



Nr. 6 (2011)

Research Area E

ISSN 2219-9268

**Firms as Bundles of Discrete Resources – Towards
an Explanation of the Exponential Distribution of
Firm Growth Rates**

Alex Coad



Firms as Bundles of Discrete Resources – Towards an Explanation of the Exponential Distribution of Firm Growth Rates*

Alex Coad[†]

January 29, 2011

Abstract

A robust feature of the corporate growth process is the Laplace, or symmetric exponential, distribution of firm growth rates. In this paper, we sketch out a class of simple theoretical models capable of explaining this empirical regularity. We do not attempt to generalize on where growth opportunities come from, but rather we focus on how firms build upon growth opportunities. We base ourselves on Penrose's (1959) description of firm growth to explain how the interdependent nature of discrete resources may lead to the triggering off of a series of additions to a firm's resources. In a

*Thanks go to Giulio Bottazzi, David Brée, Thomas Brenner, Tommaso Ciarli, Mike Dietrich, Giovanni Dosi, Davide Fiaschi, Ricardo Mamede, Luigi Marengo, Rosanna Nisticò, Bernard Paulré, Rekha Rao, Giorgio Ricchiuti, Angelo Secchi, Federico Tamagni, Ulrich Witt and participants at the 'Economic Evolution as a Learning Process' course at the Max Planck Institute, Jena, Germany, March 13-24 2006; the workshop on 'Internal Organisation, Cooperative Relationships among Firms and Competitiveness' in Lucca, January 2007; the Econophysics conference in Ancona, September 27-29, 2007; and the DIME workshop on "production theory" in Pisa, November 8-9, 2010. An anonymous EAEPE-PEPE referee provided helpful comments, for which I am grateful. Zlata Jakubovic and Katja Mehlis provided excellent research assistance. The usual caveat applies.

[†]*Corresponding Author:* Alex Coad, SPRU, University of Sussex, Freeman Centre, Falmer, Brighton, BN1 9QE. *Phone:* +44 1273 877128. *E-mail:* A.Coad@sussex.ac.uk; and Max Planck Institute of Economics, Evolutionary Economics Group, Kahlaische Strasse 10, D-07745 Jena, Germany

first formal model we consider the case of employment growth in a hierarchy, and observe that growth rates follow an exponential distribution. In a second model we include plant and capital as resources and we are able to reproduce a number of stylized facts about firm growth.

JEL codes: L1, C1

Keywords: Firm growth rates, Laplace distribution, Exponential distribution, Hierarchy, Growth autocorrelation, Self-organizing criticality

1 Introduction

It has long been known that the distribution of firm growth rates is fat-tailed. In an early contribution, Ashton (1926) considers the growth patterns of British textile firms and observes that: “In their growth they obey no one law. A few apparently undergo a steady expansion. . . With others, increase in size takes place by a sudden leap. . .” (Ashton, 1926, pp. 572-573). Little dedicates a section of his 1962 empirical study to the distribution of growth rates, and also finds that the distribution is fat-tailed. However, he concludes the section without proposing any theoretical explanation: “I do not know what plausible hypothesis explains the highly leptokurtic nature of the distributions” (Little, 1962, p. 408). Recent empirical research into industrial dynamics has discovered that the distribution of firm growth rates closely follows the Laplace distribution, also known as the symmetric exponential distribution. Using the Compustat database of US manufacturing firms, Stanley et al. (1996) and Amaral et al. (1997) observe a ‘tent-shaped’ distribution characterized by a straight line on logarithmic plots that corresponds to the Laplace density. The Laplace distribution is also found to be a rather useful heuristic when considering growth rates of firms in the worldwide pharmaceutical industry (Bottazzi et al., 2001). Giulio Bottazzi and coauthors extend these findings by considering the Laplace density in the wider context of the family of Subbotin distributions. They find that, for the Compustat database, the Laplace is indeed a suitable distribution for modelling firm growth rates, at both aggregate and disaggregated levels of analysis (Bottazzi and Secchi, 2003a). The Laplacian

nature of the distribution of growth rates also holds for other databases, such as Italian manufacturing (Bottazzi et al., 2007).¹ In addition, the Laplace distribution appears to hold across a variety of firm growth indicators, such as Sales growth, employment growth or Value Added growth (Bottazzi et al., 2007). The growth rates of French manufacturing firms have also been studied, and roughly speaking a similar shape was observed, although it must be said that the empirical density was noticeably fatter-tailed than the Laplace (Bottazzi et al., 2010).² In Figure 1, we use the Compustat database to show the heavy-tailed distribution of annual employment growth rates for large US firms.

In this paper, we argue that it would be fruitful to conceive firms as being composed of discrete, interrelated resources, that are subject to local interactions, and susceptible to containing some degree of organizational slack. We sketch out two similar theoretical models that rely on these characteristics of firms to explain a number of ‘stylized facts’ of firm growth. In section 2 we review and discuss previous models of industry growth. Section 3 contains a discussion whereby we identify some features common to firms – i.e. firms can be seen as composed of lumpy, indivisible resources that are subject to non-linear interactions. In Section 4 we present a simple model of employment growth in a hierarchical organization. Section 5 contains an extended model that is computationally rather complex, the properties of which we explore using simulation analysis. Section 6 concludes.

2 A Discussion of Previous Models

There is something of a tradition in Industrial Organization modelling to represent growth processes in purely stochastic terms. Ijiri and Simon (1977) offered an explanation of the skewed firm size distribution in terms of a random process in which the probability of a firm taking up an additional business opportunity is

¹Reichstein and Jensen (2005) investigate the growth rate distribution of Danish firms, and, unlike most previous work, they observe asymmetries in the growth rate distribution, such that the Laplace is a better fit to the upper tail than the lower tail of the growth rate distribution.

²i.e. the observed subbotin b parameter (the ‘shape’ parameter) is significantly lower than the Laplace value of 1. This highlights the importance of following Bottazzi et al. (2002) and considering the Laplace as a special case in the Subbotin family of distributions.

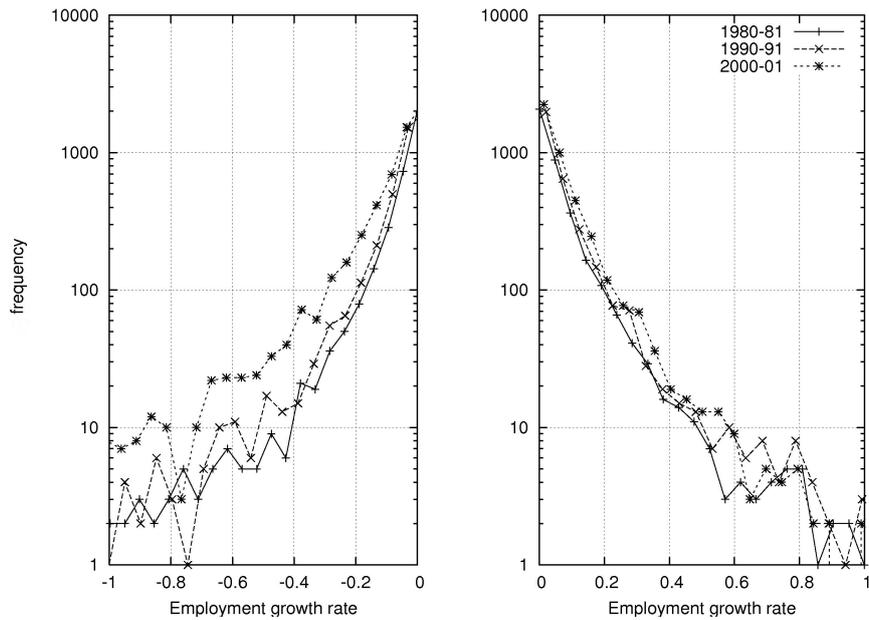


Figure 1: The empirical distribution of employment growth rates, elaborated by the author using the Compustat dataset for large US firms. Note the log scale on the y -axis. Employment growth is calculated in the usual way of taking log-differences of total employment in successive periods. Annual employment growth rates are calculated for the periods 1980-81, 1990-91 and 2000-01, obtaining 5256, 5931 and 7948 observations respectively. For each period, firms are sorted into 100 bins.

conditional upon its size. This model, dubbed the ‘island’ model because of the independent arrival of the growth opportunities, has been widely accepted, and interest in it was recently revived by Sutton (1998).

Previous models that focused on the exponential distribution of firm growth rates have also taken the approach of stochastic explanations. Amaral et al. (1997) develop a model in which the emergence of the distribution rests on a particular specification of the functional form of the stochastic growth process. However, there is little justification of the choice of such a functional form, and so one might consider that their model of the growth rate distribution appears to be more of a tautology than an explanation. The model in Bottazzi and Secchi (2003b) and Bottazzi and Secchi (2006a) also conceives of firm growth as a random process – “in our model luck is the principal factor that finally distinguishes winners from losers among the contenders” (Bottazzi and Secchi, 2006a, p. 236). The allocation of growth opportunities is governed by the ‘Polya Urn’ statistics, and as a justification of such mathematical apparatus they evoke the principle of ‘increasing returns to growth’ in the competitive process.³ As a result, their model generates ‘lumpiness’ in the growth of sales by referring to increasing returns and competition, whereas the model developed here, which can in some sense be seen as complementary, generates ‘lumpiness’ as a consequence of the indivisible nature of productive resources. Within the period of a calendar year, Bottazzi and Secchi suppose that the probability of a growth opportunity being taken up depends positively on the number of growth opportunities already taken up that year. This hypothesis, however, is difficult to reconcile with the existence of a small *negative* year-on-year autocorrelation of growth rates observed in many (but not all) empirical studies.⁴ Furthermore, given that the resulting distribution is determined by the choice of the underlying stochastic process, there are limits to how much such models can actually ‘explain’.

The choice of stochastic models to describe industrial evolution bears witness to a reluctance to generalize across firms. Firms grow for a wide variety of different

³More in the sense of Arthur (1989) rather than in that of the Kaldor-Verdoorn ‘dynamic increasing returns’.

⁴For a survey of growth rate autocorrelation, see Coad (2007a).

reasons, they are indeed heterogeneous, and it is believed that the best or only way to model growth may be by treating it as purely stochastic. To move beyond describing industry dynamics in terms of purely random shocks, we need to address the following question: “Can we generalize across firms?” Our answer is: “Yes we can, to some degree”.

Without denying the complexity of commercial organizations or the heterogeneity that exists between firms from different sectors of the economy, we maintain that there are some general features that are present in firms. (Indeed, Simon (1962) suggests that there are some broad features that appear to be common not only to all firms but to all complex systems!) The theoretical explanation proposed here is rooted in the ‘resource-based approach’, which views firms as being composed of discrete, complementary resources (Penrose, 1959). In addition, we allow for the possibility of growth being accommodated by organizational slack. Organizational slack is a widely-recognized characteristic of business firms – indeed, a firm’s resources will not be fully utilized at any given time for a number of reasons.⁵ However, managers will seek to use a firm’s resources efficiently, to have them as close as possible to ‘full utilization’. If a firm’s resources are underutilized, then growth can feed off these slack resources.⁶ On the other hand, if resources are already more or less fully employed, then growth will only be possible with the addition of new resources. In the former case, growth requires no additional investment, whilst in the latter case, firm growth will be accompanied by potentially wide-scale investment. This depiction of firm growth can be expressed in terms of self-organizing criticality. The firm can be seen as a system which tends to a ‘critical state’ of full utilization of its resources, as managers strive to organize the firms resources efficiently within the firm’s hierarchical framework.

⁵Here are a few possible examples. Slack may be present because indivisibilities of key inputs may prevent a firm from attaining perfect productive efficiency. Also, slack may creep in as the learning-by-doing effects that increase a worker’s productivity are not counterbalanced by increasing demands made of the worker. Furthermore, slack may be necessary because firms must be able to adapt and act flexibly in response to unforeseen contingencies and the changing market environment.

⁶Penrose writes “[a]t all times there exist, within every firm, pools of unused productive services and these, together with the changing knowledge of management, create a productive opportunity which is unique for each firm.” (Penrose, 1960, p. 2).

Depending upon the criticality of the system, the addition of an activity during growth will result in a (marginally) increased strain for many associated resources, thus potentially triggering off a chain reaction of subsequent growth across the whole of the organization. In this vein, Dixon comments on the criticality of a firm at a more general level: “the later addition of one person to regular activities can bring into operation a chain of reactions in the form of salaried employee increases, salary increases, and fixed asset additions” (Dixon, 1953, p. 50). Similarly, Hannan writes: “changes in one organizational feature often generate cascades of additional changes, because of the interdependence among parts of an organization” (Hannan, 2005, p. 61). Weick and Quinn put it this way: “Small changes can be decisive if they occur on the edge of chaos... in interconnected systems, there is no such thing as marginal change” (Weick and Quinn, 1999, p. 378). The ‘avalanche’ will only stop if there is sufficient slack capacity to absorb the extra workload associated with the additional resources.

To illustrate this idea, we propose models that are capable of generating heavy-tailed growth rate distributions within the time series of a single firm. In Section 4 we consider the special case of the propagation of employment growth throughout the various levels of a firm’s hierarchy. The organization of production in a hierarchy is indeed a general feature of all firms – in fact, in the Transaction-Cost-Economics literature, the words ‘firm’ and ‘hierarchy’ are used almost interchangeably. In this context, a firm may grow by adding an additional worker on the factory shopfloor, who will require the attention of a supervisor. It may occur, however, that all of the current supervisors are already too busy to take on this extra burden of supervision. With a small probability, then, the addition of this supplementary worker requires that the firm hire another supervisor. Furthermore, the addition of a supervisor may then increase the administrative workload of the central office, such that this latter also needs to hire a supplementary worker. An analogy with the classic ‘sandpile’ model⁷ can therefore be drawn, as the addition

⁷See Bak and Chen (1991); see also Bak et al. (1993) for an economic application. In the ‘sandpile’ model, grains of sand are dropped on top of each other until a sandpile is formed. “[R]andomly dropping on additional sand will result in the slope of the pile increasing to a critical slope, at which point avalanches of all sizes (limited only by the size of the pile) can occur in response to the dropping of a single additional grain of sand.” (Bak et al. (1993) p. 7).

of a supplementary worker can lead to a ‘snowball effect’ of hiring of employees at higher levels of the hierarchy. However, the organization of employees in a hierarchy is only one example of many ‘tree structures’ with multiple layers that can be used to describe the arrangement of a firm’s resources. We expand upon the basic employment hierarchy model in Section 5, so that our model consists not only of labour but also other inputs such as capital.

3 Theoretical Foundations

The path-breaking book of Penrose (1959) is a milestone for research into the theory of the firm. In this book, Penrose explains that firms are composed of ‘resources’ which are idiosyncratic assets that are essential inputs into the productive process. Although Penrose’s book is mainly concerned with human resources (in particular, the scarce resource that is managerial talent), other authors have identified other examples of resources. Brand names, in-house knowledge of technology, employment of skilled personnel, trade contracts, machinery, and efficient procedures are other such examples (Wernerfelt, 1984). Montgomery (1994) suggests that Disney’s cast of animated characters can be viewed as a resource, that has been observed to fuel diversification. Somewhat more unusual is the affirmation that even emotions such as anger and frustration can be considered to be organization-specific ‘resources’ (Feldman, 2004, p. 304). Furthermore, Winter (1995) comments on the similarity of the Penrosian concept of ‘resources’ and the evolutionary notion of ‘organizational routines’ and concludes that even routines can be considered as resources.⁸

One of the major features of these resources is their indivisible nature. This was described quite clearly by Penrose and has been recognized by many subsequent scholars. To summarize, Garnsey writes “Penrose pointed out that many of the resources required for expansion are only available in multiples that do not match

⁸Winter writes “routines clearly qualify as resources, given the expansive use of the term ‘resources’ in the literature of the resource-based view. . . . a routine in operation at a particular site can be conceived as a web of coordinating relationships connecting specific resources. . . .” (Winter, 1995, pp. 148-149).

up, as where new equipment creates excess capacity. This creates incentives to exploit unused resources through further growth.” (Garnsey, 1998, p. 539).

The indivisible resources that form the basis for a firm’s productive potential are not perfect substitutes but they need to be combined in roughly constant proportions in order for the firm to produce its output. As a consequence, firms strive to find those combinations of resources that reduce slack. Penrose describes this idea in these words: “[i]f a collection of indivisible productive resources is to be fully used, the minimum level of output at which the firm must produce must correspond to the least common multiple of the various maximum outputs obtainable from the smallest unit in which each type of resource can be acquired.” (Penrose, 1959, p. 68). It follows that “[u]nused productive services are, for the enterprising firm, at the same time a challenge to innovate [and] an incentive to expand . . .” (Penrose, 1959, p. 85). In other words, the resources in a firm are interdependent because, under circumstances where firms strive for the most efficient combination of resources, the addition of one indivisible resource may well have consequences on the desirable levels of other resources. This may lead to non-linearities in the growth process as firms add indivisible resources to arrive at an efficient level of production.

4 A Simplified Model

4.1 Intuition of the model

In the previous discussion, we argued that firms grow by adding discrete resources to a complex of interdependent resources that they already possess. In this section, we present an analytical model capable of reproducing the empirically observed functional form of the growth rate distribution by exploiting a few simplifying assumptions. Our attempts to construct a formal model, it would appear, receive the blessing of those who are concerned that the modelling of heavy-tailed phenomena has received insufficient attention in the literature (McKelvey and Andriani, 2005).

We characterize a firm as being composed of a relatively large number of hi-

erarchies.⁹ The bottom layer of the firm (i.e. the very lowest hierarchical level) is composed exclusively of production workers, whilst all of the other levels are composed of managers whose task is to supervise either production workers or subordinate managers (see Figure 2 for an illustration). A firm grows by adding a production worker. The number of managers is determined by the number of production workers and also by limits on the efficient span of control, α , which correspond to the maximum number of subordinates that a manager can effectively supervise. “At executive levels [the span of control] is seldom less than three, and seldom more than ten, and usually lies within narrower bounds – particularly if we take averages over all executives in an organization at a given level.” (Simon, 1957, p. 32). In this model, though, we do not need to attribute any specific numerical value to α and so we leave it in algebraic form. It is computationally helpful, and also theoretically meaningful, however, to assume that α is a whole number that is strictly greater than unity (i.e. $\alpha \in \mathbb{N}^+$, $\alpha > 1$). For analytical simplicity, we assume that α is a constant and does not vary either within a hierarchical level or across levels (for a discussion of the plausibility of this assumption, see Williamson (1967) p. 128). For the purposes of this model, we also must assume that adjustment of the firm’s hierarchical organization to additional production workers occurs within one time period. Finally, we assume that the firm is initially at a stable state, such that it is already efficiently organized in the sense that it is not possible for it to employ fewer managers given the number of production workers and its given value of α (i.e. the limit on the efficient span of control). The reader may notice major similarities between the model developed here and the executive compensation model of Simon (1957) and the information flows model of Williamson (1967). The fact that the same hierarchical model has been applied in quite different contexts lends credibility to its use here – indeed, we cannot be accused of having conclusions that emerge from *ad hoc* modelling assumptions.

A summary understanding can be obtained by looking at Figure 2. Two im-

⁹We do not need to define the number ‘large’ nor define what happens at the very top of the hierarchy. Also, we do not need to suppose that the number of hierarchies tends to infinity, because we only want to explain the distribution of growth rates for a certain limited range. An implication of this assumption is that this model is not suitable for describing growth processes in very small firms.

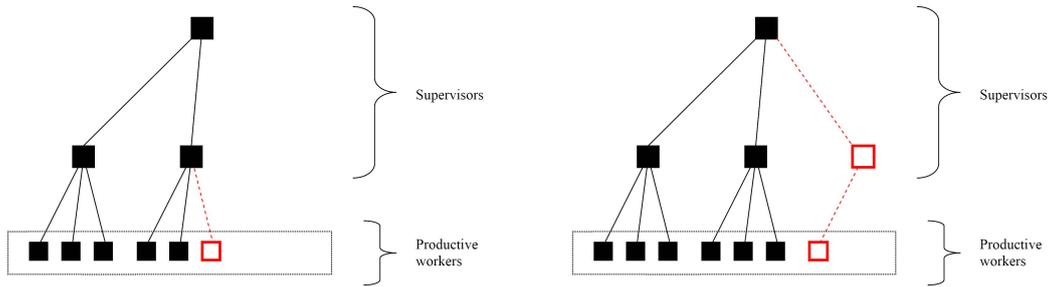


Figure 2: An illustration of the underlying intuition of the model, where the span of control is $\alpha=3$. Depending upon the ‘criticality’ of the system, the addition of a production worker may lead to an increase in the number of supervisors further up the hierarchy. If there is some slack in the system, a production worker can be added and new supervisors need not be added (see left). If, however, the attention of supervisors is already at full utilization, the addition of a production worker will require the addition of a supervisor (see right).

portant points should be emphasized. First, there is a distinction between total production n and total employment x . Total production corresponds to the number of production workers, i.e. n , while total employment corresponds to the number of both production workers and supervisors combined, i.e. x . Second, it should be noted that we do not attempt to generalize on the sources of growth opportunities, but rather we focus on how firms build upon given growth opportunities.¹⁰ We argue that the fat-tailed distribution of growth rates does not come from the distribution of opportunities available to firms, but rather on the reactions of firms to growth stimuli. The model is admittedly a gross simplification and does not take into account such factors as the interdependence of growth rates between firms, flexibility of α (the span of control parameter), liquidity constraints that limit growth, or limits on the availability of suitable workers. Nonetheless, its simplicity will make it clear to what properties we owe the emergence of the distribution.

¹⁰Our model is thus in line with the previous theoretical models of industrial structure and dynamics reviewed in Section 2, where growth opportunities are supposed to arrive by themselves and little attention is paid to their source.

4.2 Formal model

Let us begin with the simplest possible case, considering one firm that grows by adding just one production worker (i.e. $\Delta n = 1$). If new production workers can be integrated without having to add a supervisor, we have $\Delta n = \Delta x$; i.e. the number of production workers added is equal to change in total employment. It is possible, however, that all of the managers in the second hierarchical level (i.e. those that supervise the production workers) are already fully occupied. This will occur when the number of production workers (before adding the new one) is exactly a multiple of α . If this is the case, the arrival of the supplementary worker will require that one supplementary manager be hired at the next hierarchical level. This scenario will occur with probability $1/\alpha$. However, the arrival of this new manager at the second level may add to the workload of managers on the third hierarchical level, and so on. The probability that the addition of a production worker leads to at least *two* managers being hired at two successive levels is $1/\alpha \times 1/\alpha = 1/\alpha^2$. We can continue with this reasoning to end up with the following distribution of employment growth:

$$\text{Prob.}(\Delta x \geq 1 | \Delta n = 1) = 1$$

$$\text{Prob.}(\Delta x \geq 2 | \Delta n = 1) = 1/\alpha$$

$$\text{Prob.}(\Delta x \geq 3 | \Delta n = 1) = 1/\alpha^2$$

...

and so on. Formally, we have an exponential distribution with the following functional form:

$$P(\Delta x \geq \gamma | \Delta n = 1) = \alpha^{1-\gamma} \quad (1)$$

or, expressed differently,

$$P(\Delta x = \gamma | \Delta n = 1) = \alpha^{1-\gamma}(1 - 1/\alpha) \quad (2)$$

where γ is a positive integer ($\gamma \geq \Delta n$). We therefore observe that the distribution of total employment growth (Δx) of a firm that grows by adding one

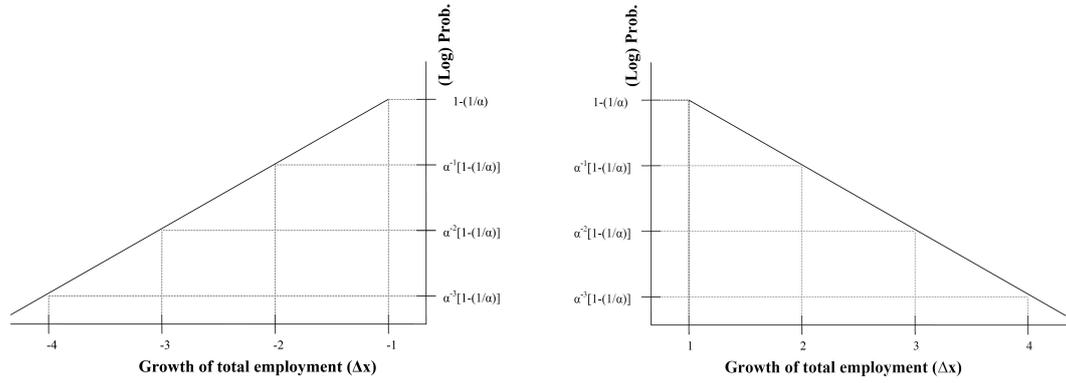


Figure 3: The distribution of growth of total employment if a firm grows by $\Delta n = 1$ (right figure, see Result 1), or if a firm shrinks by $\Delta n = 1$ (left figure, see Result 2).

production worker will follow an exponential distribution.¹¹

It is also possible to generalize for the case where a firm grows by adding $\Delta n \in \mathbb{N}^+$ production workers (with, of course, $\Delta x \geq \Delta n$). For $\Delta n < \alpha$, we obtain the following distribution:

$$\begin{aligned}
 P(\Delta x = \gamma | \Delta n) &= 1 - \Delta n / \alpha && \text{if } \gamma = \Delta n \\
 P(\Delta x = \gamma | \Delta n) &= \Delta n \cdot \alpha^{\Delta n - \gamma} (1 - 1/\alpha) && \text{if } \gamma > \Delta n
 \end{aligned} \tag{3}$$

where equation (2) corresponds to the special case where $\Delta n = 1$.

Thus, we have:

Result 1 *The distribution of growth of total employment of a firm that adds Δn production workers will follow an exponential distribution.*

An illustration is offered in Figure 3 (right).

Analogical reasoning can be applied to the case where a firm shrinks in size. Consider a firm that shrinks by Δn units – at least one supervisor will no longer

¹¹The reader may have noticed that the focus of this example is on growth *amounts*, rather than growth rates. However, our analysis investigates the potential growth increments *for a firm of a given size*. In this particular case, therefore, there is a direct correspondence between growth amounts and growth rates.

be needed when, after shrinking, the number of production workers is an exact multiple of α . Formally, we can still use Equation (3), where a firm shrinks by Δx employees as a response to shedding Δn production workers (with $\Delta n < \alpha$), i.e.:

$$\begin{aligned} P(\Delta x = \gamma|\Delta n) &= 1 - \Delta n/\alpha && \text{if } \Delta x = \Delta n \\ P(\Delta x = \gamma|\Delta n) &= \Delta n \cdot \alpha^{\Delta n - \gamma}(1 - 1/\alpha) && \text{if } \Delta x > \Delta n \end{aligned} \quad (4)$$

Result 2 *The distribution of growth of total employment of a firm that shrinks by Δn production workers will follow an exponential distribution.*

An illustration is given in Figure 3 (left).

Autocorrelation dynamics It is also possible to derive the conditional autocorrelation dynamics of the firm's growth dynamics. Consider the case where one production worker is added in each period t , i.e. $n_t = n_{t-1} + 1$. The conditional growth autocorrelation can be written as:

$$\begin{aligned} P(x_t - x_{t-1} > 1 | x_{t-1} - x_{t-2} > 1) &= 0 \\ P(x_t - x_{t-1} > 1 | x_{t-1} - x_{t-2} = 1) &= \frac{1}{\alpha - 1} \end{aligned} \quad (5)$$

If the firm experienced a growth spurt in the previous period (i.e. $x_{t-1} - x_{t-2} > 1$), it has a probability of zero of repeating this growth performance in the following period. If, however, the firm added a production worker in the previous period but this did not trigger off the addition of a supervisor, then the addition of a production worker in this period has a positive probability of leading to the further addition of a supervisor.

Given that $E(x_t - x_{t-1} | x_{t-1} - x_{t-2} > 1) = 1$ and $E(x_t - x_{t-1} | x_{t-1} - x_{t-2} = 1) > 1$, we observe negative growth autocorrelation in the case where a growth spurt of x was triggered in the previous period (i.e. when $x_{t-1} - x_{t-2} > 1$).

We will pursue our analysis of conditional autocorrelation profiles in Section 5.3.

4.3 Discussion

The model is admittedly far too simple to be realistic, yet its simplicity makes for greater visibility of the source of the emergence of the symmetric exponential distribution. The model can be seen as the simplest model in a family of possible models that view firms as coherent collections of resources that are complementary and discrete. These latter are subject to localised interactions and embedded in an organization that tends to a critical state of full utilization of its resources. In this context, a small growth stimulus working through local interaction channels can be transmitted throughout a firm to produce potentially large-scale effects. We argue that it is these properties that explain the emergence of the observed fat-tailed growth rate distributions.

The model describes the dynamics of a single, lone organization and makes no attempt to account for competitive interactions between firms. In our view, this is not a serious flaw. Other explanations of the fat-tailed growth rate distribution have emphasized the complex nature of inter-firm competition as the source of the emergence of the observed distribution (e.g. Bottazzi and Secchi (2006a), McKelvey and Andriani (2005)). Recent empirical work has nonetheless cast doubt on the importance of inter-firm competition as a factor conditioning firm growth rates. Sutton (2007) analyzes the dynamics of market shares of the largest and second largest firms in a number of Japanese industries, and finds (perhaps surprisingly) that their market share dynamics can be modeled as statistically independent. Only in the case where the combined market share of an industry's two largest firms is at least 90% of the industry total does inter-firm competition leave a detectable statistical footprint. Geroski and Gugler (2004) consider the impact of the growth of rival firms on a firm's employment growth, using a database on several thousand of the largest firms in 14 European countries. Rival firms are defined as other firms in the same 3-digit industry. In their main regression results (their Table 2) they are unable to detect any significant effect of rival's growth on

firm growth, although they do find a significant negative effect in specific industries (i.e. differentiated good industries and advertising intensive industries). In our model, it is the complex nature of interactions between the resources *within* a firm, rather than the competitive struggle between firms, that accounts for the emergence of the observed growth rates distribution.

Some caveats of the model should nonetheless be mentioned. First, the model only considers the case of employment growth and does not consider other aspects of firm growth (such as growth of sales or growth of fixed capital). Second, all of the interactions take place in a vertical direction (up or down the hierarchical channels) rather than in a horizontal direction. The model therefore places strict limits on the nature of interactions that is probably not a realistic portrayal of how a firm's resources interact. Third, the addition of an extra production worker at the bottom of the hierarchy is assumed to lead to a modification of total employment that is, if anything, instantaneous. This might not be an accurate assumption, however, considering that: “[t]here is a considerable time lag between the growth of numbers of production workers and the expansion in employment of other personnel” (McGuire, quoted in Starbuck (1971) p. 54). Fourth, the model is very simple – this is both an advantage and a disadvantage. Such a simple model cannot be an accurate portrayal of firm growth.

5 Extending the Model

The previous model had many limitations, as we have already discussed. In an attempt to improve the model, we now try to develop it into something more realistic by introducing capital as a factor of production. Whereas the previous model consisted of production workers and supervisors, here we include in an analogous fashion machines and production plants, with the restriction that each production plant can house a limited number of machines. We also introduce a further input which serves as a ‘numeraire’, which combines production labour and machines in order to produce the final good. For the sake of simplicity, we do not allow for any substitution between labour and capital in the production of the

final good, although this could be relaxed in further work.

The following model serves as a useful illustration of the growth mechanism. Due to difficulties in applying optimization algorithms to cases involving integer restrictions, however, we are prevented from firmly establishing the sensitivity of the model to different parameter settings. As such, the model is at a preliminary stage and should be seen primarily as an aid to intuition.

5.1 Introducing the Model

Consider a firm's production Q that is produced from N inputs in the following way:

$$Q = f(A, B, \dots, N) \quad (6)$$

Equation (6) implies that production of final goods requires many inputs. These inputs are lumpy and indivisible. In reality, there are many such indivisible inputs to a production process, In the present model, however, we limit ourselves to the five inputs mentioned above. In this example we envisage that A corresponds to the basic raw materials required for a 'production run', B to the number of production employees, C to the number of 'machines', D to the number of supervisors (non-production workers), and E to the number of production plants. It is also meaningful to restrict the values of these inputs to integer number values, i.e. $(A, B, C, D, E) \in \mathbb{N}^+$.

The production function has the following functional form:

$$\begin{aligned} Q &= f(A, B, C, D, E) \\ &= \text{Min}(A, B \text{ DIV } r^B, C \text{ DIV } r^C, D \text{ DIV } r^D, E \text{ DIV } r^E) \end{aligned} \quad (7)$$

where $\text{Min}(\cdot)$ is the function that selects the minimum from the list of values

in parentheses.¹² The DIV operator corresponds to integer division, whereby the result is integer and any remainder is discarded.¹³ Integer division effectively helps us to model the indivisible nature of the inputs.

The r^B coefficient (with $r^B < 1$) corresponds to the requirements of input B into the production process. Strictly speaking, r^B corresponds to the number of units of B required for one unit of A. In other words, a production run requires not only raw materials but also the labour services of an employee. Analogous interpretations exist for the r^C , r^D , and r^E coefficients.

Using a similar reasoning to that developed in Section 4, it follows that the addition of one more production run has an (unconditional) probability of r^B of resulting in the addition of a production worker, as well as an (unconditional) probability of r^C of resulting in the addition of a machine. In turn, any additional employee may lead to the addition of a supervisor, and the need for an extra machine may lead to the construction of a new production plant.

In this extended model, a firm's inputs are still organized in a hierarchical fashion, with the structure of the hierarchy being determined by the input requirement coefficients r . For instance, the addition of production employees will eventually trigger the need for an extra supervisor, once the threshold implicit in the production requirement coefficients is crossed.

It is thus apparent that the marginal costs of adding one more production run depend upon the 'criticality' of the system. In some cases, there is enough slack to accommodate the production run without any repercussions. In other cases, adding a production run results in cascading investments throughout the organization. In what follows, we will observe how profit-maximizing firms decide whether or not to satisfy marginal changes in demand with additional production runs, depending upon the 'criticality' of the organization of the firm's productive inputs. As such, the firm's marginal cost depends upon the criticality of the

¹²Our production function emphasizes complementarity in inputs, rather than substitutability. As a result, it bears some similarities to a Leontief production function (which also emphasizes complementarities) while differing from the standard Cobb-Douglas production function (which allows substitutability).

¹³In other words, the result of $B \text{ DIV } r^B$ is the integer obtained when B is divided by r^B and any remainder is thrown away. For example, $7 \text{ DIV } 2$ is the integer 3.

system.

A firm's cost function can be written as:

$$COST = g(A, B, C, D, E) = (c^A A + c^B B + c^C C + c^D D + c^E E) \quad (8)$$

Where c^A is the cost to the firm of input A , and where analogous definitions hold for c^B , c^C , c^D , and c^E .

For analytical convenience we make the usual assumption that firms seek to maximize profits:

$$\max_Q P \cdot Q - COST \quad (9)$$

subject to given demand conditions. For simplicity, we consider a downward-sloping demand curve which corresponds to the case of imperfect competition:¹⁴

$$P = h(Q) = \beta - \phi Q \quad (10)$$

where β and ϕ are parameters.

5.2 Quantitative examples

In this section we will consider the case with five inputs: A, B, C, D and E .

Substituting (10), (7) and (8) into (9), we obtain:

¹⁴If we assume that the conditions of demand are those of perfect competition (i.e. that price is given) then firms face no scale effects and the model in its present form does not rule out infinite increases in size. In this case, it would thus be necessary to impose some restriction which prevents the apparition of cases in which firms can instantly and costlessly experience implausibly large increases in size. To this end one might introduce the existence of some sort of adjustment costs. In the example developed here, however, the demand curve corresponds to the case of imperfect competition. In this paper, which focuses on one lone firm, imperfect competition constitutes the only point of contact of the firm with the rest of the economy. Given that the assumption of imperfect competition is restrictive and may limit the generalisability of the model, future work could relax this assumption.

$$\max_{A,B,C,D,E} \left\{ \begin{array}{l} \beta \cdot \text{Min}(A, B \text{ DIV } r^B, C \text{ DIV } r^C, D \text{ DIV } r^D, E \text{ DIV } r^E) \\ -\phi \cdot \text{Min}(A, B \text{ DIV } r^B, C \text{ DIV } r^C, D \text{ DIV } r^D, E \text{ DIV } r^E)^2 \\ -(c^A A + c^B B + c^C C + c^D D + c^E E) \end{array} \right\} \quad (11)$$

This problem is obviously too complex to solve analytically. To our knowledge, it is also too difficult to solve using standard optimization software, because the integer restrictions (brought on by integer division) pose serious problems to conventional optimization algorithms.

To find the maximum, it seems necessary to adopt a ‘brute force’ approach. We will attempt to cycle through all possible combinations of the (‘reasonable’) integer values for the inputs and then select out the combinations of inputs that bring about profit maximization.¹⁵ The profit-maximizing input sets are then recorded, and the time series will subsequently analyzed.

We now observe how the firm responds to changes in its market environment. This is modelled by seeing how the firm reacts to changes in the demand parameter β . Changes in β can either stem from exogenous market developments or from the firm’s own activity (e.g. advertising), and could correspond to the arrival of business opportunities assumed in the ‘islands models’ literature (Ijiri and Simon (1977), Sutton (1998), Bottazzi and Secchi (2006a)).

We therefore propose that demand follows the following simple growth process:

$$\beta_{t+1} = (1.01)\beta_t \quad (12)$$

¹⁵Cycling through all possible combinations would be extraordinarily computationally intensive. For example, if we allow variables $A - E$ to take any value from 0 to 100, we end up with $100^5 = 10^{10}$ combinations. For this reason, we restrict ourselves to ‘reasonable’ values of the inputs. In choosing the initial conditions, we begin by choosing values for the inputs that are relatively close to the values obtained from optimization when the integer constraint is relaxed. We then explore constellations of inputs around this initial value (i.e. localized search) to obtain the initial profit-maximizing input constellation. In subsequent time periods, as demand increases, we apply iterative localized search procedures, bearing in mind that in most cases the corresponding profit-maximizing set of inputs will be ‘close’ (if not equal) to those obtained in the previous step.

Table 1: Parameter settings for the three simulation scenarios: the baseline, capital-intensive and labour-intensive scenarios.

	R^b	R^c	R^d	R^e	C^a	C^b	C^c	C^d	C^e
baseline	0.14	0.09	0.022	0.003	21	26	34	74	310
K intensive	0.14	0.13	0.023	0.007	22	30	34	81	310
L intensive	0.19	0.08	0.044	0.004	17	22	39	67	336

which corresponds to a demand growth of 1% in each period t .

Given this dynamic demand, we will now observe the growth series of the profit-maximizing levels of inputs and total production. We will then observe the distribution of the resultant growth rates.

The parameter settings for the simulation models are presented in Table 1. In our baseline scenario, for example, the parameters imply that one production plant can efficiently house up to about 7 supervisors, about 30 machines, about 45 production workers and about 330 units of raw materials (corresponding to 330 production runs). To explore the sensitivity of our results to these parameters, variations on the baseline case are explored in the capital-intensive and labour-intensive scenarios.

5.3 Properties of the Model

We now investigate some properties of the model. In keeping with common practice in simulation modelling, we let the model run for a short while before taking the readings. More specifically, after running the model for 250 periods for each of the three scenarios, we discard the first 50 periods and focus on the last 200 observations only.

In keeping with empirical work on firm growth, growth rates are calculated by taking log-differences of size levels.¹⁶

¹⁶See Coad (2009) for a survey.

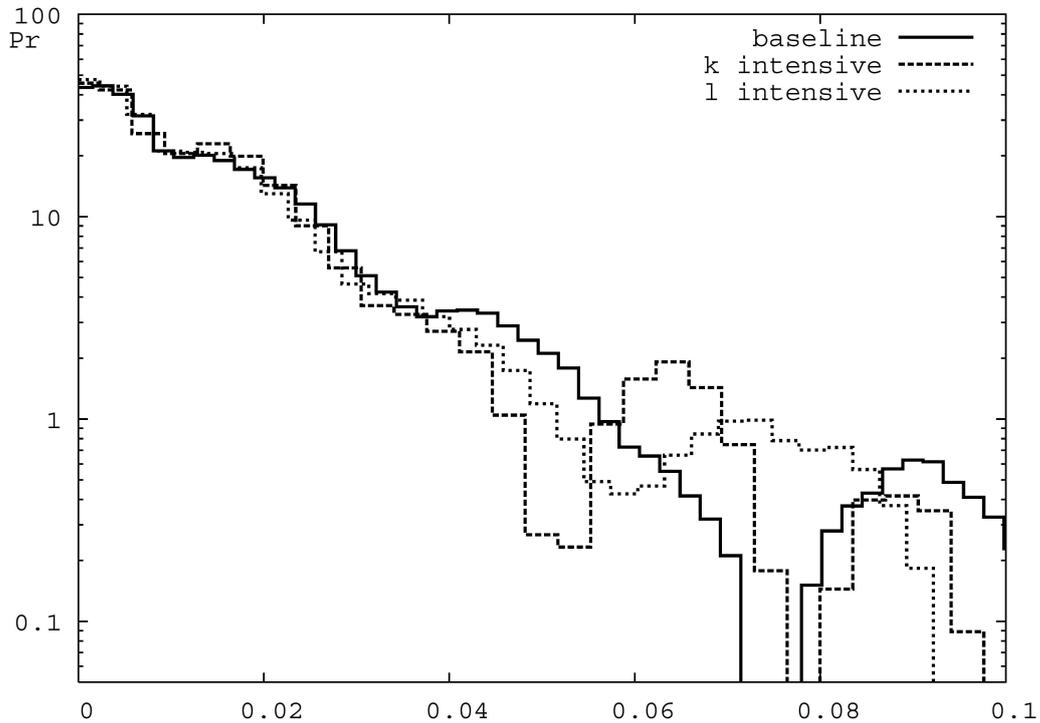


Figure 4: Kernel densities of the distribution of growth rates of output Q , for the three simulation scenarios. Kernel densities computed for 50 equispaced points using an Epanenchnikov kernel in `gbutils` 5.2. Note the log scale on the y axis.

5.3.1 Growth rates distribution

We consider first the distribution of growth rates. Table 2 presents some statistics on the cumulative distributions of some of the growth rate series, for each of the three scenarios (baseline, capital intensive, and labour intensive). These results provide evidence of fat tails in the growth rate densities of output, employment, and machines.¹⁷ Figure 4 shows the distribution of growth rates of output. This figure provides further evidence on the fat tails of the growth rates distribution.

¹⁷Note that we do not report the results for growth of profits, because this series has a strong trending component, which comes from the nature of demand growth.

Table 2: Cumulative distributions of the growth rate series (periods 51-250).

	≥ 0	≥ 0.01	≥ 0.02	≥ 0.04	≥ 0.08
Baseline					
Growth of output (Q)	200	77	41	11	2
Growth of employment (B and D combined)	200	77	42	8	2
Growth of machines (C)	200	80	40	15	2
K-intensive					
Growth of output (Q)	200	82	36	10	2
Growth of employment (B and D combined)	200	77	35	10	3
Growth of machines (C)	200	77	38	11	2
L-intensive					
Growth of output (Q)	200	79	36	11	2
Growth of employment (B and D combined)	200	69	36	9	2
Growth of machines (C)	200	68	33	15	4

5.3.2 Autocorrelation of growth rates

Many studies, although not all, have found small negative autocorrelation in firm growth rates.¹⁸ Autocorrelation coefficients from simple regressions are presented in Table 3. Regressions are performed using either OLS or least absolute deviation (LAD) techniques.¹⁹ OLS regressions suggest that, if anything, there is a mild negative autocorrelation, although the coefficients obtained from median regressions (i.e. the LAD regressions) are lower and less significant. These simulated results are thus similar to the findings in Bottazzi et al. (2010) and Coad (2006) where growth rate autocorrelation is strongly negative on average (using OLS which calculates the ‘average effect’) whilst growth rate autocorrelation for the median firm (using LAD which calculates something corresponding to the ‘median effect’) is much closer to zero.

Coad (2007a) presents a more detailed analysis of *conditional* growth rate autocorrelation using quantile autoregression techniques, which explore how growth rate autocorrelation varies over the conditional distribution of growth rates. Whilst for the average firm, growth rate autocorrelation is not so important, it is very

¹⁸See Coad (2007a) for a survey.

¹⁹OLS t -statistics that are robust to heteroskedasticity are calculated using the Huber/White/sandwich estimator (i.e. the ‘robust’ option in Stata 10). LAD t -statistics obtained after 1000 bootstrap replications.

Table 3: Autocorrelation of growth rates of output. OLS t -statistics robust to heteroskedasticity, while LAD t -statistics are obtained after 1000 bootstrap replications.

	Coefficient	t -stat	(pseudo-) R^2	Obs
Baseline				
OLS	-0.1243	-1.85	0.0155	200
LAD	-0.1213	-1.32	0.0112	200
K-intensive				
OLS	-0.1455	-2.55	0.0212	200
LAD	-0.0611	-0.69	0.0064	200
L-intensive				
OLS	-0.0816	-1.09	0.0067	200
LAD	-0.0054	-0.04	0.0001	200

important for the fast growth firms. Fast growth firms are quite unlikely to repeat their performance in the following year. As a result, the autocorrelation coefficient becomes very negative for these firms. The empirically observed autoregression profile is shown in the top-left panel in Figure 5. The three remaining panels in Figure 5 show quantile autoregression plots for our simulated data. We view the comparison of these plots to be a (small) success story. Whilst most firms experience no growth rate autocorrelation, it is the fastest growing firms (at the upper quantiles) that experience much stronger forces of negative autocorrelation. If a firm has just experienced widespread organizational growth, it is unlikely that it will grow in the following period. For such firms, growth can be accommodated by slack capacity.

5.3.3 Growth rates and size

Empirical investigations of Gibrat's Law have generally found a negative relationship between firm size and firm growth, but that above a certain size threshold firm growth appears to be relatively independent of size.²⁰ Our results are in line with these empirical findings. The results in Table 4 show that, in each of the three scenarios, growth rates are highest when the firm is smallest, but that for

²⁰For a survey, see Coad (2009, Chapter 3).

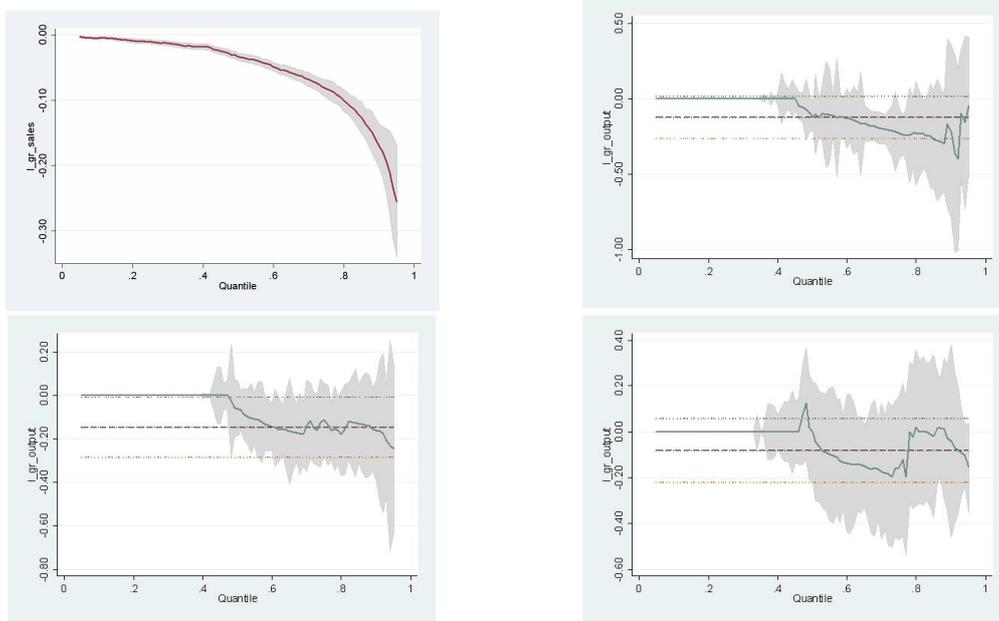


Figure 5: Quantile autoregression plots. Top left: empirical quantile autoregression profile for sales growth, for firms with non-negative sales growth rates (source: author's elaboration based on Coad (2007a, Figure 3) under the restriction of non-negative sales growth). Top right: quantile autoregression profile of output growth for the baseline simulated model. Bottom row: quantile autoregression profile of output growth for the Capital-intensive (left) and labour-intensive (right) simulations. Each plot shows the variation in the coefficient on lagged growth over the conditional quantiles of the growth rate distribution. Conditional quantiles (on the x -axis) range from 0 (for firms experiencing zero growth) to 1 (for the fastest-growing firms). Confidence intervals extend to 95% confidence intervals in either direction (confidence intervals for the simulation runs obtained after 100 bootstrap replications). Horizontal lines represent OLS estimates with 95% confidence intervals. Graphs made using the 'grqreg' Stata module (Azevedo (2004)).

the three other size groups, when the firm is larger, there no longer appears to be a clear relationship between firm size and growth.

In addition, most studies into firm growth (but not all) have observed that the variance of growth rates decreases with firm size.²¹ Table 4 presents some evidence in this direction. When we split the sample into four groups according to size, the standard deviation decreases in each of the three cases.

We suggest that our model is able to explain these empirical findings because firms grow by adding ‘lumpy’ resources, and the fixed size of these resources is relatively large (especially compared to the sizes of smaller firms). Small firms have to struggle more to internalize these discrete resources, and their growth is spurred on by slack caused by incomplete utilization of large resources, whereas larger firms are less affected by the indivisible nature of their resources.

5.3.4 The profits-growth relationship

Theoretical work into firm growth has often suggested that there is a strong link between relative financial performance and firm growth, building on the intuition that selection effects reallocate productive capacity away from inefficient firms towards more profitable firms (see among others Nelson and Winter (1982) and Jovanovic (1982)). Empirical work on the issue, however, has only found a weak link between these two variables, and this finding has been viewed as a puzzling challenge (Dosi, 2007; Coad, 2007b; Bottazzi et al., 2008).²²

Figure 6 shows some scatterplots of the relationship between growth of profits and growth of output, using the data obtained from our model. Although there is a positive relationship between these two variables, the relationship is not very clear-cut.

²¹A negative dependence of growth rate variance on size has been found in data on US manufacturing firms (Amaral et al. (1997); Bottazzi and Secchi (2003a)), for firms in the worldwide pharmaceutical industry (Bottazzi and Secchi (2006b)), and, to a lesser extent, for French manufacturing firms Bottazzi et al. (2010). In the case of Italian manufacturing firms, however, Bottazzi et al. (2007) do not observe any relationship between growth rate variability and size.

²²While most studies seem to focus on the relationship between (growth of) profit margins and firm growth (Dosi, 2007; Coad, 2007b; Bottazzi et al., 2008), other studies focus on the interrelationship between growth of profit (amounts) and firm growth (see e.g. Coad, 2010; Moneta et al., 2010). Both these streams of research have similar findings.

Table 4: Scaling of growth rate variance with firm size (size and growth measured in terms of Q).

	Mean size	Mean growth rate	Std. Dev.	Obs.
Baseline				
51-100	182.00	0.0120	0.0221	50
101-150	317.88	0.0104	0.0180	50
151-200	542.32	0.0103	0.0097	50
201-250	912.24	0.0107	0.0078	50
<i>K</i>-intensive				
51-100	175.24	0.0122	0.0306	50
101-150	314.36	0.0110	0.0162	50
151-200	535.82	0.0106	0.0126	50
201-250	904.42	0.0100	0.0068	50
<i>L</i>-intensive				
51-100	184.26	0.0119	0.0221	50
101-150	317.50	0.0101	0.0217	50
151-200	546.30	0.0107	0.0126	50
201-250	912.62	0.0101	0.0106	50

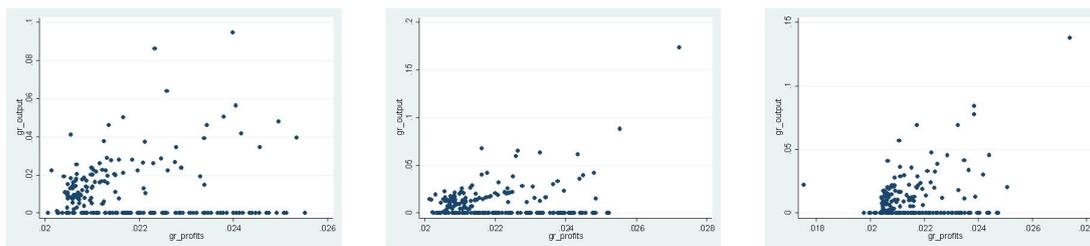


Figure 6: Scatterplots of the contemporaneous relationship between growth of profits and growth of output, for the baseline (left), capital-intensive (centre) and labour-intensive (right) simulations. Each plot shows 200 datapoints.

Table 5: Robust OLS estimation of equation (13) (coefficients and t -statistics).

	baseline	k_intensive	l_intensive
α_0	-0.0093 -0.42	-0.0570 -1.19	-0.0538 -1.21
α_1	0.9236 0.89	3.0822 1.40	2.9761 1.42
R^2	0.0060	0.0532	0.0504
obs	200	200	200

We estimate the following regression equation:

$$Gr_Q_t = \alpha_0 + \alpha_1 \cdot Gr_Profits_t + \varepsilon_t \quad (13)$$

and the results are reported in Table 5. We observe that the relationship between growth of profits and growth of output is always positive, but never statistically significant at conventional levels. The R^2 coefficient takes values of 5% or lower, which is surprisingly low.

It is indeed interesting that we only observe a weak relationship between profits and growth. Even in this deterministic model of a profit-maximizing firm, which has perfect information on its demand curve, profits are not strongly related to growth of output. Instead, the time series of output growth appears to be remarkably erratic and idiosyncratic. This model therefore provides an explanation for the lack of a strong relationship between financial performance and growth, an empirical finding that has been puzzling industrial organization scholars for quite some time now. In this model, growth is driven by the idiosyncratic configuration of indivisible resources, and as a result there is no clear relationship between growth and profits – even though the firm in this example is perfectly profit-maximising! However, other possible explanations of the absence of a strong relationship between profits and growth, such as bounded rationality on the part of firms (in particular, imperfect knowledge of the demand curve), as well as the possibility that firms do not seek to maximize profits, should also be kept in mind as additional moderating factors that weaken the relationship between financial

performance and firm growth.

6 Conclusion

We began this paper by observing that the distributions of growth rates of firms are distributed according to the Laplace distribution (also known as the symmetric exponential distribution). We then attempted to explain this emergent property. We acknowledge that there are many possible explanations for empirically observed regularities (Simon, 1968), especially for unconditional objects such as growth rate distributions (Brock, 1999). We therefore tried to base the assumptions of our model on theoretical descriptions of firm growth. While the Bottazzi and Secchi (2006a) model emphasizes competition between firms, our model focuses on the internal structure of business organizations to model the time series development of a single firm.

We drew upon insights from the resource-based view of the firm (Penrose, 1959) in order to present a class of models of firm growth. We viewed a firm as a coherent collection of complementary, discrete resources, that are subject to localised interactions, embedded in an organization that tends to a critical state of full utilization of resources. The lumpy nature of resources within a firm implies that firm expansion is characterized by nonconstant marginal costs that depend upon the degree of utilization of the firm's resources. In this context, a small stimulus working through local interaction channels can be transmitted throughout a firm to produce potentially large-scale effects.

In our first model we considered employment growth in the case of a hierarchical organization, in which a limit has been placed on the efficient span of control. We do not attempt to generalize upon where growth opportunities come from, but instead we consider how firms build upon growth opportunities. Adding an extra worker at the bottom of the hierarchy will marginally increase the workload at higher levels of the hierarchy, which may trigger off to a potentially large hiring of supervisors. It is observed that a firm's growth rate of total employment (production workers and supervisors combined) will follow an exponential distribution.

Our second model was an attempt to extend the early model, by introducing fixed capital into the model. This model was capable of explaining a number of ‘stylised facts’ about firm growth, such as the heavy-tailed growth rates distribution, the peculiar shape of the conditional autocorrelation profile, and the scaling of growth rate variance with size.

Further developments of the models presented here can also be envisaged. For example, it might be worthwhile extending the model to a population of heterogeneous firms (with heterogeneous costs and hence heterogeneous profits), instead of focusing on one lone firm. In this way, one could attempt to replicate empirical findings for the cross-section of firms. Such an extended model could presumably be driven by random fluctuations in firm-specific variables, rather than having the deterministic growth trend that features in the model in Section 5. In this model of heterogeneous firms, one might specify interaction rules such that, for example, firms with more slack are quicker at taking up new business opportunities than other firms, or that firms with slack are more open to innovation. Furthermore, it would be interesting to include a consideration of ‘time-to-build’ issues, such that large expansion plans, such as the building of a new plant, require investment over more than one time period.

References

- Amaral, L. A. N., Buldyrev, S. V., Havlin, S., Salinger, M. A., Stanley, H. E., and Stanley, M. H. R. (1997). Scaling behavior in economics: the problem of quantifying company growth. *Physica A*, 244:1–24.
- Arthur, W. B. (1989). Competing technologies, increasing returns, and lock-in by historical events. *Economic Journal*, 99(394):116–131.
- Ashton, T. S. (1926). The growth of textile businesses in the Oldham district, 1884–1924. *Journal of the Royal Statistical Society*, 89(3):567–583.
- Azevedo, J. P. W. (2004). grqreg: Stata module to graph the coefficients of a quantile regression. Technical report, Boston College Department of Economics.
- Bak, P. and Chen, K. (1991). Self-organizing criticality. *Scientific American*, January:26–33.
- Bak, P., Chen, K., Scheinkman, J., and Woodford, M. (1993). Aggregate fluctuations from independent sectoral shocks: Self-organized criticality in a model of production and inventory dynamics. *Ricerca Economica*, 47(1):3–30.
- Bottazzi, G., Cefis, E., and Dosi, G. (2002). Corporate growth and industrial structure: Some evidence from the Italian manufacturing industry. *Industrial and Corporate Change*, 11:705–723.
- Bottazzi, G., Cefis, E., Dosi, G., and Secchi, A. (2007). Invariances and diversities in the patterns of industrial evolution: Some evidence from Italian manufacturing industries. *Small Business Economics*, 29(1):137–159.
- Bottazzi, G., Coad, A., Jacoby, N., and Secchi, A. (2010). Corporate growth and industrial dynamics: Evidence from French manufacturing. *Applied Economics*, Forthcoming;DOI: 10.1080/00036840802400454.
- Bottazzi, G., Dosi, G., Lippi, M., Pammolli, F., and Riccaboni, M. (2001). Innovation and corporate growth in the evolution of the drug industry. *International Journal of Industrial Organization*, 19:1161–1187.
- Bottazzi, G. and Secchi, A. (2003a). Common properties and sectoral specificities in the dynamics of US manufacturing companies. *Review of Industrial Organization*, 23:217–232.
- Bottazzi, G. and Secchi, A. (2003b). A stochastic model of firm growth. *Physica A*, 324:213–219.

- Bottazzi, G. and Secchi, A. (2006a). Explaining the distribution of firms growth rates. *Rand Journal of Economics*, 37:234–263.
- Bottazzi, G. and Secchi, A. (2006b). Gibrat's law and diversification. *Industrial and Corporate Change*, 37(2):235–256.
- Bottazzi, G., Secchi, A., and Tamagni, F. (2008). Productivity, profitability and financial performance. *Industrial and Corporate Change*, 17(4):711–751.
- Brock, W. A. (1999). Scaling in economics: a reader's guide. *Industrial and Corporate Change*, 8(3):409–446.
- Coad, A. (2006). A closer look at serial growth rate correlation. Pisa, Sant'Anna School of Advanced Studies, LEM Working Paper series 2006/29.
- Coad, A. (2007a). A closer look at serial growth rate correlation. *Review of Industrial Organization*, 31(1):69–82.
- Coad, A. (2007b). Testing the principle of 'growth of the fitter': the relationship between profits and firm growth. *Structural Change and Economic Dynamics*, 18(3):370–386.
- Coad, A. (2009). *The Growth of Firms: A Survey of Theories and Empirical Evidence*. Edward Elgar: Cheltenham, UK.
- Coad, A. (2010). Exploring the processes of firm growth: evidence from a vector autoregression. *Industrial and Corporate Change*, 19(6):1677–1703.
- Dixon, R. L. (1953). Creep. *Journal of Accountancy*, July:48–55.
- Dosi, G. (2007). Statistical regularities in the evolution of industries: A guide through some evidence and challenges for the theory. In Malerba, F. and Brusoni, S., editors, *Perspectives on Innovation*. Cambridge University Press, Cambridge, UK.
- Feldman, M. S. (2004). Resources in emerging structures and processes of change. *Organization Science*, 15(3):295–309.
- Garnsey, E. (1998). A theory of the early growth of the firm. *Industrial and Corporate Change*, 7(3):523–556.
- Geroski, P. A. and Gugler, K. (2004). Corporate growth convergence in Europe. *Oxford Economic Papers*, 56:597–620.
- Hannan, M. T. (2005). Ecologies of organizations: Diversity and identity. *Journal of Economic Perspectives*, 19(1):51–70.

- Ijiri, Y. and Simon, H. A. (1977). *Skew Distributions and the Sizes of Business Firms*. North Holland, Amsterdam.
- Jovanovic, B. (1982). Selection and the evolution of industry. *Econometrica*, 50(3):649–670.
- Little, I. M. D. (1962). Higgledy piggledy growth. *Bulletin of the Oxford University Institute of Statistics*, 24(4):387–412.
- McKelvey, B. and Andriani, P. (2005). Why Gaussian statistics are mostly wrong for strategic organization. *Strategic Organization*, 3(2):219–228.
- Moneta, A., Entner, D., Hoyer, P., and Coad, A. (2010). Causal inference by independent component analysis with applications to micro- and macroeconomic data. Jena Economic Research Papers in Economics, Friedrich-Schiller-University Jena, Max-Planck-Institute of Economics.
- Montgomery, C. A. (1994). Corporate diversification. *Journal of Economic Perspectives*, 8(3):163–178.
- Nelson, R. R. and Winter, S. G. (1982). *An Evolutionary Theory of Economic Change*. Belknap Press of Harvard University Press, Cambridge, MA.
- Penrose, E. T. (1959). *The Theory of the Growth of the Firm*. Basil Blackwell, Oxford.
- Penrose, E. T. (1960). The growth of the firm - a case study: The Hercules powder company. *Business History Review*, 34(1):1–23.
- Reichstein, T. and Jensen, M. (2005). Firm size and firm growth rate distributions - the case of Denmark. *Industrial and Corporate Change*, 14(6):1145–1166.
- Simon, H. A. (1957). The compensation of executives. *Sociometry*, 20(1):32–35.
- Simon, H. A. (1962). The architecture of complexity. *Proceedings of the American Philosophical Society*, 106(6):467–482.
- Simon, H. A. (1968). On judging the plausibility of theories. In *Logic, methodology and philosophy of science; proceedings of the international congress for Logic, Methodology and Philosophy of Science*, pages 439–459. North-Holland, Amsterdam.
- Stanley, M. H. R., Amaral, L. A. N., Buldyrev, S. V., Havlin, S., Leschhorn, H., Maass, P., Salinger, M. A., and Stanley, H. E. (1996). Scaling behavior in the growth of companies. *Nature*, 379:804–806.

- Starbuck, W. H. (1971). Organizational growth and development. In Starbuck, W. H., editor, *Organizational Growth and Development*, chapter 1, pages 11–141. Penguin books, Harmondsworth UK.
- Sutton, J. (1998). *Technology and Market Structure: Theory and History*. MIT Press, Cambridge, MA.
- Sutton, J. (2007). Market share dynamics and the ‘persistence of leadership’ debate. *American Economic Review*, 97(1):222–241.
- Weick, K. E. and Quinn, R. E. (1999). Organizational change and development. *Annual Review of Psychology*, 50:361–386.
- Wernerfelt, B. (1984). A resource-based view of the firm. *Strategic Management Journal*, 5(2):171–180.
- Williamson, O. (1967). Hierarchical control and optimum firm size. *Journal of Political Economy*, 75(2):123–138.
- Winter, S. G. (1995). Four Rs of profitability: Rents, resources, routines, and replication. In Montgomery, C., editor, *Resource-Based and Evolutionary Theories of the Firm: Towards a Synthesis*. Kluwer Academic, Boston, MA.