Plant Level Growth Rates: A Case For Disaggregating Shocks *

Jonathan R. L. Halket†
NYU

May 8, 2006

Abstract

Plant and firm level growth rates exhibit significant fat-tailed behavior. I show that these tails still exist after conditioning on some standard explanatory variables like the age and size of a plant, and that the resultant distributions are well-described by a Laplace distribution. I then simulate a simple model which relaxes the standard central limit theorem by allowing for random numbers of random shocks.

*This paper is submitted as part of the requirements 3rd yr phd sequence. The author would like to thank Gianluca Clementi and Sydney Ludvigson for guidance and Dino Palazzo for the data. All errors are mine own.

†Dept of Economics, 110 5th Ave, New York, NY 10011; jrhalket@nyu.edu

1 Introduction

The relationships between firm growth and firm size and age has garnered much attention. Gibrat (1931), Hart & Prais (1956), Simon & Bonini (1958) and Hymer & Pashigian (1962) on through (e.g.) Evans (1987), Dunne et al. (1989) and Cabral & Mata (2003), examine the variation of firm or plant growth rates by and within industries, by size, or by age of plant, firm and industry (among others). The more recent work has taken advantage of the richer panel data sets now available.

Recent empirical evidence, first highlighted by Stanley et al. (1996), points to a tent-shape in the logarithm of the unconditional densities for firm growth. Higher moments of the distribution indicate a dramatic departure from normality in the tails and in a way which is robust to conditioning on industry and size.

I find tent-shaped densities after conditioning on the age as well as size and industry of the plant. Bottazzi & Secchi (2005) argue that a "success breeds success" type of stochastic plant growth model, based on Polya's urn can explain the tent shape in the limit. I propose a more general potential stochastic process which may be based on size and age. I then simulate a simple model using this process which is able to replicate some notable facts about the distribution.

Section 2 discusses the data and several stylized facts therein. Section 3 discusses a parametric form for the distribution called a Subbotin or generalized exponential distribution and fits the distribution to the data, conditioning on several elements. Section 4 proposes a way to generate these distributions from normal shocks and simulates the proposal. Section 5 concludes.

2 Data

The data are from the Chilean manufacturing census Encuesta Nacional Industrial Anual conducted by the Instituto Nacional de Estadisticas. All plants of 10 employees or over are surveyed. Data are from 1979 to 1996 at an annual frequency. A wide variety of industries, from bakeries to motor vehicles, are included.

Throughout, I consider several measures of firm size and growth. I consider number of employees, operating income and total wages as measures of firm size.[1][2] The data on employees is the most extensive of the three, so apart from some robustness comparisons, I use number of employees as our primary measure of size.

2.1 Growth

For many of our growth measures, it is sensible to first demean our size measures. I remove year-industry cohort means from each of our size measures[3] and then take log differences:

$$s_{i,t} = \log S_{i,t} - \frac{1}{I}\sum_{j \in I} \log S_{j,t}$$
$$g_{i,t} = s_{i,t} - s_{i,t-1}$$

where $S_{i,t}$ is the size of plant i in industry I at time t.[4]

Distributions of plant growth rates are fat-tailed relative to the normal distribution. Figure 1 plots the empirical density relative to the normal on a log scale, so as to better highlight the tails. On this scale, growth rates, evidently, are tent-shaped. Figures 2, 3, and 4 show that the tent-shape is robust across some important subsamples.

2.2 Size

Figures 5 - 11 show that the distribution of sizes is approximately a Pareto or power distribution (with the exponent given by the slope coefficient in

[1]See Appendix for exact definitions.

[2]Data on assets were not available for enough firms to be included.

[3]Henceforth, "demeaning" refers to this specific procedure.

[4]Alternative ways of demeaning (for example demeaning the growth rates instead of the levels) do not make a significant difference to the results.

Figure 1:

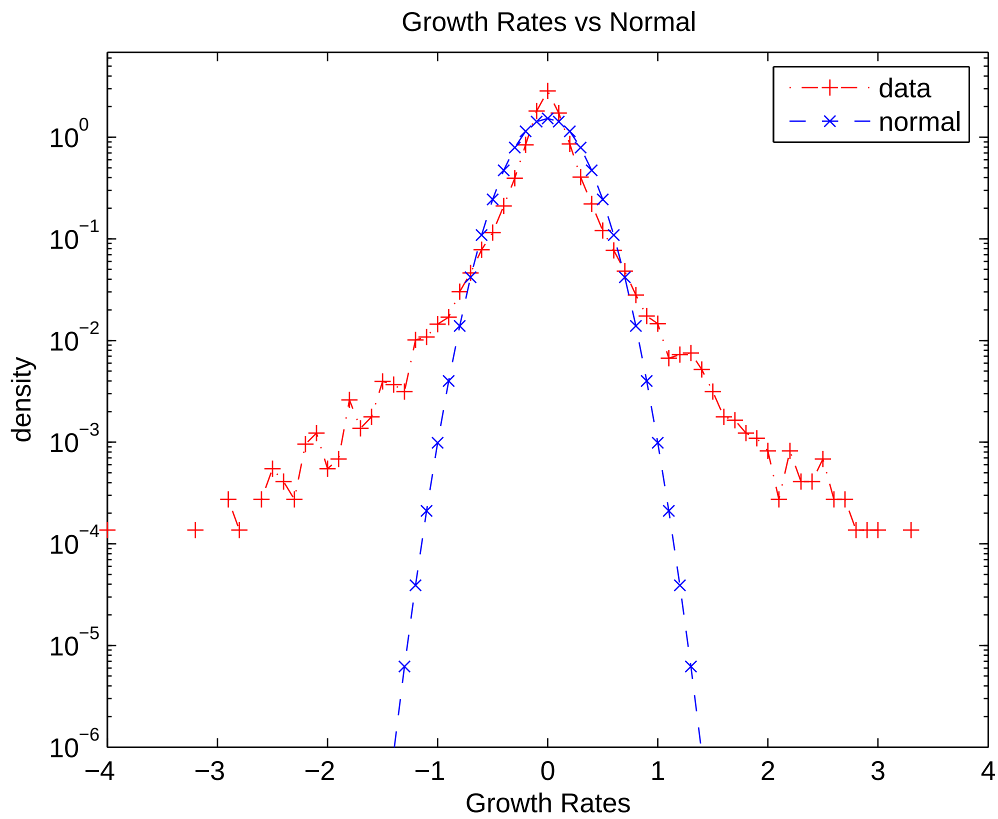

Binned empirical density of plant growth rates as measured by number of employees. Normal density has zero mean and the same variance as the included data. Data are demeaned according to above. Data for all industries and all years are used.

Figure 2:

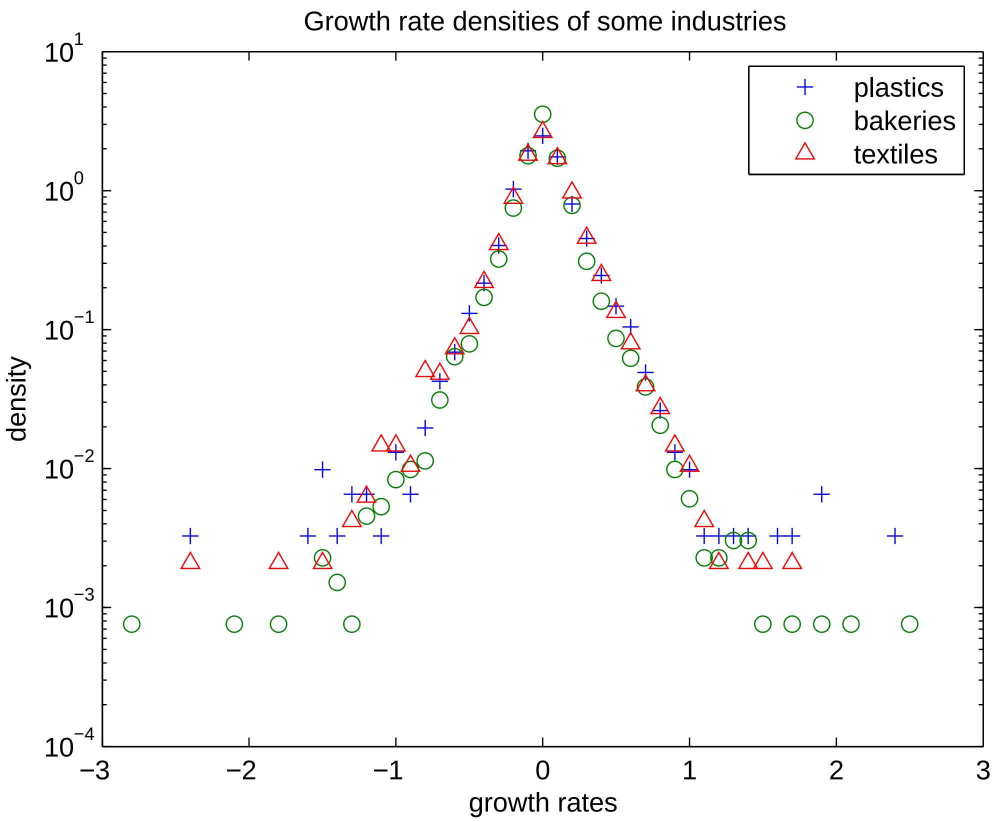

Binned empirical density of plant growth rates as measured by number of employees. Data are demeaned according to above. Data for all years are used. Industries are classified according to the United Nations Classification system ISIC Rev 2.: plastics (3560), textiles (3220), and bakeries (3117).

Figure 3:

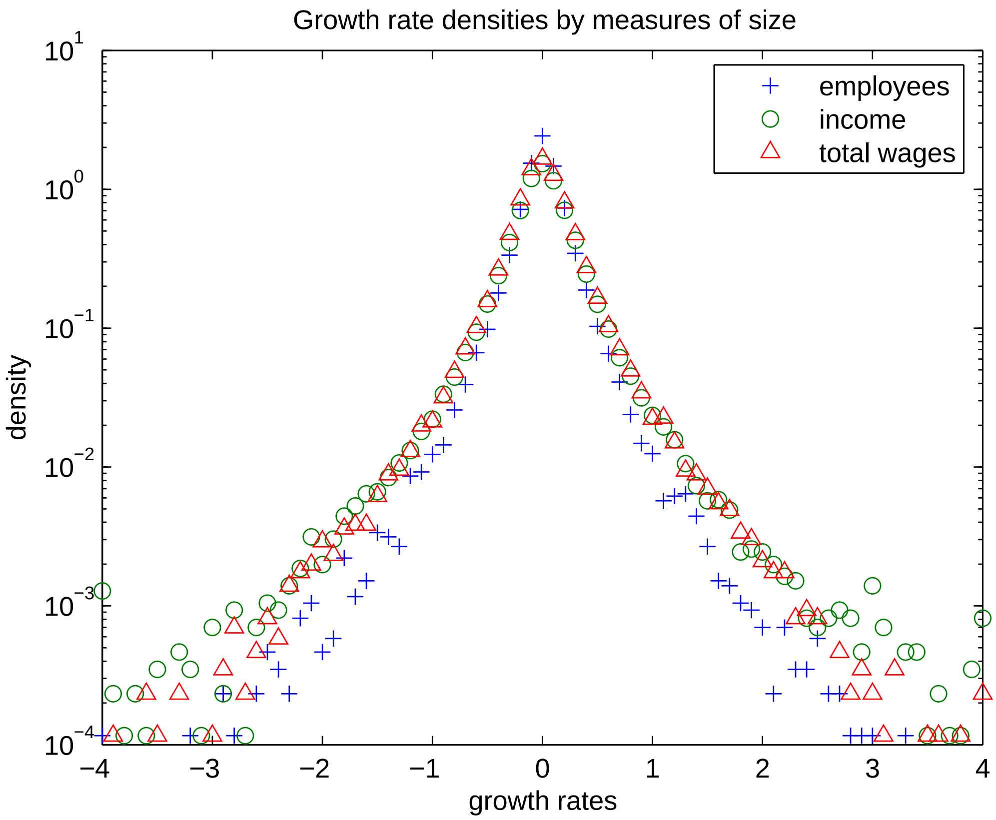

Binned empirical density of plant growth rates. Data are demeaned according to above. Data for all industries and all years are used.

Figure 4:

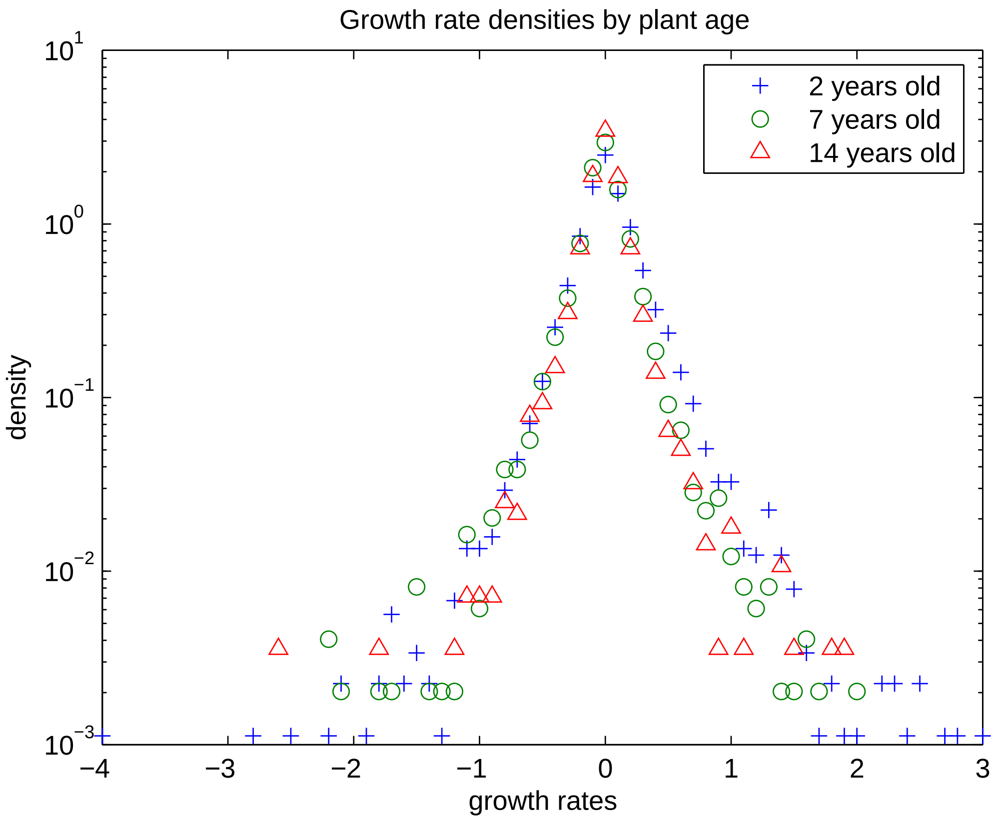

Binned empirical density of plant growth rates. Data are demeaned according to above. Data for all industries and all years are used.

Figure 5:

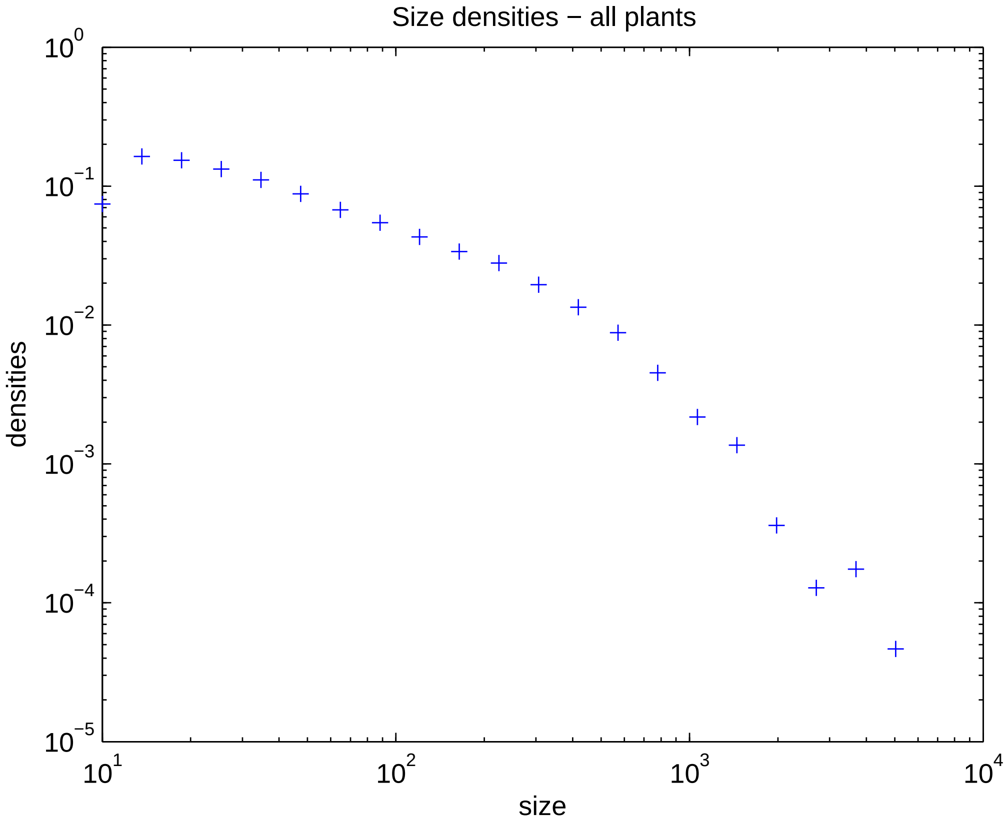

Binned empirical density of plant sizes by employees. Data are not demeaned. Data for all industries and all years are used. An OLS fit gives a slope of -1.2415 (with a standard error of .0992). The log-likelihoods of a power law distribution and log-normal distribution are -5.0865e+005 and -4.3004e+005, respectively.

Figure 6:

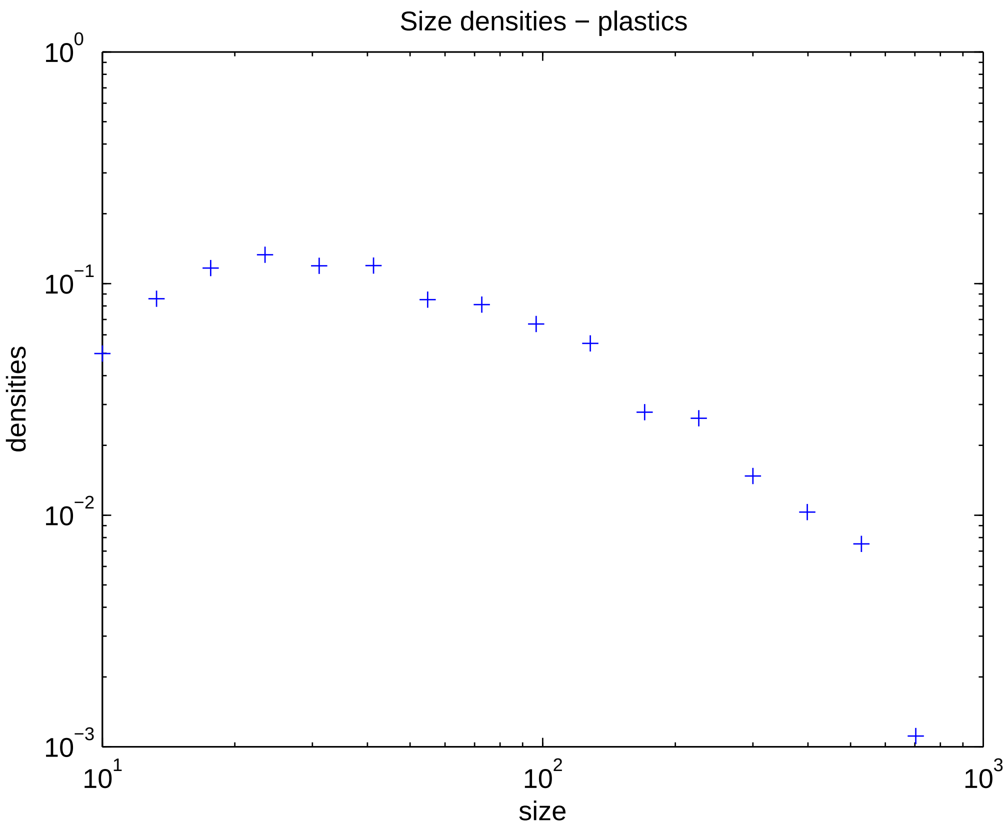

Binned empirical density of plant sizes by employees. Data are not demeaned. Data for all years are used. An OLS fit gives a slope of -0.8018 (with a standard error of .1371). The log-likelihoods of a power law distribution and log-normal distribution are -2.1750e+004 and -1.8023e+004, respectively.

Figure 7:

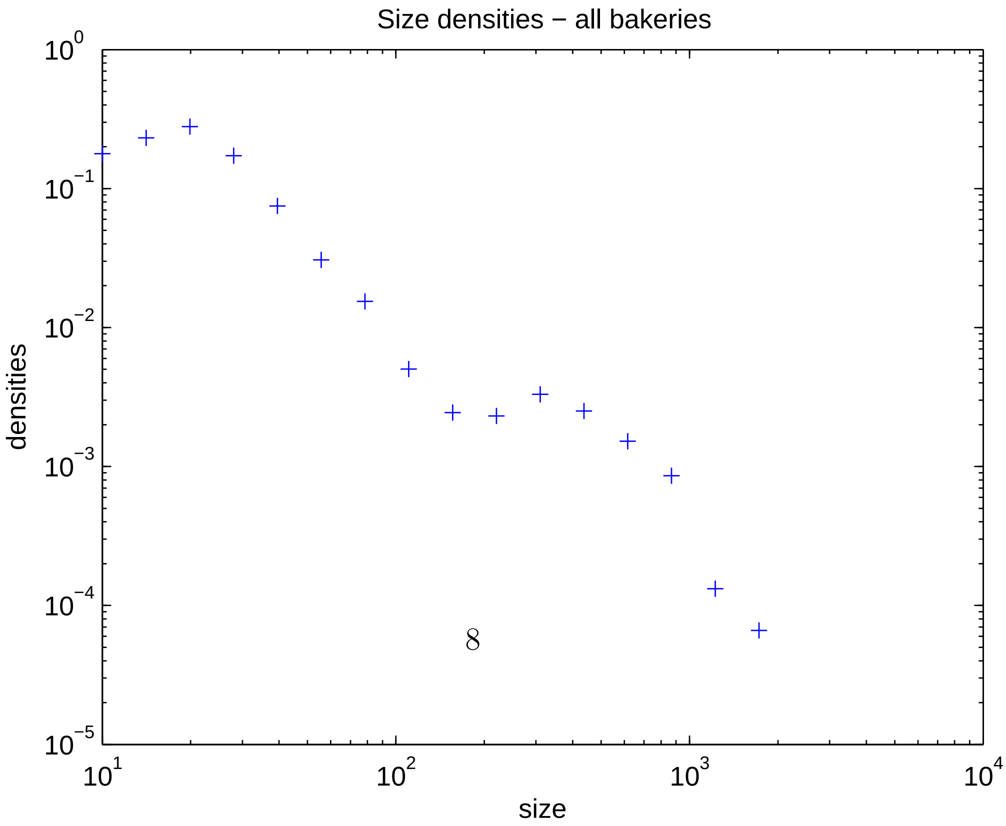

Binned empirical density of plant sizes by employees. Data are not demeaned. Data for all years are used. An OLS fit gives a slope of -1.5550 (with a standard error of .0979). The log-likelihoods of a power law distribution and log-normal distribution are -7.8439e+004 and -6.5606e+004, respectively.

Figure 8:

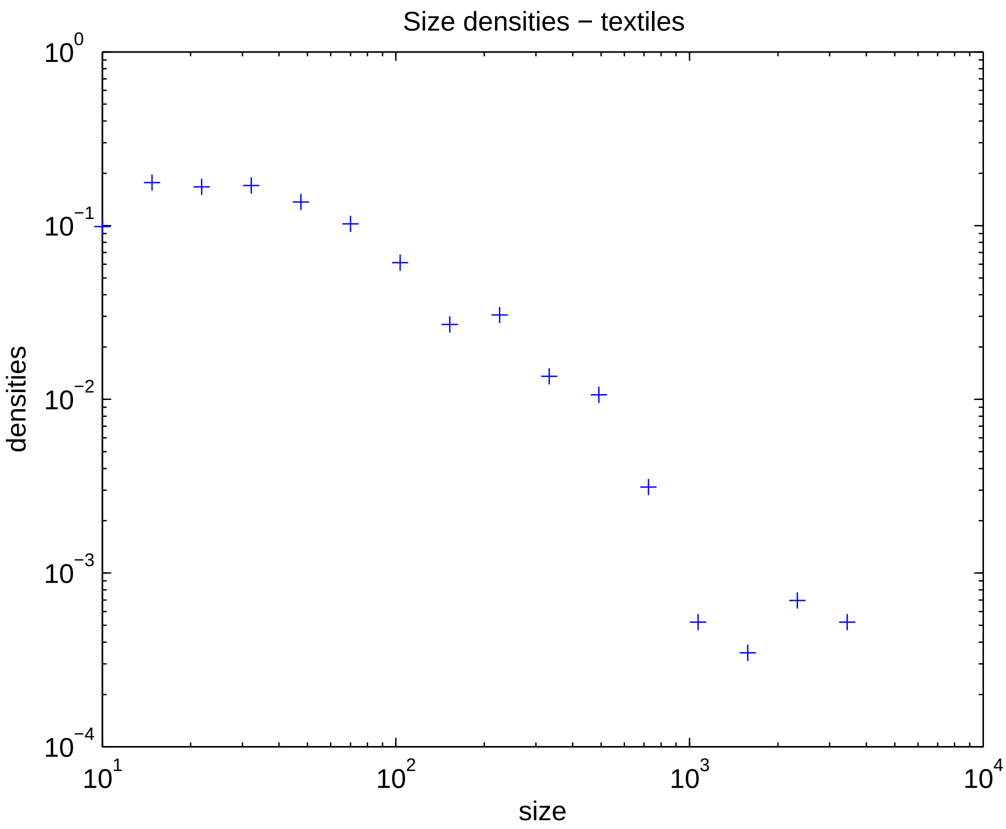

Binned empirical density of plant sizes by employees. Data are not demeaned. Data for all years are used. An OLS fit gives a slope of -1.1813 (with a standard error of .1075). The log-likelihoods of a power law distribution and log-normal distribution are -3.3795e+004 and -2.8067e+004, respectively.

Figure 9:

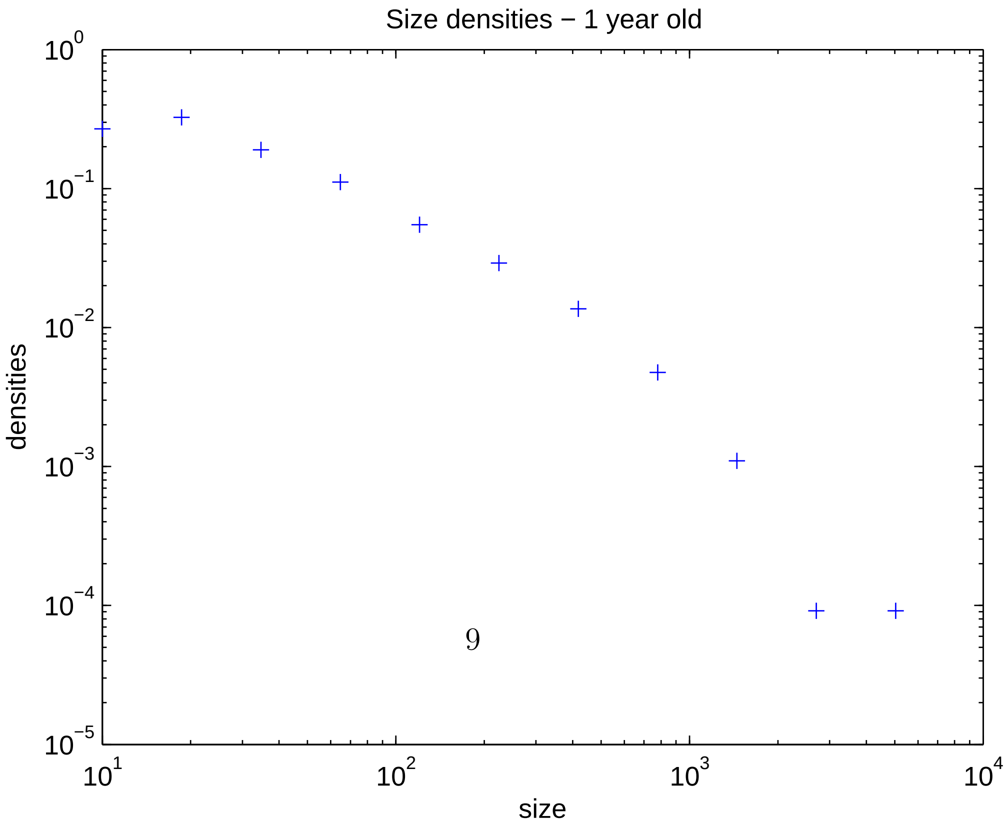

Binned empirical density of plant sizes by employees. Data are not demeaned. Data for all years and industries are used. An OLS fit gives a slope of -1.3997 (with a standard error of .1214). The log-likelihoods of a power law distribution and log-normal distribution are -6.0863e+004 and –5.0522e+004, respectively.

Figure 10:

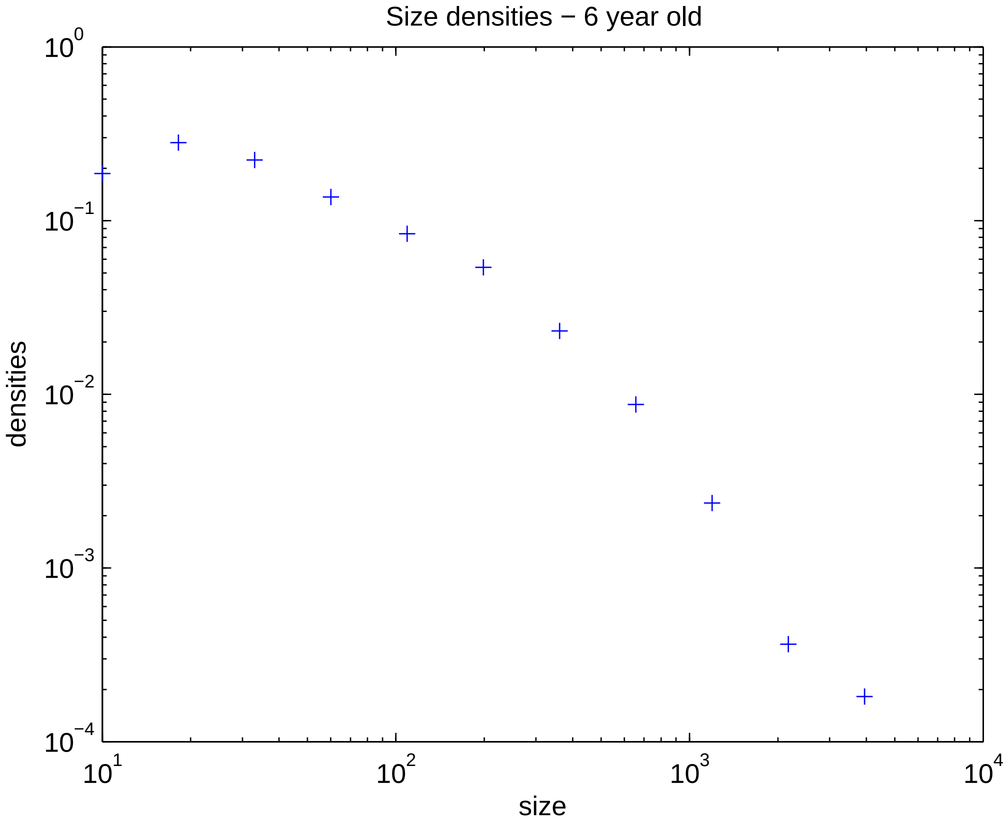

Binned empirical density of plant sizes by employees. Data are not demeaned. Data for all years and industries are used. An OLS fit gives a slope of -1.2422 (with a standard error of .1295). The log-likelihoods of a power law distribution and log-normal distribution are -3.2126e+004 and -2.6908e+004, respectively.

Figure 11:

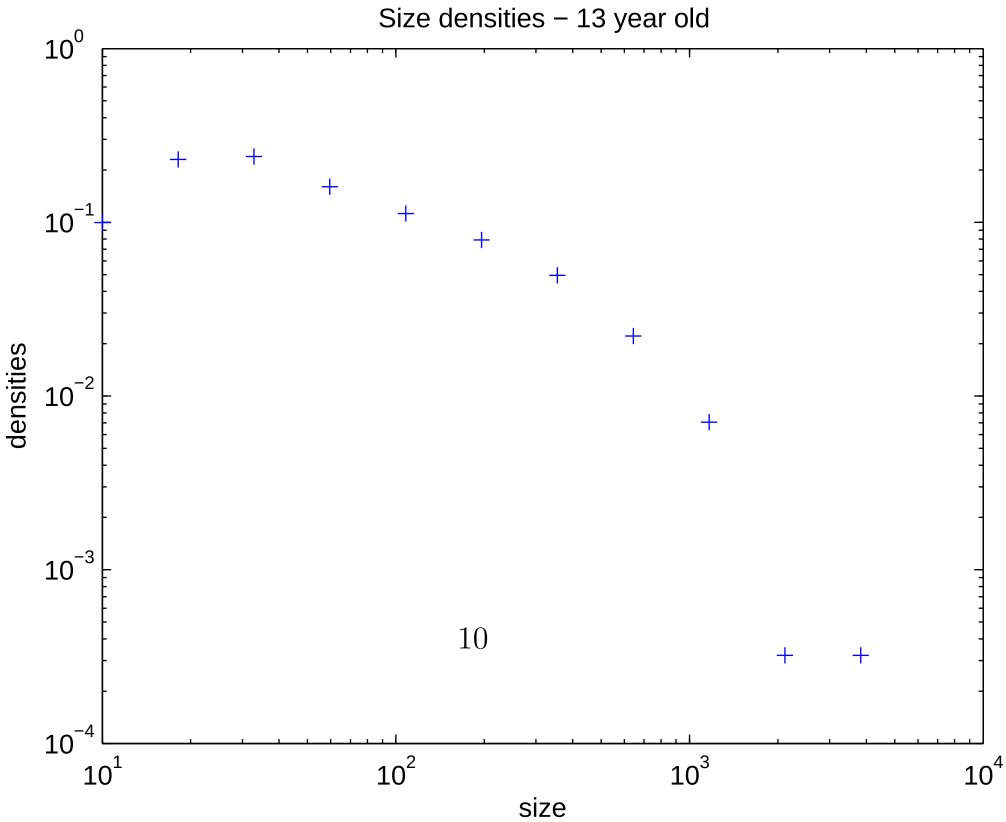

Binned empirical density of plant sizes by employees. Data are not demeaned. Data for all years and industries are used. An OLS fit gives a slope of -1.0746 (with a standard error of .1804). The log-likelihoods of a power law distribution and log-normal distribution are -1.9500e+004 and -1.6657e+004, respectively.

Figure 12:

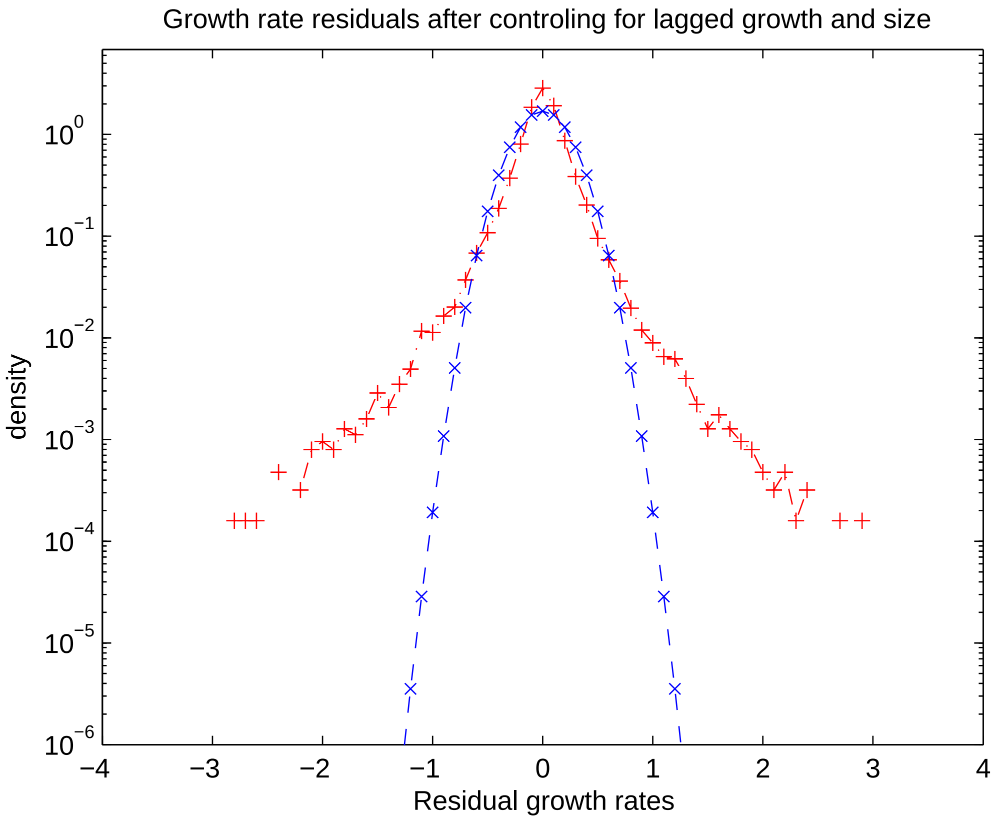

Binned empirical density of residuals of regression of demeaned growth rates on size, lagged demeaned growth rates (as measured by number of employees) and plant level fixed effects. Normal density has zero mean and the same variance as the included data. Data for all industries and all years are used.

each figure). A specific case of a power distribution is Zipf's distribution, which has a slope of -1. Comparing log-likelihoods, the empirical distribution may be better characterized by a log-normal distribution (e.g. Hart & Prais (1956)) [5]

Growth rates are significantly negatively autocorrelated and negatively correlated with size. These correlations are robust to demeaning growth rates and/or size. Figure 12 plots the residuals after controlling for these

[5] Comparing log-likelihoods of non-nested models with dissimilar numbers of parameters should be done sparingly.

Table 1: Volatility on Size and Age

$\log(\sigma)$			
employees	-.0019		.0440
	(.0073)		(.0074)
age		-.0220	-.0231
		(.0010)	(.0010)
constant	-2.4272	-2.2767	-2.3904
	(.02732)	(.0108)	(.0272)

σ is the standard deviation of the demeaned growth rates of employees. Standard deviations of coefficients included in parentheses.

correlates.

The volatility of growth rates obeys a power law $\sigma = kS^{\alpha}$. The coefficient on size is about zero (though statistically significantly positive in the last regression)[6]. If growth came from shocks that were distributed identically and independently across each incremental size (across employees, for example), then α should be $-.5$.[7] If growth were independent of size then α should $= 0$. From table 1 shows, volatility is about constant with size. Also, volatility falls with age. Older plants are more stable than young ones.

3 Fitting A Distribution

3.1 A More General Distribution

Growth rates distributions are not well described by a normal distribution. Still, they are symmetric and single peaked. A more general distribution, which nests both the normal distribution and the Laplace distribution is a generalized error distribution or Subbotin distribution, from Subbotin (1923).

[6] If I remove industry and year means from both our measure of growth and our measure of size, I get a fall in variance (with size) for the very smallest plants. Still, the coefficient on size for the whole sample remains slightly positive.

[7] This is the typical benchmark in the literature (e.g. Amaral et al. (2001) and Teitelbaum & Axtell (2005))

Figure 13: Density of the Subbotin distribution for various b

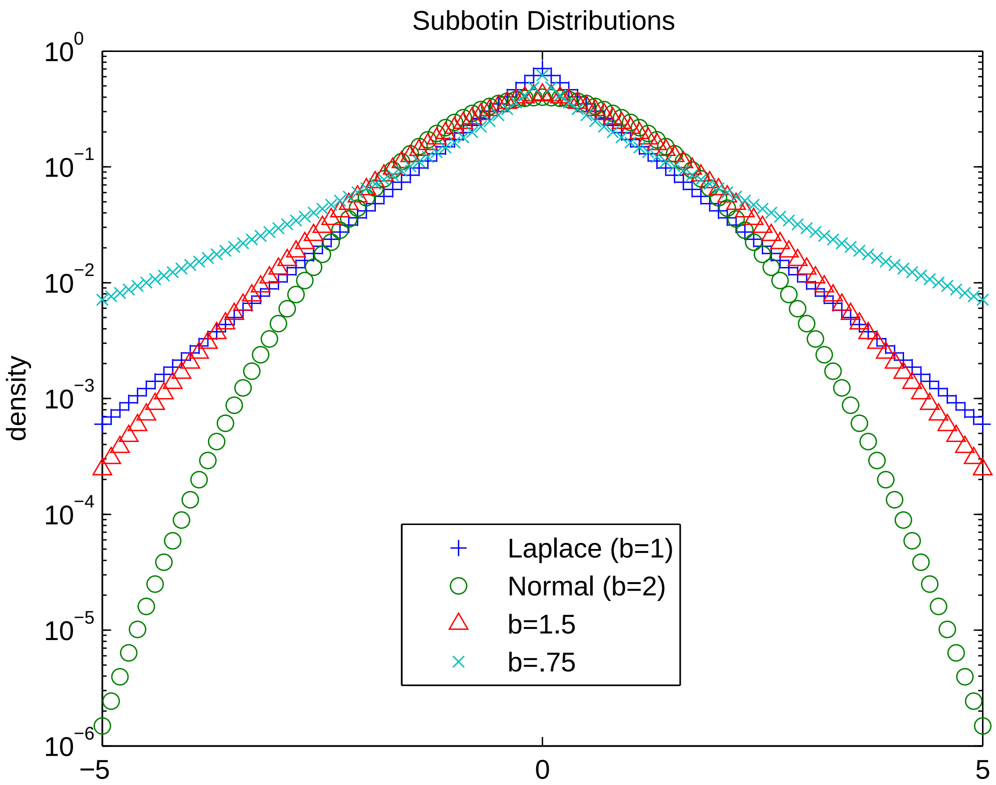

It has a pdf:

$$f_S(x) = \frac{1}{2ab^{1/b}\Gamma(1/b+1)} e^{-\frac{1}{b}|\frac{x-\mu}{a}|^b}. \tag{1}$$

b is a shape parameter. a effects the scale while μ is the mean. $\Gamma(x)$ is the gamma function $= \int\limits_0^\infty t^{-x} e^{-t} dt$.

The Laplace or double exponential distribution has pdf:

$$f_l(x) = \frac{1}{\sqrt{2}\sigma} e^{-|\frac{x-\mu}{\sqrt{2}\sigma}|}. \tag{2}$$

The Laplace distribution is a special case of the Subbotin distribution,

Table 2: Industry Estimation Results

Industry	b	Industry	b
Slaughtering	0.97	Fruit & Vegetable Canning	0.951
	(0.017)		(0.019)
Fish Canning	0.972	Grain Mill Products	0.969
	(0.016)		(0.020)
Bakeries	0.973	Wine	0.924
	(0.006)		(0.020)
Spinning & Weaving	0.982	Knitting Mills	0.979
	(0.014)		(0.014)
Manu. Wearing Apparel	0.99	Manu. Footwear	0.987
	(0.010)		(0.015)
Sawmills	0.954	Furniture Manu.	1.01
	(0.001)		(0.016)
Printing & Publishing	0.989	Plastics Manu.	0.988
	(0.013)		(0.013)
Structural Metal Manu.	0.979	Motor Vehicles	0.976
	(0.017)		(0.021)
Other Fabricated Metal	0.981	Other Machinery	0.965
	(0.017)		(0.019)

Standard deviations of coefficients included in parentheses.

with $b = 1$. The normal distribution is also a special case, with $b = 2$.

3.2 Estimation Results

Tables 2 and 3 contains results from the maximum likelihood estimates of the parameter b. For brevity's sake, I report just those industries which contained over 1000 observations in the data. Table 2 presents the same results by age of plant.

All series, whether by industry or by age, are distributed nearly Laplacean and significantly different from normal.[8]. These results are broadly similar

[8] A Kolmogorov-Smirnov test of the empirical distribution against Laplace and normal distributions also indicates that the Laplace is a much better fit for all series. The KS test, though widely used in such a manner, is not precisely appropriate for testing against fitted theoretical distributions

Table 3: Age Estimation Results

Age	b	Age	b
2	0.975	3	0.981
	(0.007)		(0.008)
4	0.984	5	0.977
	(0.009)		(0.009)
6	0.980	7	0.971
	(0.010)		(0.010)
8	0.974	9	0.979
	(0.010)		(0.011)
10	0.971	11	0.971
	(0.011)		(0.012)
12	0.963	13	0.962
	(0.012)		(0.013)
14	0.968	15	0.980
	(0.013)		(0.019)
16	0.969	17	0.971
	(0.014)		(0.016)
18	0.965		
	(0.015)		

Standard deviations of coefficients included in parentheses.

to previous estimates of b for different conditional distributions (for instance, Bottazzi & Secchi (2005) condition only on industry), though with tighter standard errors.

4 A simple model for generating a Subbotin distribution

4.1 Stochastic Process

There are a wide variety of shocks that may effect plant growth: cost shocks, demand shocks, and technology shocks, just to name a few. Moreover, in any given time period, the number as well as the size of such shocks is random. In this section I show that if the number of shocks is random and each individual shock is normal, then the sum of such shocks can attain a distribution similar to the empirical distribution of growth rates.

I use the following theorem adapted from Andrews & Mallows (1974):

Theorem 4.1 *: If X has a density function f_X which is symmetric about zero, $\exists$ independent random variables W, Z with $Z \sim N(0,1)$ and $X = Z/W$ if and only if:*

$$(-\frac{d}{dy})^k f_X(\sqrt{y}) \geq 0 \textit{ for } y > 0$$

There is no general closed form solution for the pdf of W when $X \sim f_S$, though it is generally stable.[9] For the Laplace case, $X \sim f_l$ with standard deviation σ,

$$f_W(W) = \frac{2}{\sigma^2 W^3} e^{-(\sigma^2 W^2)^{-1}}. \tag{3}$$

The following proposition is the main theoretical result of this paper:

Proposition 4.2 *: Let $x = B(v)$, where*

$$v \sim f_V(v) = \frac{2v}{\sigma^2} e^{-\frac{v^2}{\sigma^2}} \tag{4}$$

[9] West (1987) provides a good way to solve for f_W for this class of distributions.

then $x \sim f_l$ with variance σ^2.

Proof: See appendix.

Plants draw from a Brownian motion process $B(v)$ of length v. This is the continuous analogue of letting x be a draw from a random number v of standard normal shocks. It is also isomorphic to one in which a shock is drawn from a normal distribution with a random variance.[10] Still, the process above seems a much more intuitive description of plant growth. The routine assumption of a fixed number of shocks per period used in most models is with loss of generality. Underlying such an assumption is an assumption that the distribution of the number of shocks is sufficiently thin-tailed for the standard central limit theorem to hold.

4.2 Simulation

I simulate a growth process for a cohort of firms with no dynamic interaction between plants. I start each of 100000 plants at 10 employees and then simulate for 20 years at an annual frequency. I choose parameters so that the variance of growth rates, conditional on age and size, matches the results from table 1. Any firm that is not larger than 10 employees at the end of 20 years is then dropped from the sample. This is an over-simple simulation, intended as a "spin around the block".

Figures 14 and 15 show the results from this simulation. The size distribution appears to be a truncated log-normal or possibly a power distribution. The likelihood for a power distribution is larger than the likelihood for a log-normal. The slope is -6, however, which is significantly different from the empirical slopes as well as from Zipf's distribution (power distribution with slope = 1) often referred to in the literature.

[10] In fact, Subbotin distributions are useful for Monte Carlo simulations for just such a reason. See e.g. Box & Tiao (1973). The rate of decline in the tails of f_V is proportional (but not linearly so) to the shape parameter b.

Figure 14:

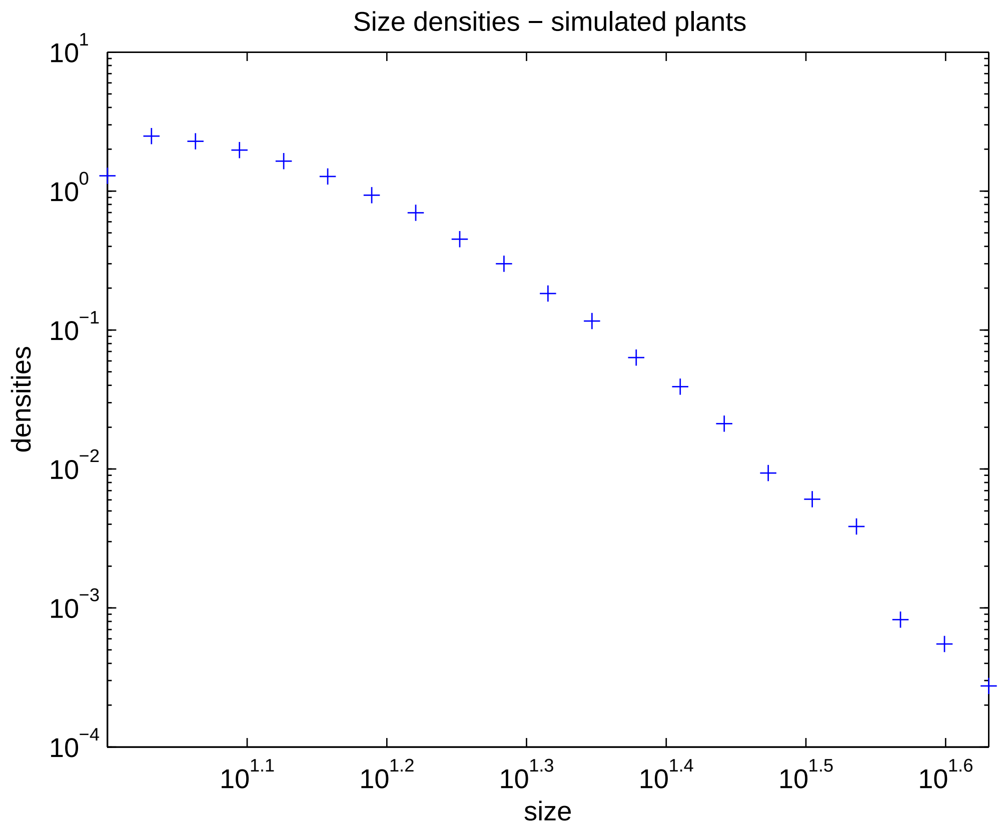
Size densities − simulated plants
densities
size
10^{1}
10^{0}
10^{-1}
10^{-2}
10^{-3}
10^{-4}
$10^{1.1}$
$10^{1.2}$
$10^{1.3}$
$10^{1.4}$
$10^{1.5}$
$10^{1.6}$

Figure 15:

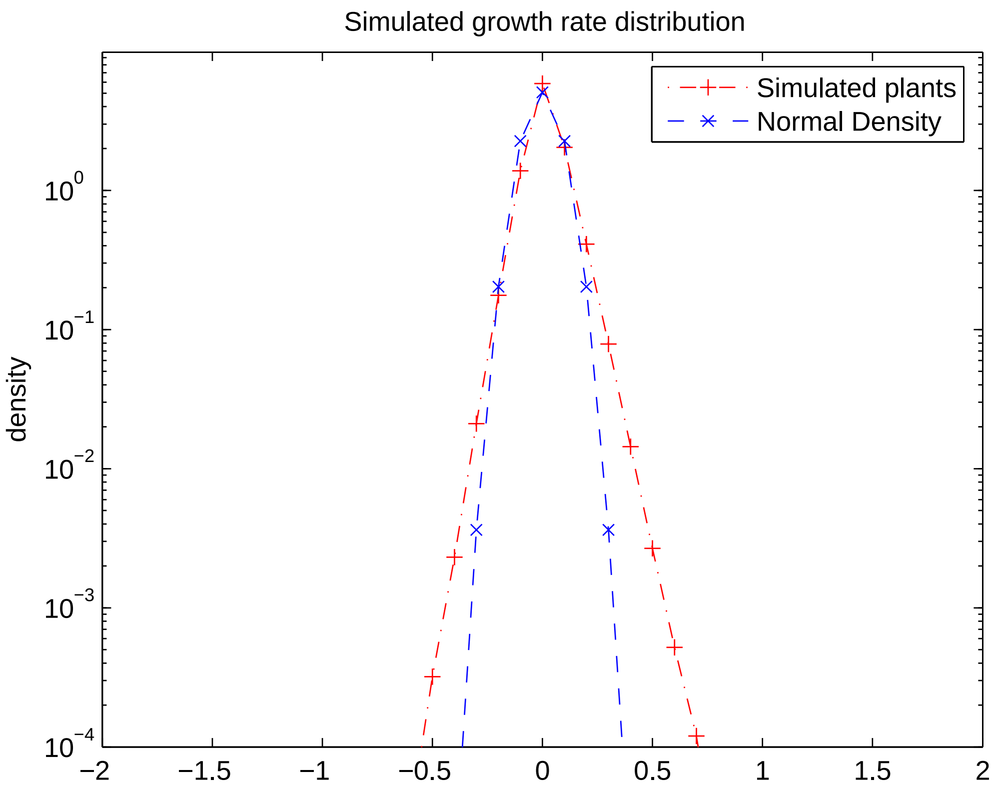
Simulated growth rate distribution
Simulated plants
Normal Density
density
10^{0}
10^{-1}
10^{-2}
10^{-3}
10^{-4}
−2
−1.5
−1
−0.5
0
0.5
1
1.5
2

5 Conclusion

Plant growth rates have fat tails; though symmetric, extreme events happen far too often for the distribution to be normal. Instead, the distribution is best characterized by a Laplace distribution. We have shown that this characterization holds up to conditioning on age and size in a simple parametric way. Nevertheless, the nondecreasing volatility with size in this data set is a new result.

I show that this growth distribution may arise from random numbers of shocks per period. That shocks are random both in their quantity as well as their size seems intuitive, but the interesting consequences this added dimension may have for many models is often overlooked or implicity assumed away.

Many further avenues of research remain. A distribution over the number of shocks should be incorporated into models of firm growth and size to examine the implications. For instance, random numbers of shocks may effect the speed of learning in a model like those of Jovanovic (1982) or Abbring & Campbell (2005). Interesting market dynamics and frictions can be included by using a framework like Ericson & Pakes (1995). Less parametric estimates of the conditional distributions should be pursued as well. Searching for evidence of the number of shocks in a period may provide a way to test our hypothesis.[11]

The discussion of fat-tails in the asset pricing literature (starting with Engle (1982)) may be relevant. That literature has followed a progressive disaggregation, from monthly or daily frequencies to tick or trade by trade level data, as the empirical evidence has indicated it was necessary. Similarly disaggregated plant data (say, at the employee level) may prove necessary as well.

[11] Indeed, the Chilean data set contains reports on the number of days lost to striking employees.

6 Appendix

6.1 Data Construction

Employees measure is the sum of the following data series: Owners (Male), Owners (Females), White Collar Production Workers (Male), White Collar Production Workers (Females), White Collar Executives (Male), White Collar Executives (Females), White Collar Administrative (Male), White Collar Administrative (Females), Blue Collar Production Workers (Males), Blue Collar Production Workers (Females), Blue Collar Nonproduction Workers (Males), Blue Collar Nonproduction Workers (Females), Workers At Home (Male), Workers At Home (Females), Salesperson On Commission (Male), Salesperson On Commission (Females).

Total wages measure is the sum of the following data series in thousands of Pesos: Wages White Collar Workers, Wages of Blue Collar Workers, Bonuses of White Collar Workers, Bonuses of Blue Collar Workers, Payroll Taxes (White Collar Workers), Payroll Taxes (Blue Collar Workers), Family Allowance Taxes (White Collar Workers) Family Allowance Taxes (Blue Collar Workers).

A firm is counted as born in the year that it first appears in the data set. A firm dies the last time it exits the data set.

6.2 Proof of Proposition 4.2

Proof The Subbotin distribution satisfies the requirements for Theorem 4.1. Let the density of x, $f(x)$ be:

$$f(x) = \frac{1}{2a} e^{-|x/a|}$$

From Theorem 1, x can be written $x = Z/W$, where $Z \sim N(0,1)$. Generalizing from West (1987), the density of W is then:

$$f_W(W) = \frac{2}{\sigma^2 W^3} e^{-(\sigma^2 W^2)^{-1}} \quad 0 \leq W \leq \infty.$$

So $x = vZ$ where:

$$f_V(v) = \frac{2v}{\sigma^2} e^{-\frac{v^2}{\sigma^2}}$$

Then:

$$f(x) = \int_0^\infty \phi(0, v) f_V(v) dW$$

where $\phi(0, v)$ is the density of a normal distribution with mean 0 and variance v. x, as a mixture of normals, can be written as Brownian motion of random length v. The proposition then follows. Q.E.D.

References

Abbring, J. H. & Campbell, J. R. (2005), 'A firm's first year', *Tinbergen Institute Discussion Paper* .

Amaral, L. A. N., Gopikrishnan, P., Pleroua, V. & Stanley, H. (2001), 'A model for the growth dynamics of economic organizations', *Physica A* **299**, 127–136.

Andrews, D. F. & Mallows, C. L. (1974), 'Scale mixtures of normal distributions', *Journal of the Royal Statistical Society* **36**(1), 99–102.

Bottazzi, G. & Secchi, A. (2005), 'Explaining the distribution of firms [sic] growth rates', *working paper* .

Box, G. E. P. & Tiao, G. C. (1973), *Bayesian inference in statistical analysis*, Addison-Wiley, New York.

Cabral, L. M. B. & Mata, J. (2003), 'On the evolution of the firm size distribution: Facts and theory', *American Economic Review* **93**(4), 1075–1090.

Dunne, T., Roberts, M. J. & Samuelson, L. (1989), 'The growth and failure of U.S. manufacturing plants', *Quarterly Journal of Economics* **104**(4), 671–698.

Engle, R. F. (1982), 'Autoregressive conditional heteroskedasticity with estimates of the variance of UK inflation', *Econometrica* **50**, 987–1008.

Ericson, R. & Pakes, A. (1995), 'Markov perfect industry dynamics: A framework for empirical work', *Review of Economics and Statistics* **62**, 53–82.

Evans, D. S. (1987), 'Tests of alternative theories of firm growth', *Journal of Political Economy* **95**(4), 657–674.

Gibrat, R. (1931), *Les Inegalites Economiques*, Sirey, Paris.

Hart, P. E. & Prais, S. J. (1956), 'The analysis of buisness concentration: A statistical approach', *Journal of the Royal Statistical Society: Series A* **119**(2), 150–191.

Hymer, S. & Pashigian, P. (1962), 'Firm size and rate of growth', *Journal of Political Economy* **70**(6), 556–569.

Jovanovic, B. (1982), 'Selection and the evolution of industry', *Econometrica* **50**(3), 649–670.

Simon, H. A. & Bonini, C. P. (1958), 'The size distribution of business firms', *American Economic Review* **48**, 607–617.

Stanley, M. H. R., Amaral, L. A. N., Buldyrev, S., Havlin, S., Leschorn, H., Maass, P., Salinger, M. A. & Stanley, H. E. (1996), 'Scaling behavior in the growth of companies', *Nature* **379**, 804–806.

Subbotin, M. T. (1923), 'On the law of frequency of errors', *Matematicheskii Sbornik* .

Teitelbaum, D. & Axtell, R. (2005), 'Firm size dynamics of industries: Stochastic growth processes, large fluctuations, and the population of firms as a complex system', *Economic Research Working Papers* .

West, M. (1987), 'On scale mixtures of normal distributions', *Biometrika* **74**(3), 646–648.