Network to Study Productivity in Canada from a Firm-Level Perspective

Productivity at the Firm Level

Salvador Navarro and David Rivers

NSPCFP Working Paper 2013-01

Productivity at the Firm Level

Salvador Navarro and David Rivers

March 26, 2013

1 Introduction and Related Literature

Interest in the estimation of production functions and the measurement of productivity has a long history in economics. The concept of productivity is a measure of the efficiency of an economy, an industry, or a firm. In this paper we will focus the discussion on firm-level productivity, but many of the ideas presented apply to more aggregate notions as well.

The easiest way of understanding firm-level productivity is to consider two firms with the same level of input usage. If firm 1 produces more output than firm 2, it is considered to be more productive. This simple notion was notoriously difficult to formalize. The first notions of productivity were focused on single-input measures such as the amount of output per worker produced by a firm, as these measures could be computed from readily available data. The problem with this notion is easy to see if one considers a generic production function:

$$Y = G\left(K, L, M\right),$$

where Y is output, K is capital, L is labor and M is intermediate inputs. For example, consider labor productivity, which is measured as output per worker: $\frac{Y}{L} = \frac{G(K,L,M)}{L}$, and only requires data on output and labor input. As the right hand side of this expression makes clear, the levels of other inputs affect this productivity measure. A notion that better matches the intuitive definition we provide above is that of total factor productivity or TFP. In its simplest form, TFP can be measured by the term A in the following expression:

$$Y = F(K, L, M)\, A.$$

It measures changes in output that are directly unrelated to changes in all inputs used by the firm. Throughout this note, we will refer to A simply as productivity.

Early work on the analysis of production efficiency focused on relating aggregate measures of value added (gross production minus expenditures on intermediate inputs) to aggregate measures of capital and labor inputs. Discussions on the appropriateness of the approach arose soon after, including theoretical discussions of the empirical model, questions concerning data quality and reliability, and econometric issues related to identification and estimation. Researchers quickly recognized that both better data and better empirical techniques and approaches were needed. Over time, data collection methods improved, both in terms of quality and the level of detail. This allowed researchers not only to improve their estimates based on existing empirical models, but also allowed for the development of new and improved models that took advantage of the new data.

This process has iterated over the long history of production function/productivity studies as new and better data is collected either as a response to researchers identifying problems with existing data sources, or because of improvements in procedures, and/or the lifting of restrictions on statistical agencies. In turn, the new datasets have allowed for new questions to be answered and for new areas of research to be identified.

The early research in this area sought to use production function estimation as a test of marginal productivity theory, to analyze the competitiveness of labor markets, and later to measure changes in aggregate productivity. More recently, the literature has tried to answer a much broader set of questions, including recent papers that have examined patterns of entry and exit of firms, the effect of research and development on firm productivity, the effect of trade policies on domestic and foreign competition, the effects of prices on measures of productivity, the size distribution of firms and its determinants, and many others.

In what follows we will outline some of the key issues in the literature. We will also discuss how access to firm-level data, and in particular, detailed data on firm activities such as pricing, participation in foreign markets, and research and development (to name just a few), are crucial for obtaining an accurate understanding of the supply side of the economy and, as a result, for developing economic policy. Finally, we will point at additional questions that can also be answered with this approach using either existing datasets or datasets that can be created with relatively small investments.

2 Importance of Detailed Firm-Level Data

In this section we outline some of the reasons why firm-level data, and in particular detailed firm-level data, are important for understanding firm behavior and outcomes. We first discuss the value of firm-level data itself, and then move on to a discussion of the benefits of detailed firm-level information about prices, foreign market participation, R&D, etc.

2.1 Firm-Level Data

Data at the firm-level has several advantages compared to industry or economy-wide aggregate information. It allows for the potentially more accurate measurement of various aggregate objects of interest like industry (and country) average productivity and average capital intensities,[1] industry trends in input usage, etc. More importantly, it also enables one to generate disaggregate objects of interest, such as detailed measurements of productivity, the measurement of firm-level heterogeneity within industries (or countries), as well as to obtain unbiased estimates of the relationship between productivity and outcomes such as exporting, importing, wages, output prices, and more.

[1] In addition to averages, micro-data also allows for the measurement of other distributional measures such as quantiles, dispersion (e.g., variance, 90/10 ratios), etc.

2.1.1 Measurement of Productivity

Early work on the estimation of production functions and the measurement of productivity relied on aggregate data from national accounts and later incorporated industry-level aggregates based on Input-Output tables. The primary interest was in estimating the relationship between aggregate measures of value added to aggregate measures of capital and labor. This early work was soon attacked on several fronts. In particular, two important sources of bias were identified with this work: aggregation bias and simultaneity bias.

Researchers questioned the existence of an aggregate production function and pointed out problems with aggregating up firm-level production functions for use with aggregate data. While Input-Output tables are a helpful way of organizing economic data and looking at inter-dependencies across industries etc., they are less useful as a means of describing the production process of an industry. The reason is that the macroeconomic models that are used to aggregate firm-level variables to the industry-level rely on assumptions about individual firms' technology (such as fixed-proportions) that is unlikely to hold in most industries.

In addition, the problem of simultaneous equations bias was identified (Mendershausen (1938) and Marschak and Andrews (1944)). The common exercise of regressing value added on inputs, with the residual measuring productivity, was shown to produce biased estimates of both the production function parameters as well as of productivity due to the correlation between inputs and unobserved productivity.

In response to these concerns, the literature shifted towards the use of micro data instead, in particular using agricultural data (Tintner (1944) and Mundlak (1961)), as initially this was one of the few sources of micro data. This directly addressed some of the theoretical and aggregation issues, and provided some help with respect to simultaneity bias as well, as some components of productivity (climate, soil quality, topology) could be controlled for with fixed effects. As firm-level data on manufacturing became available, work shifted towards the study of firm-level production.

Index number methods that could take advantage of this newly available data were then developed (see Caves, Christensen and Diewert (1982) for a review). Index models can recover productivity by using simple relationships implied by the assumed profit maximization/cost minimization of firms in perfectly competitive markets. Specifically they derive a relationship between firm-level input shares, which can be observed, and output elasticities of those inputs, which are common parameters of interest. These methods, however, require (among other non-innocuous assumptions) that all inputs into the production function be bought (or rented) in spot markets, i.e., to be static and not subject to any frictions, an assumption that the literature has found questionable.

As panel data on firms inputs and output became available, modern microeconometric techniques that take advantage of repeated measurements like simple fixed/random effects (see Mundlak (1961) for an early example and Griliches and Mairesse (1998) for a summary) and dynamic panel data methods were introduced (Arellano and Bond (1991); Blundell and Bond (1998, 2000)). These panel data methods all rely on firm-level observations in order to be able to difference out firm-level productivity and control for simultaneity bias.

As this work continued, however, researchers noticed that the simultaneity problem was not fully mitigated, and that the common approach of using panel data techniques designed to help solve the problem often led to unrealistic results. A methodological innovation came in the work of Olley and Pakes (1996) who recognized that, by bringing in additional data (investment in physical capital) to proxy for the unobserved productivity of the firm, the simultaneity problem could be addressed.

The dependence of these methods on value-added measures of output has recently been scrutinized by Gandhi, Navarro and Rivers (2011), who propose a new method that combines index methods and dynamic panel methods. All of the methods described above take advantage of firm-level information to solve problems associated with obtaining unbiased estimates of productivity, and the development of these approaches has closely followed the introduction of new data that makes these innovations possible.

2.1.2 Firm Heterogeneity

In addition to helping researchers to address estimation bias due to aggregation and simultaneity, firm-level data also allows for the measurement of objects of interest that is not otherwise feasible. It can also be used to discover dimensions of heterogeneity that are masked by the use of aggregate data. Recently, the literature has emphasized the importance of this heterogeneity in several areas.

For example, there has been substantial research into questions related to the allocative efficiency of inputs across firms within an industry. Using plant-level manufacturing data for Chile, Pavcnik (2002) shows that trade liberalization is associated with large increases in productivity at the industry level. She shows that, as opposed to existing firms simply becoming more productive via some unspecified positive effect of trade, trade is associated with a shift of productive resources within an industry from less productive to more productive firms. This reallocation is what generates a substantial increase in industry (and economy-wide) productivity. Without firm-level data and corresponding estimates of productivity, it is not possible to empirically separate the reallocative effects due to changes in trade policy from direct increases in firm-level productivity.

In the macroeconomics literature, heterogeneity in firm-level productivity has been used to measure the magnitude of the misallocation of inputs across firms due to frictions, and the effect this misallocation has on aggregate productivity and growth. For example, Hsieh and Klenow (2009) use a standard model of industry equilibrium to derive the implication that revenue total factor productivity should be equalized across firms within each industry. Using firm-level data, they measure the within-industry dispersion of productivity in China and India and find that the size of misallocation in China and India is substantial, resulting in aggregate productivity losses of up to 50% as compared to the United States. Since the key to their empirical exercise relies on estimates of productivity dispersion within industries, aggregate data is not useful as it cannot capture firm-level heterogeneity. This work has also prompted a sizable literature focused on explaining the source of these frictions, which can also be much better understood by using firm-level data.

In international trade, heterogeneity across firms (often characterized as productivity) is a key component in many models explaining foreign market participation decisions such as exporting (Eaton and Kortum (2002) and Melitz (2003)) and outsourcing and off-shoring (Antras and Helpman (2004)). These papers were written partly as a consequence of early empirical work using firm-level data that suggested a positive relationship between firm size/productivity and trade (e.g., Bernard and Jensen (1999)). The results of these papers suggest that measures of firm-level heterogeneity are necessary for understanding patterns of trade for a given set of countries or industries.

Another example is the measurement of patterns of firm-level entry and exit within industries. When U.S. census data on manufacturing firms became available, some of the first studies to employ this data were interested in analyzing these patterns. For example, the influential study by Dunne, Roberts and Samuelson (1988) found that there are different types of entrants into an industry: new firms, existing firms that open new facilities, and firms that diversify their production using existing facilities. Each of these types of entrants have very different implications in terms of investment, employment, innovation, etc.

Dunne, Roberts and Samuelson (1988) also found that entry and exit rates are highly correlated across industry, which implies that net entry rates understate the overall level of turnover, and that industry-specific factors are key in understanding these differences. In addition, they document the large failure rate of new firms entering into an industry, a key determinant of market structure.

Firm-level data is also necessary for examining questions related to productivity spillovers across firms (see Lychagin, Pinkse, Slade and Reenen (2010) for an example). With aggregate data, it is not possible to uncover the extent to which innovations at one firm lead to innovations at other firms, and how this process depends on the distance between firms. This distance can represent many factors such as closeness in physical distance, the products produced, capital-intensity, technological efficiency, and others.

2.2 Detailed Data on Firms

While firm-level information about the inputs and output of a firm has many valuable uses, there are important questions that cannot be addressed without further information about the activities of the firms aside from basic input and output decisions. This includes information about the prices of output and inputs, foreign market participation decisions (exporting, importing intermediate inputs, outsourcing, and off-shoring), expenditures on research and development, product-specific information and product mix for multi-product firms, managerial practices, innovation efforts, and more. These data also have value for improving measurement of the production process, for understanding the determinants of productivity, and for generating an understanding of how these various firm-level decisions are inter-related.

2.2.1 Measurement

Initially, production function estimation was constrained by the data available to researchers from national accounts. This data included measures of value-added, labor, and capital, but not intermediate inputs (raw materials, energy, and services). As a result, researchers were forced to make strong assumptions on the way in which intermediate inputs enter the production process in order to deal with this data limitation. In the 1960's and 1970's, several papers pointed out problems associated with the use of value-added specifications of the production function (e.g., Sims (1969) and Arrow (1972)).

Despite these concerns, and the fact that data on intermediate inputs are now a standard variable in firm-level datasets, the use of value-added specifications has continued until the present. This is partially due to the inertia generated from earlier studies that were constrained by the available data and partly due to a set of results in the 1970's showing that under a set of assumptions, including that all input markets are perfectly competitive, value added could still be used to recover a scaled version of productivity (see Bruno (1978) and references therein).

This issue was recently revisited by Gandhi, Navarro and Rivers (2011) where they show that the value-added model is unsuitable for modern applications. They find that there are theoretical reasons why estimates of productivity using value added are likely to overestimate heterogeneity. They also provide evidence from manufacturing firms for both Chile and Colombia that value-added estimates of productivity and the production function can be substantially biased relative to a gross output model that incorporates intermediate inputs, even when properly rescaled.

Recent work in the literature has also emphasized the need to control for output prices (e.g., Klette and Griliches (1996); Foster, Haltiwanger and Syverson (2008); Katayama, Lu and Tybout (2009)), as the output of the firm is typically measured in revenues rather than quantities. Without separate data on prices and quantities, it is difficult to separate supply-side from demand-side factors affecting the output, and therefore productivity, of the firm. Without these more detailed data, researchers are forced to make strong, and often unrealistic, modeling assumptions in order to obtain estimates.

A common assumption to deal with unobserved output prices is that firms behave perfectly competitively, such that they all charge the same price, and thus an aggregate price index can be used to control for prices. This is a strong assumption, and recent work using data that contains actual output prices has found evidence not only for the presence of market power in many industries, but heterogeneous degrees of market power across firms as well (De Loecker and Warzynski (2011)). For example, Smeets and Warzynski (2013) account for output price heterogeneity in their study of Danish manufacturing firms. They show that, when they deflate revenues using a constructed firm level price index, their estimates of productivity are significantly different from estimates that only use industry-level deflators.

Lack of detailed data on input prices is another source of mismeasurement of productivity. In a recent working paper Ramanarayanan (2013) provides a compelling and simple argument as to why this is the case. Consider two identical firms that produce identical amounts of output with the same inputs, i.e., they are equally productive. Now assume that the first firm has access to cheaper intermediate inputs by, for example, importing them. This firm's expenditures on intermediate inputs will be lower than the expenditures of the firm buying the more expensive domestic inputs. If a common input price deflator is employed for both firms, our measure of "real" intermediate input usage will be higher for the second firm, as their expenditures are higher, even though both firms use the exact same amount of the physical input. As a consequence, the first firm will look more productive than the second firm as it will appear as if it was producing the same output with "less" intermediate inputs than the second firm. Research by Fox and Smeets (2011) has also shown that data on the price of inputs (in their case labor) can be a good way to measure unobserved differences in the productivity of workers across firms that would otherwise lead to biased estimates.

2.2.2 Determinants of Productivity

In addition to studying the effects of firm heterogeneity, firm-level data enables the study of the sources of firm heterogeneity. Despite the fact that there exists a substantial amount of work on measuring productivity and examining how it varies across many dimensions (across firms, time, sectors, countries, different policy regimes), little is known about how and why these differences arise.

In a recent paper, Syverson (2011) surveys the literature and identifies a long list of the potential determinants of heterogeneity in firm productivity. Understanding the determinants of productivity is important as productivity is believed to be a key factor for firm survival in the long-run and, on the macroeconomics side, it is considered the engine of growth of the economy. The ability to measure how productivity heterogeneity varies across industry, country, and time, allows for researchers to better understand the sources of this heterogeneity and to evaluate its potential effects. Since the set of relevant drivers of productivity growth is quite large and diverse, depending on which factors were found to be most important for a given industry or country, there could be quite different policy implications.

Without data on potential drivers of productivity, it is very difficult to learn how it evolves, and how policy can be designed to increase overall productivity in the economy.

On even a basic level, it is important to distinguish between factors that can potentially be affected by the firm itself–the quality of inputs, managerial practices, R&D, etc. and factors shaping the environment in which the firm operates–degree of competition, pressure from imports, access to foreign markets, etc.

An additional use for these data is that many empirical models for estimating productivity involve assumptions about the process by which productivity evolves. In particular, productivity is commonly assumed to evolve in a random, possibly serially-correlated, uncontrolled fashion. However, it is very likely that there are decisions under the control of the firm that can affect the growth of productivity.

Having data on firm-level decisions beyond input usage (e.g., exporting, investment in R&D, learning-by-doing) is crucial even at a descriptive level when one is trying to identify which of these activities may be important in determining the evolution of productivity. Furthermore, formal econometric analyses need to incorporate these decisions as part of the model in order to obtain valid estimates, even if the interest of the exercise is on other aspects of the firm.

Work on understanding the determinants of productivity has been performed on data from a variety of countries. However, to date, there is no consensus as to which factors are most important, and the impact that each of these factors can have on productivity. A key reason for this is the different quality and nature of the data employed in these studies. Furthermore, many of these studies employ data from developing countries that are subject to very different restrictions (technical, political geographical, etc) than those of developed countries. As a consequence, the identified key drivers of productivity in one country may not be the same for a different country.

2.2.3 Understanding Firm-Level Decisions

Recent work taking advantage of detailed firm-level data has discovered that many dimensions of firm heterogeneity are correlated. For example, Bernard and Jensen (1995, 1999) find that exporting firms are larger and more productive on average. However, as shown in Gandhi, Navarro and Rivers (2011), this productivity differential is not robust to the use of gross output as opposed to valued added measures of output. A similar result, that exporters do not look more productive than non-exporters is also found by Smeets and Warzynski (2013) when they control for output price heterogeneity. The sensitivity of these estimated export-productivity relations highlights the importance of having access to good, detailed firm-level data that can allow researchers to relax the strong assumptions otherwise required to obtain results.

The relationship between firm-level productivity and participation in import markets is more robust. Kasahara and Rodrigue (2008) find that importers look more productive than non-importers. A similar result is obtained by both Gandhi, Navarro and Rivers (2011) and Smeets and Warzynski (2013). However, while the former find that the relation is weakened when they use gross output as opposed to value added, the latter find that importers look even more productive once they account for output price heterogeneity. To our knowledge, no results where both gross output and firm specific prices are accounted for are available.

Other firm decisions seem to be highly correlated with productivity as well. Doraszelski and Jaumandreu (2006) find that firms with more capital and firms which invested in R&D in previous periods are both more productive and more likely to engage in R&D. De Loecker and Warzynski (2011) find that exporting firms charge higher markups on average, suggesting a relationship between market power and exporting decisions. Foster, Haltiwanger and Syverson (2008) use data on firm-level output prices and quantity to disentangle the supply and demand side factors in driving entry and exit patterns. They show that firms with higher productivity are more likely to charge lower prices, but that differences in demand are more important in explaining plant survival than are differences in productivity.

Uncovering these patterns has relied on detailed firm data on the various activities of the firms. However, the benefit of this information extends beyond establishing correlations in different dimensions of firm heterogeneity. It also allows for researchers to try to uncover the source of these correlations, and in doing so we can develop a clearer understanding of firm behavior. Little work has been done in this area.

For example, Aw, Roberts and Xu (2008) find that R&D expenditures and exporting are correlated. In principle, this relationship could be driven by a number of factors. It might be due to more productive firms being more likely to engage in both activities. However, given the richness of the data available to them, they are able to identify that the source of the correlation is due in part to complementarities between exporting and R&D. This implies that any policy designed to encourage firms to export, may also lead to increases in R&D, and the other way around. Complementarities like this one increase the potency of any policy targeted to just one of these activities. On the other hand, negative complementarities can also exist that can lead to unwanted negative effects. These also need to be considered when performing policy evaluation.

2.2.4 Equilibrium Models of Firm-Level Decisions

Most of the literature we have surveyed so far tries to make one or two key points in isolation from the rest of what constitutes a firm (e.g., exporting and productivity, R&D and productivity). A few attempts have been made to integrate these into a more general model of the firm that can potentially allow for more dimensions of heterogeneity. These papers also show that accounting for multiple sources of heterogeneity is crucial for explaining patterns of firm-level behavior.

Roberts, Xu, Fan and Zhang (2011) use micro data on both trade and production for Chinese footwear manufacturers in order to analyze export demand, pricing, productivity, and market participation by destination market for these producers. By combining data on value and quantity of Chinese footwear transactions with information on the manufacturing firms themselves, they develop a model that allows for firm demand and cost (productivity) heterogeneity in order to explain the firm's input demand decisions, whether to become an exporter and which destinations to serve. They use the model to disentangle whether the export performance of these firms is due to having low costs (high productivity and access to cheap inputs) or because they have invested in order to increase the quality of their products. Roberts, Xu, Fan and Zhang (2011) find that the latter explanation plays a more important role in their export performance. This finding is key, as it highlights that these firms are not as vulnerable to other low cost competitors, such as Vietnam, as they would be if their success was just based on being low-cost producers.

Eaton, Kortum and Kramarz (2011) exploit detailed data on the exports of French firms to compare different trade theories. They begin by documenting several new stylized facts: a) the number of French firms servicing a market increases with market size, b) sales distributions are similar across very different markets, and c) sales in France increase with the number of markets being serviced by the firm. They begin with a simple general equilibrium model of trade based on fixed costs of entry and find that it fails to match the data in several dimensions (including that in a typical market there are too many firms selling small amounts and that French firms that export sell "too much" in France).

They then extend the model by adding firm (productivity) and market (demand) specific unobservables. The extended model does a much better job of matching the data and, while they find that productivity is important in explaining the probability that certain firms become exporters. This is only possible because they model the interaction between firm-level and market-level unobservables.

Access to even more detailed data, or to additional data, can also be helpful in order to extend our understanding of firm level decisions as they relate to productivity. Crozet, Head and Mayer (2012) matches the French trade and firm-level datasets with expert assessments on wine quality in order to further disentangle firm productivity from the quality that a firm produces as well as unobservable demand shocks for the Champagne producing industry in France. They find that, at least for the industry they analyze, quality is a much stronger predictor of export success than physical productivity. Furthermore, they generalize the model by allowing for pricing-to-market and find little evidence of its importance.

3 Importance for Policy

An accurate and detailed understanding of how firms interact and how they make decisions about pricing, innovation, developing new markets, etc. is important information for policy-makers to be able to make informed decisions that have their intended effect. In order to illustrate the importance of detailed firm-level data for policy, we discuss three hypothetical examples in which this data is necessary for the design of optimal policy.

3.1 Example 1: Firm Entry and Exit

Consider a case in which one examined data on the number of firms within a given industry over time, and found that the number of firms changed very little. Without more detailed information, one might conclude that these industries are quite stable. In addition, suppose that, using a separate dataset on workers, one found that workers in that industry experienced higher than normal incidence of job loss. One might then conclude that, since the industry seems to be fairly stable in terms of firm churning, firms in this industry provide unstable jobs for workers. If policy-makers were interested in reducing the number of unemployment spells for workers, they might create policies designed to incentivize firms to reduce worker turnover.

However, access to more detailed firm-level data may tell a completely different story. Suppose that, while the number of firms in the industry is stable, firm-level data confirmed that there was substantial entry and exit of firms within this industry, but in roughly equal amounts. In this situation, one might conclude that the turnover for workers is related to firm turnover instead. Therefore, if one is interested in improving worker outcomes, the correct policy should be related to firm survival rather than worker retention.

Suppose further that one could actually match the worker and firm datasets and identify at which firm each worker was employed prior to unemployment. This would allow for a direct examination of which firms were the source of worker turnover. One could then look at the characteristics of those firms as well as of the environment in which they operate, and determine which policies are most likely to have the intended effects and also which ones are most likely to have unintended negative effects. For example, a policy oriented to reduce firm turnover, while potentially beneficial in terms of worker unemployment, may also reduce the incentive for firm entry, a possibility that can only be explored with access to detailed matched firm-worker data.

3.2 Example 2: Trade Policy: Exporting and Productivity

There is a large literature in international trade that finds a strong positive correlation between firm-level productivity and exporting decisions (e.g., Bernard and Jensen (1995)). However, there is some disagreement as to the source(s) of this correlation. One story is that more productive firms self-select into exporting (Eaton and Kortum (2002) and Melitz (2003)). Another (although not mutually exclusive) explanation is that exporting actually generates increases in firm-level productivity. The idea is that, by interacting with trading partners, technologies can be transferred between countries, leading to productivity increases for exporting firms. The first possibility has little to say in terms of trade policy implications, while the second one could imply that a policy that incetivizes exporting (e.g. subsidies), may increase productivity in the industry by allowing firms to "learn by exporting".

Disentangling these two stories requires detailed data on firm productivity and exporting behavior. The literature currently supports self-selection, but is divided as to whether there is learning-by-exporting. Some authors have suggested that this disagreement stems from the different countries (and industries) studied. Since the majority of the empirical evidence for learning-by-exporting comes from developing countries, one explanation is that the transfer of technology only operates in one direction, from more to less technologically advanced. Since different firms export to different countries, the potential for learning-by-exporting could vary substantially based on trading partners. Firm-level data on export destinations would allow researchers to examine the benefits of exporting behavior. Understanding the potential benefits of exporting is an important component in evaluating various trade policies.

3.3 Example 3: R&D Subsidies

Research and development is commonly viewed as one of the key drivers of aggregate productivity growth. There is also research that suggests that R&D by one firm can have important spillover effects for the productivity of other firms (Griliches (1979) and Scherer (1982)). These spillovers create potentially large externalities that are not internalized by the firm performing the R&D. This can lead to underinvestment in R&D, and generates a role for potential government intervention. By incentivizing the R&D expenditures of some firms, the government can help reduce the effect of the externality and improve overall productivity growth.

While the benefits of R&D policies can have large effects, they can also be costly, making it important that they are carefully targeted. Research has found that R&D expenditures vary substantially across both industries and firms, and that the rate of return to R&D is correlated with firm size (e.g., Doraszelski and Jaumandreu (2006)). To the extent that the gains to R&D expenditures are correlated with other observable characteristics of firms, this should be taken into account to maximize the potential gains from the policy. Targeted subsidies to those firms most likely to benefit directly from the policy, rather than a common subsidy across all firms, could have substantially larger overall effects.

The process by which R&D enhances productivity and how this spills over to other firms is still an open question. More detailed data on what R&D expenditures are used for (e.g., hiring research workers versus purchasing research equipment), whether the expenditures are targeted to improve process or product innovation, and whether R&D is designed to create innovations specific to the firm or generally valuable to firms, would be very useful for increasing the efficacy of any related policies. Notice that this also opens the door for unintended consequences of policies unrelated to R&D. For example, a policy unrelated to R&D that negatively affects firms more likely to engage in R&D would not only affect the firms directly involved in it, but could also have an effect for the entire industry via the (now negative) spillover effects.

4 Conclusion

We have outlined several ways in which access to detailed firm-level data improves our measurement of productivity, our understanding of what drives productivity and how it affects the decisions of firms, and how together this information can aid in designing optimal policy. We have mentioned only some of the uses of these data and the reasons why this information is important for both researchers and policy-makers.

The benefits of these data are widely appreciated in the literature. This is reflected in the fact that we currently know quite a lot about firms and their activities in Colombia, Slovenia, Chile, Spain, Denmark, and India where the types of datasets described above are being made available to researchers. Without access to Canadian firm-level data, we can attempt to derive answers to questions based on more aggregate and less detailed data. However, experience has shown us that one should be hesitant to use these results, as they are subject to many sources of bias and produce at best an incomplete picture. Consequently, we are left attempting to extrapolate the results from these other countries to Canada. While this approach may be suitable for some questions, it is unlikely that the effect of trade policies in India, for example, are likely to mirror the experiences of Canadian firms.

Data availability also generates free labor for developing datasets, improving their quality, merging existing or newly created datasets, and providing recommendations for future data collection, as researchers (and their graduate students) are often happy to provide these services in the course of their work. In addition, these datasets provide unique opportunities for the training of graduate students. The ability to work directly with rich micro-data allows students to acquire valuable empirical and analytical skills, many of which can only be developed through learning-by-doing and interacting with researchers with valuable experience based on a detailed understanding of the data and policy questions of interest. Furthermore, providing researchers with access to these data will promote the exchange of ideas, state of the art models and empirical techniques, as well as detailed information about the data and collection methods between academic and government researchers.

In addition, without access to these data, researchers who would otherwise prefer to study Canadian firms and markets, are forced to focus on other countries where the necessary data are readily available, enriching the understanding of those economies and providing free analysis for their policy-makers, at the expense of Canada. The type of data employed in the studies we surveyed in this paper is only available for a handful of countries, and in many countries it does not even exist.

Fortunately, in the case of Canada, this kind of data either already exists or can be relatively easily and cheaply developed by merging and linking existing datasets available via Statistics Canada. The missing component is making these data more accessible to researchers. More importantly, the list of potentially available data via CEDR extends beyond the data used in the studies surveyed, providing a golden opportunity for better understanding answers to existing questions, as well as for the analysis of new ones.

References

Antras, Pol and Elhanan Helpman, "Global Sourcing," *Journal of Political Economy*, 2004, *112* (3), 552–580.

Arellano, Manuel and Stephen Bond, "Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations," *Review of Economic Studies*, April 1991, *58* (2), 277–97.

Arrow, Kenneth J., *The Measurement of Real Value Added*, Institute for Mathematical Studies in the Social Sciences, 1972.

Aw, Bee Yan, Mark J Roberts, and Daniel Yi Xu, "R&D investments, exporting, and the evolution of firm productivity," *The American Economic Review*, 2008, *98* (2), 451–456.

Bernard, Andrew B. and J. Bradford Jensen, "Exporters, Jobs, and Wages in U.S. Manufacturing: 1976-1987," *Brookings Papers on Economic Activity. Microeconomics*, 1995, pp. 67–119.

_ **and** _ , "Exceptional Exporter Performance: Cause, Effect or Both?," *Journal of International Economics*, 1999, *47* (1), 1–25.

Blundell, Richard and Stephen Bond, "Initial Conditions and Moment Restrictions in Dynamic Panel Data Models," *Journal of Econometrics*, August 1998, *87* (1), 115–143.

_ **and** _ , "GMM Estimation with Persistent Panel Data: An Application to Production Functions," *Econometric Reviews*, 2000, *19* (3), 321–340.

Bruno, Michael, "Duality, Intermediate Inputs, and Value-Added," in M. Fuss and McFadden D., eds., *Production Economics: A Dual Approach to Theory and Practice*, Vol. 2, Amsterdam: North-Holland, 1978, chapter 1.

Caves, Douglas W., Laurits R. Christensen, and W. Erwin Diewert, "The Economic Theory of Index Numbers and the Measurement of Input, Output, and Productivity," *Econometrica*, November 1982, *50* (6), 1393–1414.

Crozet, Matthieu, Keith Head, and Terry Mayer, "Quality Sorting and Trade: Firm-level Evidence for French Wine," *Review of Economic Studies*, 2012, *79*, 609–644.

Doraszelski, U. and J. Jaumandreu, "R&D and Productivity: Estimating Production Functions When Productivity Is Endogenous," Technical Working Paper, Harvard University 2006.

Dunne, Timothy, Mark J Roberts, and Larry Samuelson, "Patterns of firm entry and exit in US manufacturing industries," *The RAND Journal of Economics*, 1988, pp. 495–515.

Eaton, Jonathan and Samuel Kortum, "Technology, geography, and trade," *Econometrica*, 2002, *70* (5), 1741–1779.

_ , _ , and Francis Kramarz, "An Anatomy of International Trade: Evidence from French Firms," *Econometrica*, September 2011, *79* (5), 1453–1498.

Foster, Lucia, John Haltiwanger, and Chad Syverson, "Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?," *American Economic Review*, March 2008, *98* (1), 394–425.

Fox, Jeremy and Valérie Smeets, "Does Input Quality Drive Measured Differences in Firm Productivity?," *International Economic Review*, November 2011, *52.*

Gandhi, Amit, Salvador Navarro, and David Rivers, "On the identification of production functions: How heterogeneous is productivity?," *University of Western Ontario, CIBC Centre for Human Capital and Productivity Working Papers*, 2011.

Griliches, Zvi, "Issues in Assessing the Contribution of Research and Development to Productivity Growth," *Bell Journal of Economics*, Spring 1979, *10* (1), 92–116.

— **and Jacques Mairesse**, "Production Functions: The Search for Identification," in "Econometrics and Economic Theory in the Twentieth Century: The Ragnar Frisch Centennial Symposium," New York: Cambridge University Press, 1998, pp. 169–203.

Hsieh, Chang-Tai and Peter J. Klenow, "Misallocation and Manufacturing TFP in China and India," *Quarterly Journal of Economics*, 2009, *124* (4), 1403–1448.

Kasahara, Hiroyuki and Joel Rodrigue, "Does the Use of Imported Intermediates Increase Productivity? Plant-level Evidence," *The Journal of Development Economics*, 2008, *87* (1), 106–118.

Katayama, Hajime, Shihua Lu, and James R Tybout, "Firm-level productivity studies: illusions and a solution," *International Journal of Industrial Organization*, 2009, *27* (3), 403–413.

Klette, Tor Jacob and Zvi Griliches, "The Inconsistency of Common Scale Estimators When Output Prices are Unobserved and Endogenous," *Journal of Applied Econometrics*, July 1996, *11* (4), 343–361.

Loecker, J De and F Warzynski, "Markups and Firm-level Export Status.(Princeton University, Unpublished)," 2011.

Lychagin, Sergey, Joris Pinkse, Margaret E. Slade, and John Van Reenen, "Spillovers in Space: Does Geography Matter?," Technical Report, NBER Working Papers 2010.

Marschak, Jacob and William H. Andrews, "Random Simultaneous Equations and the Theory of Production," *Econometrica*, 1944, *12* (3-4), 143–205.

Melitz, Marc J, "The impact of trade on intra-industry reallocations and aggregate industry pr.," *Econometrica*, 2003, *71* (6), 1695–1725.

Mendershausen, Horst, “On the significance of Professor Douglas’ production function,” *Econometrica: Journal of the Econometric Society*, 1938, pp. 143–153.

Mundlak, Yair, “Empirical production function free of management bias,” *Journal of Farm Economics*, 1961, *43* (1), 44–56.

Olley, G. Steven and Ariel Pakes, “The Dynamics of Productivity in the Telecommunications Equipment Industry,” *Econometrica*, November 1996, *64* (6), 1263–1297.

Pavcnik, Nina, “Trade Liberalization Exit and Productivity Improvements: Evidence from Chilean Plants,” *Review of Economic Studies*, 2002, *69* (1), 245–276.

Ramanarayanan, Ananth, “Imported inputs and the gains from trade,” 2013. Unpublished Manuscript, The University of Western Ontario.

Roberts, Mark J, Daniel Yi Xu, Xiaoyan Fan, and Shengxing Zhang, “A structural model of demand, cost and export market selection for Chinese footwear producers,” 2011. Unpublished Manuscript, Duke University.

Scherer, Frederic M, “Inter-industry technology flows and productivity growth,” *The Review of Economics and Statistics*, 1982, *64* (4), 627–634.

Sims, Christopher A., “Theoretical Basis for a Double Deflated Index of Real Value Added,” *The Review of Economics and Statistics*, November 1969, *51* (4), 470–471.

Smeets, Valérie and Frederic Warzynski, “Estimating productivity with multi-product firms, pricing heterogeneity and the role of international trade,” 2013. Journal of International Economics, Forthcoming.

Syverson, Chad, “What Determines Productivity?,” *Journal of Economic Literature*, June 2011, *49* (2), 326–65.

Tintner, Gerhard, “A note on the derivation of production functions from farm records,” *Econometrica, Journal of the Econometric Society*, 1944, pp. 26–34.