Industry Dynamics and Productivity I

Paul T. Scott
Toulouse School of Economics
ptscott@gmail.com

Empirical IO
Winter 2016

Introduction

The lectures on industry dynamics and productivity will have several broad goals:

1. Overview of data used to study production and stylized facts found in them:
    - Firm size and productivity distribution
    - Entry, exit, and growth
2. Understand econometric methods for estimating production functions and measuring productivity.
3. Explore applications on the determinants of productivity

Motivation

The tools of production analysis are important for:

- Evaluating industry performance
- Understanding technological change
- Merger analysis
- Examining effects of policies on efficiency
- Note: these tools are highly relevent in other fields, especially trade, macro, and development

Some Questions

- Role of entry and exit in driving growth?
- Impact of events like trade liberalization and deregulation on productivity?
- Persistence of productivity within plant/firm?
- What are the factors driving plant-level changes in productivity and growth?

Data sources

- More and more, empirical work on production, productivity, and industry dynamics relies on national surveys of manufacturing establishments like the Census of Manufactures and Annual Survey of Manufacturing in the US.
- Researchers have also used plant-level data from Chile, China, Columbia, Denmark, France, India. Typically, these are the best data sets (and the data from certain countries have certain advantages), but they typically have barriers to access.
- Firm-level data sets like Compustat are easier to access, but they typically lack establishment-level information, less detail about inputs and outputs, and present bigger concerns with the selection of the sample.
- Some researchers have acquired data sets with detailed cost information from specific firms (e.g. Benkard (2000) on learning by doing).

The firm size distribution

- A very robust finding: the firm size distribution has a long upper tail.
- … this holds within the vast majority of industries, countries, and after conditioning on observable characteristics.
- Typically, the size distribution is approximated with a lognormal or Pareto distribbtion.

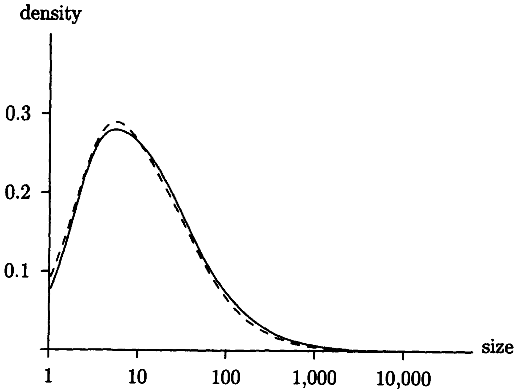

FIGURE 2. FIRM SIZE DISTRIBUTION IN 1983 (SOLID LINE) AND 1991 (DASHED LINE), BASED ON EMPLOYMENT DATA FROM THE *QUADROS DO PESSOAL* DATA SET

Source: Cabral and Mata (2003)

Gibrat's law

- Gibrat's law states that if the growth rate of a variable is independent of its size and over time, it will have a log-normal distribution in the long run.
- Let Y_{it} denote firm i's size (employment or output) in year t . Suppose it evolves according to the following process:

  $$(Y_{i,t+1} - Y_{it})/Y_{it} = \varepsilon_{it}$$

  where ε_{it} is i.i.d. across i and t
- Then, after allowing a large group of firms to evolve for a while, the cross-sectional distribution of Y_{it} will have a log-normal distribution.

"Selection and the Evolution of Industry"
Jovanovic (1982)

Motivation

- Earlier models with adjustment cost and constant returns to scale predicts growth in proportion to size.
- In contrast, micro-data shows that:
  - Smaller firms have higher and more variable growth rates.
  - Smaller firms more likely to exit.
- Model uses noisy selection to explain firm survival and growth.
- Note: this is substantively a different story than standard account of growth in the macro literature which relies on the fixity of capital (but not a mutually exclusive one).

Model: overview

- A model of a small industry with a homogenous product
- Fixed input price, and a known time path of demand and output price.
- Infinitely many small firms.
- Each firm has a cost parameter which it learns about over time. The distribution of cost parameters is known, normal.

Model: costs

- q: output
- $c(q)$: cost function which satisfies

$$c(0) = 0, \quad c'(0) = 0, \quad c'(q) > 0, \quad c''(q) > 0, \quad \lim_{q \to \infty} c'(q) = \infty$$

- Total costs for a firm are $c(q_t)\, x_t$ where x_t is a positive, continuous, and increasing function of η_t:

$$\eta_t = \theta + \epsilon_t$$

where θ is the firm's type (true cost), and ϵ_t are i.i.d. cost shocks.

Model: production

- q_t is chosen to maximize expected profits:

$$q_t = \arg\max_q p_t q - c(q) x_t^*$$

  before the realization of ν.

- x_t^* is the expectation of x_t given the information the firm has so far.
- Can write the optimal choice as $q\left(p_t / x_t^*\right)$
- From implicit function theorem and assumptions on cost function:

$$\frac{\partial q}{\partial x_t^*} = \frac{-c'}{x_t^* c''} < 0$$

Learning

- Since prior distribution of η is normal, so is the posterior distribution, and $(n, \bar{\eta})$ are sufficient statistics, where $\bar{\eta}$ is the mean of the observed history of η's, and n is the number of observations.
- Let $P(\cdot|x_t, n)$ be the posterior distribution for next period's η.

Exit

- $W > 0$ is the value of exit.
- If $V(x, n, t)$ is the value of staying in the industry, then the firms stays iff $V(x, n, t) > W$.
- The value function satisfies:

$$V(x, n, t) = \pi(p_t, x) + \beta \int \max[W, V(z, n+1, t+1)] \, P(dz|x, n)$$

- Theorem 1: a unique solution exists, and V is strictly decreasing in x.
- Define $\gamma(n, t)$ as the level of x_t^* where the firm is indifferent about exit. Note: optimal to exit for $x_t^* > \gamma(n, t)$.

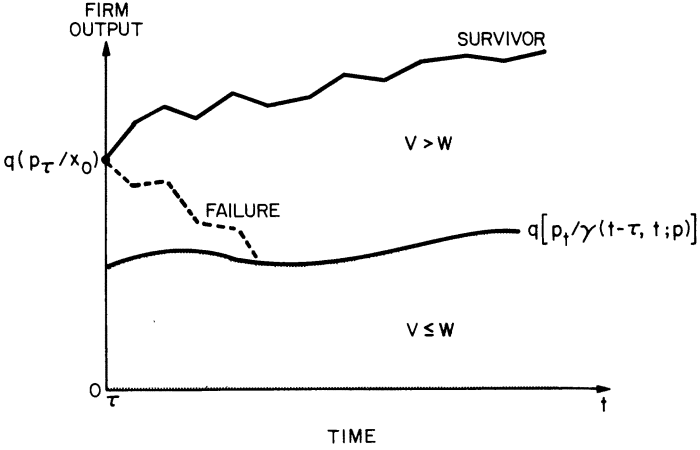

FIGURE 1

Since firms exit for high values of x, they typically exit with small values of q (recalling the monotonicity of the production decision).

Results

- Firms that exit tend to be small. x_t^* is a martingale. Thus, old firms tend to be larger than average (or new) firms.
- The variability of growth rates is largest for small firms, and old firms converge to a common growth rate since x_t^* eventually converges to a constant.
- Larger and older plants should have lower exit rates.

Results

Younger firms grow faster than large firms for two reasons:

- Selection. The firms that exit are ones which would have had low growth rates. Young firms have more exit, so they are more selected for high growth

- Note that period t output is proportional to p/x_t^*. Thus, growth is proportional to

$$x_t^*/x_{t+1}^*$$

Old firms effectively have no uncertainty in x_{t+1}^* conditional on x_t^*, so for old firms

$$E_t\left[x_t^*/x_{t+1}^*\right] \approx 1$$

but for younger firms, x_{t+1}^* is volatile and we use Jensen's inequality to conclude

$$E_t\left[x_t^*/x_{t+1}^*\right] > x_t^*/E_t\left[x_{t+1}^*\right] = 1$$

Comments

- Jovanovic's model far from a complete model of industry dynamics, but it's now widely accepted that selection are an important part of industry dynamics. Hopenhayn (1992) and Melitz (2003) are notable extentions.
- Jovanovic and previous empirical studies show that Gibrat's law does not hold, but the idea that firms face a noisy environment is also still a pervasive feature of models of industry dynamics.
- Melitz's (2003) model is one you should know, if you haven't encountered in another class. It is very elegant, and he develops an equilibrium model in which high-productivity firms select into exporting and low-productivity firms exit.

Dunne, Roberts, and Samuelson (1989)

- DRS examine patterns of employment growth and plant failure using the US Longitudinal Research Database.
- Growth rate:

$$g_t = (S_{t+1} - S_t) / S_t$$

where S_t is employment and $S_{t+1} = 0$ if the firm exits.

TABLE I
PLANT GROWTH AND EXIT RATES

Age (years)	Size (number of employees) 5–19	20–49	50–99	100–249	>250	Total
a. Mean employment growth rate of successful plants						
1–5	0.606	0.299	0.187	0.132	0.067	0.446
6–10	0.338	0.136	0.066	0.011	−0.011	0.202
11–15	0.310	0.055	−0.006	−0.015	−0.018	0.153
Total	0.519	0.226	0.130	0.077	0.026	0.353
b. Plant exit rates						
1–5	0.412	0.396	0.390	0.327	0.229	0.397
6–10	0.347	0.268	0.281	0.245	0.158	0.303
11–15	0.304	0.206	0.234	0.212	0.131	0.255
Total	0.391	0.347	0.346	0.291	0.191	0.363
c. Mean employment growth rate of all plants						
1–5	−0.056	−0.216	−0.276	−0.238	−0.178	−0.129
6–10	−0.127	−0.169	−0.234	−0.236	−0.167	−0.162
11–15	−0.089	−0.163	−0.239	−0.224	−0.147	−0.141
Total	−0.074	−0.199	−0.261	−0.236	−0.170	−0.138
d. Number of plant-year observations on successful plants/failing plants						
1–5	75,959/53,325	29,938/19,649	13,758/8,794	9,472/4,601	3,281/977	132,408/87,346
6–10	27,409/14,569	15,268/5,584	7,577/2,961	5,829/1,889	2,630/494	58,713/25,947
11–15	7,773/3,400	4,675/1,216	2,198/673	1,568/421	911/137	17,125/5,847
Total	111,141/71,294	49,881/26,449	23,533/12,428	16,869/6,911	6,822/1,608	208,246/118,690

Dunne, Roberts, and Samuelson (1989)

Main findings (controlling for industry and year interactions):

- Failure rates decline with increases in plant size and age (consistent with Jovanovic).
- Variance of growth rates declines with age (consistent with Jovanovic).
- Looking at surviving plants, mean growth rates decline with size.
  - One might conjecture that this result is mainly driven by selection, following Jovanovic's story.
  - Note that focusing on all plants and imputing $g_t = -1$ for exiting plants (as DRS do) presumably leads to bias in the other direction. We can't observe what firms would have done if they had stayed active.
  - Selection bias is always a concern when working with production data – we will discuss it in more detail later.

Productivity definitions

- The two most popular measures of productivity are labor productivity and total factor productivity (TFP).
- Labor productivity is defined as the ratio of output to labor inputs (Y_t/L_t).
- TFP is defined as the residual of a production function. For example, with the Cobb-Douglas Production function

$$Y_t = e^{\omega_t} L_t^{\alpha} K_t^{\beta},$$

which we can rewrite in logs,

$$y_t = \alpha l_t + \beta k_t + \omega_t.$$

TFP is ω_t (lower case variables represent logs of uppercase variables).

Basic concerns with productivity definitions

- Labor productivity can change due to changes in the capital-labor ratio without any changes in technology. Consequently, TFP is typically the object of choice for studies on technological change or firm performance.
- That said, TFP is not without its own conceptual and practical limitations.
  - Unlike labor productivity, TFP is defined in terms of a specifc functional form and does not have units.
  - TFP relies on measurement of capital stocks, which is typically difficult.
- However productivity is measured, it tends to have a large variance across firms, and the productivity of an individual firm tends to be highly correlated over time.

Bartelsman and Doms (2003): overview

Bartelsman and Doms (2003) review some empirical work on productivity. Stylized facts:

- Large productivity dispersion across firms.
- Within firm, productivity is highly but imperfectly persistent.
- There is considerable reallocation of labor inputs and output within industries;
  "the aggregate data belie the tremendous turmoil underneath."

What to make of these residuals?

- "I found the spectacle of economic models yielding large residuals rather uncomfortable, even when the issue was fudged by renaming them technical change and claiming credit for their "measurement." – Zvi Griliches
- Bad data could be one reason could be one reason we observe large TFP dispersion, but we observe similar levels of dispersion in developing and developed countries, and we see that measured productivities are connected to real outcomes:
  - more productive firms are less likely to exit
  - more productive firms are more likely to be exporters
  - productivities of entrants tend to be lower than average incumbents
- With most things that correlate with productivity, we should not simply include them in the TFP regression as determinants of productivity. It's natural to think more productive firms would want to become exporters as in Melitz (2003). On the other hand, it's possible that there is some causal effect.

Data limitations

- Theory is typically developed in terms of input and output quantities, but data often contain only expenditures and revenues. This means we might have to worry about output price heterogeneity.
- Rather using sales as the dependent variable, we might want to use value added (subtracting out materials and other short-term input expenditures), but there are concerns either way.
- Typically we observe labor expenditures and capital stock is just the value of assets. Both these measures are aggregates – we would like to see labor inputs at different skill levels and to know about amounts of different sorts of investments.
- Most establishments produce multiple products.

Functional forms

- While Cobb-Douglas is most common in literature, it is limiting in important ways:
  - All technological change is neutral to inputs.
  - Very restrictive on input substitutability.

Simultaneity

- $y_t = \alpha l_t + \beta k_t + \omega_t$
- Generally, we should expect input use to respond to ω_t. For example, if capital is set at $t-1$ and labor can be adjusted at t, we should expect labor to respond to the current realization of productivity.
- Input prices as instruments are a potential solution, but we often don't observe them with any variance, and if they do vary, you might question whether the variation is exogenous.

Fixed effects

- Another potential solution to the simultaneity problem is to assume $\omega_{it} = \mu_i + \varepsilon_{it}$ and that ε_{it} is uncorrelated with input decisions.
- One problem with this (and fixed effects more generally) is that they kill the signal-to-noise ratio. In particular, capital will typically have little variation within firm, and so we typically see significant downward bias in the capital coefficient when using fixed effects.
- Another problem is that productivity often doesn't seem to be time invariant in the data, and often we're interested in identifying how it responds to some change in the environment.

Selection

- We discussed before how the firms that exit are those that have low productivity draws.
- Selection will be an issue, for example, if we want to estimate how the productivity process evolves or how endogenous variables like exporter status impact productivity.