economics letters

Economics Letters 60 (1998) 335–341

Scaling the volatility of GDP growth rates

D. Canning[a,b,*], L.A.N. Amaral[d], Y. Lee[c], M. Meyer[c], H.E. Stanley[c]

[a]*Queen's University of Belfast, Belfast, UK*
[b]*Harvard Institute for International Development, 1 Eliot Street, Cambridge 02138, MA, USA*
[c]*Boston University, Boston, MA, USA*
[d]*Massachusetts Institute of Technology, Boston, MA, USA*

Received 4 December 1997; accepted 15 May 1998

Abstract

The distribution of shocks to GDP growth rates is found to be exponential rather than normal. Their standard deviation scales with GDP$^{\beta}$ where $\beta = -0.15 \pm 0.03$. These macroeconomic results place restrictions on the microeconomic structure of interactions between agents.

Keywords: Real Business Cycles; Country Size; Scale Effects; Volatility; Network Models

JEL classification: E30; 040

1. Introduction

There is a large literature on real business cycles (e.g., Kydland and Prescott, 1982; Backus et al., 1992; Quah and Sargent, 1993 and Cooper and Haltiwanger, 1996) where emphasis is placed on calibrating models to explain the variances and covariances of macroeconomic variables. The key idea is that the comovements in macroeconomic variables place restrictions on the microeconomic impulse generation and propagation mechanisms in the economy. In physics some emphasis is also placed on the study of volatility, based on the fact that theories of the microstructure of physical systems give rise to laws about how volatility varies with the scale of measurement. Studying the relationship between volatility and the scale of measurement may also give some insights into the structure of economic systems.

As we increase the level of aggregation, measured values of average behavior take averages over an increasing number of individuals. Individual specific volatility tends to be averaged out; in the aggregate, volatility depends on the strength of the correlations between people. The volatility of average behavior at different levels of aggregation can be used as an indicator of the strength of microeconomic links between agents.

For example, in network models with interactions between neighbors (see Stanley, 1971; Durlauf, 1996), if local interactions are weak then distant points are uncorrelated and the volatility of the

*Corresponding author. Tel.: +1-617-4959958; Fax: +1-617-4959706; E-mail: dcanning@hiid.harvard.edu

0165-1765/98/$19.00
PII: S0165-1765(98)00121-9

system falls with the square root of its size. If the local interactions are strong, distant points are highly correlated and volatility remains in the aggregate even as the system grows large. At the critical value dividing these regimes, volatility decays as a power law, but not as fast as the square root. Some economic models assume that agents are subject to both common and idiosyncratic disturbances; this gives rise to volatility which declines quickly with size but converges to a lower limit that depends on the volatility of the common component. Empirical analysis of the relationship between volatility and the scale of measurement can determine which model is most appropriate in a particular situation.

The reduction in the volatility of growth rates with country size is well known. Barro (1991) uses size measures to correct for possible heteroskedasticity in long run growth rates. Head (1995) calculates the Spearman rank correlation coefficient of the volatility of GDP with total GDP across countries and argues that the higher output variance of smaller countries is due to their greater openness and susceptibility to foreign shocks. We argue that there is a highly structured relationship between aggregate output shocks and the size of an economy and that microeconomic models should try to explain all of these empirical regularities.

We find that growth rate shocks are not normally distributed but are exponential and, in addition, they follow a power law; the log of their standard deviation declines linearly with log total GDP. These results severely restrict the set of microeconomic models of the economy which are compatible with the data. In particular, models which have common economy wide shocks and/or idiosyncratic shocks are inconsistent with the evidence. We suggest that the exponential form of the probability density function for the growth shocks, and the law by which their standard deviation falls as a power of GDP, appear to be such robust results that microeconomic models of growth fluctuations should be calibrated so as to be consistent with them.

Stanley et al. (1996) and Amaral et al. (1997a) examine scaling behavior in the growth of companies and also find a power law. A power law relationship is interesting because network models of interactions generically do not produce power laws for aggregate volatility; this implies that a parameter of the system is at a critical value. Bak et al. (1987) discuss how parameters controlling the dynamics of physical systems can self organize to the critical values. Schienkman and Woodford (1994) discuss the endogenous emergence of critical parameters in an economic system based on the network model with inventories in Bak et al. (1993). Buldyrev et al. (1997) and Amaral et al. (1997b) show how models with interactions between organizations with endogenous internal structures can give rise to power laws.

2. Graphical results

We calculate the growth rate of y_{it}, the GDP in country I at time t as $g_{it} = \log\,(y_{i,t+1}/y_{it})$. Data are from the Penn world tables 5.6 (see Summers and Heston, 1991). Our first problem is that in order to find the volatility of growth rates we need to measure the actual growth rate relative to its expected value, and we have to estimate the expected value. To address this problem we assume that growth rates follow a process given by

$$g_{it} = \delta_i + \psi_t + r_{it} \text{ with } E(r_{it}) = 0 \qquad (1)$$

so that expected growth is equal to a country specific constant, δ_i, plus a common disturbance across countries at each time t, ψ_t. The volatility of growth then depends on the size of the residual, r_{it}. In principle, different models of how the expected growth rate varies over time and across countries will generate different residuals. In practice, we find our results are robust to making different assumptions about expected growth. However, the correct form of the expected growth rate in Eq. (1) remains an open question. We estimate the parameters in Eq. (1) by ordinary least squares and then examine the structure of the residuals.

We begin, in Fig. 1, by plotting the empirical probability density function of the residuals after sorting the data by the total GDP of the country. For the purposes of the figure, we separate the data into three groups of equal size, low total GDP, medium total GDP and high total GDP, and plot the probability density functions for the smallest and largest groups (the graph for the middle group lies between these two). We estimate probabilities by taking the empirical frequency of observations at equally spaced growth rate intervals. The probability density is plotted on a log scale, so if the residuals were normal this graph should be described by the quadratic function log prob$(r)=a-br^2$ where a and b are constants. Instead, it appears linear in the absolute value of the residual. Note that ordinary least squares estimation of the parameters of the growth regression (1) remains consistent even if the error structure is not normal.

It is clear from Fig. 1 that the growth residuals from the low GDP countries have a higher variance than the growth residuals from the high GDP countries. We wish to investigate how the volatility of growth rates varies with the size of the country. We split the sample into groups by selecting ten equal intervals of log GDP and plot the standard deviation of the growth rate residuals for observations with total GDP in that interval; the results are shown in Fig. 2. It appears that the log of the standard deviation of the growth residual declines linearly with log GDP. The smallest countries have somewhat lower volatility than we would expect given the trend line; it may be the case that some of these small countries are in practice integrated into larger economic units. We ignore the observation

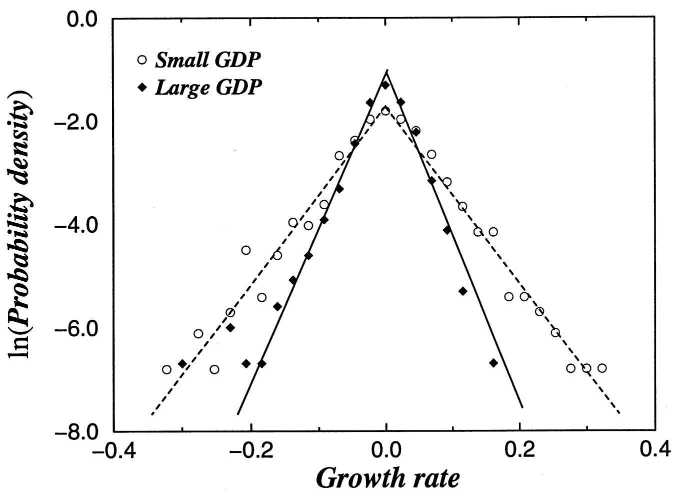

Fig. 1. Logarithm of the conditional probability density $\ln(\rho(r|y))$ as a function of the annual growth rate r. Here, the growth rate is defined as the residual in equation (1). Different symbols refer to different GDP sizes y, $6.9\times10^7<y<2.4\times10^9$ (○), and $2.2\times10^{10}<y<7.6\times10^{11}$ (◆). The lines indicate the best fit of equation (4) to the data.

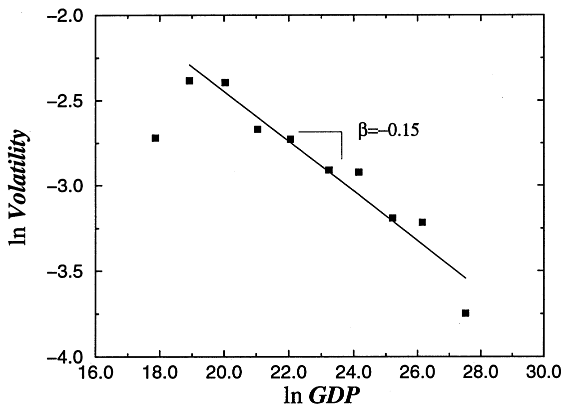

Fig. 2. The volatility $\sigma(y)$ of annual growth rate as function of the GDP size y in double-logarithmic scale. The line indicates a power-law behavior with exponent $\beta = -0.15$.

for the lowest GDP group as an outlier, and using ordinary least squares on the data in Fig. 2, we estimate $\log \sigma_r(y) = \log \alpha + \beta \log y$ with $\beta = -0.15 \pm 0.03$ where the confidence interval is twice the estimated standard deviation of the parameter estimate.

Finally, we repeat the analysis carried out in Fig. 1 but scale the growth residuals to account for the fact that the variance of growth rates varies by total GDP. We scale each growth residual by dividing by the estimated standard deviation given its GDP level according to the law $\sigma_r(y) \sim y^{-0.15}$. In Fig. 3 we graph the empirical probability density function for these scaled residuals. After scaling, the resulting empirical probability density functions appear identical for observations drawn from different populations grouped by total GDP. It appears that our power law relationship between volatility and size accounts for the heterogeneity in the distribution of the growth residuals.

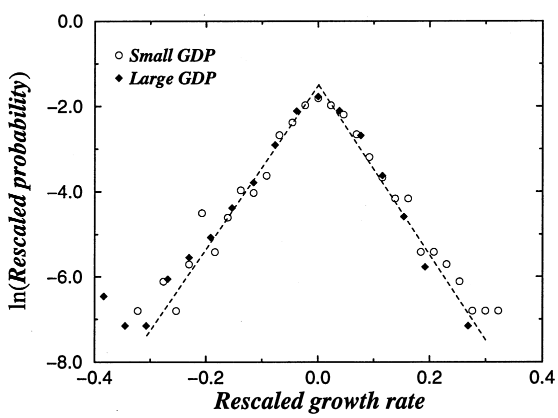

Fig. 3. Logarithm of the rescaled probability density $\ln(\sigma(y)\rho(r|y))$ as a function of the rescaled annual growth rate $r/\sigma(y)$. The symbols refer to the same GDP sizes as in Fig. 1.

3. Maximum likelihood estimates

The graphical approach used in Section 2 is required to find the basic statistical structure of the growth residuals. The advantage of this approach is that it imposes no a priori restrictions on the nature of the data. We can repeat the analysis using a maximum likelihood framework. Now, of course, we must limit the set of models we consider, though the graphical results offer a good guide. We impose the following functional form on the standard deviation of the growth residuals

$$\log \sigma_r(y) = \log \alpha + \beta \log y + \gamma(\log y)^2 \tag{2}$$

We can test the log linearity found in Fig. 2 by imposing the restriction $\gamma = 0$. If the growth residuals were normally distributed, we would have

$$P(r|y) = \frac{1}{\sqrt{2\pi}\sigma_r(y)} \exp[-r^2/2\sigma_r^2(y)] \tag{3}$$

If the linearity suggested by Figs. 1 and 3 is correct, a better model may be the exponential form given by

$$P(r|y) = \frac{1}{\sqrt{2}\sigma_r(y)} \exp[-\sqrt{2}|r|/\sigma_r(y)] \tag{4}$$

We can estimate the parameters of these models of the growth process using maximum likelihood techniques. The results are shown in Table 1.

The parameter estimates change very little under different assumptions about the error structure. What is noticeable is the improvement in the log likelihood (a measure of the goodness of fit of the model) when we move from the normal residuals to exponential residuals and when we include the power law for the standard deviation of the residuals. If we start with a prior belief that puts equal

Table 1
Maximum likelihood estimates

	Error structure	Estimate of α	Estimate of β	Estimate of γ	Log likelihood
I	Normal	0.0605			6767.7
		(0.0002)			
II	Normal	0.3801	−0.1147		6998.7
		(0.0119)	(0.0019)		
III	Exponential	0.0562			7482.5
		(0.0007)			
IV	Exponential	0.3678	−0.1159		7610.8
		(0.0417)	(0.0068)		
V	Exponential	0.1587	−0.0131	−0.0031	7611.4
		(0.1296)	(0.0991)	(0.0030)	

Number of observations = 4885.
Standard errors in parentheses.

weights on a normal or an exponential error structure, then our posterior probabilities from estimates II and IV are approximately $p(\text{normal})=1/e^{612}$, $p(\text{exponential})=1-1/e^{612}$. To put it another way, it would be hard to believe that the graphs in Fig. 1 or Fig. 3 are quadratic and that the growth shocks are normally distributed. The central limit theorem suggests that if the average growth rate is the mean over a large number individual growth rates that are independently and identically distributed it should have a normal distribution; The nonnormality suggests some other mechanism is at work.

In the log likelihood framework, adding a spurious parameter increases twice the log likelihood by a random amount that, for a sufficiently large number of observations, has a χ^2 distribution with one degree of freedom (i.e., it is $\chi^2(1)$). A small increase in the log likelihood is compatible with the additional parameter being spurious, while a large increase in the log likelihood is evidence that the additional parameter has a value different from zero. The 95% critical value for the $\chi^2(1)$ distribution is 3.84; that is, 95% of the time we would expect adding a spurious parameter to increase twice the log likelihood by less than 3.84. Looking at Table 1, the increase in the log likelihood when we move from imposing a constant variance in estimation III to include the power law term in estimation IV is very large. Twice the difference in the log likelihoods is 256.6, so we can reject the hypothesis that $\beta=0$. Taking twice the standard error of the parameter estimate to construct a confidence interval, we estimate $\beta=-0.116\pm0.014$. This figure is slightly different from that found by the graphical approach, probably due to the fact that we have not eliminated outlier observations at very low levels of total GDP.

Adding a nonlinear term in regression V does not improve the log likelihood much from that found in regression IV. Twice the gain in the log likelihood function is only 1.2; we cannot reject the hypothesis that $\gamma=0$ at the 95% confidence level. This confirms the result suggested by Fig. 2, the log of the standard deviation of growth shocks declines linearly with the log of GDP, a power law. We often regard linearity as an approximation that holds only on a limited range. Given the large number of observations, and wide range of values of total GDP in the data, it may be that the scaling function in Fig. 2 is, in fact, linear.

4. Conclusions

We can reject the hypothesis that at the microeconomic level the economy is made up of entities of equal size with independent identically distributed growth shocks. This would produce normally distributed errors with $\beta=-0.5$. Similarly we can reject a model in which shocks to agents within a country have a common component (independent of country size) and an idiosyncratic component. This would imply that the curve in Fig. 2 levels out as GDP rises since size reduces the impact of the individual disturbances at the micro level but the common component remains. The empirical regularities in the volatility of the macrodata found here place severe constraints on the structure of microeconomic interactions.

Acknowledgements

We would like to thank Jeffrey Sachs for encouraging this research and Y. Liu C.-K. Peng, G. Gopikrishnan, S.V. Buldyrev, and M.A. Salinger for their comments.

References

Amaral, L.A.N., Buldyrev, S.V., Havin, S., et al., 1997. Scaling behavior in economics: I. Empirical results for company growth. Journal de Physique I France 7, 621–633.
Amaral L.A.N., Buldyrev S.V., Havin S., Salinger M.A., Stanley H.E., 1997. Power law scaling for a system of interacting units with complex internal structure. Mimeo, Center for Polymer Studies, Boston University.
Bak, P., Chen, K., Scheinkman, J., Woodford, M., 1993. Aggregate fluctuations from independent sectoral shocks: Self organized criticality in a model with production and inventory dynamics. Richerche Economiche 47, 3–30.
Bak, P., Tang, C., Wiesenfeld, K., 1987. Self organized criticality—an explanation of $1/f$ noise. Physical Review Letters 59, 381–384.
Backus, D.K., Kehoe, P.J., Kydland, F.E., 1992. International real business cycles. Journal of Political Economy 101, 745–775.
Barro, R., 1991. Economic growth in a cross section of countries. Quarterly Journal of Economics 106, 407–444.
Buldyrev, S.V., Amaral, L.A.N., Havin, S., et al., 1997. Scaling behavior in economics: II. Modeling company growth. Journal de Physique I France 7, 635–650.
Cooper, R., Haltiwanger, J., 1996. Evidence on macroeconomic complementarities. Review of Economics and Statistics 78, 78–93.
Durlauf S.N., 1996. Statistical mechanics approaches to socioeconomic behavior. Social Systems Research Institute, University of Wisconsin at Madison, working paper No. 9617.
Head, A.C., 1995. Country size, aggregate fluctuations, and international risk sharing. Canadian Journal of Economics 28 (4b), 1096–1119.
Kydland, F.E., Prescott, E.C., 1982. Time to build and aggregate fluctuations. Econometrica 50, 1345–1370.
Quah, D., Sargent, T., 1993. A dynamic index model for large cross section. In: Stock, J., Watson, M. (Eds.), Business Cycles, Indicators and Forecasting. University of Chicago Press, Chicago.
Schienkman, J.A., Woodford, M., 1994. Self organized criticality and economic fluctuations. American Economic Review 84, 417–421.
Stanley H.E., 1971. Introduction to Phase Transitions and Critical Phenomena. Oxford University Press, New York.
Stanley, M.H.R., Amaral, L.A.N., Buldyrev, S.V., et al., 1996. Scaling behavior in the growth of companies. Nature 379, 804–806.
Summers, R., Heston, A., 1991. The Penn World Tables (Mark 5): An expanded set of international comparisons, 1950–1988. Quarterly Journal of Economics 106, 327–368.