Economics Program Working Paper Series

Exploring Innovation with Firm Level Data

Samuel Kortum

University of Chicago

January 2008

Part of the Supplemental Materials for
INNOVATION AND U.S. COMPETITIVENESS
The Conference Board report #R-1441-09-RR

About the Report: The Conference Board has recently undertaken a project on innovation and competitiveness, with funding from Microsoft Corporation. The goal of the project is to provide an overview of the current state of knowledge on the nature of innovation, and its role in stimulating economic growth and improved living standards in the U.S. The project draws on experts across the academic, corporate, and policy arenas, in addition to The Conference Board's own analysis, surveys, and focus groups of the business community. Such experts met in February 2007 to present and discuss various aspects of the innovation process and measurement thereof. Each presenter wrote a summary piece focusing on his respective area of expertise. These summary documents underpin the content in *Innovation and U.S. Competitiveness*; however the conclusions drawn are those of The Conference Board alone. These papers are retained for reference in The Conference Board Economics Program Working Paper Series.

EPWP #08 - 11

THE CONFERENCE BOARD

Economics Program
845 Third Avenue
New York, NY 10022-6679
Tel. 212-759-0900
www.conference-board.org/economics

Exploring Innovation with Firm Level Data[1]

Samuel Kortum

University of Chicago

[1] Paper prepared for the Conference Board's Workshop *Perspectives on U.S. Innovation and Competitiveness*, February 8 and 9, 2007. Contact author: Samuel Kortum.

Growth accounting is typically conducted at the aggregate national level. The major finding of growth accounting, a finding that has held up for 50 years, is that much of growth cannot be attributed to increases in the factors of production such as capital, labor, and labor quality. The remainder, Robert Solow's (1957) famous *residual*, turns out to be a substantial contributor to growth. It is naturally interpreted as the consequence of innovation and improvements in technology. But, that interpretation of the residual cries out for direct evidence about the sources of technological change in the economy, and prescriptions for how to keep this beneficial dynamic process going.

1 Innovation at the Firm Level

To dig deeper into innovation and technological change we are naturally led from the nation as the unit of observation to the firm as the unit of observation. Why not go even deeper down to the level of a plant or establishment? Because, the firm is the decision-making unit for innovation. This distinction between firms and establishments really matters because much Research and Development (R&D) takes place in large multi-establishment firms.

In part because it is the decision making unit, much of the data that we would like is available at the firm level. We observe innovative inputs, such as firm-level R&D, as well as innovative outputs, such as patents assigned to the firm, and new products developed by the firm. We also measure indirect outputs of innovation such as gains in firm productivity and size.

1.1 The Research Agenda

A major research agenda in economics, with Zvi Griliches deserving much of the credit for it, developed out of this beneficial combination of important questions and available micro data.

Much of the research was presented at the NBER, where Zvi Griliches orchestrated the NBER's Productivity Group. That was my upbringing as an economist.

The agenda involved using the micro data and econometrics to estimate relationships between (i) patents and R&D, (ii) productivity and R&D, (iii) firm value and either R&D or patents. The volume edited by Griliches (1984) is a good collection of such work.

The longer term agenda was to, aggregate back up from such firm-level findings to obtain a deeper understanding of the aggregate growth residual. That part lagged behind, but there was certainly progress in establishing some important relationships in the data.[2]

1.2 Firm-Level Findings

One of the most robust relationships is between patents and R&D at the firm level. Both across firms, and for given firms over time, patents follow R&D very closely. The paper by Hausman, Hall, and Griliches (1984) provides the basic results, and is also an excellent demonstration of the careful and innovative applied econometrics work that came out of this research agenda.

An older literature had focused on the relationship between productivity growth and R&D at the aggregate or industry level. With the new firm-level data it seemed that the relationship would be nailed down much better. If fact, while productivity and R&D are positively correlated across firms, it is very difficult to tease out a strong correlation from the within-firm dimension. Furthermore, results about the effect of R&D on productivity don't seem to be very robust. Griliches (1979) provides a nice overview of the early literature and sets out a framework for future work at the firm level. In his Presidential address before the American Economic Association, Griliches (1994) assesses what was achieved.

While the relationships above treat R&D spending as an exogenous variable, it was long understood by Griliches (1979) and others that R&D is a choice variable of the firm. It is therefore important to consider the properties and determinants of such investments.

R&D intensity, which is simply a firm's R&D as a fraction of its sales, turns out to be roughly independent of firm size (among firms that do any R&D). Another way to put it is that R&D scales up with the size of the firm. On the other hand, while R&D intensity is highly persistent within firms, it is highly skewed across firms. Some firms seem to survive quite nicely by spending almost nothing on R&D.

Not surprisingly, the NBER Productivity Group focused on firm productivity. But, elsewhere, and from long before, Herbert Simon and others (see Ijiri and Simon, 1977) were obsessed with more basic performance variables such as firm sales or employment. Their research established that firm growth is not strongly correlated with firm size among large firms, the so-called Gibrat's Law. Furthermore, the size distribution of firms is highly skewed. The work of Simon and others used stochastic models of growth to connect these two facts. See Luttmer (2007) for a modern analysis of this type.

2 The Challenge

While delving down to the level of the firm has revealed much about the process of innovation, serious problems remain. These problems are conceptual. The firm is a complicated object for economists to model. Technology is an unusual type of good. Productivity, measured at the firm level, is not what we thought it was.

2.1 Limitations of the Production Function

[2] Kortum (2004) provides a broader perspective on this agenda, including statements by others in the field.

We're trained to think of a firm as a production function, but it's not. We are tempted to apply at the firm level the production functions that served us so well in our aggregate analysis. The result has been a tension between our empirical analysis, our theoretical analysis, and reality.

Edith Penrose (1959) called them those "innovating, multiproduct, 'flesh-and-blood' organizations that businessmen call firms (p. 13)". She had it about right. Our analytic tools need to follow her vision. A problem with the production function approach is that it either puts no constraints on the size of the firm (under constant returns to scale) or else it suggests firms would never get very big (under decreasing returns). The Penrose vision was that firms could be of any size, but their growth at any point in time is constrained by the resources at their disposal.

2.2 Lessons from Endogenous Growth Theory

Another lesson came from Romer (1990), who reminded us that technology is non-rival. This property can't be ignored in applied work. In the analytic framework laid out in Griliches (1979), the production function displayed constant returns to scale in factors of production, R&D included. This specification was necessary so that Griliches could adopt a simple competitive market structure.

But, as Romer (1990) pointed out, non-rivalry of technology means that the production function should be constant returns to scale, not counting R&D. To be so, the production function needs to include all of the rival factors of production including management talent and such. But once that is done, the same technology can be used when all other factor (including the number of talented managers) are scaled up.

This logic forces us to work with an imperfectly competitive market structure. Then R&D investments are recouped out of the profits derived under imperfect competition. Romer (1990) went on to construct a very tractable model of endogenous technological change with imperfect competition.

Aghion and Howitt (1991) and Grossman and Helpman (1991) built on that structure to capture the force of creative destruction in their endogenous growth models. Their approach is the starting point for what I do below.

2.3 Elusive Firm-Level Productivity

Forcing an analogy between productivity for a country and productivity for a firm can lead to confusion.

A country's technological advances show up in productivity, the value of output per worker (or per unit of input). It is for this reason that growth accountants attributed residual growth to technological change.

A firm's technological advances, on the other hand, show up in firm size. The firm's value of output per worker need not be higher due to a technological advance. The firm will simply expand until its value of output per worker is equated to the wage plus any markup of price over marginal cost. Thus we may be on a hopeless quest in trying to uncover the firm-level productivity effects of R&D investments. Instead, we should be looking for how R&D investments allow the firm to enter new markets or to expand in other ways.

Technological advances are a win-win situation for a country. Technological advances are more of a zero-sum game among competing firms. The technological advance of one firm is likely to come at the expense of its competitors. Yet the nation as a whole gains even if the firm's competitor is hurt.

Getting a quantitative handle on the aggregate production function has been a corner-stone of modern macroeconomics. I do not think the same can be said for firm-level estimation of production functions. While estimating firm-level production functions is a small industry, it is not clear that such estimates get us any closer to answering questions about the sources of innovation. In what follows I will describe a model, developed with Tor Jakob Klette, that takes a very different approach, much closer to Penrose and Romer.

3 Innovative Firms and Aggregate Innovation

The model that we developed, Klette and Kortum (2004), tries to incorporate the firm-level findings of Griliches and others into an aggregate theory of growth. The model is highly stylized. It has been extended in several directions by Lentz and Mortenson (2005, 2007), who use it in their analysis of Danish firms. Luttmer (2007b) has developed it further to better match facts about the size distribution of firms.[3]

It has yet to be used for its original purpose: to connect firm-level observations on innovative activity with economy-wide issues. Here, I sketch out the basic idea of the model, with equations kept to a minimum.

3.1 Aggregate Setting

The economy consists of many differentiated products, each yielding a given amount of revenue. An innovation in this setting means a better product, an ι% improvement on what was available before. Head-to-head competition between firms leads to the better product displacing an existing one.

[3] Sekar (2007) extends the model along several dimensions to confront data on Chilean producers.

A steady stream of innovations, arriving at rate μ, leads to technological change, and hence growth in the standard of living, at rate $\mu\iota$. The aggregate model is essentially the Quality Ladders model of Grossman and Helpman (1991).

At this point it is useful to get a mental picture of the model. Imagine the computer game Tetris. The blocks come down from the top of the computer screen at rate μ. The blocks begin piling up from the bottom, representing technological change. Rather than the variety of shapes in Tetris, imagine each block as just a single square. Rather than the 10 columns of blocks in Tetris, imagine many columns, representing the many differentiated goods available in the economy.

3.2 A Firm

A firm consists of an integer n of products, a subset of all the products in the economy. By investing in R&D, the firm may transit to size $n + 1$. It's chance of moving up a notch is $I = G(R, n)$. This innovation production function combines R&D of the firm with the firm's knowledge capital n. The function G(., .) is constant returns to scale in its two arguments, with diminishing returns to R holding n fixed. Since I can be interpreted as the Poisson parameter generating patented inventions (as in Hausman, Hall and Griliches, 1984) the specification captures the relationship between patents and R&D.

A firm is subject to having another firm innovate on one of its products. If that happens, the firm transits to size $n - 1$. Since each of the firm's n products is subject to such an event, the overall hazard of the firm moving down a notch is μn.

Returning to the Tetris analogy, imagine many colors of blocks, with a firm consisting of a particular color. The size of a firms is the number of blocks of its color along the top of the pile of blocks (the number of columns occupied by that color). The firm shrinks if a different colored block lands on top of one of its blocks. The firm expands if one of its color of block lands on a new column.

3.3 Value of a Firm

Each of a firm's products yields revenue of 1 (e.g. $1 million) and profits of $\pi < 1$. The value of a size *n* firm is $V(n)$. At an interest rate *r*, the return on a size *n* firm is $rV(n)$ which represents profits on current products less R&D expenditures plus capital gains from new innovations less capital losses via innovations by competitors. This logic is captured in the Bellman equation:

$$rV(n) = \max_R \{\pi n - R + G(R,n)[V(n+1) - V(n)] - \mu n[V(n) - V(n-1)]\}$$

The solution to the Bellman equation is simply $V(n) = \upsilon n$, $I(n) = \lambda n$, and $R(n) = c(\lambda)n$, all proportional to *n*. As part of the solution we get expressions for the terms υ and λ, expressed as functions of r, μ, π, and the parameters of the innovation production function *G*. Note the prediction that R&D scales with firm size, consistent with observations of firm R&D investments.

The value of the firm decomposes into the value of its production division and the value of its research division. The value of the production division is simply the future profits that can be earned on the firms existing products before they become obsolete. The value of the research division comes from the value of knowledge capital in the firm that makes it possible to develop new innovations in the future.

With some manipulation, the solution captures Griliches' idea of a firm's stock of knowledge:

$$E[n_t \mid R_s] = \alpha \int_{t_0}^{t} e^{-\mu(t-s)} R_s \, ds$$

The appropriate stock of knowledge is n which is expected to equal the firm's past R&D depreciated at the rate at which new innovations come along (thus making existing ones obsolete).

3.4 Firm Dynamics

The size of a given firm follows a simple stochastic process. Starting from size $n = 1$, the size of a firm fans out into a geometric distribution with ever greater variance, conditional on the firm surviving. Notice that a big firm gets that way by a string of good luck in its innovative outcomes.

A firm can also be hit by a string of bad luck. The firm dies when it loses its last product and becomes size $n = 0$. The chance that the firm will have died approaches 1. A young or small firm is more likely to die.

An old firm is more likely to be large, hence less likely to die. Firm growth is independent of size, consistent with Gibrat's Law. This prediction requires that we do not condition on firm survival. The model is thus consistent with findings on the dynamics of firm size.

3.5 Innovators and Imitators

We assume that when they are born firms acquire a certain value of an innovative parameter. A large value of it allows the firm to make larger improvements, but the firm must spend more on R&D. If they can, firms will spend more on R&D to make big product improvements. Those that can't will spend little on R&D thus making marginal product improvements. Firms of either type grow at the same rate.

Big-step firms will appear more productive since they wield more market power. The model can thus capture the correlation between R&D and productivity in the cross-section. Yet, since all firms grow at the same rate there will still be no correlation between R&D intensity and size. The model, albeit with some special assumptions about the relationship between R&D cost and the size of innovations, is consistent with both the heterogeneity in research intensity as well as with the scaling of R&D with size for any given firm.

3.6 Micro to Macro

A powerful feature of this model, and one justification for keeping it so stylized, is the ease with which it aggregates from the level of individual firms to the aggregate national economy.

We can count up firms using the notation $M_n(t)$, the measure of size n firms at date t. Since there is a unit measure of goods in the economy, we have

$$\sum_{n=1}^{\infty} nM_n(t) = 1$$

It follows from the solution to the Bellman equation that the total number of innovations by incumbent firms is

$$\sum_{n} M_n(t) I(n) = \lambda$$

What's left is innovation by entrants:

$$\eta = \mu - \lambda$$

Entrants are launched with one new product, thus starting out as size $n = 1$. From there they grow or shrink like any other firm of size 1.

Eventually the economy settles down to a stationary size distribution of firms M_n. In an economy with more entry, there will be fewer very big firms. As in the work of Simon, the model connects entry and firm dynamics with the size distribution.

In general equilibrium, the rate of technological change $\mu\iota$ can be expressed as a function of the parameters of the model (in particular the parameters of the innovation production function *G*) thus completing the connection between firm-level observations and the residual from growth accounting.

4 Conclusion

Since Solow's "discovery" of the growth-accounting residual in the late 1950's, economists have attempted to estimate the economic forces driving technological change. Firm-level datasets are a valuable source of evidence for this purpose. But, the typical emphasis on estimating firm-level productivity, while a natural analog to growth accounting, is misguided. The problem is that a productive firm has an incentive to expand until its productivity advantage is dissipated. I've described one conceptual framework for firm-level analysis that redirects the emphasis away from firm-level productivity and toward firm innovation and growth through expansion into new product lines. This framework makes sense of the firm-level findings of Griliches and others while linking them to aggregate technological change.[4]

Recent work, most notably Corrado, Hulten, and Sichel (2005, 2006), has shown how to measure intangible investments to bring them into the growth accounting framework.[5] What I'd like to see is a marriage between that work at the aggregate level and an analytic approach along the lines of what I've laid out here at the firm level.

[4] The approach outlined here is admittedly overly stylized along many dimensions. For example, there is a continuum of firms. Work in the Industrial Organization literature (e.g. Ericson and Pakes (1995)) shows how to proceed with a small number of firms.
[5] That approach has recently been extended into the international domain by McGrattan and Prescott (2007).

References

Aghion, Philippe and Peter Howitt (1992), "A Model of Growth Through Creative Destruction," Econometrica, 60: 323-351.

Corrado,Carol,Charles Hulten,and Daniel Sichel (2005), "Measuring Capital and Technology: An Expanded Framework," in Carol Corrado, John Haltiwanger, and Daniel Sichel, eds. Measuring Capital in the New Economy, University of Chicago Press.

Corrado,Carol,Charles Hulten,and Daniel Sichel (2006), "Intangible Capital and Economic Growth," NBER Working Paper No. 11948.

Ericson, Richard and Ariel Pakes (1995), "Markov-Perfect Industry Dynamics: A Framework for Empirical Work," Review of Economic Studies, 62: 53-82.

Griliches, Zvi (1979), "Issues in Assessing the Contribution of R&D to Productivity Growth," Bell Journal of Economics, 10: 92-116.

Griliches, Zvi (1984), R&D, Patents, and Productivity, University of Chicago Press.

Griliches, Zvi (1994), "Productivity, R&D, and the Data Constraint," American Economic Review, 84: 1-23.

Grossman, Gene M. and Elhanan Helpman (1991), Innovation and Growth in the Global Economy. Cambridge, MA: MIT Press.

Hausman, Jerry, Bronwyn H. Hall, and Zvi Griliches (1984), "Econometric Models for Count Data with an Application to the Patents-R&D Relationship," Econometrica 52: 909-938.

Ijiri, Yuji and Herbert Simon (1977), Skew Distributions and the Sizes of Business Firms, Amsterdam: North-Holland.

Klette, Tor Jakob and Samuel Kortum (2004), "Innovating Firms and Aggregate Innovation," Journal of Political Economy, 112: 986-1018.

Kortum, Samuel (2004), "An R&D Roundtable," Economics of Innovation and New Technology, 13: 349-363.

Lentz, Rasmus and Dale Mortenson (2005), "Productivity Growth and Worker Reallocation," International Economic Review, 46: 731-751.

Lentz, Rasmus and Dale Mortenson (2007), "An Empirical Model of Growth Through Product Innovation," mimeo University of Wisconsin-Madison.

Luttmer, Erzo G.J. (2007a), "Selection, Growth, and the Size Distribution of Firms," Quarterly Journal of Economics, 122: 1103-1144.

Luttmer, Erzo G.J. (2007b), "On the Mechanics of Firm Growth," Minnesota Economics Research Report, No. 2007-04.

McGrattan, Ellen and Edward Prescott (2007), "Technology Capital and the U.S. Current Account," Federal Reserve Bank of Minneapolis Working Paper, No. 646.

Penrose, Edith. (1959), The Theory of the Growth in the Firm, Oxford: Blackwell.

Romer, Paul M. (1990), "Endogenous Technical Change," Journal of Political Economy, S71-S102.

Sekar, Murat (2007), "A Structural Model of Establishment and Industry Evolution: Evidence from Chile." mimeo, University of Minnesota.

Solow, Robert (1957), "Technical Change and the Aggregate Production Function," Review of Economic Studies, 39: 312-330.