. *Contemp Phys* 4:323–351
- Okuyama K, Takayasu M, Takayasu H (1999) Zipf's law in income distribution of companies. *Phys A* 269(1):125–131
- Pascoal R, Augusto M, Monteiro AM (2016) Size distribution of Portuguese firms between 2006 and 2012. *Phys A* 458:342–355

- Pellegrino B (2019) Product differentiation, oligopoly, and resource allocation. UniCredit Foundation Working Paper Series 133:1–95
- Robinson J (1969) The economics of imperfect competition. Springer, Berlin
- Ruppert D (2014) Statistics and finance: an introduction. Springer, Berlin
- Schlag KH (1998) Why imitate, and if so, how?: a boundedly rational approach to multi-armed bandits. *J Econ Theory* 78(1):130–156
- Schlichte F, Junge S, Mammen J (2019) Being at the right place at the right time: does the timing within technology waves determine new venture success? *J Bus Econ* 89(8):995–1021
- Silvestre J (1993) The market-power foundations of macroeconomic policy. *J Econ Lit* 31(1):105–141
- Skilton PF, Bernardes E (2015) Competition network structure and product market entry. *Strateg Manag J* 36(11):1688–1696
- Smirnov N (1948) Table for estimating the goodness of fit of empirical distributions. *Ann Math Stat* 19(2):279–281
- Sokol DD (2014) Tensions between antitrust and industrial policy. *Geo Mason Law Rev* 22:1247
- Stumpf MP, Porter MA (2012) Critical truths about power laws. *Science* 335(6069):665–666
- Subbotin MT (1923) On the law of frequency of error. *Matematicheskii Sbornik* 31(2):296–301
- Sugden R (2009) Credible worlds, capacities and mechanisms. *Erkenntnis* 70(1):3–27
- Tesfatsion L (2017) Modeling economic systems as locally-constructive sequential games. *J Econ Methodol* 24(4):384–409
- Wilensky U (1999) NetLogo. Northwestern University, Evanston
- Will O, Hegselmann R (2014) Modeling Hume's theory of norm emergence. In: Conte R, Andrighetto G, Campenni M (eds) *Minding norms*. Oxford University Press, Oxford, pp 68–80
- Zhang J, Chen Q, Wang Y (2009) Zipf distribution in top Chinese firms and an economic explanation. *Phys A* 388(10):2020–2024
- Zipf GK (1949) *Human behavior and the principle of least effort*. Addison-Wesley Press, Boston
- Zollman KJS (2010) The epistemic benefit of transient diversity. *Erkenntnis* 72(1):17–35

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Authors and Affiliations

Jan Schulz¹ · Daniel M. Mayerhoffer¹

✉ Daniel M. Mayerhoffer
daniel.mayerhoffer@uni-bamberg.de

Jan Schulz
jan.schulz@uni-bamberg.de

¹ University of Bamberg, Feldkirchenstr. 21, 96052 Bamberg, Germany