

Coad, Alex

Working Paper

A closer look at serial growth rate correlation

LEM Working Paper Series, No. 2006/29

Provided in Cooperation with:

Laboratory of Economics and Management (LEM), Sant'Anna School of Advanced Studies

Suggested Citation: Coad, Alex (2006) : A closer look at serial growth rate correlation, LEM Working Paper Series, No. 2006/29, Scuola Superiore Sant'Anna, Laboratory of Economics and Management (LEM), Pisa

This Version is available at:

<https://hdl.handle.net/10419/89396>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



Laboratory of Economics and Management
Sant'Anna School of Advanced Studies

Piazza Martiri della Libertà, 33 - 56127 PISA (Italy)
Tel. +39-050-883-343 Fax +39-050-883-344
Email: lem@sssup.it Web Page: <http://www.lem.sssup.it/>

LEM

Working Paper Series

A Closer Look at Serial Growth Rate Correlation

Alex COAD

MATISSE, Université Paris I Panthéon-Sorbonne, Paris (France) and
LEM, Scuola Superiore Sant'Anna, Pisa (Italy)

2006/29

November 2006

A Closer Look at Serial Growth Rate Correlation *

Alex Coad *†‡

★ MATISSE, Univ. Paris 1 Panthéon-Sorbonne, Paris, France

‡ LEM, Sant'Anna School of Advanced Studies, Pisa, Italy

October 30, 2006

Abstract

Serial correlation in annual growth rates carries a lot of information on growth processes – it allows us to directly observe firm performance as well as to test theories. Using a 7-year balanced panel of 10 000 French manufacturing firms, we observe that small firms typically are subject to negative correlation of annual growth rates, whereas larger firms display positive correlation. Furthermore, we find that those small firms that experience extreme positive or negative growth in any one year are unlikely to repeat this performance in the following year.

JEL codes: L11, L25

Keywords: Serial correlation, firm growth, quantile regression, French manufacturing, fast-growth firms

*Thanks go to Carlo Bianchi, Giulio Bottazzi, Domenico Delli Gatti, Giovanni Dosi, Bernard Paulré, Rekha Rao and participants at the LEM conference on 'Reappraising Economics as an Empirically Disciplined Science' in Volterra in September 2006 for helpful comments. I am also indebted to Larry White (the editor) and two anonymous referees for many interesting remarks and suggestions that led to considerable modification of an earlier draft (which had previously circulated as a Paris 1 working paper (i.e. Coad (2006b))). The usual disclaimer applies.

† *Corresponding Author:* Alex Coad, MATISSE, Maison des Sciences Economiques, 106-112 Bd. de l'Hôpital, 75647 Paris Cedex 13, France. *E-mail:* a.coad@sssup.it.

1 Introduction

A lot of information on the processes of firm growth can be obtained by studying serial correlation in growth rates. At first glance, it allows us to directly observe the evolution of industries by better understanding patterns of year-on-year growth at the firm-level. Such research may have policy implications if, for example, it is desirable to prevent large firms from experiencing cumulative growth, or if one should want to investigate the ability of small firms to generate durable employment, i.e. jobs that have not disappeared by the following year.

Another more subtle motivation for studying serial correlation is that it allows us to judge between theories by comparing the hypothetical predictions with the empirically-observed regularities. First of all, if it were observed to be significant, the existence of serial correlation would lead us to reject Gibrats law of proportionate effect and the associated stochastic models of industry evolution. This strand of the literature treats firm growth as a purely stochastic phenomenon in which a firm's size at any time is simply the product of previous growth shocks. Following Sutton (1997), we define the size of a firm at time t by x_t , and represent growth by the random variable ε_t (i.e. the 'proportionate effect') to obtain:

$$x_t - x_{t-1} = \varepsilon_t \cdot x_{t-1}$$

whence:

$$x_t = (1 + \varepsilon)x_{t-1} = x_0(1 + \varepsilon_1)(1 + \varepsilon_2) \dots (1 + \varepsilon_t) \quad (1)$$

According to equation (1), a firm's size can be seen as the simple multiplication of independent growth shocks. This simple model has become a popular benchmark for modelling industrial evolution because, among other properties, it is able to generate the observed log-normal firm-size distribution, and also the proposition that expected growth is independent of size does find empirical support (roughly speaking). However, such a model would be inappropriate if the assumption of serial independence of growth rates does not find reasonable empirical support.

Second, the notion of a firm- or industry-specific 'optimal size' and the related 'adjustment cost' hypothesis of firm growth can be rejected by looking at the characteristics of serial growth correlation. The traditional, static representation of the firm considered it as having an 'optimal size' determined in a trade-off between production technology and decreasing returns to bureaucratization. This conceptualization of firms having an 'optimal size' was then extended to the case of growing firms. According to this approach, firms have a target size that they tend towards, but the existence of non-linear adjustment costs prohibits them from instantly attaining their ideal size. Instead, they grow gradually by equating at the margin the gains from having a larger size and the costs of growing. If this theory is to be believed, we should expect to find a positive autocorrelation in growth rates as firms approach their 'optimal size'. However, in reality we do not always observe positive autocorrelation in annual growth rates which leads us to doubt the validity of this theory.

Third, looking at autocorrelation statistics will allow us to judge between the different models that attempt to explain the heavy-tailed distribution of annual firm growth rates. The explanation offered by Bottazzi and Secchi (2006) hinges on the notion of increasing returns in the growth process, which would lead us to expect positive autocorrelation in annual growth rates. The explanation offered by Coad (2006a), however, considers that firms grow by the addition of lumpy resources. It follows from the discrete and interdependent nature of these resources that the required additions in any one year are occasionally rather large. In this case, we would expect a negative autocorrelation of annual growth rates.

Another motivation for this study is to observe what happens to those firms that grow extremely fast. Indeed, a robust ‘stylised fact’ that has emerged only recently is that annual firm growth rates distributions are remarkably fat-tailed and can be approximated by the Laplace distribution (Stanley *et al.* 1996, Bottazzi and Secchi 2003, Bottazzi *et al.* 2005, Bottazzi *et al.* 2006). A considerable proportion of employment creation takes place within just a handful of fast-growing firms. Conventional regression techniques that focus on what happens to the ‘average firm’, and that dismiss extreme events as ‘outliers’, may thus be inappropriate. In this study we therefore include semi-parametric regression techniques (i.e. quantile regression) to tackle this issue.

This paper provides several novel results. In particular, we observe that autocorrelation dynamics vary with firm size, such that whilst large firms experience positive feedback in year-to-year growth rates, the growth of smaller firms is marked by an erratic, ‘start-and-stop’ dynamics. Indeed, small and large firms appear to operate on different ‘frequencies’. For those small firms that experience extreme growth in one year, significant negative correlation indicates that they are quite unlikely to repeat this performance in the following year. Larger firms undergoing extreme growth events, however, do not experience such strong negative autocorrelation.

Section 2 reviews the previous literature relating to this subject, and section 3 presents the database. In section 4, we begin with some summary statistics and results using conventional regressions, and then apply quantile regression techniques in Section 5. Section 6 concludes with a discussion of our findings.

2 Literature review

The relevant empirical questions in this section are the sign, the magnitude, and also the time-scale of serial correlation in the growth rates of firms.

Early empirical studies into the growth of firms measured serial correlation when growth was measured over a period of 4 to 6 years. Positive autocorrelation of 33% was observed by Ijiri and Simon (1967) for large US firms, and a similar magnitude of 30% was reported by Singh and Whittington (1975) for UK firms. However, much weaker autocorrelation was later reported in comparable studies by Kumar (1985) and Dunne and Hughes (1994).

More recently, availability of better datasets has encouraged the consideration of annual autocorrelation patterns. Indeed, persistence should be more visible when measured over shorter time horizons. However, the results are quite mixed. Positive serial correlation has often been observed, in studies such as those of Chesher (1979) and Geroski *et al.* (1997) for UK quoted firms, Wagner (1992) for German manufacturing firms, Weiss (1998) for Austrian farms, Bottazzi *et al.* (2001) for the worldwide pharmaceutical industry, and Bottazzi and Secchi (2003) for US manufacturing. On the other hand, negative serial correlation has also been reported – some examples are Boeri and Cramer (1992) for German firms, Goddard *et al.* (2002) for quoted Japanese firms, Bottazzi *et al.* (2003) for Italian manufacturing, and Bottazzi *et al.* (2005) for French manufacturing. Still other studies have failed to find any significant autocorrelation in growth rates (see Almus and Nerlinger (2000) for German start-ups, Bottazzi *et al.* (2002) for selected Italian manufacturing sectors, Geroski and Mazzucato (2002) for the US automobile industry, and Lotti *et al.* (2003) for Italian manufacturing firms). To put it mildly, there does not appear to be an emerging consensus.

Another subject of interest (also yielding conflicting results) is the number of relevant lags to consider. Chesher (1979) and Bottazzi and Secchi (2003) found that only one lag was

significant, whilst Geroski *et al.* (1997) find significant autocorrelation at the 3rd lag (though not for the second). Bottazzi *et al.* (2001) find positive autocorrelation for every year up to and including the seventh lag, although only the first lag is statistically significant.

It is perhaps remarkable that the results of the studies reviewed above have so little in common. It is also remarkable that previous research has been so little concerned with this question. Indeed, instead of addressing serial correlation in any detail, often it is ‘controlled away’ as a dirty residual, a blemish on the ‘natural’ growth rate structure. The baby is thus thrown out with the bathwater. In our view, the lack of agreement would suggest that, if there are any regularities in the serial correlation of firm growth, they are more complex than the standard specification would be able to detect (i.e. that there is no ‘one-size-fits-all’ serial correlation coefficient that applies for all firms). We therefore consider how serial correlation changes with two aspects of firms – their size, and their growth rate – and our results, though preliminary, are nonetheless encouraging. In a nutshell, our results suggest that the discrepancies between autocorrelation coefficients in previous studies can be explained by the different firm-size compositions of these databases.

3 Database

This research draws upon the EAE databank collected by SESSI and provided by the French Statistical Office (INSEE).¹ This database contains longitudinal data on a virtually exhaustive panel of French firms with 20 employees or more over the period 1989-2002. We restrict our analysis to the manufacturing sectors. Since data reporting norms changed over the period, we maintain statistical consistency by only using the period 1996-2002 and we consider only continuing firms over this period. Firms that entered midway through 1996 or exited midway through 2002 have been removed. Since we want to focus on internal, ‘organic’ growth rates, we exclude firms that have undergone any kind of modification of structure, such as merger or acquisition. Because of limited information on restructuring activities and in contrast to some previous studies (e.g. Bottazzi *et al.*, 2001), we do not attempt to construct ‘super-firms’ by treating firms that merge at some stage during the period under study as if they had been merged from the start of the period. Firms are classified according to their sector of principal activity.² To start with we had observations for around 22000 firms per year for each year of the period.³ In the final balanced panel constructed for the period 1996-2002, we arrive, somewhat serendipitously, at exactly 10 000 firms for each year.

Our analysis is based on precisely the same database used by Bottazzi *et al.*’s (2005) study addressing the size distribution of firms, Gibrat’s law, and statistical properties of growth rates. Readers interested in such topics are referred to this paper.

4 Analysis

4.1 Summary statistics

We begin by looking at some summary statistics of firms in our database (see Table 1). First, in keeping with the elementary ‘stylized facts’ of industry structure, we observe that the firm-size

¹The EAE databank has been made available to the author under the mandatory condition of censorship of any individual information.

²The French NAF classification matches with the international NACE and ISIC classifications.

³22319, 22231, 22305, 22085, 21966, 22053, and 21855 firms respectively

Table 1: Summary statistics of the firm size distribution

year	obs.	mean	std. dev.	skewness	kurtosis	median	1%	99%
SALES (FF '000)								
1996	10000	93622	324276	14.80	316.24	29660	5640	1149076
1997	10000	98792	364255	19.44	618.02	30815	5665	1179687
1998	10000	104734	383413	19.38	611.57	33117	5960	1227714
1999	10000	107321	381536	17.15	456.54	34080	6042	1318392
2000	10000	117369	424978	17.27	473.09	36617	6044	1429880
2001	10000	121774	445042	17.33	463.25	37845	6009	1548911
2002	10000	120637	456510	18.56	515.27	37091	5638	1502079
EMPLOYMENT								
1996	10000	97.07	225.30	14.78	398.26	44	19	885
1997	10000	97.40	223.16	14.54	386.69	45	20	868
1998	10000	98.41	222.91	14.47	385.01	45	19	889
1999	10000	99.20	222.55	14.32	376.92	45	20	894
2000	10000	101.41	224.41	13.54	328.31	46	20	909
2001	10000	103.47	230.70	13.29	307.61	47	19	925
2002	10000	102.55	233.19	13.97	339.49	46	19	922
SALES GROWTH								
1997	10000	0.0359	0.2337	0.1410	25.74	0.0317	-0.6858	0.7689
1998	10000	0.0665	0.2163	0.0311	22.33	0.0592	-0.6030	0.7516
1999	10000	0.0257	0.2155	-0.0998	23.40	0.0278	-0.6206	0.6695
2000	10000	0.0647	0.2160	0.8329	27.06	0.0569	-0.5604	0.7088
2001	10000	0.0308	0.2114	-0.5049	24.65	0.0351	-0.6677	0.6045
2002	10000	-0.0252	0.2206	-0.3137	22.09	-0.0091	-0.7216	0.5864
EMPLOYMENT GROWTH								
1997	9990	0.0094	0.1359	0.1746	21.79	0.0000	-0.4199	0.4394
1998	9989	0.0143	0.1368	0.1673	18.54	0.0000	-0.4387	0.4199
1999	9989	0.0083	0.1464	4.3251	181.57	0.0000	-0.4162	0.3895
2000	9999	0.0225	0.1451	-0.8492	59.09	0.0118	-0.3947	0.4595
2001	10000	0.0116	0.1380	-0.1701	31.84	0.0000	-0.4480	0.4055
2002	10000	-0.0147	0.1386	-0.7697	32.14	0.0000	-0.4788	0.3460




Figure 1: Distribution of sales growth rates (source: Bottazzi *et al.*, 2005)




Figure 2: Distribution of employment growth rates (source: author's elaboration)

distribution is right-skewed (compare the mean and the median, look also at the skewness and kurtosis statistics). Second, whilst the distribution appears to be roughly stationary, a closer inspection reveals subtle shifts in the sample characteristics over time, with firms on average getting bigger but at a decreasing rate. Indeed, one caveat of working with balanced panels is that the characteristics of the firms may change slightly as we move from the beginning to the end of the period under consideration.

Our two measures of size and growth are sales and number of employees, which are highly correlated with each other.⁴ Figures 1 and 2 present the distributions of sales and employment growth rates, where these growth rates are cleaned of size dependence, serial correlation and heteroskedasticity effects according to the procedure described in Bottazzi *et al.* (2005). The main point of interest here is that the distribution is fat-tailed and resembles the Laplace (i.e. it appears to be approximately ‘tent-shaped’ on logarithmic axes). This testifies that relatively large growth events in any year occur not altogether infrequently. It also indicates that regression estimators based on the assumption of normally distributed errors (such as OLS) may be unreliable.

4.2 Regression analysis

To begin with, we apply regression analysis to obtain point estimates for autocorrelation coefficients, although the main results of this paper come from the quantile regressions.

In keeping with previous studies, we define our dependent variable *GROWTH* as the log-difference of size:

$$GROWTH_{i,t} = \log(SIZE_{i,t}) - \log(SIZE_{i,t-1}) \quad (2)$$

for firm i at time t , where $SIZE_{i,t}$ is measured either in terms of sales or employees. We then estimate the following regression equation:

$$GROWTH_{i,t} = \alpha_0 + \alpha_1 \log(SIZE_{i,t-1}) + \sum_{k=1}^K \beta_k GROWTH_{i,t-k} + \epsilon_{i,t}. \quad (3)$$

Given that the Gibrat Law literature has identified a dependence of growth rates upon firm size, we introduce lagged size as a control variable.

⁴The correlation between sales and number of employees is 0.8404 (with $N=70\ 000$), and the correlation between sales growth and employment growth is 0.3903 (with $N = 59\ 967$; taking logs of employment we lose firms who at some point in time had 0 employees). Both are very highly significant.

Table 2: OLS estimation of equation (3), taking 3 lags. Coefficients significant at the 5% level appear in bold.

t	α_1	β_1	β_2	β_3
SALES				
2000	0.0026	-0.2136	-0.0995	-0.0231
(SE)	(0.0017)	(0.0239)	(0.0195)	(0.0172)
2001	-0.0055	-0.2119	-0.0533	0.0029
(SE)	(0.0017)	(0.0237)	(0.0192)	(0.0149)
2002	0.0016	-0.2523	-0.1294	-0.0357
(SE)	(0.0018)	(0.0294)	(0.0238)	(0.0168)
EMPL				
2000	-0.0105	-0.1110	0.0361	0.0452
(SE)	(0.0015)	(0.0286)	(0.0163)	(0.0172)
2001	-0.0017	-0.1185	0.0174	0.0430
(SE)	(0.0015)	(0.0373)	(0.0139)	(0.0148)
2002	-0.0055	-0.1042	-0.0084	0.0448
(SE)	(0.0015)	(0.0253)	(0.0285)	(0.0180)

Table 3: MAD estimation of equation (3), taking 3 lags. Coefficients significant at the 5% level appear in bold.

t	α_1	β_1	β_2	β_3
SALES				
2000	0.0066	-0.0501	0.0018	0.0207
(SE)	(0.0012)	(0.0066)	(0.0068)	(0.0062)
2001	-0.0028	-0.0530	0.0180	0.0359
(SE)	(0.0012)	(0.0064)	(0.0064)	(0.0063)
2002	0.0025	-0.0568	-0.0336	0.0082
(SE)	(0.0014)	(0.0076)	(0.0076)	(0.0074)
EMPL				
2000	-0.0015	0.0123	0.0588	0.0476
(SE)	(0.0014)	(0.0076)	(0.0088)	(0.0088)
2001	0.0039	0.0045	0.0109	0.0163
(SE)	(0.0004)	(0.0025)	(0.0025)	(0.0026)
2002	-0.0004	0.0003	0.0007	0.0008
(SE)	(0.0000)	(0.0001)	(0.0001)	(0.0001)

To begin with, we estimate equation (3) by OLS but, since the residuals are known to be approximately Laplace-distributed, OLS is likely to perform relatively poorly. Similarly, many other estimators, including the Binder-Hsiao-Pesaran (2005) short-panel VAR estimator, require normality of residuals and are thus inappropriate in this particular case. Instead, we follow Bottazzi *et al.* (2005) and prefer the results obtained by Minimum Absolute Deviation (MAD) estimation of equation (3). The MAD estimator is to be preferred on theoretical grounds because it provides reliable results for Laplace-distributed residuals. Regression results are reported in Tables 2 and 3. When growth is measured in terms of sales, we observe a small negative autocorrelation for the first lag, in the order of -5%. The second lag is smaller, sometimes significant, but variable across the three years; and the third lag is small and positive. Regarding employment growth, we observe a small yet positive and statistically significant correlation coefficient for the average firm, for each of the first three lags.

However, it has previously been noted that one calendar year is an arbitrary period over which to measure growth (see the discussion in Geroski, 2000). We will now look at growth rate autocorrelation over periods of two and three years, by MAD estimation of equation (3). The results are presented in Tables 4 and 5. When we measure growth over periods of two or three years, we obtain quite different results. Regarding autocorrelation of sales growth, we obtain a positive and significant coefficient when growth is measured over a three-year interval, which contrasts with the results presented in Table 3 for annual data. In addition, the coefficients for employment growth autocorrelation are much larger when growth is measured over two or three years. In showing these results, we are not trying to confuse the reader by showing that the autocorrelation coefficients vary wildly for different specifications. Rather, we are trying to demonstrate that an autocorrelation coefficient is only ever meaningful when it refers to a specific time period.

These results highlight some important features that should be kept in mind when investigating serial correlation. First, both the magnitude and even the sign of the observed

Table 4: MAD estimation of equation (3), with sales growth measured over different periods. Coefficients significant at the 5% level appear in bold.

t	α_1	β_1	β_2
<i>00-02</i>	0.0023 (0.0013)	<i>98-00</i> 0.0043 (0.0055)	<i>96-98</i> 0.0062 (0.0055)
<i>99-01</i>	-0.0029 (0.0012)	<i>97-99</i> 0.0306 (0.0047)	
<i>98-00</i>	0.0062 (0.0013)	<i>96-98</i> 0.0205 (0.0055)	
<i>99-02</i>	0.0024 (0.0013)	<i>96-99</i> 0.0126 (0.0045)	

Table 5: MAD estimation of equation (3), with employment growth measured over different periods. Coefficients significant at the 5% level appear in bold.

t	α_1	β_1	β_2
<i>00-02</i>	-0.0006 (0.0000)	<i>98-00</i> 0.0010 (0.0001)	<i>96-98</i> 0.0005 (0.0001)
<i>99-01</i>	0.0038 (0.0005)	<i>97-99</i> 0.0135 (0.0021)	
<i>98-00</i>	-0.0016 (0.0014)	<i>96-98</i> 0.0522 (0.0065)	
<i>99-02</i>	-0.0005 (0.0000)	<i>96-99</i> 0.0006 (0.0001)	

autocorrelation coefficients are sensitive to the accounting period used. We should be reluctant to speak of ‘mean reversion’ in the growth process generally, for example, if we observe negative autocorrelation in annual growth rates, because these findings may not be robust to changes in time periods. Second, the conventional accounting period of one year is arbitrary and does not correspond to any meaningful duration of economic activity. Given these important qualifications, our following analysis is nonetheless able to provide useful insights into the growth process because it explores systematic variation in serial correlation patterns, conditional on firm size and conditional on growth rates.

4.3 Does autocorrelation vary with firm size?

As firms grow, they undergo many fundamental changes (Greiner, 1998). Whilst smaller firms are characteristically flexible, larger firms are more routinized, more inert and less able to adapt. In large firms, everything takes place on a larger scale, there is less reason to fear a ‘sudden death’, and the time-scale of strategic horizons extend much further than for a smaller counterpart. Larger firms may well have longer-term investment projects that unfold over a period of several years, whereas smaller firms can adjust much more rapidly. It is therefore meaningful to suppose that differences in the behavior of large firms and smaller firms will also be manifest in their respective growth processes. It has previously been conjectured that large and small firms operate on a different ‘frequency’ or time-scale, and respond to different stimuli (Hannan and Freeman, 1984).⁵ However, to my knowledge, no empirical study has

⁵To be precise, Hannan and Freeman write about: “the proposition that time-scales of selection processes stretch with size. . . One way to visualize such a relationship is to consider environmental variations as composed of a spectrum of frequencies of varying lengths - hourly, daily, weekly, annually, etc. Small organizations are more sensitive to high-frequency variations than large organizations. For example, short-term variations in the availability of credit may be catastrophic to small businesses but only a minor nuisance to giant firms. To the extent that large organizations can buffer themselves against the effects of high-frequency variations, their viability depends mainly on lower-frequency variations.” Hannan and Freeman, 1984:161




Figure 3: Autocorrelation of annual sales growth




Figure 4: Autocorrelation of annual employment growth

explicitly considered this relationship. The results in Dunne and Hughes (1994: Table VI) and in Wagner (1992: Table II) would appear to lean in this direction, but the authors fail to comment upon this possibility. The aim of this section is thus to compare growth rate autocorrelation among firms of different sizes.

We sort firms into 20 equipopulated bins according to their Sales in 1996,⁶ and calculate their growth rate autocorrelation by MAD estimation of equation (3). The evidence presented in Figures 3 and 4 would seem to support the hypothesis that annual growth rate autocorrelation varies with size, being negative, on average, for small firms and positive for larger ones. Further evidence in support of this hypothesis will also be presented in what follows.

We should be careful how we interpret these results. It may not be meaningful to say that large firms have positive feedback and smaller firms have negative feedback in their growth dynamics because, as discussed previously, it is possible that the magnitudes and signs of the autocorrelation coefficients would change if we were to measure growth over a different time period. However, one thing that we can infer from these results is that large firms and small firms operate on different time scales.

5 Quantile regression analysis

In this section we begin by explaining why we believe quantile regression techniques to be a useful tool to this study. First we describe the intuition of quantile regression analysis, and then we present the quantile regression model in a few introductory equations. We then present the results.

5.1 An introduction to quantile regression

Standard least squares regression techniques provide summary point estimates that calculate the average effect of the independent variables on the ‘average firm’. However, this focus on the average may hide important features of the underlying dynamics. As Mosteller and Tukey explain in an oft-cited passage: “What the regression curve does is give a grand summary

⁶The issue of ascribing growing firms to different size classes is not as easy as one could imagine. A drawback of sorting firms in this way is that their size in the initial period could be a poor proxy for their longer-term characteristics (this is commonly known as the problem of the ‘regression fallacy’ – see Friedman (1992) for a discussion). In the Appendix we develop an alternative methodology for sorting firms according to size (by taking their mean size over the 7-year period), and we obtain qualitatively similar results.

for the averages of the distributions corresponding to the set of x 's. We could go further and compute several regression curves corresponding to the various percentage points of the distributions and thus get a more complete picture of the set. Ordinarily this is not done, and so regression often gives a rather incomplete picture. Just as the mean gives an incomplete picture of a single distribution, so the regression curve gives a correspondingly incomplete picture for a set of distributions" (Mosteller and Tukey, 1977:266). Quantile regression techniques can therefore help us obtain a more complete picture of the underlying dynamics of firm growth processes.

In our case, estimation of linear models by quantile regression may be preferable to the usual regression methods for a number of reasons. First of all, we know that the standard least-squares assumption of normally distributed errors does not hold for our database because growth rates follow a heavy-tailed distribution. Whilst the optimal properties of standard regression estimators are not robust to modest departures from normality, quantile regression results are characteristically robust to outliers and heavy-tailed distributions. In fact, the quantile regression solution $\hat{\beta}_\theta$ is invariant to outliers that tend to $\pm \infty$ (Buchinsky, 1994). Another advantage is that, while conventional regressions focus on the mean, quantile regressions are able to describe the entire conditional distribution of the dependent variable. In the context of this studies, high growth firms are of interest in their own right, we don't want to dismiss them as outliers, but on the contrary we believe it would be worthwhile to study them in detail. This can be done by calculating coefficient estimates at various quantiles of the conditional distribution. Finally, a quantile regression approach avoids the restrictive assumption that the error terms are identically distributed at all points of the conditional distribution. Relaxing this assumption allows us to acknowledge firm heterogeneity and consider the possibility that estimated slope parameters vary at different quantiles of the conditional growth rate distribution.

The quantile regression model, first introduced by Koenker and Bassett (1978), can be written as:

$$y_{it} = x'_{it}\beta_\theta + u_{\theta it} \quad \text{with} \quad \text{Quant}_\theta(y_{it}|x_{it}) = x'_{it}\beta_\theta \quad (4)$$

where y_{it} is the growth rate, x is a vector of regressors, β is the vector of parameters to be estimated, and u is a vector of residuals. $Q_\theta(y_{it}|x_{it})$ denotes the θ^{th} conditional quantile of y_{it} given x_{it} . The θ^{th} regression quantile, $0 < \theta < 1$, solves the following problem:

$$\min_{\beta} \frac{1}{n} \left\{ \sum_{i,t:y_{it} \geq x'_{it}\beta} \theta |y_{it} - x'_{it}\beta| + \sum_{i,t:y_{it} < x'_{it}\beta} (1 - \theta) |y_{it} - x'_{it}\beta| \right\} = \min_{\beta} \frac{1}{n} \sum_{i=1}^n \rho_\theta u_{\theta it} \quad (5)$$

where $\rho_\theta(\cdot)$, which is known as the 'check function', is defined as:

$$\rho_\theta(u_{\theta it}) = \begin{cases} \theta u_{\theta it} & \text{if } u_{\theta it} \geq 0 \\ (\theta - 1)u_{\theta it} & \text{if } u_{\theta it} < 0 \end{cases} \quad (6)$$

Equation (5) is then solved by linear programming methods. As one increases θ continuously from 0 to 1, one traces the entire conditional distribution of y , conditional on x (Buchinsky, 1998). More on quantile regression techniques can be found in the surveys by Buchinsky (1998) and Koenker and Hallock (2001); for applications see Coad and Rao (2006) and also the special issue of *Empirical Economics* (Vol. 26 (3), 2001).




Figure 5: Regression quantiles for sales (left) and employment (right) autocorrelation coefficients, with 95% confidence intervals.

Table 6: Quantile regression estimation of equation (7) for the 10%, 25%, 50%, 75% and 90% quantiles, allowing for only one lag in serial correlation. Coefficients significant at the 5% level appear in bold.

	10%	25%	50%	75%	90%
Sales gr.					
β_1	-0.1354	-0.0725	-0.0449	-0.0596	-0.1267
(<i>t</i> -stat)	-12.35	-17.79	-15.63	-14.15	-10.49
Pseudo- R^2	0.0294	0.0259	0.0189	0.0237	0.0294
Empl. gr.					
β_1	-0.0924	-0.0206	0.0000	0.0034	-0.0547
(<i>t</i> -stat)	-8.45	-5.11	0.00	0.7100	-4.0000
Pseudo- R^2	0.0091	0.0041	0.0007	0.0121	0.0237

5.2 Quantile regression results

The regression equation that we estimate is:

$$GROWTH_{i,t} = \alpha_0 + \alpha_1 \log(SIZE_{i,t-1}) + \beta_1 GROWTH_{i,t-1} + y_t + \epsilon_{i,t}. \quad (7)$$

where y_t are yearly dummies. The quantile regression results are presented in Table 6, and a summary representation is provided in Figure 5. The coefficients can be interpreted as the partial derivative of the conditional quantile of the dependent variable with respect to particular regressors (Yasar *et al.*, 2006). Evaluated at the median, we observe that there is only slight negative autocorrelation in sales growth and totally insignificant autocorrelation in employment growth. (In fact, the median quantile regression corresponds to the MAD regression estimate.) The story does not end here, however, because the serial correlation coefficient estimates vary considerably across the conditional growth rate distribution. For firms experiencing dramatic losses in sales or employment at time t , the sharply negative coefficient implies that in the previous period $t - 1$ these firms were probably experiencing above-average growth. Similarly, for those fastest-growing firms at time t , the negative coefficient estimate indicates that these firms probably performed relatively poorly in the previous period $t - 1$. It would appear then that, although in any one year there are some firms that undergo signif-




Figure 6: regression quantiles for sales growth autocorrelation coefficients across the 10 size groups (group ‘1’ = smallest group)

icant growth events, these firms are unlikely to repeat this performance.⁷ According to this evidence, it would appear that the better analogy would probably be that of the ‘hare and tortoise’ rather than notions of cumulative ‘snowball effect’ dynamics or even serial independence of growth rates.

5.3 Robustness across size groups

Are the previous results robust across size? Or is the relationship displayed in Figure 5 just the result of aggregating firms of different sizes – where smaller firms are the extreme growers and it is these same firms that experience the negative autocorrelation? It does not appear, for this dataset, that growth rate variance decreases dramatically with size (compare Bottazzi *et al.* (2005) with Bottazzi and Secchi (2003)). Nevertheless, in this section we will investigate possible heterogeneity across size classes by applying quantile regression analysis to different size groups. We sort and split the firms into 10 size groups according to their initial size (sales in 1996). We then explore the regression quantiles for each of these 10 groups. Results are presented in Table 7 and Figures 6 (sales growth) and 7 (employment growth).

The results are reasonably consistent whether we consider sales growth or employment growth. For the larger firms, the results support the previous finding that, on average, these firms experience a slightly positive autocorrelation in annual growth rates. Even as we move

⁷One potential problem that we thought deserved investigation was the possibility of data entry errors. Despite the INSEE’s reputation for providing high-quality data, we were concerned that there could be cases of omitted numbers in which a firm’s sales (or employees) were observed to shrink by tenfold in one year and grow by tenfold in the next. Where we found such cases, we checked for consistency with other corresponding variables (e.g. value added, employees etc). As it happens, the database appeared consistent under scrutiny and we are pleased to acknowledge that our suspicions were a waste of time.

Table 7: Quantile regression estimation of equation (7) for the 10%, 25%, 50%, 75% and 90% quantiles for 10 size groups (**1** = smallest), allowing for only one lag in serial correlation. The size groups are sorted according to sales in 1996. Coefficients significant at the 5% level appear in bold.

	10%	25%	50%	75%	90%
Sales gr.					
1: β_1	-0.1694	-0.1420	-0.0847	-0.0810	-0.1427
(<i>t</i> -stat)	-4.74	-10.98	-8.92	-6.23	-4.17
2: β_1	-0.1507	-0.1295	-0.0942	-0.1147	-0.1601
(<i>t</i> -stat)	-4.12	-8.58	-8.21	-7.47	-4.26
3: β_1	-0.1337	-0.0989	-0.0630	-0.0718	-0.1447
(<i>t</i> -stat)	-4.05	-6.96	-4.62	-4.19	-3.00
4: β_1	-0.0375	-0.0336	-0.0456	-0.0701	-0.1381
(<i>t</i> -stat)	-1.06	-2.09	-3.71	-3.93	-2.95
5: β_1	-0.1356	-0.0832	-0.0876	-0.0971	-0.2052
(<i>t</i> -stat)	-3.79	-6.24	-9.89	-6.85	-4.62
6: β_1	-0.0984	-0.0808	-0.0706	-0.1248	-0.1960
(<i>t</i> -stat)	-2.72	-6.16	-7.15	-7.42	-5.15
7: β_1	-0.1366	-0.0584	-0.0574	-0.0997	-0.2218
(<i>t</i> -stat)	-3.38	-3.86	-5.52	-7.26	-6.22
8: β_1	-0.0797	-0.0292	-0.0287	-0.0795	-0.1689
(<i>t</i> -stat)	-2.21	-2.35	-3.84	-5.59	-3.07
9: β_1	-0.0150	-0.0009	-0.0018	-0.0476	-0.1271
(<i>t</i> -stat)	-0.41	-0.07	-0.19	-3.79	-2.95
10: β_1	0.0662	0.0786	0.0641	0.0526	-0.0055
(<i>t</i> -stat)	1.71	5.36	6.07	3.35	-0.13
Empl. gr.					
1: β_1	-0.1837	-0.1051	0.0000	-0.0390	-0.0918
(<i>t</i> -stat)	-5.61	-6.80	0.00	-3.64	-1.90
2: β_1	-0.2002	-0.1016	-0.0196	-0.0354	-0.1544
(<i>t</i> -stat)	-5.72	-5.45	-5.88	-4.10	-3.24
3: β_1	-0.1039	-0.0296	0.0000	-0.0247	-0.0624
(<i>t</i> -stat)	-4.17	-2.66	0.00	-1.19	-1.19
4: β_1	-0.0709	-0.0128	0.0000	0.0006	-0.0184
(<i>t</i> -stat)	-1.92	-1.01	0.00	0.04	-0.46
5: β_1	-0.1355	-0.0734	-0.0248	-0.0660	-0.1367
(<i>t</i> -stat)	-3.67	-5.27	-6.75	-4.08	-3.13
6: β_1	-0.1067	-0.0331	0.0000	-0.0002	-0.0740
(<i>t</i> -stat)	-3.21	-2.46	0.00	-0.01	-2.44
7: β_1	0.0141	0.0381	0.0000	0.0131	-0.0673
(<i>t</i> -stat)	0.49	2.77	0.00	0.79	-1.47
8: β_1	-0.0395	0.0078	0.0063	0.0149	-0.0679
(<i>t</i> -stat)	-1.09	0.59	2.07	1.10	-1.71
9: β_1	0.0528	0.0550	0.0668	0.0848	0.0582
(<i>t</i> -stat)	1.29	3.95	9.77	5.37	1.51
10: β_1	0.1094	0.1361	0.1890	0.1975	0.1780
(<i>t</i> -stat)	2.64	11.92	33.19	18.30	7.70




Figure 7: regression quantiles for employment growth autocorrelation coefficients across the 10 size groups (group ‘1’ = smallest group)

to the extremes of the conditional distribution, the autocorrelation coefficient does not change too dramatically. This may be because diversification has a stabilizing effect on growth rates. Smaller firms, however, typically experience negative correlation which is moderate near the median but quite pronounced towards the extreme quantiles. This is in line with previous observations on “the prevalence of interruptions to growth” for small firms (Garnsey and Heffernan (2005: 675)). For these firms, prolonged periods of high growth are quite unusual.

5.4 Robustness to temporal disaggregation

Up until now, we have pooled together the observations from all of the years of the panel database. However, it might not be a valid methodology to pool together observations from

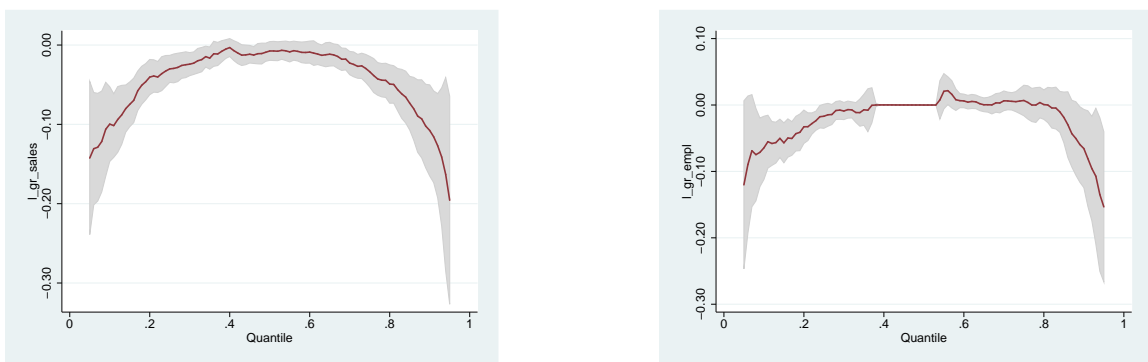


Figure 8: Regression quantiles for sales (left) and employment (right) autocorrelation coefficients for $t=1999$, with 95% confidence intervals.




Figure 9: Regression quantiles for sales (left) and employment (right) autocorrelation coefficients for $t=2002$, with 95% confidence intervals.

different years if the autocorrelation structure varies with time (e.g. over the business cycle). We now check how our results stand up to temporal disaggregation by estimating the quantile regressions for individual years. In a sense, we are moving from a panel dataset to five yearly cross-sections. However, our findings appear to be robust to temporal disaggregation. The quantile regression plots for the years 1999 and 2002 (for sales and employment growth) are shown in Figures 8 and 9.

5.5 Robustness to sectoral disaggregation

Rigorous empirical methodology requires us to also ensure that these results are not due to aggregation over heterogeneous industries. In this section, we report quantile regression results for 20 2-digit industries. Summary information on these sectors is provided in Table 8 and the results are presented in Table 9.

Generally speaking, the properties that were visible at the aggregate level are also visible for 2-digit industries. Firms near the median experience only moderate autocorrelation (either positive or negative), whereas firms at the extreme quantiles of the conditional growth rate distribution experience much stronger forces of negative autocorrelation. Although sectoral disaggregation does not qualitatively change our key findings, there are a few sectors in which the results are quite ‘messy’. This may be because we aggregate over firms from the same industry but of different sizes. One interpretation would be that, in determining autocorrelation in growth processes, the most relevant dimension is size and conditional growth rate, rather than sector of activity.

Table 8: Description of the 2-digit manufacturing sectors studied (source: Bottazzi *et al.* 2005)

ISIC class	Description	No. obs. (1996)	Mean size (€'000 in 2002)
17	Manuf. of textiles	730	9703
18	Manuf. of wearing apparel, dressing and dyeing of fur	498	9623
19	Tanning and dressing of leather, manuf. of luggage, handbags, ...	205	14629
20	Manuf. of wood and products of wood and cork, except furniture; ...	314	9083
21	Manuf. of paper and paper products	364	22428
22	Publishing, printing and reproduction of recorded media	820	13745
23	Manuf. of coke, refined petroleum products and nuclear fuel	19	73547
24	Manuf. of chemicals and chemical products	496	52378
25	Manuf. of rubber and plastics products	685	17964
26	Manuf. of other non-metallic mineral products	426	21624
27	Manuf. of basic metals	265	34411
28	Manuf. of fabricated metal products, except machinery and equipment	2276	8041
29	Manuf. of machinery and equipment n.e.c.	987	19343
30	Manuf. of office, accounting and computing machinery	23	39850
31	Manuf. of electrical machinery and apparatus n.e.c.	357	26740
32	Manuf. of radio, television and communication equipment and apparatus	218	25159
33	Manuf. of medical, precision and optical instruments, watches and clocks	354	12452
34	Manuf. of motor vehicles, trailers and semi-trailers	280	49195
35	Manuf. of other transport equipment	137	68192
36	Manuf. of furniture; manufacturing n.e.c.	546	14411

Table 9: Quantile regression estimation of equation (7) for the 10%, 25%, 50%, 75% and 90% quantiles for 20 2-digit sectors (17-36), allowing for only one lag in serial correlation. Coefficients significant at the 5% level appear in bold.

	10%	25%	50%	75%	90%	10%	25%	50%	75%	90%
	Sales gr.					Empl. gr.				
17: β_1	0.0588	0.0221	0.0003	-0.0094	-0.0303	-0.0217	0.0128	0.0003	0.0294	-0.0191
(<i>t</i> -stat)	2.24	1.69	0.02	-0.57	-0.66	-0.50	0.92	1.33	1.85	-0.45
18: β_1	-0.1438	-0.0392	-0.0284	-0.05	-0.1170	-0.0907	-0.0312	0.0000	-0.0699	-0.1097
(<i>t</i> -stat)	-3.30	-2.56	-2.21	-2.38	-2.10	-2.13	-1.86	0.00	-4.52	-1.82
19: β_1	-0.1227	-0.0296	-0.0032	-0.0967	-0.1718	-0.2684	-0.1256	-0.0379	-0.0727	-0.1550
(<i>t</i> -stat)	-2.56	-1.42	-0.22	-4.95	-2.08	-3.08	-5.19	-2.98	-2.00	-1.67
20: β_1	-0.0458	-0.0188	0.0144	0.0383	-0.1454	-0.1578	-0.0531	-0.0019	-0.0319	-0.0896
(<i>t</i> -stat)	-1.09	-0.98	0.68	1.63	-1.86	-2.92	-1.80	-0.79	-0.88	-1.12
21: β_1	-0.2197	-0.0961	-0.0563	-0.1180	-0.2096	-0.1027	-0.0304	-0.0001	0.0018	-0.0531
(<i>t</i> -stat)	-5.86	-4.39	-4.01	-5.80	-3.46	-1.80	-1.74	-1.29	0.12	-0.82
22: β_1	-0.1691	-0.0574	-0.0194	-0.0374	-0.0903	-0.1506	-0.0753	0.0000	-0.0507	-0.1484
(<i>t</i> -stat)	-4.12	-4.95	-2.85	-3.57	-2.48	-2.84	-5.32	0.00	-3.20	-3.72
23: β_1	0.0961	0.1317	0.1883	0.1523	0.1098	-0.0899	0.0683	0.1788	0.1926	0.0678
(<i>t</i> -stat)	0.20	1.80	9.09	5.16	0.12	-0.22	0.80	2.52	1.98	0.27
24: β_1	-0.0667	-0.0155	-0.0100	-0.0272	-0.0558	0.0634	0.0550	0.0587	0.0818	0.0068
(<i>t</i> -stat)	-1.46	-1.01	-0.84	-1.48	-1.14	1.34	3.40	7.66	4.79	0.10
25: β_1	-0.1346	-0.0582	-0.0246	-0.0440	-0.1126	-0.1009	-0.0262	0.0000	0.0091	-0.0361
(<i>t</i> -stat)	-2.50	-3.93	-2.34	-2.87	-3.05	-2.06	-1.50	0.00	0.46	-0.82
26: β_1	-0.0685	-0.3990	-0.0239	-0.0686	-0.1429	-0.1252	-0.0216	-0.0003	-0.0689	-0.1483
(<i>t</i> -stat)	-1.30	-2.31	-2.20	-3.31	-2.60	-1.77	-1.13	-0.56	-3.34	-2.45
27: β_1	-0.1552	-0.1052	-0.0189	0.0094	0.0096	-0.0744	0.0194	0.0115	0.0455	0.0176
(<i>t</i> -stat)	-2.42	-5.02	-0.79	0.29	0.19	-3.05	2.02	3.24	3.58	0.60
28: β_1	-0.1801	-0.1219	-0.1003	-0.1153	-0.1748	-0.1262	-0.0485	-0.0005	-0.0353	-0.1284
(<i>t</i> -stat)	-7.71	-13.85	-14.57	-12.46	-7.07	-6.97	-5.73	-6.37	-3.14	-4.27
29: β_1	-0.2043	-0.1438	-0.1062	-0.1354	-0.1874	-0.0909	-0.0084	0.0006	-0.0028	-0.0758
(<i>t</i> -stat)	-4.91	-8.98	-12.20	-9.37	-4.34	-2.91	-0.77	2.21	-0.18	-1.87
30: β_1	0.0074	-0.0959	0.0194	0.0593	-0.1050	-0.3232	-0.0403	0.1071	0.1416	0.2377
(<i>t</i> -stat)	0.30	-0.70	0.14	0.33	-0.31	-0.86	-0.20	1.49	1.23	0.22
31: β_1	-0.0551	-0.0749	-0.0437	-0.0602	-0.1216	-0.0490	0.0227	0.0172	0.0735	0.1041
(<i>t</i> -stat)	-1.00	-2.69	-2.91	-2.39	-2.11	-0.87	1.09	1.26	3.07	1.57
32: β_1	-0.1094	-0.0904	-0.0610	-0.0228	-0.0538	-0.0106	-0.0138	0.0384	0.0692	0.1172
(<i>t</i> -stat)	-1.31	-2.42	-1.88	-0.93	-0.91	-0.14	-0.62	1.66	1.87	1.16
33: β_1	-0.1573	-0.1179	-0.0763	-0.0779	-0.1457	-0.1160	0.0053	0.0000	0.0895	0.0967
(<i>t</i> -stat)	-2.60	-5.47	-5.33	-2.63	-1.87	-2.03	0.22	0.00	3.45	1.37
34: β_1	-0.0696	-0.0472	-0.0193	-0.0386	-0.0438	-0.1472	-0.0298	0.0245	0.0592	0.0715
(<i>t</i> -stat)	-1.09	-1.98	-0.92	-1.83	-0.51	-2.51	-1.12	1.09	1.75	0.83
35: β_1	-0.2325	-0.1190	-0.1097	-0.1439	-0.2830	-0.0719	0.0036	0.0443	0.0228	-0.1280
(<i>t</i> -stat)	-2.52	-3.53	-4.51	-3.74	-1.90	-0.59	0.15	2.09	0.61	-0.89
36: β_1	0.0052	-0.0304	-0.0172	-0.0588	-0.1527	0.0113	0.0087	0.0039	-0.0482	-0.1278
(<i>t</i> -stat)	0.17	-2.51	-1.53	-3.25	-2.60	0.32	0.61	1.40	-1.87	-2.48

6 Summary and Conclusions

We began by exploring serial correlation in annual growth rates using standard regression techniques, and detected a statistically significant influence of past growth even for the third lag. When sales growth was considered, the coefficient on the first lag was typically around 5%, whereas for employment growth it was generally positive although smaller in magnitude. We also found evidence that growth rate autocorrelation varied with firm size, consistent with the hypothesis that small firms operate on a different time scale (i.e. a shorter ‘frequency’) than larger ones. In the case of annual growth rates, we obtained negative coefficients for groups of smaller firms and positive ones for larger firms. This systematic variation of autocorrelation coefficients across firm size helps explain why previous studies using different databases (reviewed in Section 2) found such inconclusive results.

An important recent discovery in the industrial organization literature is that firm growth rates are fat-tailed and follow closely the Laplace density. This means that we can expect that, in any given year, a significant proportion of turbulence in market share or employment is due to just a handful of fast-growing firms. Although small in number, these firms are of special interest to economists. What are the characteristics of these firms? Standard regression techniques, that focus on the ‘average firm’, are of limited use in this case. Instead, we apply quantile regression techniques that explicitly recognise heterogeneity between firms, and present results from various quantiles of the conditional growth rates distribution. Although we find a small negative annual autocorrelation at the aggregate level, there exist more powerful autoregressive forces for those firms that matter the most - the extreme-growth firms. These firms may grow a lot in one period, but it is unlikely that the spurt will last long. We also observed an interaction between the characteristics of the extreme-growth firms and size. Whilst smaller fast-growth firms are much more prone to dramatic negative autocorrelation, larger firms seem to have much smoother growth dynamics.

Our results can be related to several well-known theories in the industrial organization literature. The models of passive and active learning in the evolution of industries (as proposed by Jovanovic (1982) and Ericson and Pakes (1995), respectively) appear to be supported by our findings, because the growth paths of small firms are quite erratic and noisy whereas those of larger firms are relatively smoothed. Our results also have implications for Gibrat’s Law. On the basis of our findings, this ‘law’ would be rejected because, in many cases, growth rates in consecutive years are not independent.

It is, of course, far too early to speak of the possibility of ‘stylized facts’, but since our findings are reasonably robust and also theoretically meaningful, we anticipate that future research will corroborate our results. We also consider that more should be done in way of investigation of the characteristics of extreme high-growth firms (see also Coad and Rao, 2006). These firms are just a small proportion in the number of firms but account for a great proportion of employment growth or market share turbulence. Conventional regression techniques are of limited use in this respect. Quantile regression techniques are far more appropriate, although perhaps future work on high-growth firms should also consider an approach by case studies.

Appendix

In this Appendix we provide further evidence of the robustness of our findings. In particular, we check the robustness of our findings by using an alternative technique for sorting firms according to size.

Up until now, we have sorted the firms according to their size in the first time period of our panel dataset, i.e. 1996. However, this could give misleading results. If we sort firms according to size in any one year, there is a danger that the size of some firms in that particular year will not be representative of their size in the other years. For example, suppose a firm experiences a temporary negative shock to its size in one year, but next year it returns to its ‘usual’ size. Such a firm will thus be erroneously classified as a ‘fast-growing small firm’ if it is put into a size class during the year that it is small. Conversely, the year before it would perhaps have been classified as a ‘fast-shrinking medium-sized firm’. If we classify firms according to their size in any one particular year (such as the initial year), we may have a tendency to exaggerate the growth of small firms and underestimate the growth of larger firms. There may also be implications for the relationship between autocorrelation and firm size.

This statistical problem of sorting growing entities according to size is commonly known as the ‘regression fallacy’. In his 1992 discussion of this issue, Milton Friedman suggests that firms should be allocated to size classes according to their average size for the whole period of analysis, rather than attributing them to a size class according to their size in any one year. Therefore, in this Appendix, we classify firms according to their mean size over the whole period (more precisely, the mean number of employees 1996-2002).

We begin by examining whether autocorrelation coefficients do indeed vary with firm size according to this new size-classification scheme. The evidence is shown in Figures 10 and 11. Again, we see that autocorrelation does vary with firm size.

We also repeat the quantile regression analysis by sorting firms according to their mean size rather than initial size. The plots are shown in Figures 12 and 13, and results are reported in Table 10. We conclude that our findings are qualitatively similar to those obtained by classifying firms according to their initial size.




Figure 10: Autocorrelation of annual sales growth




Figure 11: Autocorrelation of annual employment growth




Figure 12: regression quantiles for sales growth autocorrelation coefficients across the 10 size groups (group '1' = smallest group)




Figure 13: regression quantiles for employment growth autocorrelation coefficients across the 10 size groups (group '1' = smallest group)

Table 10: Quantile regression estimation of equation (7) for the 10%, 25%, 50%, 75% and 90% quantiles for 10 size groups (1 = smallest), allowing for only one lag in serial correlation. The size groups are sorted according to their mean size (employees) 1996-2002. Coefficients significant at the 5% level appear in bold.

	10%	25%	50%	75%	90%
Sales gr.					
1: β_1	-0.1984	-0.1380	-0.1084	-0.0980	-0.1775
(<i>t</i> -stat)	-5.65	-8.57	-10.94	-7.08	-6.36
2: β_1	-0.2706	-0.1566	-0.0974	-0.1324	-0.2147
(<i>t</i> -stat)	-6.63	-12.98	-10.28	-9.54	-5.22
3: β_1	-0.2229	-0.1426	-0.1022	-0.1236	-0.2071
(<i>t</i> -stat)	-5.36	-12.11	-13.68	-8.61	-4.99
4: β_1	-0.1742	-0.1144	-0.0708	-0.1049	-0.1928
(<i>t</i> -stat)	-6.25	-10.33	-8.63	-6.62	-4.71
5: β_1	-0.0650	-0.0239	-0.0482	-0.0785	-0.1626
(<i>t</i> -stat)	-2.01	-2.04	-4.69	-5.88	-4.23
6: β_1	-0.0966	-0.0497	-0.0261	-0.0345	-0.1153
(<i>t</i> -stat)	-3.36	-3.59	-2.45	-2.44	-3.38
7: β_1	-0.1440	-0.0701	-0.0447	-0.0469	-0.1197
(<i>t</i> -stat)	-4.24	-4.98	-5.70	-3.47	-3.28
8: β_1	-0.0285	0.0057	0.0083	-0.0132	-0.0503
(<i>t</i> -stat)	-0.88	0.39	0.83	-0.90	-1.13
9: β_1	-0.1103	-0.0291	0.0255	0.0194	-0.0400
(<i>t</i> -stat)	-2.92	-2.01	2.53	1.43	-1.04
10: β_1	0.0255	0.0920	0.0862	0.0709	0.0779
(<i>t</i> -stat)	0.64	6.93	7.96	3.81	2.11
Empl. gr.					
1: β_1	-0.1414	-0.0436	0.0000	-0.0645	-0.1182
(<i>t</i> -stat)	-3.36	-10.28	0.00	-10.03	-3.47
2: β_1	-0.2482	-0.1412	0.0000	-0.0973	-0.2019
(<i>t</i> -stat)	-6.96	-10.68	0.00	-7.72	-5.14
3: β_1	-0.2384	-0.1528	-0.0820	-0.1584	-0.3008
(<i>t</i> -stat)	-5.60	-9.69	-13.70	-9.85	-6.23
4: β_1	-0.1135	-0.0673	0.0000	-0.0855	-0.1966
(<i>t</i> -stat)	-4.19	-5.28	0.00	-4.43	-3.97
5: β_1	-0.0536	-0.0172	0.0000	-0.0603	-0.1660
(<i>t</i> -stat)	-1.54	-1.90	0.00	-6.92	-3.98
6: β_1	-0.0663	-0.0051	0.0000	0.0032	-0.1066
(<i>t</i> -stat)	-2.52	-1.04	0.00	0.26	-2.81
7: β_1	-0.0924	0.0157	0.0137	0.0514	0.0154
(<i>t</i> -stat)	-2.69	2.59	5.97	3.75	0.40
8: β_1	0.0410	0.0645	0.0755	0.1165	0.0944
(<i>t</i> -stat)	1.60	4.86	9.79	8.41	2.63
9: β_1	0.0450	0.0449	0.0932	0.1700	0.1811
(<i>t</i> -stat)	1.37	3.48	13.01	11.6	4.95
10: β_1	0.0888	0.1194	0.1770	0.1996	0.1943
(<i>t</i> -stat)	1.82	9.60	32.09	18.93	6.92

References

- Almus, M. and E. Nerlinger (1992), 'Testing 'Gibrat's Law' for Young Firms - Empirical Results for West Germany', *Small Business Economics*, **15**, 1-12.
- Binder, M., C. Hsiao and C. H. Pesaran (2005), 'Estimation and inference in short panel vector autoregressions with unit roots and cointegration', *Econometric Theory*, **21**, 795-837.
- Boeri, T. and U. Cramer (1992), 'Employment growth, incumbents and entrants: evidence from Germany', *International Journal of Industrial Organization*, **10**, 545-565.
- Bottazzi, G., G. Dosi, M. Lippi, F. Pammolli and M. Riccaboni (2001), 'Innovation and Corporate Growth in the Evolution of the Drug Industry', *International Journal of Industrial Organization*, **19**, 1161-1187.
- Bottazzi, G., E. Cefis and G. Dosi (2002), 'Corporate Growth and Industrial Structure: Some Evidence from the Italian Manufacturing Industry', *Industrial and Corporate Change*, **11**, 705-723.
- Bottazzi, G. and A. Secchi, (2003) 'Common Properties and Sectoral Specificities in the Dynamics of U. S. Manufacturing Companies', *Review of Industrial Organization*, **23**, 217-232.
- Bottazzi, G., A. Coad, N. Jacoby and A. Secchi (2005), 'Corporate Growth and Industrial Dynamics: Evidence from French Manufacturing', Pisa, Sant'Anna School of Advanced Studies, LEM Working Paper Series 2005/21.
- Bottazzi, G., E. Cefis, G. Dosi and A. Secchi (2006), 'Invariances and Diversities in the Evolution of Manufacturing Industries', *Small Business Economics*, forthcoming.
- Bottazzi, G. and A. Secchi (2006), 'Explaining the Distribution of Firms Growth Rates', *The RAND Journal of Economics*, forthcoming.
- Buchinsky, M. (1994) 'Changes in the U.S. Wage Structure 1963-1987: Application of Quantile Regression', *Econometrica* **62**, 405-458.
- Buchinsky, M. (1998) 'Recent Advances in Quantile Regression Models: A Practical Guide for Empirical Research' *Journal of Human Resources* **33** (1), 88-126.
- Chesher, A. (1979), 'Testing the Law of Proportionate Effect', *Journal of Industrial Economics*, **27** (4), 403-411.
- Coad, A., (2006a), 'Towards an Explanation of the Exponential Distribution of Firm Growth Rates', Cahiers de la Maison des Sciences Economiques No. 06025 (Série Rouge), Université Paris 1 Panthéon-Sorbonne, France.
- Coad, A., (2006b), 'Understanding the Processes of Firm Growth – a Closer Look at Serial Growth Rate Correlation', Cahiers de la Maison des Sciences Economiques No. 06051 (Série Rouge), Université Paris 1 Panthéon-Sorbonne, France.
- Coad, A. and R. Rao (2006), 'Innovation and Firm Growth in High-Tech Sectors: A Quantile Regression Approach', Pisa, Sant'Anna School of Advanced Studies, LEM Working Paper Series 2006/18.

- Dunne, P. and A. Hughes (1994), 'Age, Size, Growth and Survival: UK companies in the 1980s', *Journal of Industrial Economics*, **42** (2), 115-140.
- Ericson, R. and A. Pakes (1995), 'Markov-perfect industry dynamics: a framework for empirical work', *Review of Economic Studies*, **62**, 53-82.
- Friedman, M. (1992), 'Do Old Fallacies Ever Die?', *Journal of Economic Literature*, **30** (4), 2129-2132.
- Garnsey, E. and P. Heffernan (2005), 'Growth Setbacks in New Firms', *Futures*, **37** (7), 675-697.
- Geroski, P. A. (2000), 'The growth of firms in theory and practice', in N. Foss and V. Mahnke (eds), *New Directions in Economic Strategy Research*, Oxford: Oxford University Press.
- Geroski, P., S. Machin and C. Walters (1997), 'Corporate Growth and Profitability', *Journal of Industrial Economics*, **45** (2), 171-189.
- Geroski P. and M. Mazzucato (2002), 'Learning and the Sources of Corporate Growth', *Industrial and Corporate Change*, **11** (4), 623-644.
- Goddard, J., J. Wilson and P. Blandon (2002), 'Panel tests of Gibrat's Law for Japanese Manufacturing', *International Journal of Industrial Organization*, **20**, 415-433.
- Greiner, L. (1998), 'Evolution and revolution as organizations grow', *Harvard Business Review*, May-June, 55-67.
- Hannan, M. and J. Freeman (1984), 'Structural Inertia and Organizational Change', *American Sociological Review*, **49** (2), 149-164.
- Ijiri, Y. and H. A. Simon (1967), 'A Model of Business Firm Growth', *Econometrica*, **35** (2), 348-355.
- Jovanovic, B. (1982), 'Selection and the evolution of industry', *Econometrica*, **50**, 649-670.
- Koenker, R. and G. Bassett (1978) 'Regression Quantiles' *Econometrica* **46**, 33-50.
- Koenker, R. and K. F. Hallock (2001), 'Quantile Regression', *Journal of Economic Perspectives*, **15** (4), 143-156.
- Kumar, M. S. (1985), 'Growth, Acquisition Activity and Firm Size: Evidence from the United Kingdom', *Journal of Industrial Economics*, **33** (3), 327-338.
- Little, I. M. D. (1962) "Higgledy Piggledy Growth" *Bulletin of the Oxford University Institute of Statistics* **24** (4), November, 387-412.
- Lotti, F., E. Santarelli and M. Vivarelli (2003), 'Does Gibrat's Law hold among young, small firms?', *Journal of Evolutionary Economics*, **13**, 213-235.
- Mosteller, F. and J. Tukey (1977), 'Data Analysis and Regression', Addison-Wesley, Reading, MA.
- Singh A. and G. Whittington (1975), 'The Size and Growth of Firms', *Review of Economic Studies*, **42** (1), 15-26.

- Stanley, M. H. R., L. A. N. Amaral, S. V. Buldyrev, S. Havlin, H. Leschhorn, P. Maass, M. A. Salinger and H. E. Stanley (1996), 'Scaling behavior in the growth of companies', *Nature*, **379**, 804-806.
- Sutton, J. (1997), 'Gibrat's Legacy', *Journal of Economic Literature*, **35**, 40-59.
- Wagner, J. (1992), 'Firm Size, Firm Growth, and Persistence of Chance: Testing Gibrat's Law with Establishment Data from Lower Saxony, 1978-1989', *Small Business Economics*, **2** (4), 125-131.
- Weiss, C. (1998), 'Size, Growth, and Survival in the Upper Austrian Farm Sector', *Small Business Economics*, **10**, 305-312.
- Yasar, M., C. H. Nelson and R. M. Rejesus (2006), 'Productivity and Exporting Status of Manufacturing Firms: Evidence from Quantile Regressions', *Weltwirtschaftliches Archiv*, forthcoming.