

Economic Growth: Lecture 3

Doug Hanley

The Schumpeterian framework that was introduced in the last lecture provides a good framework in which to think about technological progress and how it relates to creative destruction. It also has fairly precise implications regarding the life cycle dynamics of the firm, that is, how they are created and destroyed and how they might expand or contract over time.

Going down this path will bring us closer to certain industrial organization topics, such as firm entry and exit, competition, and market structure. These topics are interesting and worthy of understanding in their own right, but in broadening our scope, we also gain a better understanding of the overall growth picture.

There is a wealth of empirical literature on US firm level data. Much of this is made possible by the extensive data available from the Census in the form of the Longitudinal Business Database. Additionally, for growth related topics, data on patenting is invaluable. In the case of publicly traded firms, one can also consult the more widely available Compustat database.

A seminal work on this topic in terms of scope is Davis et al. (1998). One should also see Bartelsman and Doms (2000) for a shorter and slightly more recent overview. These papers attempt to shed light on the linkage between productivity growth at the aggregate level and productivity growth that can be observed at the firm level.

1 Theory of the Firm

In this lecture, we will be focusing initially on theory. The primary source will be an important paper in this strand of the literature, Klette and Kortum (2004). At its core, the model featured in this paper inherits many features with the quality ladder model presented in lecture 2. The primary difference is the introduction of the notion of a firm, as well as the distinction between entrant and incumbent innovation.

At any given time, a firm denoted by $f \in [0, F]$ owns some collection of n_f product lines $\mathcal{J}_f = \{j_f^1, \dots, j_f^{n_f}\} \in [0, 1]^{n_f}$. At any given time, each product line j is owned by some firm f . Incumbent firms can generate a flow rate of innovations X_f using research labor C_f according to the production function

$$X_f = G(n_f, C_f) = \nu n_f^{1-\gamma} C_f^\gamma$$

if we denote variables that are normalized by the number of products n_f with lower case, this implies

$$x_f = g(c_f) = \nu c_f^\gamma$$

Thus an equivalent representation would be that each product line has an associated research lab that can generate innovations at rate x_f using research labor $g(c_f)$. Inverting the above, the cost in terms of research labor for a particular flow rate x_f is

$$c_f = c(x_f) = \left(\frac{x_f}{\nu}\right)^{\frac{1}{\gamma}}$$

Now let's attempt to find the value of owning a particular product line. For a

firm with n product lines, the value function satisfies

$$rV_n - \dot{V} = n\pi - n\tau [V_{n-1} - V_n] + n \max_x \{-wc(x) + x[V_{n+1} - V_n]\}$$

where by construction $V_0 = 0$. We can now posit a per-product line value of the form $\tilde{V}_n = n\tilde{V}$, which yields

$$(r + \tau)V - \dot{V} = \pi + \max_x \{-wc(x) + xV\} \equiv \pi + \Omega$$

where Ω is referred to as the option value of innovation. Normalizing by the aggregate growth rate g , we then find

$$(r + \tau - g)\tilde{V} = \tilde{\pi} + \max_x \{-\tilde{w}c(x) + x\tilde{V}\} = \tilde{\pi} + \tilde{\Omega}$$

The first order condition for the above is then simply

$$\tilde{w}c'(x) = \tilde{V}$$

Recall the primary results from the production side equilibrium from the previous lecture

$$\tilde{\pi} = 1 - \lambda^{-1} \quad \text{and} \quad \ell = \frac{\lambda^{-1}}{\tilde{w}}$$

Innovation is undertaken by both incumbent firms and entrants. Suppose that there is a pool of entrants with a linear technology for producing innovation, whereby they can achieve an aggregate rate e by spending χe in terms of research labor. A successful entrant steals one product line from an existing incumbent firm. Thus we should have the condition

$$\tilde{V} = \tilde{w}\chi$$

Since there is a unit mass of product lines and the per-product innovation rate of incumbents is x , this is also the aggregate innovation rate by incumbents. Thus the overall aggregate innovation rate is

$$\tau = x + e$$

Finally, there is the labor market clearing condition with $L = 1$

$$1 = P + R = \frac{\lambda^{-1}}{\tilde{w}} + c(x) + \chi e$$

2 Firm Size Distribution

We can now begin to characterize the firm size distribution resulting from this model. We will find the distribution over the number of products n . However, because everything is identical across firms and products at the product line level, these same results will hold proportionally for observables such as revenue, income, profit, and employment.

For any given incumbent firm with n product lines, it will gain products at the rate nx and lose products at rate $n\tau = n(x + e)$. Thus in expectation each incumbent firm is shrinking at the rate

$$\frac{\dot{n}}{n} = \frac{nx - n\tau}{n} = -e$$

because some are being replaced by new entrants. Let μ_n be the mass of firms with n product lines. Note that when a firm loses its last product line and reaches $n = 0$, it is assumed without loss of generality that they exit, as they have no research capacity or sales. Thus the μ distribution should satisfy the

flow equations

$$\text{Inflows} = \text{Outflows}$$

$$e + 2\tau\mu_2 = x\mu_1 + \tau\mu_1$$

$$(n-1)x\mu_{n-1} + (n+1)\tau\mu_{n+1} = nx\mu_n + n\tau\mu_n \quad \text{for } n > 1$$

In addition, we will also have $\sum_{n=1}^{\infty} n\mu_n = 1$.

Now let's guess that μ takes the form

$$\mu_n = \frac{AB^{-n}}{n}$$

Plugging this in to the $n > 1$ clause and canceling, we find

$$xB + \tau B^{-1} = x + \tau$$

which has the roots $B = 1$, which we reject as non-summable, and $B = \tau/x$. Using this and the $n = 1$ clause, we can then find that $A = e/x$. Thus full expression is

$$\mu_n = \frac{1}{n} \left(\frac{e}{x}\right) \left(\frac{x}{e+x}\right)^n = \frac{\hat{e}}{n} \left(\frac{1}{1+\hat{e}}\right)^n$$

where $\hat{e} = e/x$ is the relative entry rate. This can be verified to satisfy the unit sum condition. We can also find the total mass of firms, which is simply $F = \sum_{n=1}^{\infty} \mu_n$. Here we use the identity

$$\sum_{n=1}^{\infty} \frac{z^n}{n} = \ln\left(\frac{1}{1-z}\right)$$

which can be obtained by integrating both sides of the more familiar identity

$$\sum_{n=0}^{\infty} z^n = \frac{1}{1-z}$$

In our case $z = x/(e+x) = 1/(1+\hat{e})$, meaning the total mass of firms is equal to

$$F = \hat{e} \ln(1 + \hat{e}^{-1}) \in [0, 1]$$

Note that F is thus increasing in \hat{e} . Because there is a fixed unit mass of products, the average firm size is thus $\bar{n} = 1/F$ and is decreasing in \hat{e} . Using this, we can now find the true distribution of firms (rather than the mass) with $\tilde{\mu}_n = \mu_n/F$. This yields

$$\tilde{\mu}_n = \frac{(1 + \hat{e})^{-n}}{n \ln(1 + \hat{e}^{-1})}$$

The net effect of higher relative entry on the distribution is to concentrate more mass in one product line firms, while lower entry spreads the distribution out. Section 2 depicts this.

3 Goals of the Paper


This paper is nice because it clearly sets out what trends in the data it wishes to explain.¹ Furthermore, it manages to capture the bulk of them using a fairly simple and analytically tractable model. Some of the most relevant ones for are purposes are

1. **Research intensity is independent of firm size.**

This is simply the ratio of research investment to revenue, which is con-

¹These are often called "stylized facts," but I think this term is falling out of favor.

Figure 1: Product Line Distributions



stant as both values scale linearly with number of products.

2. **The distribution of research intensity is highly skewed, and a considerable fraction of firms report zero research.**

A highly skewed distribution of number of products yields the former, while the latter is not incorporated.

3. **Differences in research intensity across firms are highly persistent.**

The authors extend the model to include exogenous and persistent differences in innovation step size and cost. This gets the job done, but is somewhat *ad hoc*.

4. **Firm research investment follows essentially a geometric random walk.**

This arises because firms grow and shrink symmetrically by gaining and losing products, and this linearly affects research spending.

5. **The size distribution of firms is highly skewed.**

One could argue that this is true, however the distribution is still not fat tailed. Empirical studies have found the size distribution of firms to approximately Pareto.

6. **Smaller firms have a lower probability of survival, but those that survive tend to grow faster than larger firms. Among larger firms, growth rates are unrelated to past growth or to firm size.**

The first clause of this statement holds true because firms exit only by losing their last product line. Additionally, conditioning on not exiting preferentially selects for firms that have not lost products, meaning they have in net gained more products than average. This does not arise with large firms, where we are far away from exit through product loss.

One dimension where this model falls a bit short is in capturing the true extent of skewness in the observed firm size distribution, at least in the US. We might be able to approximate this using very low entry rates, but these would be implausibly so. Fundamentally, the resulting distribution does not have the right shape.

Luttmer (2011) notes that in order to achieve the proper level of skewness, one needs a different generating process as well as persistent heterogeneity in firm growth rates. In this model, firm growth rates are invariant to firm size, a result known as **Gibrat's Law**. This seems to be true in the data as well. However, growth rates are also not correlated over time for specific firms. This doesn't seem to be the case in the data, particularly when looking at firms that are now very large such as Walmart or GE.

Another object which we have sort of lost track of at this point is the distribution of firm level productivity. Recall from lecture 2 that the aggregate productivity index is

$$Q = \exp \left(\int_0^1 \log(q_j) dj \right)$$

Suppose that we define a normalized productivity $\hat{q}_j \equiv q_j/Q$. This helps us look only at the distribution, rather than the level which we know to be growing steadily over time. Over short time intervals, this evolves according to

$$\hat{q}_j(t + \Delta) = \begin{cases} \frac{\lambda q_j(t)}{1 + \Delta g} & \text{w.p. } \Delta\tau \\ \frac{q_j(t)}{1 + \Delta g} & \text{w.p. } 1 - \Delta\tau \end{cases}$$

Thus it follows an asymmetric random walk. The implication is then that the variance, both in terms of expectation and across product lines, increases without bound over time.

This is not a problem in this specific setting because of the log-log aggregation. However, for general elasticities of substitution, this would be an issue. What is required is some kind of resetting process. For instance, there might be a lower bound on productivity below which firms either cease producing or are dragged along with the economy. Alternatively, new or existing products could draw their productivity from some average of the current distribution, as is done in Acemoglu and Cao (2010).

4 Firm Heterogeneity

In addition to *ex post* heterogeneity in firms induced differential research outcomes, there can be other sources of firm growth and size variability. In particular, the above discussion treated all product lines as identical (both *ex ante* and *ex post*) as far sales, employment, and profitability go. This may very well not be the case.

Klette and Kortum (2004) introduce persistent differences at the firm level in terms of innovation step size. Certain firms simply produce larger innovations but must also employ more researchers in order to so. This yields persistent differences in innovation intensities, while avoiding persistent differences in growth rates (which would otherwise lead to larger firms being more research intensive

and profitable, a trend that we do not observe in the data).

Nonetheless, we do observe persistent differences in firm growth rates. Imagine, for example, the differences between a firm like Google and one like Xerox. However, squaring this with the fact that larger firms are on average no more profitable or research intensive than small firms requires a somewhat nuanced explanation. Roughly speaking though, one can imagine a firm life cycle in which some start out highly innovative, achieving rapid growth early on, but over time revert to innovative capacities more in line with or below the average firm.

It is difficult to write down a closed-firm solution to many models featuring further firm heterogeneity, but we can get a bit of analytic traction by going down that path. To do so, we need to move away from the unit elasticity ($\varepsilon = 1$) setting. Recall from the previous lecture that the production level chosen by a monopolist with productivity q_j is

$$x_j^M = \left[\left(\frac{\varepsilon - 1}{\varepsilon} \right) \frac{q_j}{w} \right]^\varepsilon Y$$

Now consider the limit pricing scenario. Here the firm will charge a price equal to the marginal cost of their competitor, namely $p_j = \lambda w/q_j$. This results in production of

$$x_j^L = \left[\left(\frac{1}{\lambda} \right) \frac{q_j}{w} \right]^\varepsilon Y$$

The question arises then, which of these is larger? The monopoly solution is valid only when the other firm is sure not to enter. Thus the firm will choose the maximum of the two. The condition for $x_j^M > x_j^L$ is simply

$$\lambda > \frac{\varepsilon}{\varepsilon - 1}$$

Let's assume that this holds so the firm will in fact always charge the monopoly

price. We could also define $\bar{\lambda} = \min\{\lambda, \varepsilon/(\varepsilon - 1)\}$ and proceed accordingly. Note that we are implicitly assuming stationary strategies here (and throughout). The question of equilibrium when allowing for multi-period strategies is quite a bit more complicated.

Now we can compute some observable outcomes at the product line level. In particular, the production level of the firm was computed to be

$$x_j = \left[\left(\frac{\varepsilon - 1}{\varepsilon} \right) \frac{q_j}{w} \right]^\varepsilon Y$$

Plugging this into the final goods production function, we find the identity

$$\frac{w}{Q} = \frac{\varepsilon - 1}{\varepsilon}$$

where Q is a productivity aggregate defined as

$$Q^{\varepsilon-1} = \int_0^1 q_j^{\varepsilon-1} dj$$

This will also satisfy $Y = QP$. Now let the relative productivity be $\hat{q}_j = q_j/Q$. By construction, we then have

$$1 = \int_0^1 \hat{q}_j^{\varepsilon-1} dj$$

We can further show that

$$\pi(\hat{q}) = \left(\frac{1}{\varepsilon} \right) \hat{q}^{\varepsilon-1} Y \quad \text{and} \quad w\ell(\hat{q}) = \left(\frac{\varepsilon - 1}{\varepsilon} \right) \hat{q}^{\varepsilon-1} Y$$

When a firm successfully innovates, it improves the productivity of a random product line by a factor λ .² The *relative* productivity from then on is falling

²In reality, firms are likely to be able to target their research efforts. This assumption is

over time as other products in the economy are improved upon. Lets write down a value function that can capture this dynamic, letting g be the growth rate of Q (and hence Y)

$$\begin{aligned} rV(\hat{q}) &= \pi(\hat{q}) + \max_x \{-wc(x) + x\bar{V}\} - \tau V(\hat{q}) - g\hat{q}V_{\hat{q}}(\hat{q}) + \dot{V} \\ &= \pi(\hat{q}) + \Omega - \tau V(\hat{q}) - gV_{\hat{q}}(\hat{q}) + \dot{V} \end{aligned}$$

where $\bar{V} = \mathbb{E}_{\hat{q}}[V(\hat{q})]$ is the expected return from innovation. Now suppose the value function takes the form

$$V(\hat{q}) = [A + B\hat{q}^{\varepsilon-1}] Y$$

This should satisfy

$$(r - g + \tau) [A + B\hat{q}^{\varepsilon-1}] = \left(\frac{1}{\varepsilon}\right) \hat{q}^{\varepsilon-1} + \tilde{\Omega} - g(\varepsilon - 1)B\hat{q}^{\varepsilon-1}$$

Equating term by term yields

$$A = \frac{\tilde{\Omega}}{r - g + \tau} \quad \text{and} \quad B = \frac{1/\varepsilon}{r - g + \tau + (\varepsilon - 1)g}$$

Note that $\tilde{\Omega}$ is a function of $\tilde{V} = \bar{V}/Y$. Thus there is one more layer to solve there. Assuming a constant elasticity research cost function, one can solve for \tilde{V} in closed form. From here the innovation rate x can be found. The remaining task is to calculate the growth rate

$$\begin{aligned} Q^{\varepsilon-1}(t + \Delta) &= \int_0^1 [\Delta\tau(\lambda q_j)^{\varepsilon-1} + (1 - \Delta\tau)q_j^{\varepsilon-1} dj] \\ &= Q^{\varepsilon-1}(t)(1 + \Delta\tau[\lambda^{\varepsilon-1} - 1]) \end{aligned}$$

made purely for tractability.

Thus the growth rate of Q is simply

$$g = \left[\frac{\lambda^{\varepsilon-1} - 1}{\varepsilon - 1} \right] \tau$$

At this point we run into issues of the existence of the distribution over \hat{q} , as discussed earlier. Nonetheless, we can write down flow equations for this. Let the distribution be denoted by $F(\cdot)$. This should then satisfy

$$F(\hat{q}, t + \Delta) = \Delta \tau F\left(\frac{\hat{q}(1 + \Delta g)}{\lambda}, t\right) + (1 - \Delta \tau) F(\hat{q}(1 + \Delta g), t)$$

Taking the limit as $\Delta \rightarrow 0$, we find a sort in inflow-outflow differential equation

$$g \hat{q} f(\hat{q}) = \tau [F(\hat{q}) - F(\hat{q}/\lambda)]$$

For a very detailed version of the above model, with the addition of an *ex ante* distribution of product line shares, see Lentz and Mortensen (2008). Note that in this work, the authors actually match the distribution of \hat{q} to the data rather than use the steady state implied by the model, which doesn't exist.

References

- D. Acemoglu and D. V. Cao. Innovation by entrants and incumbents. Technical report, National Bureau of Economic Research, 2010.
- E. J. Bartelsman and M. Doms. Understanding productivity: lessons from longitudinal microdata. *Journal of Economic literature*, pages 569–594, 2000.
- S. J. Davis, J. C. Haltiwanger, and S. Schuh. Job creation and destruction. *MIT Press Books*, 1, 1998.
- T. J. Klette and S. Kortum. Innovating firms and aggregate innovation. *Journal of Political Economy*, 112(5):986, 2004.

- R. Lentz and D. Mortensen. An empirical model of growth through product innovation. *Econometrica*, 76(6):1317–1373, 2008.
- E. G. Luttmer. On the mechanics of firm growth. *Review of Economic Studies*, 78(3):1042–1068, 2011.