On the Autocorrelation of Growth Rates

Evidence for Micro, Small and Large Firms from the Austrian Service Industries, 1975–2004

Alex Coad · Werner Hölzl

Received: 6 November 2008 / Revised: 28 January 2009 /

Accepted: 6 February 2009 / Published online: 18 March 2009

© The Author(s) 2009. This article is published with open access at Springerlink.com

Abstract This paper studies the serial autocorrelation of annual growth rates in employment for selected Austrian service industries over a 30-year period using quantile regression techniques. The autocorrelation of growth rates provides important information on firms growth processes. We find that the growth patterns of micro firms are strikingly different from the growth patterns of small, medium-sized and larger firms. First, we do find a positive dependency of growth on size for growing micro firms, while this relationship is negative for the other size groups. Second, growing micro firms are subject to negative autocorrelation of annual growth rates making sustained growth a very rare occurrence, while larger growing firms usually display a positive autocorrelation suggesting that high growth episodes of larger firms stretch over a longer time horizon. This indicates that the growth of micro firms in particular is characterized by a rather lumpy growth profile. Furthermore, we find that the autocorrelation patterns are asymmetric with regard to decline and growth.

Keywords serial correlation · firm growth · service sector · quantile regression · Austria

JEL Classification L11 · L25

A. Coad

Evolutionary Economics Group, Max Planck Institute of Economics,
Kahlaische Strasse 10, 07745 Jena, Germany
e-mail: coad@econ.mpg.de

W. Hölzl (✉)

Austrian Institute of Economic Research (WIFO),
P.O. Box 91, 1103 Vienna, Austria
e-mail: werner.hoelzl@wifo.ac.at

1 Introduction

This paper studies the serial correlation of growth rates for firms in Austrian service industries. The serial correlation of growth rates provides important information on the processes of firm growth. It allows to study the persistence of firm growth processes. The persistence in employment growth—are new jobs likely to be discontinued 1 year later or is the growth process continuing—is of special interest, as in recent years there is a growing interest in the job-creating potential of small firms. The study of serial correlation for different size classes helps to get a clearer picture of the differences between small and large firm growth. In addition the study of the serial correlation of growth rates allows to assess theories of firm growth by comparing theoretical predictions with observed regularities. For instance, if serial correlation were observed to be significant this would lead from a strictly methodological viewpoint to a rejection of Gibrat's Law of proportionate effect and the associated stochastic models of industry evolution. The assumption that firm growth is a purely stochastic and the product of independent growth shocks is not appropriate if growth rates are not serial independent (Chesher 1979). Also the nature of adjustment costs can be studied by looking at the serial correlation of growth rates. Convex adjustment costs prevent firms from immediately attaining their chosen size and lead to a gradual partial adjustment over time. The path to the desired new equilibrium level is a smooth, partial adjustment-like path. We should therefore observe a positive autocorrelation of growth rates. Non-convex adjustment costs in contrast are more easily reconciled with the empirical evidence that employment change is non-smooth but lumpy (Hamermesh and Pfann 1996; Caballero et al. 1997). If non-convex adjustment costs play an important role we should expect to find zero or even negative autocorrelation especially for high growth rates. Firms expand at one point in time and wait for the next expansion in order to economize on fixed adjustment costs. The distribution of employment should be characterized by a high proportion of extreme events. The non-convex adjustment cost hypothesis predicts that sharp adjustment or employment should be followed by relatively long periods of inaction, that is no change in the level of employment.

We use a dataset that allows us to identify the growth processes of micro firms. Usually national statistical offices gather data on firms above a certain size threshold (e.g. above 20 employees). The neglect of micro firms is a strong limitation especially if one considers service industries where the average firm size is smaller than in manufacturing. Micro firms face particular difficulties in their very early growth phase, presumably because of high fixed (non-convex) adjustment costs related to their small size.¹ The growth patterns of fast-growing small firms are thus particularly erratic (Garnsey et al. 2006; Santarelli and Vivarelli 2007). In addition, research suggests that most micro firms do not display any changes in employment for a long time (Hölzl and Huber 2008). In our dataset, 65% of firms in the 1–9 employees category stay at the same size from 1 year to the next. Although researchers often focus on high-growth firms (often referred to as 'gazelles'), we should acknowledge that most micro firms don't grow. We also need to know more about the growth history of these gazelles, instead of focusing on their growth in one single year.

¹For example, if a firm with one employee hires an additional employee this leads to a growth rate of 100%.

Our research concentrates on the service sector. Research into firm growth has focused on the manufacturing sector at the expense of the service sector (Delmar 1997). As the major part of economic activity in modern economies takes place in the service sector this constitutes a shortcoming of many studies. The situation is improving, however—in recent years some studies have focused on firm growth in the services sector (Audretsch et al. 1999b, 2004; Lotti 2007; Teruel-Carrizosa 2009). Most studies find that there is a negative relationship between firm size and expected growth rate, indicating that most service sectors have something like a minimum efficient scale of operation (Variyam and Kraybill 1992; Johnson et al. 1999). In addition, the service sector is interesting for the present research as the adjustment costs are most likely related to costly changes in the level of employment. Adjustment costs due to capital investment are likely less important than in manufacturing industries. Thus the finding of a negative autocorrelation would indicate primarily adjustment costs to labor. Third there is not much available evidence for Austria. Weiss (1998) looks at growth dynamics of farms in Upper Austria. He finds that the most dynamic farms are the smallest ones. In addition he observes positive autocorrelation. However, this finding is based on a very specific sector of the economy.

Using a 30 year unbalanced panel of Austrian firms, this paper finds that autocorrelation dynamics vary with firm size, such that larger firms display a positive autocorrelation and smoother growth processes than smaller firms. For micro firms we record negative autocorrelation that paints a picture of erratic ‘start-and-stop’ growth dynamics. Indeed, small and large firms seem to operate on different ‘frequencies’. High-growth micro firms are very unlikely to repeat their growth performance the following year, while larger firms experience a positive feedback, that leads to sustained growth.

The paper is organized as follows: Section 2 provides a review of the literature. Section 3 presents the database and descriptive statistics on the growth process in the Austrian service sector. Sections 4 and 5 presents the empirical analysis. In Section 5 we start with the presentation of the size and the growth rate distribution and provide evidence that the growth rate distribution exhibits heavy tails. This has implication for the choice of estimation technique. Section 5 presents results for the autocorrelation of growth rates using quantile regression techniques using both at the aggregated data and disaggregating according to time periods and firm size. Section 6 concludes the paper.

2 Background

The first empirical studies that considered the serial correlation of growth rates considered growth over a period of 4 to 6 years. Ijiri and Simon (1967) and Singh and Whittington (1975) observed positive autocorrelation around 30% for large US and UK firms, respectively. However, Kumar (1985) and Dunne and Hughes (1994) report much weaker autocorrelation in comparable studies. The substantial heterogeneity of findings is even more evident in studies that use annual data (Coad 2007). The persistence of firm growth should more visible when measured over shorter time horizons. However, the results are quite mixed. Positive autocorrelation has been found in studies as different 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 firms or Bottazzi and Secchi (2003) for US manufacturing firms. On the other hand, negative serial correlation has been reported by Boeri and Cramer (1992) for German firms, Goddard et al. (2002) for quoted Japanese firms, and Bottazzi et al. (2007, 2009) for Italian and French manufacturing firms. Other studies failed to find any significant autocorrelation in growth rates, e.g. Bottazzi et al. (2002) for selected Italian manufacturing sectors or Geroski and Mazzucato (2002) for the US automobile industry. Overall there is no clear pattern emerging regarding the autocorrelation of firm growth rates. If one takes into account the details of the adopted samples, especially the coverage of smaller firms then there is some hint that smaller firms may be characterized by negative autocorrelation. However, Wagner (1992) and Bottazzi and Secchi (2003) observe positive autocorrelation in samples that include also smaller firms. Thus to put it mildly, there does not appear to be an emerging consensus. Even more remarkable is that previous research has been so little concerned with this question. In none of this studies any attempt was made to uncover and investigate why this could be. In most empirical studies of firm growth serial correlation is not addressed in any detail, often even controlled away as dirty influence that affects the ‘natural’ structure of firm growth rates. However, the analysis of growth autocorrelation in itself is of much interest—“serial correlation in firm growth rates ... is of considerable economic interest and deserves to be examined in its own right” (Singh and Whittington 1975). Coad (2007) is an exception, as he studies explicitly the autocorrelation structure. In his study of French manufacturing firms Coad (2007) establishes that small firms typically are subject to negative correlation of annual growth rates, whereas larger firms display positive serial correlation. In addition, serial correlation is strongly negative for (small) firms that have just experienced a large growth event in the recent past. However, there are still a number of limits to Coad’s analysis. He analyses a dataset of firms with over 20 employees in the French manufacturing sector. This leaves out a large part of industry. This is true also for the broader picture. Table 1 shows that in all EU countries micro firms make up more than 83% of firms and account for at least 19% of employment. There is still room for considerable differences across countries, for example, the employment structure in Italy is much more dominated by micro firms than the employment structure in other countries.

Table 1 Size distribution of firms and employment across selected countries, 2003

Country	Number of enterprises				Employment			
	1–9	10–49	50–249	250+	1–9	10–49	50–249	250+
Austria	86.7%	11.3%	1.7%	0.3%	25.5%	23.9%	19.0%	31.7%
Germany	83.0%	14.2%	2.3%	0.5%	19.6%	21.9%	18.7%	39.8%
France	92.2%	6.5%	1.1%	0.2%	23.3%	20.7%	16.9%	39.2%
Italy	94.5%	4.9%	0.5%	0.1%	47.2%	21.9%	12.4%	18.4%
Sweden	93.4%	5.5%	0.9%	0.2%	27.8%	19.2%	16.2%	36.8%
EU 27	91.3%	7.3%	1.2%	0.2%	28.7%	20.4%	17.5%	33.4%

Notes: Size classes are measured in number of employees. Numbers refer to sections C, D, E, F, G, H, I, J and K. For the distribution of employment sections C and E are excluded for Austria as data is missing

Source: Hölzl and Reinstaller (2009)

This study seeks to complement the existing literature on growth rate autocorrelation by focusing on the dynamics of micro firms. We consider this to be an important topic considering the distinctive character that small firm growth possesses. While for large firms selecting an expansion strategy is very much a matter of firm strategy or taste, small firms struggle to grow in order to attain a minimum efficient scale and to increase their chances of survival (Phillips and Kirchoff 1989). There is clear evidence that firm size influences the survival probabilities. Audretsch et al. (1999a), Agarwal and Audretsch (2001), Arauzo-Carod and Segarra-Blasco (2005) and Kaniovski and Peneder (2008) document that start-up size is an important determinant of post-entry performance and firm survival. This suggests that growth in terms of employment is quite appropriate as indicator of success for small firms, whereas this is less adapt for larger firms, where profitability and productivity are considered to be better measures of firm success (Davidsson et al. 2008).

3 Database description

The data we use to measure firm level employment stem from the Austrian Social Security files. These data include information on all employers and employees in the Austrian private sector for the time period from 1974 to 2004. They contain a daily calendar of the starting date of an employment relationship at the individual basis. From this data we construct a data set which reports yearly employment stocks for all private sector firms with at least one employee for the time period from 1974 to 2004. Relative to the data used in most of the literature our data have the advantage of a wide coverage as they include also all micro enterprises. We have available information on business units from most sectors (except for public services). This, however, comes at the price of limited information on the firms included in the data. We lack information on firms (such as productivity or sales) other than employment, (NACE-3-digit) industry affiliation and region of operation. The entities can be enterprises or establishments. The anonymous firm identifiers in the social security files are administrative accounts only, and it is left to discretion of firms whether it chooses to report at the enterprise or establishment level. For this reason the data have been cleaned, using a series of plausibility checks to ensure that business units are properly defined. Stiglbauer (2003) argues that the vast majority of observations are at the enterprise level, since having one account reduces administrative burdens when reporting social security contributions.

Regarding the sector coverage it needs to be noted that earlier data are more unreliable than more recent data, as the number of firms without sector classification decreased substantially over time. Therefore, especially results reported for the first 10 years (1975–1984) should be considered with some caution. We include them primarily for robustness purposes. However, as we see later in this paper there is not much difference in the results between the sub periods.

This administrative data set has been widely used in empirical research, especially for labour market research (e.g. Winter-Ebmer 2003), but also to study the job creation and destruction of entry and exit (Hölzl et al. 2007), the comparison of the duration of new jobs in new and old firms (Winter-Ebmer and Böheim 2006), to study the evolution of the firm size distribution (Huber and Pfaffermayer 2007) and the survival determinants of Austrian firms (Kaniovski and Peneder 2008; Kaniovski

Table 2 Share in GDP and employed persons for NACE sections H, J and K

	Share in GDP			Share in employed persons		
	H	J	K	H	J	K
1984	5.0%	4.8%	12.4%	4.3%	3.4%	4.6%
1994	4.4%	5.8%	14.9%	4.8%	3.7%	6.0%
2004	4.2%	5.0%	17.3%	5.2%	3.6%	9.6%

Source: National accounts and Austrian Institute of Economic Research (WIFO)

et al. 2008). Similar data has been used to study firm growth. For example, Lotti et al. (2003) use Italian administrative data to study Gibrat's law.

We restrict our sample to selected Austrian service industries, as there is not much research into firm growth for the service sector. In addition the service sector is interesting for the present research as adjustment costs are primarily related to labor adjustment costs. We use Sections H (Hotels and Restaurants), Section J (Financial intermediation) and Section K (Real estate, renting and business activities).² Table 2 reports the share in GDP and the share in dependent employment for the NACE sections in the sample for the years 1984, 1994 and 2004. Dependent employment is the appropriate indicator as our sample covers firms with at least one dependent employee. Especially section K experienced an expansion of both its share in GDP and its share in employment from 1984 to 2004. Hotels and restaurants (section H) saw a decrease in its share in GDP but also an increase of employment of almost one percentage point between 1984 and 2004.

4 Firm size and growth rate distributions

4.1 Firm size distributions

Our measure of firm size and firm growth is the number of employees. Figure 1 presents the size distribution of firms for the years 1984, 1994 and 2004. The firm size distribution has almost a straight line negative slope over most of its support, which is a feature that is common to both lognormal and Pareto distributions (Mitzenmacher 2003). This is largely in line with previous work on firm size distributions (e.g. Axtell 2001; Marsili 2005; de Wit 2005). In addition Table 3 presents the size distribution of firms for all years (1974–2004) and for the year 2000. One can see that the distribution is quite similar. Over 90% of all firms are micro firms with 1 to 9 employees. Small firms with 10 to 49 employees account for 5 to 6 percent of all firms. Medium-sized and large sized firms account for the remaining 1% of firms. Most of the firms in the sample we consider are thus very small firms.

²This restriction of the data is also due to computational issues. Even with modern computers the estimation of quantile regression is extremely time consuming if applied to large datasets.

Fig. 1 Kernel estimates of the density of firm size (employees) for our sample of Austrian firms, 1984, 1994 and 2004

Notes: Kernel densities computed using an Epanechnikov kernel for 50 equipaced bins (using gbutils 5.2). Note the log scale on the y-axis. The negative domain at the beginning of the distribution is an artefact of the smoothing algorithm and the bandwidth selection of the kernel density estimation. No individual firm has a negative value for (log) size.

4.2 Growth rate distributions

In keeping with previous studies, we define firm growth (GR) as the log-difference of size:

$$GR_{i,t} = \log(S_{i,t}) - \log(S_{i,t-1}), \quad (1)$$

for firm i at time t , where S is measured in terms of employment and we consider only annual growth rates.

Table 4 provides some preliminary information on the distribution of growth rates. It depicts the employment change of firms for four different size classes: micro firms (1–9 employees), small firms (10–49 employees), smaller medium sized firms (50–99 employees) and larger medium sized and large firms (100+ employees). The allocation of firms to size classes is based on the average size of the firm,

$$AVS_{i,t} = \frac{S_{i,t} + S_{i,t-1}}{2}, \quad (2)$$

for firm i at time t . This measure has the advantage that the allocation is not biased toward smaller size classes when $S_{i,t-1}$ would have been used, nor it is biased towards

Table 3 Size distribution of firms: size classes

Size group	All years			Year 2000		
	Freq.	Percent	Cum.	Freq.	Percent	Cum.
1–9	2,324,625	93.4	93.4	96,196	92.8	92.8
10–49	138,264	5.6	98.9	6,252	6.0	98.8
50–99	15,132	0.6	99.5	695	0.7	99.5
100+	11,853	0.5	100.0	546	0.5	100.0
Total	2,489,874	100.0		103,689	100.0	

Table 4 Growth rate distribution: size classes

	1 to 9 employees		10 to 49 employees		50 to 100 employees		100+ employees	
	Freq.	Percent	Freq.	Percent	Freq.	Percent	Freq.	Percent
< -50			351	0.28	225	1.56	432	3.74
-20 to -50			1,609	1.26	314	2.17	712	6.17
-10 to -19	2,615	0.16	1,684	1.32	688	4.76	945	8.19
-5 to -9	9,941	0.60	5,506	4.32	1,411	9.77	1,032	8.94
-3 to -4	28,019	1.68	9,608	7.54	1,206	8.35	658	5.70
-2	54,284	3.26	11,029	8.66	842	5.83	368	3.19
-1	248,197	14.92	16,605	13.03	1,034	7.16	388	3.36
0	1,092,439	65.68	24,530	19.25	1,099	7.61	400	3.47
1	161,785	9.73	18,485	14.51	1,059	7.33	463	4.01
2	41,753	2.51	13,452	10.56	1,027	7.11	408	3.54
3 to 4	18,513	1.11	12,683	9.95	1,550	10.73	722	6.26
5 to 9	5,121	0.31	8,109	6.36	1,997	13.83	1,454	12.60
10 to 19	604	0.04	2,665	2.09	1,217	8.43	1,543	13.37
20 to 49			978	0.77	608	4.21	1,340	11.61
> 50			114	0.09	167	1.16	675	5.85
Total	1,663,271	100	127,408	100	14,444	100	11,540	100

larger size classes when $S_{i,t}$ would have been used.³ We see that more than 65% of the micro firms (almost one third) do not change their employment over 1 year. If we add those firms which create or destroy one job we arrive at approximately 90% of all micro firms, while only 46% of small firms, 22% of smaller medium sized firms and 10% of large firms are inactive or change their employment by one employee. The distribution of job creation and destruction clearly shows that the probability of a larger absolute change in employment increases with firm size.

Figure 2 presents the growth rate distributions. These figures are based on the log growth rates and are not cleaned of size dependence, serial correlation and heteroskedasticity effects. The distribution of growth rates is clearly fat-tailed and resembles the Laplace distribution with its characteristic ‘tent-shape’. The Laplace distribution has been shown to have a good fit with the empirical growth rates distribution for US data (Stanley et al. 1996; Bottazzi and Secchi 2003), French data (Bottazzi et al. 2009), Italian data (Bottazzi et al. 2007) and Community Innovation Survey (CIS) data (Hölzl and Friesenbichler 2008), although it appears that, in our sample, the fat-tailed nature of the growth rate distribution is even more extreme.⁴ This is indicated by the fact that the two tails of the distribution (i.e. the sides of the ‘tent’) are not straight lines but appear to be convex to the origin.

Our observed growth rate distribution is particularly reminiscent of the theoretical model and empirical evidence in Fu et al. (2005), who observe that a Laplace distribution of growth rates is suitable for large multiproduct firms, while the distribution

³For more on the statistical problems of sorting growing entities into size classes (and the ‘regression fallacy’ in particular) see Friedman (1992).

⁴One should be aware, however, that detailed comparisons of our results to previous results are not entirely warranted because of differences in methodology. For example, Bottazzi et al. (2009) estimate the parametric form of the growth rate distribution after cleaning the growth rate series of dependence on size and lagged growth, and also of any influence of size on growth rate variance.

Fig. 2 Kernel estimates of the density of firm growth (employees) for our sample of Austrian firms, 1984, 1994 and 2004

Notes: Kernel densities computed using an Epanenchnikov kernel for 50 equispaced bins (using gbutils 5.2). Note the log scale on the y-axis.

of growth rates for small firms has fatter tails and tends towards a power law. An analysis of the growth rate distribution of micro firms has not been duly investigated in the existing literature, however.

In order to investigate the possible departures from the Laplace, we estimate the parametric form of the growth rate distribution in the context of the Subbotin family of distributions (also known as the asymmetric exponential power densities).

The Subbotin distribution is formally presented by the following equation:

$$f_S(x) = \frac{1}{2ab^{1/b} \Gamma(1/b + 1)} e^{-\frac{1}{b} \left| \frac{x-\mu}{a} \right|^b}, \quad (3)$$

where $\Gamma(x)$ corresponds to the Gamma function. The distribution has three parameters—the location parameter μ , the dispersion parameter a and the shape parameter b . As the shape parameter b decreases in value, the tails of the density become fatter. The density is leptokurtic for $b < 2$. That means that it has a higher probability than a normally distributed variable of values near the mean and a higher probability than a normally distributed variable of extreme values (fat tails). It is platokurtic for $b > 2$. That means that the distribution has a lower probability than a normally distributed variable of values near the mean and a lower probability than a normally distributed variable of extreme values (thin tails). Two notable special cases of the Subbotin distribution are the Gaussian distribution (for which $b = 2$) and the Laplace distribution (with $b = 1$).

The values of the fitted parameters are presented in Table 5. Our estimates are obtained using the program Subbotools 0.9.8.1 developed by Bottazzi (2004) and assuming that the location parameter μ is equal to 0, which is the median value of the growth rate distribution. Let us consider only the period 1995–2004. When the full sample is taken, the shape parameter b takes a value of 0.1974, which is very low compared to the Laplace case of $b = 1$ or the Gaussian case of $b = 2$. The growth rate distribution of the firms in our sample is indeed very heavy-tailed. Most micro firms do not grow in any 1 year, while a minority of firms will experience

Table 5 Estimates of the growth rate distribution parameters 1975–1984, 1985–1994 and 1995–2004 using the Subbotin distribution

		b	SE	a	SE	Obs.
75–84	Full sample	0.3355	0.0008	0.0374	0.0001	398,676
	10+ empl	0.4358	0.0024	0.0718	0.0004	30,077
	15+ empl	0.4647	0.0027	0.0753	0.0004	20,216
	20+ empl	0.4760	0.0029	0.0775	0.0005	14,624
	50+ empl	0.5030	0.0033	0.0846	0.0005	5,390
85–94	Full sample	0.2950	0.0006	0.0466	0.0001	536,938
	10+ empl	0.4244	0.0019	0.0924	0.0004	46,470
	15+ empl	0.4587	0.0023	0.1011	0.0005	29,600
	20+ empl	0.4752	0.0024	0.1066	0.0005	21,744
	50+ empl	0.5073	0.0028	0.1197	0.0006	8,139
95–04	Full sample	0.1974	0.0004	0.0345	0.0001	661,023
	10+ empl	0.3441	0.0012	0.0912	0.0004	64,490
	15+ empl	0.3917	0.0015	0.1052	0.0005	41,554
	20+ empl	0.4114	0.0016	0.1123	0.0005	30,780
	50+ empl	0.4578	0.0020	0.1302	0.0006	11,681

Notes: In all estimations we restrict μ to 0. All other settings are the defaults in Subbotools 0.9.8.1. The data are unbalanced panels of annual growth rates. We pool the observations over the 10-year periods

rapid growth, that is all the more impressive when their growth is expressed as a proportional growth rate. When we focus on firms above a certain size threshold, however, the growth rate distribution becomes closer to the Laplace. When we consider firms with 10 or more employees,⁵ the b parameter takes a value of 0.3441, and when we look at firms with 50 or more employees the b parameter rises to 0.4578. We observe a similar pattern for the time periods 1875–1984 and 1985–1994. The steepest growth rate distributions are recorded for the time period 1995–2004.

The heavy-tailed nature of the growth rate distribution has two practical implications for our analysis of growth rate autocorrelation. First, the non-Gaussian nature of the growth rate distribution provides a strong warning that least squares estimators, which assume normally-distributed residuals, will perform poorly. Instead, we use quantile regressions that are robust to outliers. Second, it is of little interest to focus on a regression coefficient that corresponds to ‘the average effect for the average firm’. The growth rate distributions suggest that the median and the average firm do not grow or have a marginal growth rates, while there is a minority of firms that experience very rapid growth. Quantile regressions will allow us to investigate the autocorrelation structure across the entire distribution of growth rates.

5 Analysis

We begin by discussing why quantile regression is a useful technique in our present case (Section 5.1), before conducting our analysis at an aggregated level (Section 5.2).

⁵ As explained before, firms are sorted into size classes using average size following Friedman (1992).

We then investigate how growth autocorrelation varies for firms of different sizes in Section 5.3. In addition we repeat the analysis for several periods and observe that our results are generally stable over time.

5.1 Introducing quantile regression

Estimation of linear models by quantile regression may here be preferable for a number of reasons. First of all, most econometric techniques aim at identifying average behaviour. The great advantage of a quantile regression is that it enables us to consider the entire distribution of firm growth. Thus, quantile regression is able to provide a more “complete” story of the relationship between variables. As the name quantile regression suggests, it is not limited to regression against averages, and hence it is not limited in its explanatory value, since it also uses information that it obtains from the underlying distribution of the dependent variable (Koenker 2005). In the present study we think that high growth firms are of interest in their own right—we don’t want them to consider as outliers. We think that it is worthwhile to study them in detail. Quantile regression allows us to analyze the growth of these firms in more details and to investigate whether these high-growth firms have durable growth or whether the jobs they create are likely to disappear in the next period.

In addition, we know that the standard least-squares assumption of normally distributed errors does not hold for our data because the growth rates follow a fat-tailed distribution (as was demonstrated earlier in Fig. 2 and Table 5). While standard regression estimators are not robust to departures from normality, the quantile regression estimator is characteristically robust to outliers on the dependent variable that tend to $\pm \infty$ (Buchinsky 1998).⁶

Finally, quantile regression relaxes the restrictive assumption that error terms are identically distributed at all points of the conditional distribution. Avoiding this assumption allows to analyze differences in the relationship between the endogenous and exogenous variables at different points of the conditional distribution of the dependent variable. Quantile regression allows us to examine how the partial correlation changes across the quantiles. This provides an understanding of the entire shape of the distribution and how it may be shaped by the explanatory variables. For more on quantile regression see Appendix and Koenker (2005).

5.2 Aggregate quantile regression analysis

The non-Gaussian nature of the growth rates distribution provides a strong warning that least squares estimators, which assume normally-distributed residuals, will perform poorly, while median regression (least absolute deviation) regression is more appropriate to study the ‘average’ firm. However, the growth rates distributions suggest in addition that the average firm does not grow, while there is a minority of firms that experience very rapid growth or rapid decline. Thus methodologies

⁶Conventional least-squares estimators minimize the squares of the residuals, while quantile regression is a generalization of the median estimator that minimizes the absolute deviations. One analogy might be that least-squares estimates correspond to the mean whereas quantile regression estimates correspond to the quantiles. As such, quantile regression estimates are more appropriate for the analysis of heavy-tailed phenomena.

Table 6 Quantile regression estimation of Eq. 4 for the 10%, 25%, 50%, 75% and 90% quantiles, 1975–2004

	10%	25%	50%	75%	90%
α_1	-0.208 (0.000)	-0.042 (0.000)	0.000 (0.000)	0.032 (0.000)	0.159 (0.001)
β_1	0.000 (0.000)	-0.072 (0.003)	0.000 (0.000)	-0.058 (0.001)	-0.238 -(0.002)
β_2	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	-0.016 (0.001)	-0.079 (0.001)
Pseudo- R^2	0.175	0.032	0.000	0.020	0.077

Notes: Number of observations: 1,395,795. Standard errors appear in brackets. Coefficients significant at the 5% level appear in bold. All observations are pooled together over the 30-year period

focusing on the average behavior are of minor interest compared to econometric methods that allow to study the whole distribution of growth rates at different points. Therefore we use quantile regression. In the context of firm growth quantile regression was used by Fotopoulos and Louri (2004), Coad and Rao (2008), and Reichstein et al. (2009).

We estimate the following regression equation:

$$GR_{i,t} = \alpha_0 + \alpha_1 \log(S_{i,t-1}) + \sum_{k=1}^K \beta_k GR_{i,t-k} + y_t + \varepsilon_{i,t}, \quad (4)$$

for firm i at time t , with y_t corresponding to yearly dummies that control for common macroeconomic phenomena. S is firm size measured in employment and GR is log growth. As the dependence of growth rates upon size emerged as stylized fact from the Gibrat's Law literature, we introduce lagged size as control variable. Given the evidence on the growth rate distribution—and the knowledge that their residuals will be approximately Laplace-distributed, OLS is likely to perform poorly in estimating Eq. 4 and so we prefer quantile regression which, at the 50% quantile, corresponds to the median regression estimator (also known as the least absolute deviation (LAD) estimator). We experimented with the number of relevant lags to consider, and observed that a model with two lags provided the best representation of the growth process. A third lag was never statistically significant. Most studies found that one lag is statistically significant (e.g. Chesher 1979; Bottazzi and Secchi 2003; Coad 2007), while (Geroski et al. 1997) find significant autocorrelation at the third lag. In our regressions of Eq. 4 we therefore restrict ourselves to the two-lag model (i.e. $K = 2$).⁷

The results are presented in Table 6. Figure 3 presents a summary representation for the coefficients on lagged growth. The coefficients can be interpreted as the partial derivative of the conditional quantile of the dependent variable with respect to the particular regressors.

Concerning size, we observe that growth is independent of size when firm growth is evaluated at the median. This result is of little interest, however, considering that the most firms located around the median of the growth rate distribution do not grow. Instead, it is of interest to consider the effect of size on growth at the

⁷This reduces the sample size from 1,772,436 observations to 1,252,750 observations, as short-lived firms with less than 2 observations of growth rates are dropped from the sample.

Fig. 3 Regression quantiles for employment growth autocorrelation coefficients

extreme quantiles of the growth rate distribution. At the lower quantiles, we observe a negative coefficient for α (of magnitude -0.208 at the 10% quantile) which testifies that, among declining firms, larger firms tend to have faster decline. At the upper quantiles, we observe a positive coefficient for α (of magnitude 0.159 at the 90% quantile) which provides a contrasting result that it is the larger firms that are more likely to experience fast growth. The results in Table 6 are thus in contrast with previous findings that smaller firms grow faster than larger firms. We will check the robustness of this finding in the next section, where we differentiate across size groups and thus check for possible composition bias due to the large number of micro firms in the sample.

Evaluated at the median, we observe that there is no autocorrelation in employment growth. Again this result is likely due to the fact that a large number of micro firms does not grow at all (cf. Table 4). The focus on the median does not tell the whole story. The serial correlation coefficient estimates are very different across the conditional growth rate distribution. For firms experiencing a dramatic decline in employment at time t we do not observe a significant autocorrelation of growth rates. In fact as Fig. 3 shows, we observe a striking asymmetry of the autocorrelation of growth rates for high positive and negative growth rates. For the fastest-growing firms at time t , the negative coefficient estimates indicates that these firms probably performed relatively poorly in the two previous periods. The coefficient for first-order autocorrelation is always larger than the coefficient for second-order autocorrelation.

The asymmetry in the quantile regression plot deserves further comment. For the fastest growing firms at the upper quantiles, the coefficient turns sharply negative at the first lag and is also negative at the second. This suggests that these fast-growing firms are likely to have had very low growth rates in the previous periods. The quantile regression coefficients at the lower quantiles, however, are much closer to zero, indicating that declining firms experience no such negative correlation.

In order to provide a first robustness check of our result we consider whether the pattern of autocorrelation changed over the past 30 years in our set of Austrian service industries. We group the growth rates into three 10-year periods (1975–1984, 1985–1994, 1995–2004) and apply quantile regression techniques to the three sub periods. Table 7 reports the results. The results are quite similar to the results over

Table 7 Quantile regression estimation of Eq. 4 for three different decades

	10%	25%	50%	75%	90%
time period 1995–2004: 524,815 observations					
α_1	-0.226 (0.000)	-0.052 (0.001)	0.000 (0.000)	0.029 (0.000)	0.065 (0.002)
β_1	0.000 (0.000)	-0.039 (0.004)	0.000 (0.000)	-0.047 (0.002)	-0.281 (0.004)
β_2	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	-0.018 (0.015)	-0.050 (0.007)
Pseudo- R^2	0.129	0.038	0.000	0.015	0.040
time period 1985–1994: 458,066 observations					
α_1	-0.208 (0.000)	-0.041 (0.001)	0.000 (0.000)	0.034 (0.000)	0.161 (0.000)
β_1	0.000 (0.000)	-0.070 (0.001)	0.000 (0.000)	-0.082 (0.004)	-0.260 (0.004)
β_2	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	-0.026 (0.002)	-0.099 (0.004)
Pseudo- R^2	0.184	0.029	0.000	0.023	0.088
time period 1977–1984: 305,776 observations					
α_1	-0.151 (0.001)	-0.031 (0.001)	0.000 (0.000)	0.035 (0.001)	0.151 (0.001)
β_1	-0.112 (0.005)	-0.053 (0.005)	0.000 (0.000)	-0.075 (0.006)	-0.212 (0.010)
β_2	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	-0.025 (0.003)	-0.061 (0.009)
Pseudo- R^2	0.216	0.022	0.000	0.031	0.158

Notes: Standard errors appear in brackets. Coefficients significant at the 5% level appear in bold. All observations are pooled together over the 10-year period

the whole time period. We do not observe striking differences in the coefficients, except that for the time period 1975–1984 we observe a statistical negative first order autocorrelation coefficient β_1 for the 10% quantile that is not present for the other two time periods.

The size coefficient α_1 shows the same surprising pattern as earlier: a negative association of size to the growth rate for declining firms and a positive association of size to the growth rate for growing firms. The first order autocorrelation of growth rates is negative for growing firms as is the second order autocorrelation. For declining firms we observe a statistically significant negative first order autocorrelation for the 25% quantile. Thus the result are surprisingly robust across decades.

5.3 Does autocorrelation vary with firm size?

The evidence presented suggests that firms that undergo high-growth events are unlikely to repeat this performance in the next time periods and growth processes are characterized by substantial non-convex adjustment costs. However, are these results robust across size classes, or is the relationship displayed in Fig. 3 just the result of aggregating firms of different sizes—a large number of micro firms and a smaller number of small, medium and large firms? Our sample is characterized by a large number of micro firms thus it is likely that the results presented so far are primarily

due to micro firms. Coad (2007) has shown that there is a marked difference between the growth experiences of small and large firms in French manufacturing industries.

Firm growth leads to fundamental changes in the organization of firms. Micro and small firms are characteristically more flexible and more labor intensive—larger firms more routinized, more capital intensive but also more inert and less able to adapt. Thus it is of special interest to compare growth rate autocorrelation among firms of different sizes. We will now investigate the possible heterogeneity across firm size classes applying quantile regression analysis separately to four size groups (1–9, 10–49, 50–99 and 100+ employees). The allocation to size classes is based on the average firm size over two consecutive years (AVS_{it}). The results for the time period 1995–2004 is in Table 8. Figure 4 reports the associated quantile plots. Table 9 reports the quantile regression results for the time periods 1975 to 1984 and 1985 to 1994.

Table 8 Quantile regression estimation of Eq. 4 for four size classes, 1995–2004

	10%	25%	50%	75%	90%
size class 1–9 employees: 463,108 observations					
α_1	-0.369 (0.000)	-0.138 (0.000)	0.000 (0.000)	0.049 (0.001)	0.241 (0.003)
β_1	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	-0.098 (0.003)	-0.333 (0.008)
β_2	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	-0.049 (0.002)	-0.130 (0.007)
Pseudo- R^2	0.286	0.074	0.000	0.004	0.090
size class 10–49 employees: 50,693 observations					
α_1	-0.115 (0.007)	-0.045 (0.002)	-0.023 (0.002)	-0.113 (0.001)	-0.184 (0.004)
β_1	-0.122 (0.012)	-0.085 (0.009)	-0.008 (0.002)	0.015 (0.005)	0.056 (0.010)
β_2	-0.047 (0.007)	-0.012 (0.004)	0.018 (0.004)	0.047 (0.006)	0.099 (0.009)
Pseudo- R^2	0.018	0.007	0.001	0.045	0.083
size class 50–99 employees: 6,092 observations					
α_1	-0.520 (0.033)	-0.202 (0.011)	-0.156 (0.009)	-0.231 (0.017)	-0.503 (0.034)
β_1	-0.002 (0.011)	0.015 (0.009)	0.046 (0.010)	0.080 (0.015)	0.083 (0.012)
β_2	-0.007 (0.007)	-0.002 (0.008)	0.021 (0.010)	0.065 (0.015)	0.049 (0.013)
Pseudo- R^2	0.086	0.041	0.038	0.078	0.179
size class 100+ employees: 4,922 observations					
α_1	-0.031 (0.009)	-0.010 (0.002)	-0.012 (0.001)	-0.024 (0.002)	-0.054 (0.004)
β_1	-0.002 (0.013)	0.007 (0.004)	0.055 (0.015)	0.124 (0.026)	0.233 (0.027)
β_2	-0.018 (0.026)	0.012 (0.008)	0.034 (0.013)	0.097 (0.028)	0.125 (0.028)
Pseudo- R^2	0.003	0.002	0.013	0.040	0.082

Notes: Standard errors appear in brackets. Coefficients significant at the 5% level appear in bold. All observations are pooled together over the 10-year period

Fig. 4 Regression quantiles for employment growth autocorrelation coefficients for four size classes, 1995–2004

Our results clearly indicate differences between micro firms (1–9 employees) and larger firms allocated to other size classes. The aggregate results are mostly due to micro firms. Only for micro firms we observe a negative first and second order autocorrelation of growth rates for growing firms. Also only for micro firms we observe a positive size dependency for growing firms. Let us now discuss the results in turn.

With regard to size (α_1) we see that larger micro firms grow on average faster than smaller micro firms, while a large firm size is associated with faster decline as the negative coefficient for the 10% and the 25% quantiles shows. Note that the fact that the (log) growth rate for a firm moving from 2 to 1 employee is higher in magnitude than the growth rate of a firm with 9 employees moving to 7 employees would suggest a different pattern of α_1 for both growing and declining firms. Thus it is not possible to explain the result for growing firms on the basis of the measurement of growth rates. A possible explanation is that the probability of adjustment increases with firm size as non-convex adjustment costs that give rise to lumpy adjustment patterns are more important for smaller firms. In line with this explanation we obtain coefficients of quite similar magnitude for micro firms (1–9 employees) for all three sub periods (cf. Tables 8 and 9). For the other size classes we observe generally a negative size dependency for both growing and declining firms. This means that larger firms in the other size classes are less likely to grow and more likely to decline. In fact, in most cases we observe that the (negative) coefficient α_1 has its lowest value in absolute terms at the median growth rate (50% quantile). Interestingly the

effect of the size dependency is strongest—when measured by the magnitude of the coefficient—for the size class covering 50–99 employees. It is weaker for the size class 10–49 employees and for the firm in the largest size class (100+ employees).

Thus our results shed new light on the well-known result of a negative dependence of growth on size. This large literature (surveyed in Coad 2009) observes that, on average, smaller firms have faster growth than larger firms. Our results add a nuance to this line of research. For micro firms in the service sectors we used in this paper this seems not to be the case, while the a negative size dependency—in line with most of the literature—can be observed for small, medium-sized and larger firms.

If we move to the autocorrelation coefficients β_1 and β_2 we observe for the micro firms (1–9 employees) that the both autocorrelation coefficients are 0 and insignificant for declining firms and strongly negative and increasing in absolute magnitude for growing firms. This implies that growing micro firms had likely a negative growth experience in the past, while for declining firms we do not observe such a pattern. Decline seems to be uncorrelated with previous growth experiences. This finding confirms on the one hand the erratic character of growth processes of micro firms and on the other hand it indicates that micro firms are subject to considerable non-convex adjustment costs that makes sustained employment growth very difficult.

For the other size classes we do observe statistically significant negative autocorrelation only for declining firms, except for the first order correlation coefficient β_1 for the second size class 10–49 employees for the time periods 1985–1994 and 1975–1984 (see Table 9). However, for the last time period 1995–2004 we observe also for this size class a statistically significant positive first order autocorrelation. For the larger size classes we observe generally positive autocorrelation coefficients that are higher for firms in higher size classes. Figure 4 shows for the third size class (50–99) employees that the autocorrelation coefficient is again decreasing for the fastest growing firms for the time period 1995–2004. However it still remains positive. The largest magnitude of positive autocorrelation coefficients (both for β_1 and β_2) is recorded for the highest size class (100+ employees) suggesting that high-growth episodes of larger firms are longer lived than high growth episodes of small and micro firms. In addition this suggests that for larger firms non-convex adjustment costs seem not to play an important role, this is reinforced by the size and significance of coefficient β_2 . Also the second order Autocorrelation is generally higher in size classes with larger firms. If anything the adjustment costs seem to be convex for large firms giving rise to a gradual adjustment to a desired firm size. This finding is in line with the results by (Acs et al. 2008) who find that over a larger time period only large firms are characterized by persistence of high growth, while for small firms high growth appears to be episodic.

Except for the second size class (10–49 employees) we do not observe statistically significant autocorrelation for declining firms across size classes across time periods. In this size class the first order autocorrelation is statistically significant and negative. In addition we find negative and statistically significant autocorrelation for the 25% quantile for the time periods 1975–1984 and 1985–1994, and statistically significant positive autocorrelation for the 25% quantile in the largest size class (100+ employees) for the time period 1985–1994.

Table 9 Quantile regression estimation of Eq. 4 for four size classes, 1985–1994 and 1975–1984

	1985–1994 s				1975–1984					
	10%	25%	50%	75%	90%	10%	25%	50%	75%	90%
	size class 1–9 employees 411,882 observation				size class 1–9 employees 245,579 observations					
α_1	−0.369 (0.000)	−0.146 (0.001)	0.000 (0.000)	0.067 (0.001)	0.262 (0.000)	−0.317 (0.000)	−0.079 (0.001)	0.000 (0.000)	0.072 (0.001)	0.254 (0.004)
β_1	0.000 (0.000)	−0.062 (0.005)	0.000 (0.000)	−0.172 (0.006)	−0.392 (0.006)	0.000 (0.000)	−0.128 (0.010)	0.000 (0.000)	−0.159 (0.013)	−0.334 (0.015)
β_2	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	−0.067 (0.001)	−0.130 (0.006)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	−0.072 (0.005)	−0.080 (0.015)
Pseudo- R^2	0.313	0.117	0.000	0.016	0.165	0.304	0.043	0.000	0.023	0.226
	size class 10–49 employees 38,263 observations				size class 10–49 employees 20,153 observations					
α_1	−0.055 (0.005)	−0.035 (0.003)	−0.007 (0.006)	−0.099 (0.001)	−0.143 (0.004)	−0.036 (0.009)	−0.031 (0.005)	0.003 (0.000)	−0.101 (0.002)	−0.134 (0.004)
β_1	−0.174 (0.016)	−0.132 (0.009)	−0.017 (0.016)	−0.018 (0.006)	0.010 (0.009)	−0.102 (0.022)	−0.093 (0.010)	0.006 (0.000)	−0.013 (0.005)	−0.006 (0.007)
β_2	−0.044 (0.007)	−0.001 (0.006)	0.009 (0.008)	0.046 (0.006)	0.086 (0.006)	0.00 (0.011)	0.01 (0.007)	0.00 (0.000)	0.02 (0.004)	0.03 (0.008)
Pseudo- R^2	0.015	0.011	0.000	0.038	0.066	0.006	0.006	0.000	0.036	0.064

	size class 50–99 employees 4,305 observations			size class 50–99 employees 2,263 observations		
α_1	-0.315 (0.039)	-0.125 (0.011)	-0.094 (0.006)	-0.146 (0.024)	-0.080 (0.014)	-0.132 (0.013)
β_1	-0.066 (0.051)	-0.010 (0.016)	0.033 (0.015)	-0.060 (0.018)	-0.006 (0.032)	0.072 (0.030)
β_2	0.005 (0.014)	0.012 (0.016)	0.049 (0.013)	0.003 (0.039)	0.003 (0.012)	0.062 (0.026)
Pseudo- R^2	0.049	0.023	0.028	0.026	0.012	0.055
	size class 100+ employees 3,616 observations			size class 100+ employees 1,874 observations		
α_1	0.007 (0.004)	0.000 (0.001)	-0.006 (0.001)	0.001 (0.005)	-0.001 (0.003)	-0.014 (0.001)
β_1	0.026 (0.016)	0.033 (0.013)	0.068 (0.029)	-0.014 (0.019)	0.015 (0.025)	0.144 (0.041)
β_2	-0.013 (0.014)	0.008 (0.012)	0.029 (0.023)	-0.017 (0.028)	0.012 (0.039)	0.117 (0.026)
Pseudo- R^2	0.002	0.002	0.011	0.001	0.000	0.031
						0.053

Notes: Standard errors appear in brackets. Coefficients significant at the 5% level appear in bold. All observations are pooled together over the 10-year period

Table 10 Quantile regression estimation of Eq. 4 for four size classes 1995–2004 and 1985–1994 allowing for two lags in serial correlation and allocation to size classes based on base size

	1995–2004				1985–1994					
	10%	25%	50%	75%	90%	10%	25%	50%	75%	90%
	size class 1–9 employees 461,350 observations					size class 1–9 employees 410,751 observations				
α_1	–0.369 (0.000)	–0.114 (0.003)	0.000 (0.000)	0.065 (0.001)	0.259 (0.001)	–0.369 (0.000)	–0.092 (0.000)	0.000 (0.000)	0.080 (0.001)	0.262 (0.000)
β_1	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	–0.113 (0.003)	–0.334 (0.008)	0.000 (0.000)	–0.052 (0.001)	0.000 (0.000)	–0.173 (0.011)	–0.381 (0.008)
β_2	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	–0.049 (0.003)	–0.126 (0.010)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	–0.074 (0.005)	–0.119 (0.008)
Pseudo- R^2	0.255	0.054	0.000	0.013	0.097	0.301	0.041	0.000	0.027	0.173
	size class 10–49 employees 52,486 observations					size class 10–49 employees 39,404 observations				
α_1	0.094 (0.005)	0.042 (0.004)	0.000 (0.000)	–0.019 (0.001)	–0.036 (0.003)	0.089 (0.005)	0.040 (0.002)	0.000 (0.000)	–0.016 (0.001)	–0.033 (0.003)
β_1	–0.166 (0.014)	–0.124 (0.008)	0.000 (0.000)	0.016 (0.005)	0.063 (0.008)	–0.226 (0.018)	–0.152 (0.009)	0.000 (0.000)	–0.025 (0.006)	0.017 (0.007)
β_2	–0.079 (0.012)	–0.028 (0.005)	0.000 (0.000)	0.055 (0.005)	0.112 (0.005)	–0.060 (0.012)	–0.016 (0.006)	0.000 (0.000)	0.044 (0.004)	0.091 (0.007)
Pseudo- R^2	0.018	0.009	0.000	0.006	0.017	0.028	0.017	0.000	0.005	0.010

	size class 50–99 employees 6,076 observations			size class 50–99 employees 4,322 observations		
α_1	0.021 (0.029)	–0.003 (0.009)	0.005 (0.010)	–0.026 (0.018)	0.061 (0.022)	–0.008 (0.007)
β_1	–0.003 (0.020)	0.002 (0.009)	0.088 (0.017)	0.140 (0.028)	0.039 (0.010)	0.016 (0.004)
β_2	– 0.065 (0.014)	–0.010 (0.006)	0.055 (0.017)	0.080 (0.019)	–0.016 (0.023)	0.005 (0.015)
Pseudo- R^2	0.003	0.000	0.015	0.027	0.012 (0.005)	0.047 (0.013)
	size class 100+ employees 4,903 observations			size class 100+ employees 3,589 observations		
α_1	0.016 (0.004)	0.004 (0.001)	– 0.010 (0.002)	– 0.018 (0.005)	0.017 (0.003)	– 0.000 (0.001)
β_1	–0.019 (0.020)	0.004 (0.005)	0.099 (0.016)	0.206 (0.034)	0.006 (0.019)	0.028 (0.022)
β_2	– 0.095 (0.036)	–0.008 (0.010)	0.057 (0.013)	0.117 (0.022)	–0.011 (0.033)	0.025 (0.018)
Pseudo- R^2	0.004	0.000	0.020	0.051	0.003	0.005

Notes: Standard errors appear in brackets. Coefficients significant at the 5% level appear in bold. All observations are pooled together over the 10-year period. Size class allocation is made on the basis of S_{t-1}

Table 11 Quantile regression estimation of Eq. 4 for four size classes 1995–2004 allowing for two lags in serial correlation

	1995–2004					1985–1994				
	10%	25%	50%	75%	90%	10%	25%	50%	75%	90%
	size class 1–9 employees 463,108 observations					size class 1–9 employees 411,882 observations				
β_1	-0.318 (0.005)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	-0.369 (0.001)	-0.422 (0.005)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	-0.309 (0.003)
β_2	-0.142 (0.007)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	-0.106 (0.005)	-0.181 (0.008)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	-0.041 (0.007)
Pseudo- R^2	0.054	0.000	0.000	0.000	0.051	0.068	0.000	0.000	0.000	0.051
	size class 10–49 employees 50,693 observations					size class 10–49 employees 38,263 observations				
β_1	-0.114 (0.015)	-0.068 (0.010)	0.000 (0.000)	0.023 (0.006)	0.047 (0.010)	-0.160 (0.011)	-0.110 (0.008)	0.000 (0.000)	-0.019 (0.006)	0.005 (0.006)
β_2	-0.051 (0.006)	-0.010 (0.004)	0.000 (0.000)	0.070 (0.005)	0.109 (0.008)	-0.040 (0.009)	0.003 (0.005)	0.000 (0.000)	0.061 (0.008)	0.089 (0.008)
Pseudo- R^2	0.006	0.003	0.000	0.005	0.012	0.012	0.008	0.000	0.004	0.007

	size class 50–99 employees 6,092 observations		size class 50–99 employees 4,305 observations	
β_1	-0.002 (0.015)	0.042 (0.015)	0.147 (0.031)	-0.103 (0.034)
β_2	-0.053 (0.014)	-0.005 (0.007)	0.110 (0.020)	-0.012 (0.059)
Pseudo- R^2	0.002	0.001	0.028	0.004
	size class 100+ employees 4,922 observations		size class 100+ employees 3,616 observations	
β_1	-0.005 (0.016)	0.055 (0.017)	0.261 (0.031)	0.023 (0.196)
β_2	0.015 (0.029)	0.011 (0.013)	0.140 (0.036)	-0.014 (0.117)
Pseudo- R^2	0.005	0.001	0.057	0.001
			0.029	0.002
			0.133 (0.035)	0.032 (0.014)
			0.137 (0.033)	0.072 (0.023)
			0.144 (0.027)	0.031 (0.019)
			0.183 (0.050)	0.098 (0.027)
			0.146 (0.356)	0.028 0.009
			0.149 (0.021)	0.023 0.032

Notes: Standard errors appear in brackets. Coefficients significant at the 5% level appear in bold. All observations are pooled together over the 10-year period. Allocation to size classes is made on the basis of average size

5.4 Robustness

Firm size is a critical variable in our empirical analysis. On the one hand the allocation of firms to size classes may lead to different results. On the other hand the control variable lagged firm size in Eq. 4 may lead to high correlation with firm growth in $t - 1$. Thus it is important to assess whether our results are affected by these econometric issues. We present two robustness results.

First consider the issue of the method of allocating firms to size classes. Table 10 reports the results for two time periods (1995–2004 and 1985–1994) using $S_{i,t-1}$ to allocate firms to the size classes. The results for the autocorrelation coefficients β_1 and β_2 are quite similar to the results in Tables 8 and 9, while the size coefficient α_1 is—except for the smallest size class—lower, as expected. This shows that the results are reasonably robust regarding the method of allocation of firms to size classes. Next we consider the issue of multicollinearity. Table 11 reports the quantile regression results for a specification without size control. The autocorrelation coefficients β_1 and β_2 are largely of the same magnitude and sign as the estimates reported in Tables 8 and 9, except for the smallest size class (1–9 employees). For this size class we record substantial differences with regard to the 75% quantile where we do not record any autocorrelation and the 10% quantile where we record substantial negative autocorrelation when excluding the size control. However, these differences do not invalidate or affect our general results. The robustness results confirm that larger growing firms experience a positive autocorrelation in annual growth rates while growing micro firms typically experience negative autocorrelation that becomes insignificant near the median but quite pronounced towards the upper extreme quantiles. For micro firms sustained employment growth is quite unusual. Interestingly, the pattern of autocorrelation is asymmetric and negative or positive autocorrelation can be observed for growing but not for declining firms, except for the two smallest size classes where we observe negative autocorrelation for rapidly declining firms.

6 Summary and conclusions

Our analysis explored serial correlation in annual growth rates for Austrian service sectors (Hotels and restaurants, Financial intermediation and real estate, Renting and business activities) that includes a large number of micro firms with employment from 1 to 9 employees. In fact, the firm size distribution is dominated by micro firms but shows a pattern that is in line with most of the previous work on firm size distributions and seems to be characterized by a lognormal or a Pareto distribution.

A recent discovery in the industrial organization literature is that the firm growth rates are fat-tailed and follow closely the Laplace density. Our study of the firm growth rate distribution showed that the distribution is in fact ‘tent-shaped’, however that the tail of the distribution are even more extreme than the Laplace distribution. Using a quite general approach we find that if we exclude micro and smaller firms the growth rate distribution becomes closer to the Laplace distribution. This is a first indication that micro firms are characterized by lumpy adjustment patterns. In addition, this finding has important implications for our econometric analysis of the serial correlation of growth rates. This finding implies that a significant extent

of turbulence in employment generation and destruction is generated due to just a handful of fast-growing and fast-declining firms. These firms, although small in number, are of special interest.⁸ However, standard econometric techniques, which focus on the ‘average firm’, are not useful in this case. Therefore, we apply quantile regression that explicitly recognize that firms are heterogeneous and present results from various quantiles of the conditional growth rate distribution. Our aggregate results suggest that there is a significant negative autocorrelation for growing firms. However, once we disaggregate across size classes we find that the growth pattern of high-growth micro firms is strikingly different from the growth patterns of larger high-growth firms. First, only high-growth micro firms are prone to dramatic negative autocorrelation of growth rates, whilst larger high-growth firms have much smoother growth pattern characterized by positive autocorrelation. Second, we observe the negative dependence of growth on size only for small, medium-sized and larger firms. For growing micro firms we observe a positive size dependency. This is in line with previous observations on the lumpiness of growth for micro firms. Our results are quite robust, as we do not find striking differences across time periods. In contrast to Coad (2007) we do find a strong asymmetry in the serial correlation of growth rates of growing and declining firms except for micro firms when we do not control for firm size. While growing firms seem to be characterized by substantial positive (small, medium-sized and larger firms) or negative (micro firms) autocorrelation, for declining firms we do not find any meaningful autocorrelation of growth rates except for small firms. This suggests that at least in the service sectors we study firm growth and firm decline seem to follow quite different patterns.

Our results can be related to some well-known theories in industrial organization. Our results appear to support the passive learning model of the evolution of industries, as proposed by Jovanovic (1982), because of the quite erratic growth paths of micro firms, whereas the growth paths of larger firms are relatively smooth. On the basis of our results Gibrat’s Law would be rejected because, in many cases growth rates in consecutive years are not independent. In addition our results suggest that lumpy adjustment due to fixed adjustment costs (known in the economics literature as non-convex adjustment costs) is especially important for micro-firms, while lumpy adjustments seem not to characterize growth paths of larger firms. If anything, larger firms seem to adjust their size gradually to a desired level at least in the short run. Our finding of a asymmetry of serial correlation of growth patterns suggests that decline is not growth with a reversed sign and connects well to models that explicitly model declining industries (Ghemawat and Nalebuff 1990).

Our findings are reasonably robust and theoretically meaningful. We anticipate that further research will corroborate some of our findings. Sectoral disaggregation and research covering other sectors and countries is needed to establish our finding that the growth patterns of micro firms are that different from small, medium and larger firms as robust empirical regularity.

Acknowledgements We are grateful to Peter Huber, Andreas Reinstaller, two anonymous referees and the managing editor Michael Peneder for useful comments and suggestions on a previous version of the paper. Of course, all errors and omissions are ours.

⁸Recent studies on high growth firms are Acs et al. (2008) and Hözl (2009).

Open Access This article is distributed under the terms of the Creative Commons Attribution Noncommercial License which permits any noncommercial use, distribution, and reproduction in any medium, provided the original author(s) and source are credited.

Appendix: Quantile regression

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

$$y_{it} = x'_{it}\beta_{\theta} + u_{\theta it} \quad \text{with} \quad \text{Quant}_{\theta}(y_{it}|x_{it}) = x'_{it}\beta_{\theta} \quad (5)$$

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

$$\begin{aligned} \min_{\beta} \frac{1}{n} \left\{ \sum_{i,t: y_{it} \geq x'_{it}\beta} \theta |y_{it} - x'_{it}\beta| + \sum_{i,t: y_{it} < x'_{it}\beta} (1 - \theta) |y_{it} - x'_{it}\beta| \right\} \\ = \min_{\beta} \frac{1}{n} \sum_{i=1}^n \rho_{\theta}(u_{\theta it}) \end{aligned} \quad (6)$$

where $\rho_{\theta}(\cdot)$ is defined as:

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

Equation 6 is then solved by linear programming methods. As one increases θ from 0 to 1, one can trace the entire conditional distribution of y condition on x . More on quantile regression techniques can be found in Koenker (2005).

References

- Acs, Z., Parsons, J., and Tracy, W., High-Impact Firms: Gazelles Revisited. SBA Office of Advocacy: Washington DC, 2008.
- Agarwal, R. and Audretsch, D., “Does entry size matter? the impact of the life cycle and technology on firm survival,” *Journal of Industrial Economics*, vol. 49, pp. 21–43, 2001.
- Arauzo-Carod, J. and Segarra-Blasco, A., “The determinants of entry are not independent of start-up size: Some evidence from spanish manufacturing,” *Review of Industrial Organization*, vol. 27, pp. 147–165, 2005.
- Audretsch, D., Santarelli, E., and Vivarelli, M., “Start up size and industrial dynamics: Some evidence from Italian manufacturing,” *International Journal of Industrial Organization*, vol. 17, pp. 965–983, 1999a.
- Audretsch, D.B., Klomp, L., Santarelli, E., and Thurik, A.R., “Gibrat’s law: Are the services different?” *Review of Industrial Organization*, vol. 24, pp. 301–324, 2004.
- Audretsch, D.B., Klomp, L., and Thurik, R., “Do services differ from manufacturing?” in Audretsch, D.B. and Thurik, R. (eds.), *Innovation, Industry Evolution and Employment*. Cambridge (Cambridge University Press), pp. 230–252, 1999b.
- Axtell, R.L., “Zipf distribution of US firm sizes,” *Science*, vol. 293, pp. 1818–1820, 2001.

- Boeri, T. and Cramer, U., “Employment growth, incumbents and entrants: Evidence from Germany,” *International Journal of Industrial Organization*, vol. 10, pp. 545–565, 1992.
- Bottazzi, G., *Subbotools User's Manual*. LEM Papers Series 2004/14, Laboratory of Economics and Management (LEM), Sant'Anna School of Advanced Studies: Pisa, Italy (2004).
- Bottazzi, G., Cefis, E., and Dosi, G., “Corporate growth and industrial structure: Some evidence from the Italian manufacturing industry,” *Industrial and Corporate Change*, vol. 11, pp. 705–723, 2002.
- Bottazzi, G., Cefis, E., Dosi, G., and Secchi, A., “Invariances and diversities in the patterns of industrial evolution: Some evidence from Italian manufacturing industries,” *Small Business Economics*, vol. 29(1), pp. 137–159, 2007.
- Bottazzi, G., Coad, A., Jacoby, N., and Secchi, A., “Corporate growth and industrial dynamics: Evidence from French manufacturing,” *Applied Economics*, forthcoming, 2009.
- Bottazzi, G. and Secchi, A., “Common properties and sectoral specificities in the dynamics of US manufacturing companies,” *Review of Industrial Organization*, vol. 23, pp. 217–232, 2003.
- Buchinsky, M., “Recent advances in quantile regression models: A practical guide for empirical research,” *Journal of Human Resources*, vol. 33(1), pp. 88–126, 1998.
- Caballero, R., Engle, E., and Haltiwanger, J., “Aggregate employment dynamics: Building from microeconomic evidence,” *American Economic Review*, vol. 89, pp. 921–946, 1997.
- Chesher, A., “Testing the law of proportionate effect,” *Journal of Industrial Economics*, vol. 27(4), pp. 403–411, 1979.
- Coad, A., “A closer look at serial growth rate correlation,” *Review of Industrial Organization*, vol. 31(1), pp. 69–82, 2007.
- Coad, A., *The Growth of Firms: A Survey of Theories and Empirical Evidence*. Edward Elgar: Cheltenham UK, 2009.
- Coad, A. and Rao, R., “Innovation and firm growth in high-tech sectors: A quantile regression approach,” *Research Policy*, vol. 37(4), pp. 633–648, 2008.
- Davidsson, P., Steffens, P., and Fitzsimmons, J., “Growing profitable or growing from profits: Putting the horse in front of the cart?” *Journal of Business Venturing*, forthcoming, 2008.
- de Wit, G., “Firm size distributions: An overview of steady-state distributions resulting from firm dynamics models,” *International Journal of Industrial Organization*, vol. 23(5–6), pp. 423–450, 2005.
- Delmar, F., “Measuring growth: Methodological considerations and empirical results,” in Donckels, R. and Miettinen, A. (eds.), *Entrepreneurship and SME Research: On its Way to the Next Millennium*. Aldershot, VA, (Avebury), pp. 190–216, 1997.
- Dunne, P. and Hughes, A., “Age, size, growth and survival: UK companies in the 1980s,” *Journal of Industrial Economics*, vol. 42(2), pp. 115–140, 1994.
- Fotopoulos, G. and Louri, H., “Firm growth and fdi: Are multinationals stimulating local industrial development?” *Journal of Industry, Competition and Trade*, vol. 4, pp. 163–189, 2004.
- Friedman, M., “Do old fallacies ever die?” *Journal of Economic Literature*, vol. 30(4), pp. 2129–2132, 1992.
- Fu, D., Pammolli, F., Buldyrev, S., Riccaboni, M., Matia, K., Yamasaki, K., and Stanley, H., “The growth of business firms: Theoretical framework and empirical evidence,” *Proceedings of the National Academy of Sciences*, vol. 102(52), pp. 18801–18806, 2005.
- Garnsey, E., Stam, E., Heffernan, P., “New firm growth: Exploring processes and paths,” *Industry and Innovation*, vol. 13(1), pp. 1–20, 2006.
- Geroski, P., Machin, S., and Walters, C., “Corporate growth and profitability,” *Journal of Industrial Economics*, vol. 45(2), pp. 171–189, 1997.
- Geroski, P. and Mazzucato, M., “Learning and the sources of corporate growth,” *Industrial and Corporate Change*, vol. 11(4), pp. 623–644, 2002.
- Ghemawat, P. and Nalebuff, B., “The devolution of declining industries,” *Quarterly Journal of Economics*, vol. 105, pp. 167–186, 1990.
- Goddard, J., Wilson, J., and Blandon, P., “Panel tests of Gibrat's law for Japanese manufacturing,” *International Journal of Industrial Organization*, vol. 20(3), pp. 415–433, 2002.
- Hamermesh, D. and Pfann, G., “Adjustment costs in factor demand,” *Journal of Economic Literature*, vol. 34, pp. 1264–1292, 1996.
- Hölzl, W., “Is the r&d behaviour of fast growing smes different? Evidence from CIS 3 data for 16 countries,” *Small Business Economics*, in press, 2009.
- Hölzl, W. and Friesenbichler, K., *Gazelles. Final Report for the Europa Innova Systematic Project*. WIFO - Austrian Institute of Economic Research, 2008.
- Hölzl, W. and Huber, P., *An Anatomy of Employment Dynamics over the Business Cycle: The Case of Firm Size*. Mimeo, Austrian Institute of Economic Research, 2008.

- Hölzl, W., Huber, P., Kaniovski, S., and Peneder, M., "Wifo-Weissbuch: Gründungen, Schliessungen und Entwicklung von Unternehmen. Evidenz für Österreich," WIFO Monatsberichte, vol. 80, pp. 233–247, 2007.
- Hölzl, W. and Reinstaller, A., "Market structure: Sector indicators," in Peneder, M. (ed.), *Sectoral Growth Drivers and Competitiveness in the European Union*. European Communities, Luxembourg, 2009.
- Huber, P. and Pfaffermayer, M., *The Anatomy of the Firm Size Distribution: The Evolution of its Variance and Skewness*. WIFO Working Paper 295, 2007.
- Ijiri, Y. and Simon, H.A., "A model of business firm growth," *Econometrica*, vol. 35(2), pp. 348–355, 1967.
- Johnson, P., Conway, C., and Kattuman, P., "Small business growth in the short run," *Small Business Economics*, vol. 12(2), pp. 103–112, 1999.
- Jovanovic, B., "Selection and the evolution of industry," *Econometrica*, vol. 50(3), pp. 649–670, 1982.
- Kaniovski, S. and Peneder, M., "Determinants of firm survival: A duration analysis using the generalized gamma distribution," *Empirica*, vol. 35, pp. 41–58, 2008.
- Kaniovski, S., Peneder, M., and Smeral, E., "Determinants of firm survival in the Austrian accommodation sector," *Tourism Economics*, vol. 14, pp. 527–544, 2008.
- Koenker, R., *Quantile Regression*. Cambridge University Press: Cambridge, 2005.
- Koenker, R. and Bassett, G., "Regression quantiles," *Econometrica*, vol. 46(1), pp. 33–50, 1978.
- Kumar, M., "Growth, acquisition activity and firm size: Evidence from the United Kingdom," *Journal of Industrial Economics*, vol. 33(3), pp. 327–338, 1985.
- Lotti, F., "Firm dynamics in manufacturing and services: A broken mirror?," *Industrial and Corporate Change*, vol. 16(3), pp. 347–369, 2007.
- Lotti, F., Santarelli, E., and Vivarelli, M., "Does Gibrat's law hold among young, small firms?," *Journal of Evolutionary Economics*, vol. 13(3), pp. 213–235, 2003.
- Marsili, O., "Technology and the size distribution of firms: Evidence from Dutch manufacturing," *Review of Industrial Organization*, vol. 27(4), pp. 303–328, 2005.
- Mitzenmacher, M., "A brief history of generative models for power law and lognormal distributions," *Internet Mathematics*, vol. 1(2), pp. 226–251, 2003.
- Phillips, B. and Kirchhoff, B., "Formation, growth and survival: small firm dynamics in the US economy," *Small Business Economics*, vol. 1(1), pp. 65–74, 1989.
- Reichstein, T., Dahl, M.S., Ebersberger, B., and Jensen, M.B., "The devil dwells in the tails—a quantile regression approach to firm growth," *Journal of Evolutionary Economics*, in press, 2009.
- Santarelli, E. and Vivarelli, M., "Entrepreneurship and the process of firms' entry, survival and growth," *Industrial and Corporate Change*, vol. 16(3), pp. 455–488, 2007.
- Singh, A. and Whittington, G., "The size and growth of firms," *Review of Economic Studies*, vol. 42(1), pp. 15–26, 1975.
- Stanley, M.H.R., Amaral, L.A.N., Buldyrev, S.V., Havlin, S., Leschhorn, H., Maass, P., Salinger, M.A., and Stanley, H.E., "Scaling behavior in the growth of companies," *Nature*, vol. 379, pp. 804–806, 1996.
- Stiglbauer, A., *Job and Worker Flows in Austria 1978–1998*. PhD thesis, University of Linz, 2003.
- Teruel-Carrizosa, M., "Gibrat's law and the learning process," *Small Business Economics*, in press, 2009.
- Variyam, J.N. and Kraybill, D.S., "Empirical evidence on determinants of firm growth," *Economics Letters*, vol. 38, pp. 31–36, 1992.
- Wagner, J., "Firm size, firm growth, and persistence of chance: Testing Gibrat's law with establishment data from Lower Saxony, 1978–1989," *Small Business Economics*, vol. 4(2), pp. 125–131, 1992.
- Weiss, C., "Size, growth, and survival in the Upper Austrian farm sector," *Small Business Economics*, vol. 10(4), pp. 305–312, 1998.
- Winter-Ebmer, R., "Benefit duration and unemployment entry: A quasi-experiment in Austria," *European Economic Review*, vol. 47(2), pp. 259–73, 2003.
- Winter-Ebmer, R. and Böheim, R., "Firmengründung und Beschäftigungsstabilität," *Wirtschaftspolitische Blätter*, vol. 53, pp. 205–211, 2006.