What Determines Productivity at the Micro Level?*

Chad Syverson
University of Chicago Booth School of Business
and National Bureau of Economic Research
chad.syverson@chicagobooth.edu

August 2009

* I thank Roger Gordon, John Haltiwanger, and anonymous referees for helpful comments. Financial support from the NSF (SES-0519062 and SES-0820307) is gratefully acknowledged.

1. Introduction

Thanks to the massive infusion of detailed production activity data into economic study over the past couple of decades, researchers in many fields have learned a great deal about how firms turn inputs into outputs. Productivity, the efficiency with which this conversion occurs, has been a topic of particular interest. The particulars of these studies have varied depending on the researchers' specific interests, but there is a common thread. They have, virtually without exception, documented enormous and persistent measured productivity differences across producers, even within narrowly defined industries.

The magnitudes involved are striking. 90^{th} percentile to 10^{th} percentile total factor productivity (TFP) ratios of 3 to 1 and more are not unusual across producers within 4-digit SIC (6-digit NAICS) industries. For example, I found in Syverson (2004a) that in the U.S. Census of Manufactures, the average 90-10 logged TFP difference across plants within 4-digit industries was 0.651. This corresponds to a TFP ratio of $e^{0.651} = 1.92$. To emphasize just what this number implies, it says that the plant at the 90^{th} percentile of the productivity distribution makes almost *twice* as much output with the *same measured inputs* as the 10^{th} percentile plant. This is the average 90-10 range. The range's standard deviation across 4-digit manufacturing industries is 0.173, so many industries see even larger productivity differences among their producers. U.S. manufacturing is not exceptional in terms of productivity dispersion. Indeed, if anything, it is small relative to the productivity variation observed elsewhere. Hsieh and Klenow (2009), for example, find even larger productivity differences in China and India: with 90-10 TFP ratios of over 5:1 in China and India.[2]

These productivity differences across producers are not fleeting, either. Regressing a producer's current TFP on its one-year-lagged TFP yields autoregressive coefficients on the order of 0.6 to 0.8 (see, e.g., Ábrahám and White (2007) and Foster, Haltiwanger, and Syverson (2008)). Put simply, some producers seem to have figured out their business (or at least are on

---

[2] These figures are for revenue-based productivity measures; i.e., where output is measured using plant revenues (deflated across years using industry-specific price indexes). TFP measures that use physical quantities as output measures rather than revenues actually exhibit even *more* variation than do revenue-based measures, as documented in Foster, Haltiwanger, and Syverson (2008). Hsieh and Klenow (2009) also find greater productivity dispersion in their TFP measures that use quantity proxies to measure output (actual physical quantities are not available for most producers in their data). Even though it is only a component of revenue-based TFP (the other being the producer's average price), quantity-based TFP can be more dispersed because it tends to be negatively correlated with prices. That is, more efficient producers sell at lower prices. Thus revenue-based productivity measures, which combine quantity-based productivity and prices, tend to understate the variation in producers' physical efficiencies.

their way), while others are woefully lacking. Far more than bragging rights are at stake here: another robust finding in the literature—virtually invariant to country, time period, or industry—is that higher productivity producers are more likely to grow and survive than their less efficient industry competitors. Productivity is quite literally a matter of survival for businesses.

1.1. How Producer-Level Productivity Variation and Persistence Has Influenced Research

The discovery of the ubiquity of large and persistent productivity differences has shaped research agendas in a number of fields. Here are some examples of this influence (though by no means is it meant to be a comprehensive accounting). They speak to the breadth of the impact that answers to this paper's title question would have.

Macroeconomists are dissecting aggregate productivity growth—the source of almost all per capita income differences across countries—into various micro-components, with the intent of better understanding the sources of such growth. Foster, Halitwanger, and Krizan (2001), for example, overview the substantial role of reallocations of economic activity toward higher productivity producers (both among existing plants and through entry and exit) in explaining aggregate productivity growth. Hsieh and Klenow (2009) ask how much larger the Chinese and Indian economies would be if they achieved the same efficiency in allocating inputs across production units as does the U.S. Models of economic fluctuations driven by productivity shocks are increasingly being enriched to account for micro-level patterns, and are estimated and tested using plant- or firm-level productivity data rather than aggregates. Micro productivity data have also been brought to bear on issues of long-run growth, income convergence, and technology spillovers. They offer a level of resolution unattainable with aggregated data.

In industrial organization, research has linked productivity levels to a number of features of technology, demand, and market structure. Examples include the effect of competition (Syverson (2004b), Schmitz (2005)), the size of sunk costs (Collard-Wexler (2008)), and the interaction of product market rivalry and technology spillovers (Bloom, Schankerman, and Van Reenen (2007)). Another line of study has looked at the interaction of firms' organizational structures with productivity levels (e.g., Maksimovic and Phillips (2002), and Schoar (2002), and Hortaçsu and Syverson (2007, 2009)).

Labor economists have explored the importance of workers' human capital in explaining productivity differences (Abowd et al. (2003), Fox and Smeets (2009)), and the productivity effects of incentive pay (Lazear (2000)), other various human resources practices (Ichniowski and Shaw (2003)), and managerial talent and practices (Bloom and Van Reenen (2007)). There has also been a focus on the role of productivity-driven reallocation on labor market dynamics via job creation and destruction (Haltiwanger, Scarpetta, and Schweiger (2008)).

Perhaps in no other field have the productivity dispersion patterns noted above had a greater influence on the trajectory of the research agenda than in the trade literature. Theoretical frameworks using heterogeneous-productivity firms like Melitz (2003) and Eaton and Kortum (2003) are now the dominant conceptual lenses through which economists view trade impacts. In these models, the trade impacts vary across producers, and depend on their productivity levels in particular. Aggregate productivity gains come from improved selection and heighted competition that trade brings. A multitude of empirical studies have accompanied and been spurred by these theories (e.g., Pavcnik (2002); Bernard, Jensen, and Schott (2006); and Verhoogen (2008)). They have confirmed many of the predicted patterns, and raised questions of their own.

1.2. The Question of "Why?"

Given the important role that productivity differences play in these disparate literatures, the facts above raise obvious and crucial questions. *Why* do firms (or factories, stores, offices, or even individual production lines, for that matter) differ so much in their abilities to convert inputs into output? Is it dumb luck, or instead something—or many things—more systematic? Can producers control the factors that influence productivity, or are they purely external products of the operating environment? What supports such large productivity differences in equilibrium?

A decade ago, when Bartelsman and Doms (2000) penned the first survey of the micro-data productivity literature for this journal, researchers were just beginning to ask the "Why?" question. Much of the work to that point had focused on establishing facts like those above—the "What?" of productivity dispersion. Since then, the literature has focused more intensely on the reasons why productivity levels are so different across businesses. There has definitely been progress. But we've also learned more about what we *don't* know, and this is guiding the ways in which the productivity literature will be moving. This article is meant to be a guide to and comment on this research.

I begin by setting some boundaries. I have to. A comprehensive overview of micro-founded productivity research is neither possible in this format nor desirable. There are simply too many studies to allow adequate coverage of each. I will focus on empirical work. This is not because I view it as more important than theory. Rather, it affords a deeper coverage of this important facet of a giant literature, and it better reflects my expertise. That said, I will occasionally bring up specific theoretical work with particularly close ties to the empirical issues discussed. Furthermore, for obvious reasons, I will focus on research that has been done since the Bartelsman and Doms (2000) was written.

Even within these boundaries, there are more studies than can be satisfactorily described individually. I see this article's role as filtering the broader lessons of the literature through the lens of a subset of key studies. The papers I focus on here are not necessarily chosen because they are the first or only good work on their subject matter, but rather because they had an archetypal quality that lets me weave a narrative of the literature. I urge readers whose interests have been piqued to more intensively explore the relevant literatures. There is far more to be learned than I can convey here.

A disclaimer: some of my discussion contains elements of commentary. These opinions are mine alone and may or may not be the consensus of researchers in the field.

I organize this article as follows. The next section sketches the conceptual background: what productivity is, and how it is often measured in practice. Section 3 looks at influences on productivity that operate primarily within the business. This can be at the firm level, plant level, or even on specific processes within the firm. Many of these influences may potentially be under the control of the economic actors inside the business. In other words, they may be "levers" that management or others have available to impact productivity. Section 4 focuses on the interaction of producers' productivity levels and the markets in which they operate. These are elements of businesses' external environments that can affect productivity levels. This impact might not always be direct, but they can induce producers pull some of the levers discussed in Section 3, indirectly influencing observed productivity levels in the process. They may also be factors that limit the amount of productivity dispersion that can be sustained in equilibrium, and thus influence observed productivity differences through that channel. Section 5 discusses what I see as the big questions about business-level productivity patterns that still need to be answered. A short section that follows this concludes.

2. Productivity—What It Is, and How It's Measured

This section briefly reviews what productivity is conceptually and how it is measured in practice. Deeper discussions on the theory of productivity indexes can be found in Caves, Christensen, and Diewert (1982) and the references therein. More detail on measurement issues can be found in the large literature on the subject; see, for example, Griliches and Mairesse (1998), Olley and Pakes (1996), and Levinsohn and Petrin (2003).

Productivity in Concept. Simply put, productivity is efficiency in production: how much output is obtained from a given set of inputs. As such, it is typically expressed as an output-input ratio. Single-factor productivity measures reflect units of output produced per unit of a particular input. Labor productivity is the most common measure of this type, though occasionally capital or even materials productivity measures are used. Of course, single-factor productivity levels will be affected by the intensity with which the excluded inputs are used. Two producers may have quite different labor productivity levels even though they use the same production technology, if one happens to use capital much more intensively (say because they face different factor prices).

Because of this, researchers often use a productivity concept that is invariant to the intensity of use of observable factor inputs. This measure is called total factor productivity, or TFP (it's also sometimes called multifactor productivity, MFP). Conceptually, TFP differences reflect shifts in the isoquants of a production function: variation in output produced from a fixed set of inputs. Higher-TFP producers will produce greater amounts of output with the same set of observable inputs than lower-TFP businesses, and hence have isoquants that are shifted up and to the right. Factor price variation that drives factor intensity differences does not affect TFP, because it induces shifts *along* isoquants rather than shifts *in* isoquants.

TFP is most easily seen in the often-used formulation of a production function where output is the product of a function of observable inputs and a factor-neutral shifter:

$$Y_t = A_t F\left(K_t, L_t, M_t\right),$$

where Y_t is output, $F(\cdot)$ is a function of observable inputs capital Y_t, labor L_t, and intermediate materials M_t, and A_t is the factor-neutral shifter. In this type of formulation, TFP is A_t. It captures variations in output not explained by shifts in observable inputs (which act through $F(\cdot)$).[3]

Productivity is, at its heart, a residual. As with all residuals, it is in some ways a measure of our ignorance: it is the variation in output that cannot be explained based on observable inputs. So it's fair to interpret the work discussed in this survey as an attempt to "put a face on" that residual (or more accurately, "put faces on," given the multiple sources of productivity variation). The literature has made progress when it can explain systematic influences on output across production units that do not come from changes in observable inputs like standard labor or capital measures.

Measuring Productivity. While productivity is straightforward in concept, a host of measurement issues arise when constructing productivity measures from actual production data.

The first set of issues regards the output measure. Many businesses produce more than one output. Should these be aggregated to a single output measure, and how if so? Further, even detailed producer microdata do not typically contain measures of output quantities. Revenues are typically observed instead. The standard approach has been to use revenues (deflated to a common year's real values using price deflator series) and hope for the best. This is not driven by apathy toward mismeasurement but rather the realities of data limitations. Recent work has begun to dig deeper into the consequences of assuming single-product producers and using revenue to measure output. I'll discuss this more below. In the mean time, I will go forward for the time being assuming deflated revenues accurately reflect the establishment's or firm's output.

The second set of measurement issues regards inputs. For labor, there is the choice of whether to use number of employees, employee-hours, or some quality-adjusted labor measure (the wage bill is often used in this role, based on the notion that wages capture marginal products of heterogeneous labor units). Capital is typically measured using the establishment or firm's book value of its capital stock. This raises several questions. How good of a proxy is capital stock for the flow of capital services? Should the stock be simply the producer's reported book

[3] I use a multiplicatively separable technology shift to make exposition easy, but TFP can be extracted from a general time-varying production function $Y_t = G_t(A_t,K_t,L_t,M_t)$. Totally differentiating this production function gives:

$$dY_t = \frac{\partial G}{\partial A} dA_t + \frac{\partial G}{\partial K} dK_t + \frac{\partial G}{\partial L} dL_t + \frac{\partial G}{\partial M} dM_t .$$

Without loss of generality, we can choose units to normalize $\partial G/\partial A = 1$. Thus when observed inputs are fixed ($dK_t = dL_t = dM_t = 0$), differential shifts in TFP, dA_t, create changes in output dY_t.

value, and what are the deflators? Or should the stock be constructed using observed investments and the perpetual inventory method (and what to assume about depreciation)? With regard to measuring intermediate materials, there is an issue similar to the revenue-as-output matter above: typically all that is available are the producer's total expenditures on inputs, not quantities. And more fundamentally, how should intermediate inputs be handled? Should one estimate a gross output production function and include intermediate inputs directly, or should they simply be subtracted from output so as to deal with a value-added production function?

The third set of measurement concerns involves aggregating multiple inputs in a TFP measure. As described above, TFP differences reflect shifts in output while holding inputs constant. To construct the output-input ratio that measures TFP, a researcher must weight the individual inputs appropriately when constructing a single-dimensional input index. The correct weighting is easiest to see when the production function is Cobb-Douglas:

$$TFP_t = A_t = \frac{Y_t}{K_t^{\alpha_k} L_t^{\alpha_l} M_t^{\alpha_m}}.$$

In this case, the inputs are aggregated by taking the exponent of each factor to its respective output elasticity. It turns out that this holds more generally as a first-order approximation to any production function. The input index in the TFP denominator can be constructed similarly for general production functions.[4]

Even after solving the problem of how to construct the index, however, one must determine how the output elasticities $\alpha_j, j \in \{k,l,m\}$ are measured. Several approaches are common in the literature. One is to build upon assumptions of cost-minimization to construct the elasticities directly from observed production data. A cost-minimizing producer will equate an input's output elasticity with the product of that input's cost share and the scale elasticity. If cost shares can be measured (obtaining capital costs are usually the practical sticking point here), and the scale elasticity either estimated or assumed, then the output elasticities α_j can be directly constructed. If a researcher is willing to make some additional though not innocuous assumptions (namely, perfect competition and constant returns to scale), then the elasticities equal the inputs' payments shares of revenues. This makes constructing the α_j simple.

[4] While Cobb-Douglas-style approaches are probably the most common in the literature, many researchers also use the translog form (see Caves, Christensen, and Diewert 1982), which is a second-order approximation to general production functions, and as such is more flexible, though more demanding of the data. There is also an entirely nonparametric approach called data envelopment analysis that is used in certain, somewhat distinct circles of the literature. See Cooper, Seiford, and Tone (2006) for an overview of DEA methods.

Materials' and labor's shares are typically straightforward to collect with the wage bill and materials expenditures data at hand, and capital's share can be constructed as the residual, obviating the need for capital cost measures.

A separate approach is to estimate the elasticities *α* by estimating the production function. In this case, (logged) TFP is simply thet estimated sum of the constant and the residual. In the Cobb-Douglas case (which again, recall, is a first-order approximation to more general technologies), the estimated equation is:

$$\ln Y_t = \alpha_0 + \alpha_k \ln K_t + \alpha_l \ln L_t + \alpha_m \ln M_t + \omega_t \,.$$

Hence the TFP estimate would be $\hat{\alpha}_0 + \hat{\omega}_t$, where the first term is common across production units in the sample (typically the technology is estimated at the industry level), and the second is idiosyncratic to a particular producer.

This approach raises econometric issues. As first pointed out by Marschak and Andrews (1944), input choices are likely to be correlated with the producer's productivity ω_t (more efficient producers are, all else equal, likely to hire more inputs). There is also potential selection bias when a panel is used, since less efficient producers (those with low ω_t) are more likely to exit from the sample. (As will be discussed below, the positive correlation between productivity and survival is one of the most robust findings in the literature.) A substantial literature has arisen to address these issues; see Griliches and Mairesse (1998) and Van Biesebroeck (2008) for overviews. There is debate as to which method is best. In the end, as I see it, choosing a method is a matter of asking oneself which assumptions one is comfortable making. Certainly one cannot escape the fact that *some* assumptions must be made when estimating the production function.

Fortunately, despite these many concerns, many of the results described in this paper are likely to be quite robust to measurement peculiarities. When papers have tested robustness directly, they typically find little sensitivity to measurement choices. The inherent variation in establishment- or firm-level microdata is typically so large as to swamp any small measurement-induced differences in productivity measures. Simply put, a high-productivity producer will tend to look efficient regardless of the specific way that their productivity is measured. I usually use cost-share-based TFP index numbers as a first pass in my own work; they are easy to construct and offer the robustness of being a nonparametric first-order approximation to a general production function. That said, it is always wise to check one's results for robustness to specifics of the measurement approach.

3. Productivity and the Plant/Firm

This section discusses factors that impact productivity at the micro level by operating within the plant or firm. They are potential "levers" that management or others have available to impact the productivity of their business. The second part of the discussion, which will be covered in the next section, focuses on influences external to the firm: elements of the industry or market environment that can induce productivity changes or support productivity dispersion.

I've broken up the discussion of direct productivity impacts by category for the sake of exposition. However, it's good to keep in mind that there can often be spillovers across these categories, and multiple mechanisms can act in concert. I'll point out many of these across-category links as the discussion goes along.

3.1. Managerial Practice/Talent

Researchers have long proposed that managers drive productivity differences.[5] Whether sourced in the talents of the managers themselves or the quality of their practices, this is an appealing argument. Managers are conductors of an input orchestra. They coordinate the application of labor, capital, and intermediate inputs in an efficient manner. And, just as a poor conductor can lead to a cacophony rather than a symphony, one might expect poor management to lead to discordant production operations.

Still, perhaps no potential driving factor of productivity differences has seen a higher ratio of speculation to actual empirical study. Data limitations have been the stumbling block.

---

[5] I mean *long* proposed: Walker (1887) posits that managerial ability is the source of differences in surplus across businesses: "The excess of produce which we are contemplating comes from directing force to its proper object by the simplest and shortest ways; from saving all unnecessary waste of materials and machinery; from boldly incurring the expense—the often large expense—of improved processes and appliances, while closely scrutinizing outgo and practising a thousand petty economies in unessential matters; from meeting the demands of the market most aptly and instantly; and, lastly, from exercising a sound judgment as to the time of sale and the terms of payment. It is on account of the wide range among the employers of labor, in the matter of ability to meet these exacting conditions of business success, that we have the phenomenon in every community and in every trade, in whatever state of the market, of some employers realizing no profits at all, while others are making fair profits; others, again, large profits; others, still, colossal profits." It's impressive how Walker's description closely matches (albeit with the flowing prose typical of the time) the viewpoints of researchers over 120 years later. We finally are becoming able, with the growing availability of broad-based production microdata, to test such hypotheses on a comprehensive basis.

As great an increase in detail has the proliferation of production microdata afforded, such data rarely contains detailed information on any aspect of managerial inputs. Sometimes there may be some distinction made between blue- and white-collar or production and nonproduction employees, but that's usually it. The identity, much less the characteristics, practices, or time allocation of individual managers is rarely known. Furthermore, managerial inputs can be very abstract. It's not just time allocation that matters, but what the manager does with their time, like how they incentivize workers or deal with suppliers.

A recent set of papers has made considerable efforts to close this measurement gap. These studies have often focused on single-industry or even single-firm case studies by necessity, given the detail required in the data. More comprehensive efforts that cover a broader cross section of economic activity are underway, however.

Bloom and Van Reenen (2007)—hereafter BvR—offer one of the most comprehensive studies relating management practices (though not managers per se) to productivity. They and their team surveyed managers from over 700 medium-sized firms in the US, UK, France, and Germany. (They surveyed plant managers, so the measured practices revolve around day-to-day and close-up operations rather than the broader strategic choices made at the executive level.)

Surveys were conducted over the phone, by a questioner with the same native language as the respondent. Information was probed on 18 specific management practices in four broad areas: operations, monitoring, targets, and incentives. The interviewers used these responses to score the firm on its practices. Given the inherent subjective element of this measurement process, BvR took several steps to ensure as accurate and consistent measurement as possible. Managers were not told they were being scored. Questions on practices were open-ended (e.g., "Can you me how you promote your employees?" rather than "Do you promote your employees based on tenure?"). Financial performance was not discussed. The firms were small enough so that the interviewers would not already be aware of the performance of the firms they surveyed. Each interviewer conducted dozens of interviews, allowing BvR to control for interviewer fixed effects when relating management scores to outcomes. Further, over 60 firms were surveyed twice by different interviewers; the correlation between the separate management practice scores for the same firms was 0.73.

Much of what was scored as "best practice" management in the interviews was based on the recommendations of the management consulting industry. This raises concerns about whether these practices are actually related to performance, or just the management fad of the moment. Importantly, therefore, BvR document that higher-quality management practices (and higher scores) are correlated with several measures of productivity and firm performance, including labor productivity, TFP, return on capital, Tobin's Q, sales growth, and the probability of survival.

The correlation between a firm's management practice score and its total factor productivity is statistically strong and economically non trivial. Spanning the interquartile range of the management score distribution, for example, corresponds to a productivity change of between 3.2 and 7.5 percent. This is between 10 and 23 percent of TFP's 32 percent interquartile range in their sample.

BvR show two factors are important predictors of the quality of management practice in a firm. More intense competition in the firm's market, measured in several ways, is positively correlated with best-practice management. Additionally, management practice scores are lower when the firm is family-owned *and* primogeniture determined the current CEO's succession—i.e., he is the eldest son of the firm's founder. (I will discuss the competition-productivity link extensively in the following section. Interestingly, primogeniture's tie to productivity is not about family ownership *per se*—in fact, family ownership is positively correlated with good management in isolation.) These two factors are responsible for explaining most of the difference between the country-level average management scores in the sample. The variation in these averages are largely the result of the UK and France having a left tail of poorly managed firms, both countries with tax laws that favor primogeniture.

Disentangling whether these correlations are causal is more challenging. Perhaps management consultancies base their recommendations off of observed practices at successful firms, but some excluded factor is driving both management practice and performance. BvR are aware of this issue. They estimate a specification using competition and primogeniture measures to instrument for management scores, based on the notion that the competitive and legal environments are orthogonal to other factors that drive management practices, at least in the short run. The estimated effect of management practices on TFP remains statistically significant and is in fact larger than the OLS case. This may suggest that unobserved third factors have a modest role if any and that BvR's management practice scores are noisy measures of true managerial acumen.

BvR have since expanded their management practice survey program to gain greater coverage of business practices across economies. Bloom, Dorgan, Dowdy, Rippin, and Van Reenen (2007) describe results from an extension of this survey program to over 4000 firms in eleven countries, including fast-growing China and India. The broader results echo those above. A particularly interesting pattern emerging from the early analysis is that the much lower average management practice scores in China and India are driven not so much by lower productivity across the board (though this is present to some extent), but in particular by a large left tail of very poorly managed firms. This has obvious implications for how growing trade and its assorted competitive pressures might impact productivity evolution in these and other countries. (More about Chinese and Indian firms' TFP levels below.) One potential expansion of the survey program that would be extremely useful is for it to incorporate a panel element. This would allow one to see how firms' management practices change when their external environment does.

This research line is studying managerial actions and policies that are often at levels below the executive suite. Other work has focused on how those at the apexes of corporate hierarchies influence performance. Bertrand and Schoar (2003) study top executives (e.g., CEOs, CFOs, Presidents, etc.) who manage at least two firms for three years each during their 1969-1999 sample period. Following managers across multiple firms lets them test if individual executives can explain variation in firms' performance measures. While they don't measure productivity specifically, they do find that the individual manager fixed effects (particularly for CEOs) have significant explanatory power over firms' returns on assets. Adding these fixed effects to a regression of returns on firm and year fixed effects raises the adjusted R^2 from 0.72 to 0.77.

These results reflect performance differences that can be explained by the identity of the managers. This still leaves open the question of what the managers *do* or *know* that affects performance. Bertrand and Schoar don't have the sort of detailed management practice data of BvR, but they do regress their estimated manager fixed effects on two variables they observe for the executives in their data: age and MBA status. They find that while age is not a significant factor, managers with MBAs have significantly higher ROA effects (by roughly 1 percent, off of a mean of 16 percent).

A series of studies using detailed data collected within individual firms have also looked at the management-productivity relationship. Bushnell and Wolfram (2007) find that power plant operators have nontrivial impacts on the thermal efficiency of power plants. The best can boost their plant's fuel efficiency by over three percent, saving millions of dollars of fuel costs per year. Unfortunately, the data are less clear about what particular actions or attributes predict good plant management.

Other within-firm work has suggested that the human resources components of management, in particular, can affect productivity. This research—see for example Ichniowski, Shaw, and Prennushi (1997); Lazear (2000); Hamilton, Nickerson, and Owan (2003); and the papers cited in Ichniowski and Shaw (2003)—uses highly detailed, production-line-specific data to tie non-standard HR management practices like pay-for-performance schemes, work teams, cross-training, and routinized labor-management communication to productivity growth. These papers have suggested some interesting details about the productivity effects of these practices. For instance, there is some evidence that the practices are complements: while they may have little or modest impact on productivity when implemented in isolation, their total impact is larger than the sum of its parts when the practices are used in conjunction with one another. Further, these practices are likely to have heterogeneous effects across production lines, even in the same plant, if different lines produce product variants of differing complexity. Boning, Ichniowski, and Shaw (2007), for example, find an interaction between the complexity of the production process and the usefulness of more innovative HR management in raising productivity.

Mas (2008) shows in a vivid case study that poor management-labor relations can have productivity effects. He looks at the resale values of equipment made at plants and times where Caterpillar was experiencing labor disputes during the 1990s. Compared to otherwise identical products made at plants or times without unrest, these products had about 5 percent lower resale values. This substantial productivity impact due to the implied reduction in the equipment's quality-adjusted service flows totaled $400 million.

With these and other studies, the evidence the management and productivity are related is starting to pile up. Further, some of this work strongly suggests that this relationship is casual. Still, establishing causality definitively is still a key remaining issue for research. Bloom, Eifert, Mahajan, McKenzie, and Roberts (2009) are attempting to establish as much by using what many consider to be the gold standard for establishing causality: a randomized field experiment.

They are providing management consulting to a random set of Indian firms and will compare productivity growth in this treatment group to that observed in a set of control firms not receiving the intervention. This study could go a long way toward establishing whether or not a causal link exists. (Though given what known about the left tail of Indian firms, if management consulting were to be effective anywhere, it would be in India. So the experiment might offer the upper bound causal effect of management practices.)

3.2. Higher-Quality General Labor and Capital Inputs

Management is an unmeasured input in most production functions, and hence is embodied in the productivity measure. Similarly, the productive effects of inputs like (non-management) labor and capital can also enter productivity if there are input quality differences that standard input measures do not capture.[6]

There is of course an enormous literature on human capital, far too large to cover here, that has tied several factors to labor quality, including education, training, overall experience, and tenure at a firm. Much of this work in labor economics has focused on wages at the outcome of interest. A smaller set of work has looked at human capital's impact on productivity.

Newer work using matched employer-employee datasets, which allow individual workers to be tracked across plants or firms over time, have offered evidence on the importance of labor quality. Abowd et al. (2005) offer a broad survey of the early evidence from these types of datasets, which tend to be newly constructed and therefore have short panel histories. Their applicability for studying productivity, while limited now, will greatly increase over time. Still, some progress has been made with such data. Ilmakunnas, Maliranta, and Vainiomäki (2004), for example, use Finnish matched worker-plant data to show that (not surprisingly) productivity is increasing in workers' education and age.

As great a potential as such data may hold, the results in Fox and Smeets (2009) suggest that matched employer-employee data will not answer all of the literature's burning questions. They use matched employer-employee records from the Danish economy to control for worker

[6] Attempts to capture labor quality differences in labor measures rather than productivity are the impetus behind using the wage bill to measure labor inputs rather than the number of employees or employee-hours. The notion is that market wages reflect variations in workers' contributions to production; firms with more productive workers will have a higher wage bill per employee. Of course, there are problems with this approach: wage variation might reflect the realities of local labor markets, or causation could be in the other direction, if more productive producers earn rents that are shared with or captured by employees. Hence, more direct labor-quality measures are needed to definitively pin down labor quality's productivity contribution.

education, gender, experience, and industry tenure in production function estimation. While these labor quality measures have significant coefficients in the production function, accounting for their influence only decreases the average within-industry 90-10 percentile productivity ratio from 3.74 to 3.36. There's plenty of productivity variation left to be explained. (By the way, using the decline in productivity dispersion as a metric of a newly measured factor's importance in explaining productivity—or an R^2-type measure as Bertrand and Schoar use—is a good idea. All studies seeking to explain productivity dispersion should strive to conduct and report similar exercises.)

Capital can also vary in quality in ways that may not be captured with standard capital measures. If capital vintages differ from one another in how much technological progress they embody, the common book-value-based capital stock measures will tend to miss variations in average capital vintages across producers. Several studies have tried to address this issue by measuring the rate of capital-embodied technological progress by carefully constructing measures of the distribution of capital vintages within plants or firms. Sakellaris and Wilson (2004) do exactly this using the annual investment histories of plants in the U.S. Annual Survey of Manufactures and industry-year-specific depreciation measures. They estimate a production function that is standard in all respects except that, rather than measuring capital inputs with as single dollar-valued stock, they use a weighted sum of the plant's past investments. The weights are a combination of the cumulative depreciation of a particular vintage's investment, as well as a technological progress multiplier that they estimate. In essence, they assume that productive capital, measured in efficiency units, grows at a constant rate per year. They estimate this rate to be between 8 and 17 percent per year, depending on the specification. These numbers are striking in their implications about how much productivity growth can be driven by investment alone. (Note that, unlike simple capital deepening effects of investment that only affects labor productivity, capital-embodied technological progress like this also impacts TFP.) Other studies using different methodologies (e.g., Cummins and Violante 2002) have found somewhat smaller values, on the order of five percent per year. This seems to be an area desperate for further evidence, given its potential importance.

Van Biesebroeck (2003) measures the productivity impact of auto assembly plants shifting to "lean" technologies, which in that context involves new capital as well as a host of complementary practices (teamwork, just-in-time ordering, etc.). This is clearly related to the managerial practice discussion earlier. Van Biesebroeck finds that both the entry of new lean plants and the transformation of older-style plants are responsible for the industry's acceleration of labor productivity growth during the late 1980s and early 1990s. Interestingly, his estimates of each technology's parameters suggest that capital-augmenting productivity is the primary driver of labor productivity growth under lean processes, while Hicks-neutral TFP-type productivity drives growth in the traditional technology plants.

3.3. IT and R&D

While the research described above indicates that input heterogeneity matters, the productivity effects of a particular type of capital input—information technology (IT)—has been the subject of intense study. This is rightly so; many have hypothesized that IT was behind the resurgence in U.S. aggregate productivity growth in the mid-1990s after 20 years of sluggish performance. Given the sheer size of economic output variation that can be driven by even a modest change in aggregate productivity growth over a sustained period, it is no surprise that sources of such changes receive considerable research attention. Because of this attention, I will overview the work done on this particular capital input separately here.

An overview of IT capital's broad impacts on productivity, particularly in driving the resurgence, can be found in Jorgenson, Ho, and Stiroh (2005, 2008) and Oliner, Sichel, and Stiroh (2007). These studies document that IT-related productivity gains—both spectacular productivity growth in IT-*producing* industries and in non-IT industries increasingly using IT to improve the efficiency of their own operations—play an important role in explaining aggregate U.S. productivity growth over the past couple of decades.

At the same time, van Ark, O'Mahony, and Timmer (2008) show that the comparably sluggish productivity growth over the same period in the EU is explained in large part by the later emergence and smaller size of IT in European economies. Bloom, Sadun, and Van Reenen (2007) suggest that it is not geography *per se* that matters, but rather the location of the owning firm. They show U.S.-based multinationals operating in the U.K. are more productive than their U.K. counterparts, and this productivity advantage is primarily derived from IT capital. Further supporting this theory is the fact that U.K. plants bought by U.S. multinationals see IT-related productivity gains that plants purchased by non-U.S. multinational firms do not.

These broad patterns raise the question of what specific micro mechanisms actually underlie the aggregate relationship between IT and productivity growth. Bartel, Ichniowski, and Shaw (2007) dig deep into the production process in one particular industry to try to answer this question. They show that more intensive use of IT has several benefits in their sample of valve manufacturing lines. Better computer numerically controlled (CNC) machining centers—automated devices that manipulate and shape parts from raw material stock—reduce production time. This increased speed has multiple sources, including reduced setup times (for example, programming the particular machining process into the CNC machine), a faster production run (the CNC machine's shaping of the raw stock with various tools), and even quicker inspections. The appealing element of the study's empirical approach is that both the products and the production process, with the exception of the particular pieces of IT capital whose contribution is of interest, remain constant across observations. The paper also shows IT offers gains beyond simply reduced production times for the firms' current product set. IT-intensive product design tools (e.g., computer-aided design packages) make it easier to design customized parts, and lower setup times make multiple production runs less costly. Offering a broader array of parts allows the firms to better match their production capabilities to their customers' desires, increasing the surplus of their sales.

The gain in surplus from product specialization raises an important point about productivity measurement. The improved customization due to IT can raise the firms' average sales prices. Measures of productivity in physical units of output (e.g., number of valves per unit input) may therefore not fully capture the surplus gained. Revenue-based output measures, which are often the *only* output measure available in producer microdata, do allow these gains to be captured. On the other hand, many other factors can cause price variation across producers (e.g., local market power or demand shocks) that researchers would not want to be included in productivity. Yet these cannot be purged from revenue-based productivity measures. This is a conceptually tricky issue that we will return to below.

Some have suggested a link between firm organization, decentralization specifically, and the ease with which productive new technologies can be adopted. This is the explanation favored by Bloom, Sadun, and Van Reenen (2007) for European firms' laggardness, and is also the subject of Acemoglu, Aghion, Lelarge, Van Reenen, and Zilibotti (2007). Evidence on the matter tends to be suggestive but indirect, however, and this is an area where careful work in measuring firm structures (not an inherently easy task) could pay big dividends. I will discuss broader connections between firm structure and productivity below.

There is a long literature linking R&D and productivity, and recent additions to it have focused on exploring the ties at the micro level. As with many input-based stories of productivity differences (such as the management practices studies above), the difficulty is in separating correlation from causation. There are many reasons why more productive firms might do more R&D, suggesting that some of the causation may go the other way.

Doraszelski and Jaumandreu (2009) attempt to endogenize firm productivity growth as the consequence of R&D expenditures with uncertain outcomes. Estimating the model using a panel of Spanish firms, they find that R&D does appear to explain a substantial amount of productivity growth. However, firm-level uncertainty in the outcome of R&D is considerable: in some industries, over half of the variance in productivity is explained by productivity innovations that are orthogonal to observed R&D activity.

Aw, Roberts, and Xu (2009) highlight the bidirectional causality between R&D and productivity in their study of Taiwanese electronics exporters. They find that firms that select into exporting tend to already be more productive than their domestic counterparts (more on this in the trade section below), but that the exporting decision is often joint with a firm making large R&D investments. These R&D investments in turn raise exporters' productivity levels further. Therefore, they highlight evidence of both selection and causal effects that tie productivity to R&D.

Of course, R&D is simply one of the more observable components of firms' overall innovative efforts. Many firms undertake both process and product innovation without formally reporting R&D spending. This limits the ability of this literature to give a comprehensive look into the relationships between productivity and innovation. Still, it is certainly a very useful start, and the mechanisms the R&D literature highlights are likely to overlap in many cases into the effects of unmeasured innovative spending. (I will also discuss product innovation's ties to productivity differences in further detail below.)

Griffith, Harrison, and Van Reenen (2007) make an interesting point about R&D-productivity connections that also foreshadows the discussion below about technological spillovers. They show that the geographic location of a firm's R&D activity matters. Using patent data to pin down the historical locations where U.K. firms conduct R&D (pre-sample locations are used to minimize endogeneity of the location of R&D), they find that the productivity growth of U.K. firms with a greater U.S. research presence is faster overall and has a higher correlation with U.S. productivity growth. It thus appears that having a research presence in the U.S. makes it easier for firms to tap into to knowledge base of the U.S. economy, which tends to be at the technological lead of most industries. The precise mechanism through which this technology tapping occurs is unclear, and would be an interesting area for research.

3.4. Learning-by-doing

The very act of operating can increase productivity. Experience allows producers to identify opportunities for process improvements. This productivity growth, often called learning-by-doing, has a long and rich history of study in the literature, but has recently been investigated in more detail given newly available unit-level production data.

Benkard (2000) looks at the precipitous drop in the labor hours Lockheed needed to assemble its L-1011 TriStar wide-body aircraft. The first few units off the line required more than one million person hours (equivalent to three shifts a day of 2500 workers each for fifty work days). This was cut in half by about the 30th plane, and halved again by the 100th unit. Benkard estimates both the learning rate—how fast past production increases productivity (decreases unit labor requirements)—and the "forgetting" rate, which is how fast the knowledge stock built by learning depreciates. Forgetting is quantitatively important in this setting: Benkard estimates that almost 40 percent of the knowledge stock depreciates each year. This may not be literal forgetting, but could instead primarily reflect labor turnover. An additional factor in "forgetting" was the shift to a new variant of the plane after about 130 units. This new variant was different enough that the imperfect substitutability of the knowledge stock between the original and new variants led to a temporary but substantial increase in labor requirements.

Thornton and Thompson (2001) investigate what *types* of experience matter in productivity growth from learning by doing. Their data includes unit labor requirements for several design variants of 4000 Liberty ships produced by multiple shipyards during World War II. The multi-design/multi-yard nature of the data lets them estimate the relative productivity contributions of four different measures of past production experience: the yard's past production experience with a particular design, the same yard's past production of other designs, other yards' experience with the particular design, and other yards' production of other designs. Not surprisingly, a yard's past production of a particular model matters most for productivity growth in that same model. After that comes the yard's experience with other ship designs, which is about 60 percent the size of the own-design effect. Cross-yard spillovers are considerably smaller—only about five to ten percent of the own-yard, own-design learning impact. These cross-plant learning effects, while relatively modest here, do show that producers may become more productive by learning from other businesses. I discuss this mechanism more below.

Kellogg (2009) looks at oil and gas drilling in Texas to study how learning occurs when an upstream and downstream producer work together. He follows the efforts of pairs of producers and drillers. The former are companies actively involved in exploring for and extracting oil, while the latter firms specialize in boring out the wells that the producers hope will yield oil. Since producers typically work with multiple drillers and vice versa during his sample, and work in different fields, Kellogg is able to separately measure the relationships between experience of producers alone (i.e., regardless of the drilling firms they work with), drillers alone, and within particular producer-driller pairs. He finds that accumulated experience between a producer-driller pair increases productivity above and beyond that of the firm's overall experience levels. This relationship-specific experience is a type of capital that is lost if the firms split up, giving them incentives to preserve their contracting environment.

3.5. Product Innovation

Innovations in product quality may not necessarily raise the quantity of output (measured in some physical unit) per unit input, but they can increase the product price and therefore the firm's revenue per unit input. If one thinks about productivity as units of quality delivered per unit input, then product innovation enhances productivity. This productivity effect is captured in standard revenue-based productivity measures since they embody price variations across plants or firms in an industry. (Though as mentioned above and discussed further below, revenue productivity can also be misleading when price variation due to variations in market power across producers exists.) Product innovation can be aimed at entering new markets, or to refocus a firm's efforts on growing demand segments, as documented in Acemoglu and Linn (2004).

Product innovation's productivity effects have been studied in several recent papers. As touched on above, one of the mechanisms behind IT-based productivity growth that Bartel, Ichniowski, and Shaw (2007) point to is an improved ability to customize products. Other inputs mentioned above, like R&D and higher-quality employees, can also spur innovation.

Lentz and Mortensen (2008) use Danish firm-level data to estimate a model of firms' product innovation efforts in the vertical-quality-ladder style of Klette and Kortum (2004). They find that about 75 percent of aggregate productivity growth comes from reallocation of inputs (employment in their setup) to innovating firms. About one-third of this comes from entry and exit channels. The other two-thirds occurs as inputs move toward growing firms (and hence innovating firms, as seen through the lens of their model) from firms that lose market share when they fall behind the quality frontier.

Balasubramanian and Sivadasan (2009) link detailed and broad-based data on firms' patenting activities (they merge the NBER patent database with the U.S. Census Business Register) to see what happens when a firm patents. They find clear evidence that new patent grants are associated with increases in firm size (by any one of a number of measures) and scope (the number of products it makes), and increases in TFP (though the evidence is weaker here). Whether these correlations reflect the causal effects of patents is not clear; patenting activity could be just one part of a firm's coordinated push into new markets. Nevertheless, given the breadth of the study's coverage and its result that correlations exist, more research in this area would be worthwhile.

Bernard, Redding, and Schott (forthcoming) show that a firm's TFP is positively correlated with the number of products it produces. This holds both in the cross section and within firms over time. At the very least, these results indicate that productivity growth accompanies expansion of the variety of products a firm offers. It's less clear whether innovative activity drives both productivity and product-variety growth, or whether firms experiencing general productivity shocks "strike while the iron is hot," expanding their product offerings in response. The role of changes in product scope in firm size and productivity growth is one that is just beginning to get the attention it deserves in research agendas.

3.6. Firm Structure Decisions

A lot of the micro productivity literature uses the establishment (e.g., factory, store, or office) as the unit of analysis. This is in part data driven: many surveys are conducted at this level. Plus, plants often embody the smallest indivisible unit of a production process, and as such are a natural level at which to study technologies. But it's also clear that sometimes firm-level factors, and in particular, the organizational structure of the firm's production units—the industries they operate in, their vertical and horizontal linkages, their relative sizes, and so on—will sometimes be related to the productivity levels of the firm's component plants.

Hortaçsu and Syverson (2009) examine the productivity of plants in vertically structured firms. We use the Longitudinal Business Database, which contains most private non-agricultural establishments in the U.S., to determine which industries firms operate in. Information on inter-industry goods and services flows from the Benchmark Input-Output tables indicates which firms own vertically linked plants. We show that vertically integrated plants have higher productivity levels than non integrated plants in the same industry. Most of the difference reflects selection of high-type plants into vertical structures, rather than any change in type associated with a formerly unintegrated plant becoming integrated. Surprisingly, these productivity advantages—and indeed the very choice of a vertical structure—usually have little to do with facilitate transfers of goods along the production chain, as is often presumed. Shipments from firms' upstream units to their downstream units are surprisingly low, relative to both the firms' total upstream production and their downstream needs. Roughly one-third of upstream plants report no shipments to their firms' downstream units, and half ship less than three percent of their output internally. These patterns primarily reflect selective sorting of high plant types into large firms; once we account for firm size, vertical structure per se matters much less. We propose an alternative explanation for vertical ownership: rather than moderating goods transfers down production chains, integration instead allows more efficient transfers of intangible inputs (e.g., managerial oversight) within the firm.

Maksimovic and Phillips (2002) and Schoar (2002) both investigate the productivity of plants within conglomerate (read: multiple two- or three-digit SIC industry) firms. Their work was spurred on in part by the extensive finance literature on the "diversification discount," the term for the oft-measured negative correlation between a firm's financial returns and the number of business lines it operates. Both papers leverage U.S. manufacturer microdata to convincingly argue that the diversification discount is not about low productivity (or even, in one case, any sort of underperformance). They differ, however, in their explanations.

Maksimovic and Phillips (2002) make a selection argument. Firms that choose to specialize are likely to have idiosyncratically high productivity draws in a particular line of business, but considerably weaker draws outside this segment. Firms that choose conglomerate structures, on the other hand, are likely to have high draws in several industries, but not exceptionally high draws in any particular industry. Thus if one simply compares the performance of all a conglomerate's segments to the focused and highly productive segments of a specialist, the latter would expectedly be higher. This result does not rely on the previous literature's favored explanations of management overreach, cross subsidization of weak segments, or other agency problems at conglomerates. It simply reflects the optimal allocation of resources within a business given the firm's inherent abilities. They support their efficient allocation argument by showing that conglomerate firms' most productive plants are in their largest segments, and segments of a given rank are more productive in larger firms. Furthermore, conglomerates expand on their strongest margins: their largest, most productive segments are more sensitive to demand shifts than their smaller, less efficient lines of business.

Schoar (2002) notes that in her sample, plants in conglomerates have, if anything, higher permanent productivity levels. The observed discount reflects the temporary impacts of the very act of diversifying by buying plants in new lines of business. She shows that when a conglomerate diversifies, the plants it buys actually experience productivity growth, suggesting that they are in fact being reallocated to more capable management (there will be more on the reallocation of productive inputs below). At the same time, however, the conglomerate's existing plants suffer productivity drops. Since conglomerates have on average many more existing plants than acquired ones, average productivity in the firm falls for a period. Schoar attributes these productivity changes to a "new toy" effect: managers (over-) concentrate their efforts on integrating the new plants and business lines at the expense of existing ones. She also finds evidence that the firms' wages absorb any performance rents, also leading to a bifurcation between performance as measured by productivity and by stock market returns.

4. External Drivers of Productivity Differences

The previous section discussed productivity drivers that operate within the firm and which, at least in theory, producers have some degree of control. Plants and firms have some scope to choose their management practices, input mix, level of innovative efforts, and so on. This section focuses instead on how the producers' operating environments can influence productivity levels and growth. These elements may not operate directly, but they can affect the incentives producers have to apply the within-firm factors discussed in the previous section. By their nature, they are also the most closely tied to government policies. Hence policies that create environments that look like those discussed below will foster faster productivity growth. Therefore understanding these drivers merits special attention when one wants to consider the productivity implications of market interventions.

Before discussing the specific external drivers, it's worth taking a minute to discuss why the operating environment should affect the productivity level of a given business. The most basic producer theory, after all, says any profit-minimizing firm will minimize its costs of producing its chosen quantity. This prediction is invariant to the structure of the market in which the firm operates.

One possibility is that spillovers are present. If another firms' production practices somehow influence the productivity level of a business, obviously external factors will affect the productivity level of any firm, cost-minimizing or not. We first consider these sorts of situations in this section.

But more complex cases exist when there are no direct productivity spillovers. Why would we expect the market environment to affect producers' productivity levels if firms cost-minimize? The reason is that this simple world of the standard, static cost-minimizing firm model is an inadequate description of the processes of technology adoption. A richer model like that in Holmes, Levine, and Schmitz (2008)—who build off Arrow's (1962) seminal work—points out additional channels through which a firm's market environment (and the competitive structure in particular) shifts producers' incentives to increase productivity. Suppose adopting the productivity-enhancing practice involves disruption costs: a temporary period where costs are actually *higher* than before any technological changes were made. These costs could be due to installation issues, fine-tuning new technology, retraining workers, etc. If there are adoption costs, producers facing less competition have less incentive to adopt the new technology. The reason is that the higher per-unit profits that monopoly power brings raise the opportunity cost of adopting new production techniques.[7]

The reality of production is also much more complex than in the standard model. Most technologies, even if routinized, are intricate, multifaceted processes that require considerable

[7] Of course, these adoption costs should be included as inputs when computing productivity, at least in concept. In practice, they can be missed if they aren't part of the typical investment and labor inputs values firms report in typical producer microdata. Nonetheless, Holmes, Levine, and Schmitz's (2008) point remains: for a given adoption cost (properly counted or not), the probability of adoption decreases with the firm's market power.

coordination. They are constantly being buffeted by shocks to input costs and availability and demand-driven shifts in capacity requirements. Cost-minimizing production practice is really therefore a moving target, a constantly shifting ideal combination of operations. Elements of a firm's market environment can affect the incentives it has to chase that moving target.

4.1. Productivity Spillovers

It's sensible to think that laggard producers' productivity growth might be related to that of the productivity leaders in their industries. Certainly the importance of productivity in explaining business growth and survival indicates that firms have considerable incentive to emulate industry leaders. Further, some productivity-enhancing practices may not be completely excludable; businesses cannot always keep every facet of their production process secret.

The evidence on the persistence of productivity differences within industries suggests that any such emulation/spillover process is far from perfect, however. It's clear there are frictions that keep less efficient producers from fully replicating the best practices of their industry leaders. Still, productivity spillovers seem likely to exist. A crucial research question is, if they do, how large are they, and what are the mechanisms though which these productivity spillovers occur? These questions have been a focus of inquiry of several studies.

Moretti (2004) matches information from the 1980 and 1990 U.S. Population Censuses with the 1982 and 1992 Census of Manufactures at the city-industry level to explore productivity spillovers among plants in a metro area. He measures the impact of local human capital spillovers by including in a plant-level production function the share of the metro area's workers in *other* industries who have completed some college. The estimated marginal product of this outside educated labor is interpreted as the productivity spillover. The results are that plants in dense human capital markets have higher productivity levels than observably similar plants in thin human capital markets. Furthermore, he shows that measured spillovers are stronger across plants that are "close," both in geographic and technological senses. Moretti also shows that these productivity spillovers are reflected in workers' wages.

Bloom, Schankerman, and Van Reenen (2007) make the point that spillovers can cut two ways: there can be technological spillovers that benefit everyone, but there can also be market-stealing effects on the product market side. They separately identify these two effects by comparing the effects of firms' R&D activities on other firms at varying technological and product market distances. Technological distance is measured by correlations in firms' patenting patterns, while product market distance is reflected in the correlation in their sales across business segments. Since these two distances are not perfectly correlated across firms, the separate impacts of R&D can be measured. They find that both types of spillovers are a factor, but that the technological spillover quantitatively dominates, creating a net positive externality.

Bartelsman, Haskel, and Martin (2008) make an interesting distinction between the global and economy-specific technology frontiers. They show using microdata from numerous countries that convergence is stronger toward a plant's national frontier than it is toward the global frontier. That is, for a given productivity gap between the plant's own and the frontier productivity level, the rate at which such gaps are closed is higher for the domestic leader than for the global leader. A second interesting result is that if the plant is sufficiently behind the global frontier, any pull toward convergence disappears. However, the convergence to the national frontier remains no matter the size of the gap. Thus plants are always converging toward national leaders in their industries (conditional on survival, of course), and if they are not too far behind, toward the global frontier as well.

Crespi et al. (2008) also look at cross-border productivity convergence, with a particular eye toward identifying the sources of information flows that could be the source of this convergence. They combine production microdata with survey information on where firms gain information to use in their efforts at technological innovation. They find that, not surprisingly, "nearby" firms (e.g., suppliers and competitors, though less so buyers) are primary sources; that much of this information, particularly from competitors, is free; and that multinational presence aids these flows.

Keller and Yeaple (forthcoming) study the role that international movement in capital flows play in increasing productivity. They look at the ties between productivity growth among publicly traded U.S. firms and foreign direct investment in those firms sectors by foreign-owned multinationals. They find that FDI spillovers account for a substantial portion of productivity growth, and are particularly strong in high-tech sectors.

These papers and others like them suggest that spillovers are present and can operate through various mechanisms, though again the observed productivity dispersion also makes clear that substantial frictions remain. Policies meant to increase such spillovers must be careful, however, to not destroy firms' incentives to innovate. If spillover-enhancing policies make it too hard for firms to appropriate the benefits of their innovation, they could do more damage in the long run.

4.2. Competition

Pressures from threatened or actual competitors can affect productivity levels within an industry. Competition drives productivity increases through two key mechanisms. The first is through efficiency increases within plants or firms. Heightened competition induces firms to take costly productivity-raising actions that they may otherwise not. The second mechanism is selection among producers with heterogeneous productivity levels. Competition will drive market share toward more efficient producers, shrinking relatively high-cost firms/plants (sometimes forcing their exit) and opening up room for more efficient producers. This section discusses research that investigates competition's productivity effects operating through both of these mechanisms.

Because of the substantial literature that has been built on the productivity impacts of trade competition, I will discuss it in a separate subsection below. I first cover general competitive effects.

4.2.1. Intra-market Competition

An indicator that product-market competition is enhancing productivity is a positive correlation between productivity and producer growth and survival. Such correlations have been a robust finding in the literature; Foster, Haltiwanger, and Krizan (2001) offer a broad-based overview, for example.[8] Several recent studies have looked at particular mechanisms through which competition leads to a Darwinian selection process.

Syverson (2004b) investigates the connection between competition and productivity in a case study of the ready-mixed concrete industry, which is well suited for this type of investigation. The industry's physically homogeneous product and very high transport costs make spatial differentiation paramount. Differences in competitiveness across markets should

---

[8] Foster, Haltiwanger, and Syverson (2008) point out that these results linking selection to productivity actually reflect selection on *profitability*, since intra-industry price variation caused by idiosyncratic demand differences across plants is buried in standard revenue-based productivity measures. They show that such demand variation is extremely important—perhaps even dominant—in explaining plant survival patterns, even in their sample of plants in homogeneous-product industries. This broader interpretation of the evidence to include demand-side factors will be discussed further in the following section.

therefore be related to the density of concrete producers in the market. Inefficient concrete producers should therefore find it more difficult to be profitable in dense markets; if they charge the high prices necessary to cover their costs, consumers can easily shift purchases to their more efficient competitors. This implies the productivity distribution of ready-mixed plants should be truncated from below as density rises. This is indeed what holds in the data. Markets with denser construction activity (the construction sector buys almost all of the ready-mixed industry's output and, because of concrete small cost share in construction, is likely exogenous to variations in concrete market competitiveness across markets) have higher lower-bound productivity levels, higher average productivity, and less productivity dispersion. Syverson (2007) shows that these patterns of competition-driven selection on costs are also reflected in ready-mixed prices.

Outside of manufacturing, Foster, Haltiwanger, and Krizan (2006) find that aggregate productivity growth in the U.S. retail sector is almost exclusively through the exit of less efficient single-store firms and by their replacement with more efficient national chain store affiliates. This evokes the stories surrounding the successful growth of discount retailers like Wal-Mart and Target over the past two decades.

The mechanism driving productivity growth in these papers is Darwinian selection: more competition weeds out the less efficient, raising average productivity levels. But competition could induce innovation among existing producers too. This is sometimes referred to as the "within effect" in the literature, in contrast to the "selection effect" described above.[9]

Schmitz (2005) offers an example of productivity growth in an industry that is driven almost entirely by within-effect efficiency improvements. He follows U.S. iron ore mining during the period the industry was first facing competition from foreign producers. (The specific competitors in this case are Brazilian mines. We will discuss more examples of trade-induced productivity change in a separate section below.) The case study shows how competition can drive existing industry firms to take steps to improve their productivity.

[9] Many studies attempt to quantify the relative contributions of the within and selection effects by decomposing aggregate productivity growth into terms that reflect the separate effects. Petrin and Levinsohn (2008) have recently raised caveats about the robustness of these type of "accounting decompositions." They advocate the use of a method that aggregates the shifts in the gaps between the estimated social marginal benefits and costs caused by the movement of inputs across producers. While distinct in theory and empirical practice, such "gap methods" hold the same conceptual goal: to separately measure how much aggregate productivity growth comes from businesses becoming more efficient themselves and how much comes from reallocation of economic activity to more efficient producers.

The industry had been protected by competition from abroad by the considerable transport costs involved in moving ore from its other sources on the globe (e.g., Australia and Brazil). By 1980, however, increased production from low-cost mines in Brazil brought delivered prices for Brazilian ore in the Great Lakes region in line with delivered prices from the Mesabi Range in northern Minnesota, the major ore-producing area of the U.S. Facing competition from abroad for the first time, U.S. producers attempted to lower costs by making drastic changes in their production procedures. Schmitz shows most of these changes were centered on the strict work rules that the U.S. mines had in place. For instance, mine managers previously had very little flexibility in their ability to assign different workers to different tasks. The introduction of serious competition allowed the mines to gain back flexibility in new contacts, raising the mines' utilization of available labor and enabling them to shed unneeded overhead workers. It was extremely successful. While for decades preceding 1980, the average labor productivity of the industry was roughly constant at two tons of ore per hour, average industry productivity doubled to four tons per hour within five years. As a result, the mines remained competitive even in the face of continuing falling prices for Brazilian ore.

Other recent studies have shown these detailed case studies appear emblematic of much broader competitive effects acting in numerous industries and economies. For example, Syverson (2004a) looks at the entire U.S. manufacturing sector. Disney, Haskel, and Heden (2003a and 2003b) and the studies described in UK Office of Fair Trading (2007) show similar results in the U.K. And Nicoletti and Scarpetta (2005) overview of evidence across OECD countries.

4.2.2. Trade Competition

As indicated in Schmitz's results for the iron ore industry, the presence—or even just the threat—of imports from abroad is another form of competitive pressure. This phenomenon is the focus of a burgeoning line of research, driven in part by the recent theoretical trade literature focusing on heterogeneous-productivity producers and their response to trade; see especially Eaton and Kortum (2002) and Melitz (2003).

Pavcnik (2002) shows how trade liberalization during the 1970s impacted productivity growth among Chilean manufacturers. The paper demonstrates that sectors facing new competition from imports saw higher productivity growth over the 1979-86 sample period than sectors producing primarily non-tradable goods. Pavcnik goes on to show that these industry-level gains came both from existing producers raising their productivity levels (the within effect) and from the reallocation of activity away from—and sometimes, the exit of—less efficient, formerly protected producers (the selection effect).

Multiple studies using producer microdata have since found comparable results in other settings. Examples include Eslava et al. (2004); Muendler (2004); Bernard, Jensen, and Schott (2006); Fernandes (2007); and Verhoogen (2008). The specific mechanisms through which trade-oriented competition is postulated to increase productivity do vary across the papers, from quality upgrading within plants to heighted selection across plants. Amiti and Konings (2007) highlight a separate mechanism through which trade can increase productivity: the expansion of the set (or declines in the effective price) of inputs into production, as imported inputs become more available. I will discuss the input-market channel further below.

Interestingly, despite the strong correlation between the average productivity level of an industry's plants and that industry's exposure to trade, there less evidence of strong productivity impacts on the domestic plants that begin exporting when trade barriers fall. That is, exporters are almost inevitably (much) more productive than their industry counterparts who do not export, but most studies to this point have found that this correlation largely reflects selection rather than a causal impact of exporting on productivity. In other words, plants that end up choosing to export were *already* more productive when trade barriers were higher.

That said, Van Biesebroeck (2005) and De Loecker (2007) find cases where exporters' productivity advantage grows *after* entry into the export market (aka the "learning-by-exporting" hypothesis). Both are in somewhat atypical settings, which might explain in part why they find post-export productivity growth while many others have not. In Van Biesebroeck's (2005) sample of Subsaharan African exporters, their post-exporting growth appears create productivity gains by reducing credit and contract enforcement constraints. This type of mechanism raises the question of whether it would apply to any firm that chooses to export (if so, why wouldn't everyone?), or whether these effects, while causal, reflect heterogeneous treatment effects, with those firms the most apt to benefit choosing to export. De Loecker (2007) finds that Slovenian firms that begin exporting during the post-transition period saw productivity growth after entering foreign markets. Interestingly, firms exporting to higher-income regions saw greater productivity growth. This suggests that the structure of the export market, and not just the exporter itself, matters. Further, there are interesting selection issues regarding which markets firms choose to export to, even conditional on the decision to export in the first place.

4.3. Deregulation or Proper Regulation.

Poorly regulated markets can create perverse incentives that reduce productivity. Deregulating or reformatting to smarter forms of regulation can reverse this.

Bridgman, Qi, and Schmitz (2007) show how regulations in place for decades in the U.S. sugar market destroyed incentives to raise productivity. The U.S. Sugar Act, passed in 1934 as part of the Depression-era restructuring of agricultural law, paid sugar beet farmers a per-ton subsidy financed by a tax on downstream sugar production. Sugar refiners were compensated for this tax by quota protection from imports and government-imposed limits on domestic competition (don't ask where antitrust law was on this one). This transfer scheme obviously led to the standard price distortions, but it also distorted incentives for efficient production. Since farmers received a flat payment per ton of sugar contained in their beets, their optimal response was to simply grow the largest beets they could. The problem is that larger beets have lower sugar-to-pulp ratio, so more energy is required to refine a given amount of sugar from them. At the same time, given the restraints on competition in the refined sugar market, refiners had less incentive to improve sugar extraction on the margin. The combined result of these incentives is readily apparent in the data. When the Sugar Act was passed, a ton of beets yielded an average of 310 pounds of refined sugar, a yield that had been steadily rising from 215 pounds per ton in 1900. But this trend suddenly reversed after 1934. Yields dropped to 280 pounds per ton by 1950 and 240 pounds by 1974, the year the Act was repealed. Lo and behold, yields began to climb again immediately after repeal, to about 295 pounds per ton by 2004. (It's a testimony to the Act's productivity distortions that yields 70 years after the act were still lower than when it was passed.)

Knittel (2002) and Fabrizio, Rose, and Wolfram (2007) examine how power plant operations change when the regulatory structure they operate under does. Both studies involve a moving plants away from a traditional-style, cost-plus regulated monopoly structure into alternative forms. Knittel (2002) studies the implementation of "incentive regulation" programs, where regulators explicitly tie the earnings operators can make to the achievement of particular efficiencies in operations. Fabrizio et al. look the effect of electrical market reforms that occurred in many regions in the U.S. during the 1990s. Both studies note that plants saw operating efficiency gains after the shift in the regulatory environment. Fabrizio et al. also show that, in line with what one would expect, the productivity gains were largest among investor-owned utilities and smallest in municipally-operated utilities.

Beyond these case studies, recent work has also taken a broader look at how product market regulations impact productivity at the micro level. For example, Greenstone, List, and Syverson (2009) show how environmental regulations (the U.S. Clean Air Act Amendments specifically) reduce manufacturing plants' productivity levels. Nicoletti and Scarpetta (2005) and Arnold, Nicoletti, and Scarpetta (2008) discuss product-market regulations' productivity effects in OECD economies. A related, yet distinct effect of legal structure on productivity is the effect of privatization of formerly state owned firms. One of the more comprehensive studies of these effects is Brown, Earle, and Telegdy's (2006) study of formerly state owned enterprises in several Eastern European countries. They document broad-based productivity increases at the plant-level after privatization, though they also note considerable variation in the size of the impact across countries, with more than 15 percent average TFP growth in Romania, but a slightly negative impact in Russia.

4.4. Flexible markets

As discussed above, competition increases productivity. If one thinks of competition as flexibility in product markets (in more competitive markets, it's easier for consumers to shift their purchases from one producer to another), it's logical to think that flexible *input* markets might also effect productivity levels. Just as flexible product markets are better at allocating resources toward more productive firms, factor market flexibility serve the same role on the supply side. Indeed, there's almost surely complementarities between product market and input market flexibility. If consumers want to reallocate their purchases across producers, those firms who experience the increase in demand will need to hire the additional inputs required to meet that demand. The easier it is for productive resources to move toward higher-productivity businesses, the faster and more smoothly the allocation mechanism works. This section discusses recent research tying factor market flexibility to productivity.

Maksimovic and Phillips (2001) look at the market for U.S. manufacturing plants as productive assets. Specifically, they ask a plant's productivity changes when it is sold from one firm to another. They find that, on average, a plant's productivity rises after it is sold to another firm. That is a reassuring result: the market tends to allocate inputs in a way that increases efficiency, rather than in response to the ambitions of empire-building managers or other inefficient motives. Also consistent with this efficiency-enhancing role of inputs markets is the fact that the plants that are sold tend to come from the selling firm's less productive business lines. The sellers are in essence moving away from activities that they are less good at.

Petrin and Sivadasan (2006) look at labor market flexibility's productivity effects using a novel approach. They measure the difference between Chilean plants' marginal products of labor (as derived from industry-level production functions that they estimate) and their average wages. Such gaps can be caused by any one of a number of market distortions, like market power, taxes, or the firing costs that are the object of the study. Allocative efficiency is achieved, at least in the cross section, if this gap is equated across plants. (Though of course overall inefficiencies still exist unless these gaps are all zero.) This means that efficiency increases if labor inputs are reallocated on the margin from low- to high-gap plants, because the net change in marginal product caused by the input shift outstrips the change in labor costs. Their results indicate that reducing firing costs did appear to increase allocative efficiency, since labor moved from lower to higher gap plants on average.

Several recent papers have taken these ideas and asked whether, more broadly speaking, economies efficiently allocate inputs across heterogeneous production units. Hsieh and Klenow (2009) use the measured TFP dispersion across firms in the Chinese and Indian economies to infer the size of producer-level distortions that act together to reduce aggregate productivity. Their methodology is similar in concept Petrin and Sivadasan's; it is also based on the notion that measured deviations from competitive-market first order conditions reflect wedges that prevent producers from equating marginal social products and costs. After measuring these plant-level distortions, they compare them to the analogous distribution measured in U.S. microdata. (This is used as the comparison rather than the first-best outcomes because it is a more realistic comparison. The U.S. data contain, and hence can be used to control for, gaps that reflect adjustment costs and measurement error that may be immutable to policy action.) Hsieh and Klenow find that Chinese aggregate TFP could increase by 30-50 percent and Indian TFP by 40-60 percent by achieving the U.S. level of allocative efficiency with their existing resources.

Bartelsman, Haltiwanger, and Scarpetta (2008) look at the success of allocation across several countries. Rather than using a gap-type methodology like Hsieh and Klenow, they measure efficiency using the correlation between a plant's share of industry output and its productivity level. The logic of this metric is straightforward and similar to that discussed at the beginning of the competition section. If markets are functioning well, production should be reallocated to the most productive plants, leading to a positive correlation between output share and productivity. An additional advantage of the metric is that it is easy to compute. Its limitation is that it is an accounting decomposition, and as such is not directly tied to welfare theory the way gap-type measures are. However, Bartelsman, Haltiwanger, and Scarpetta do show using a simple model how the various types of producer-level distortions do in fact lead to reductions in the output-productivity correlation within an industry.

5. Big Questions

That is a brief summary of what we know about the causes of productivity differences at the micro level, and why we would want to know those causes. I want to emphasize that while the discussion draws out major themes of that body of knowledge, it really only just scratches the surface of the literature.

I think it's a fair reading of the discussion above to say that we've learned a lot about productivity since the Bartelsman and Doms (2000) survey. At the same time, it's hardly time to declare victory and go home. Many pressing issues and open questions remain. Until we can reach the point where we can get $R^2 = 1$ when we estimate production functions, and be able to explain all the things that determine how producers use which inputs, there will be more to learn. In this section, I will briefly lay out what I see to be the major questions about productivity that the research agenda should address.[10]

What is the importance of demand? Productivity is typically thought of as a supply-side concept. As discussed in Section 2, it is the component of the production function that is not related to observable labor, capital, and intermediate inputs. But productivity as actually measured in producer microdata is generally reflects more than just supply-side forces. Because producer-specific prices are unobserved in most business-level microdata, output is typically measured by

[10] Conversations with John Haltiwanger were very helpful in writing this section.

revenue divided by an industry-level deflator. This means that within-industry price differences are embodied in output and productivity measures. If prices reflect in part idiosyncratic demand shifts or market power variation across producers—a distinct likelihood in many industries—then high "productivity" businesses may not be particularly technologically efficient. Much of the literature described above therefore documents the joint influence of productivity *and* demand factors that show up in within-industry price variation.

A new strand of research has begun to extend the productivity literature to explicitly account for such idiosyncratic demand effects as well. These new frameworks—see Das, Roberts, and Tybout (2007); Eslava et al. (2008); Foster, Haltiwanger, and Syverson (2008, 2009), and De Loecker (2009), for example—allow an additional and realistic richness in the market forces that determine producers' fates. The work to this point indicates that demand factors are indeed important. They exert a considerable influence on business growth and survival. And while many of the basic results above that have been checked while accounting for the supply-demand dichotomy have been robust, the results do suggest some reinterpretations of the effects of productivity as inferred from standard measures. The scope of issues that this new line of research is still small, however. Demand could play an important role in many more settings that have been hidden to this point due to measurement issues. Unwinding this knot is a top priority.

What is the role (or hope for) government policies that encourage productivity growth? Clearly, many of the productivity drivers discussed above can be influenced by government policies. This is especially true of the "external" drivers in the previous section—the elements of the market environment that can induce business to take actions to raise their productivity or that affect the Darwinian selection process that whittles out inefficient producers.

Several policy-related questions are prime targets for research. There have been many policy reforms (particularly in trade policy and market regulation design) that had plausibly productivity-enhancing effects. Many studies have evaluated specific reforms in isolation, taking the policy change as given. But a policy change, even one that moves in the right direction, may not necessarily be optimal. Alternative reforms, either in size or magnitude, might be more cost effective. Research has typically compared the effects of policy reforms to a null of no reform, but perhaps an equally important comparison is among possible reform alternatives. What type of reform is most effective for a given type of market or friction? What is the optimal size and timing of policy changes? These are the next set of questions the literature should chase in this area.

A related issue is why reforms, even if they are welfare enhancing in their productivity effects, don't always happen. There could be economic reasons for this. Established interests could be earning rents in the unreformed environment, and may be able to stave off reform, especially if its benefits are diffuse while its losses are concentrated. Characterizing the nature of these barriers to aggregate productivity gains—who wins, who loses, and by how much, for example—could be fruitful.

Which productivity drivers matter the most? The research described above has framed *which* factors might explain variation in productivity levels. The relative quantitative importance of each, however, is still unclear. Summarized succinctly, if we could easily measure these factors and add them to the production function, which would have the largest R^2?

Of course, it's quite likely that the quantitative impact of factors varies across industries or markets. A concomitant question, then, is which factors matter most in what sectors? Research that ties observable attributes of the industry's technology or demand structure to the quantitative importance of productivity-influencing factors would be an incredible advance in our ability to explain productivity growth.

What factors determine whether selection or within-producer growth is more important in a market/sector/industry? In many settings above, there was a prominent distinction between aggregate productivity growth coming from "within" (productivity growth at a given plant or firm) and "between" (reallocation-based selection across existing businesses or entry and exit) sources. Just as the relative quantitative contribution of various influences in producer-level efficiencies still needs to be characterized, so too does the relative importance of the within and between components in explaining aggregate productivity growth.

We do know some patterns already—that (at least in the U.S.) aggregate productivity growth in the retail sector seems to be almost exclusively from reallocation, for example. But of course nowhere near the full span of sectors and economies has yet been covered. More importantly, we do not yet have a good model of what sectoral features (again on either the supply or demand side) might determine the relative importance of each. Why is within-store productivity growth so small on average in retail, but not manufacturing, for example? Answering questions like this would go a long way to developing our understanding of how micro productivity differences drive the aggregate productivity movements.

What is the role of misallocation as source of variation in emerging economies? Productivity differences explain much of the per capita income variation across countries. As seen above, recent research with producer microdata is building the case that a substantial portion of these productivity gaps arise from poor allocation of inputs across production units in developing countries.

In some ways, this is a hopeful finding: these countries could become substantially more productive (and raise their incomes) by simply rearranging the inputs they already have. Not everything hangs on some unattainable technologies that are out of reach.

On the other hand, the result is depressing. While research has identified misallocation as a source of the problem, it hasn't really pinned down exactly what distortions create gaps between the social marginal benefits and costs of inputs across production units. It is hard to implement policies that close these gaps and the variation between them (i.e., reallocate inputs more efficiently) without knowing the nature of the gaps in the first place.

That said, there has been some early progress on this front. Witness the efforts to tie misallocation to various labor market policies. Much remains to be done, however, and this is a ripe area for further work.

Can we predict innovation based on market conditions? Here I speak of innovation broadly—product and process innovation, measured or unmeasured by formal R&D numbers. This question is in some ways a corollary to the one above about quantifying and predicting the split between within-producer and between-producer productivity growth. Within-productivity growth is in many cases not simply the passive accumulation of efficiency; it comes from active innovative activity on the part of producers as well. What market or technological factors determine how large innovative activity will be? Can we predict whether product or process innovation will dominate, based on market features?

The nature of intangible capital. Many of the primary drivers of productivity naturally create persistence in productivity levels at plants and firms. This can include learning-by-doing; innovative efforts; and in many cases investment in higher quality managerial, labor, or capital inputs. Explaining such persistence is easy under an interpretation of such productivity-building factors as investments by producers in a stock of intangible capital—know-how about their businesses that is embodied in the organization.

Understanding how such intangible capital stocks are built and sustained will shed light on many productivity-related issues, as well as speak toward active literatures on the subject in macroeconomics and finance. How much uncertainty is inherent in intangible capital investment? What is the distribution of rates of return across producers, and what predicts them? Is intangible capital fully excludable, or are there spillovers to other firms? How well do R&D measures capture investment in intangibles? Are there other proxies that could augment such measures?

Management vs. Managers. We know more about the role of management than before, but what about *managers*? Some good work on CEOs aside, we don't really know if good managerial practices enough to attain productivity gains, or whether they are complementary to the skills of those who implement them. If there are complements, what skills matter? Are they built by experience, tenure in the industry or on the job, education, or something else? Understanding these issues might also help pin down the causal nature of management practices. If good management practices reflect in large part the fact that they are what good managers do, then the causal impact might be limited. On the other extreme, if managers don't seem to matter at all, then it is quite likely that managerial practices have a strong causal impact on productivity.

6. Conclusion

The research into the productivity differences across businesses has come a long way since Bartelsman and Doms (2000) surveyed the literature a decade ago. We know more about what causes the measured differences in productivity, and how factors both internal and external to the plant or firm shape the distribution. These insights have been applied to research questions in numerous fields.

That said, there is still plenty to be learned. Fortunately, I see no sign that the rate at which researchers accumulate knowledge in this area is slowing. I am excited to see what the next several years bring in this research agenda, as the content of 2020's survey unfolds.

References

Abowd, John, John Haltiwanger, Ron Jarmin, Julia Lane, Paul Lengermann, Kristin McCue, Kevin McKinney, and Kristin Sandusky. 2005. "The Relationship Between Human Capital, Productivity and Market Value: Building Up From Microeconomic Evidence." In *Measuring Capital in the New Economy* (eds. Carol Corrado, John Haltiwanger and Daniel Sichel), Chicago: University of Chicago Press.

Ábrahám, Árpád and T. Kirk White. 2007. "The Dynamics of Plant-level Productivity in U.S. Manufacturing." Working paper, University of Rochester.

Acemoglu, Daron, Philippe Aghion, Claire Lelarge, John Van Reenen, and Fabrizio Zilibotti. 2007. "Technology, Information, and the Decentralization of the Firm." *Quarterly Journal of Economics*, 122(4): 1759-99.

Acemoglu, Daron and Joshua Linn. 2004. "Market Size in Innovation: Theory and Evidence from the Pharmaceutical Industry." *Quarterly Journal of Economics*, 119(3): 1049-90.

Amiti, Mary and Jozef Konings. "Trade Liberalization, Intermediate Inputs, and Productivity: Evidence from Indonesia." *American Economic Review*, 97(5): 1611-38.

Arnold, Jens, Giuseppe Nicoletti, and Stefano Scarpetta. 2008. "Regulation, Allocative Efficiency and Productivity in OECD Countries: Industry and Firm-level Evidence." OECD Economics Department Working Paper 616.

Arrow, Kenneth. 1962. "Economic Welfare and the Allocation of Resources for Inventions." In *The Rate and Direction of Inventive Activity*, R. Nelson, ed. Princeton, N.J.: Princeton University Press.

Aw, Bee Yan, Mark J. Roberts, and Daniel Yi Xu. 2009. "R&D Investment, Exporting, and Productivity Dynamics." Working paper.

Balasubramanian, Natarajan and Jagadeesh Sivadasan. 2009. "What Happens When Firms Patent? New Evidence from U.S. Economic Census Data." Working Paper.

Bartel, Ann, Casey Ichniowski and Kathryn Shaw. 2007. "How Does Information Technology Affect Productivity? Plant-Level Comparisons of Product Innovation, Process Improvement and Worker Skills." *Quarterly Journal of Economics*, *122*(4): 1721-1758.

Bartelsman, Eric and Mark Doms. 2000. "Understanding Productivity: Lessons from Longitudinal Microdata." *Journal of Economic Literature*, *38*(3): 569-594.

Bartelsman, Eric, John Haltiwanger and Stefano Scarpetta. 2008. "Cross Country Differences in Productivity: The Role of Allocative Efficiency." Working paper.

Bartelsman, Eric J., Jonathan Haskel, and Ralf Martin. 2008. "Distance to Which Frontier? Evidence on Productivity Convergence from International Firm-level Data." CEPR Discussion Paper 7032.

Benkard, C. Lanier. 2000. "Learning and Forgetting: The Dynamics of Aircraft Production." *American Economic Review*, 90(4): 1034-54.

Bernard, Andrew, J. Bradford Jensen and Peter Schott. 2006. "Trade Costs, Firms and Productivity." *Journal of Monetary Economics*, 53(5): 917-37.

Bernard, Andrew, Stephen Redding, and Peter Schott. Forthcoming. "Multiple-Product Firms and Product Switching." *American Economic Review*.

Bertrand, Marianne and Antoinette Schoar. 2003. "Managing with Style: The Effect of Managers on Firm Policies," *Quarterly Journal of Economics*, *118*(4): 1169–1208.

Bloom, Nick, Benn Eifert, Aprajit Mahajan, David McKenzie, John Roberts. 2009. "Management as a Technology: Evidence from India." Working Paper.

Bloom, Nick, Stephen Dorgan, John Dowdy, and John Van Reenen. 2007. "Management Practice and Productivity: Why They Matter." Center for Economic Performance and McKinsey policy release.

Bloom, Nick, Rafaella Sadun, and John Van Reenen. 2007. "Americans Do I.T. Better: US Multinationals and the Productivity Miracle." CEP Discussion Paper No. 788.

Bloom, Nick, Mark Schankerman and John Van Reenen. 2007. "Identifying Technology Spillovers and Product Market Rivalry." NBER Working Paper no. 13060.

Bloom, Nick and John Van Reenen. 2007. "Measuring and Explaining Management Practices Across Firms and Countries." *Quarterly Journal of Economics*, 122(4): 1351-1408.

Boning, Brent, Casey Ichniowski, and Kathryn Shaw. 2007. "Opportunity Counts: Teams and the Effectiveness of Production Incentives." *Journal of Labor Economics*, *25*(4): 613-50.

Bridgman, Benjamin, Shi Qi, and James A. Schmitz, Jr. 2007. "Does Regulation Reduce Productivity? Evidence from Regulation of the U.S. Beet-Sugar Manufacturing Industry during the Sugar Acts, 1934-74." Federal Reserve Bank of Minneapolis, Staff Report: 389.

Brown, J. David, John S. Earle, and Almos Telegdy. 2006. "The Productivity Effects of Privatization: Longitudinal Estimates from Hungary, Romania, Russia, and Ukraine." *Journal of Political Economy*, 114(1): 61-99.

Bushnell, James B. and Catherine Wolfram. 2007. "The Guy at the Controls: Labor Quality and Power Plant Efficiency." NBER Working Paper 13215.

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

Collard-Wexler, Allan. 2008 "Demand Fluctuations in the Ready-Mix Concrete Industry." Working Paper.

Cooper, William W., Lawrence M. Seiford, and Kaoru Tone. 2006. *Introduction to Data Envelopment Analysis and Its Uses*. New York: Springer.

Crespi, Gustavo, Chiara Criscuolo, Jonathan E. Haskel, and Matthew Slaughter. "Productivity Growth, Knowledge Flows, and Spillovers." NBER Working Paper No. 13959

Cummins, Jason G. and Giovanni L. Violante. 2002. "Investment-Specific Technical Change in the United States (1947-2000): Measurement and Macroeconomic Consequences." *Review of Economic Dynamics*, *5*(2): 243-84.

Das, Sanghamitra, Mark J. Roberts, and James R. Tybout. "Market Entry Costs, Producer Heterogeneity, and Export Dynamics." *Econometrica*, *75*(3), 2007, 837-73.

De Loecker, Jan. 2007. "Do Exports Generate Higher Productivity? Evidence from Slovenia." *Journal of International Economics*, *73*(1), 69-98.

De Loecker, Jan. 2009. "Product Differentiation, Multi-Product Firms and Estimating the Impact of Trade Liberalization on Productivity." Working Paper.

Disney, Richard, Jonathan Haskel, and Ylva Heden. 2003a. "Entry, Exit and Establishment Survival in UK Manufacturing." *Journal of Industrial Economics*, *51*(1): 91-112.

Disney, Richard, Jonathan Haskel, and Ylva Heden. 2003b. "Restructuring and Productivity Growth in UK Manufacturing." *Economic Journal*, *113*(489): 666-94.

Doraszelski, Ulrich and Jordi Jaumandreu. 2009. "R&D and Productivity: Estimating Endogenous Productivity." Working paper.

Eaton, Jonathan and Samuel Kortum. 2002. "Technology, Geography, and Trade." *Econometrica*, *70*(5): 1741-79.

Eslava, Marcela, John C. Haltiwanger, Adriana Kugler and Maurice Kugler. 2004. "The Effects of Structural Reforms on Productivity and Profitability Enhancing Reallocation: Evidence from Colombia." *Journal of Development Economics*, *75*(2): 333-72.

Eslava, Marcela, John C. Haltiwanger, Adriana Kugler and Maurice Kugler. 2008. "Plant Survival, Market Fundamentals and Trade Liberalization." Working Paper.

Fabrizio, Kira, Nancy Rose and Catherine Wolfram. 2007. "Do Markets Reduce Costs? Assessing the Impact of Regulatory Restructuring on U.S. Electric Generation Efficiency." *American Economic Review*, *97*(4), 1250-1277.

Fernandes, Ana M. 2007. "Trade Policy, Trade Volumes and Plant-Level Productivity in Colombian Manufacturing Industries." *Journal of International Economics*, 71(1): 52-71.

Foster, Lucia, John Haltiwanger and C.J. Krizan. 2001. "Aggregate Productivity Growth: Lessons from Microeconomic Evidence." *NBER Studies in Income and Wealth, vol. 63: New Developments in Productivity Analysis*. Chicago and London: University of Chicago Press, 303-63.

Foster, Lucia, John C. Haltiwanger and C.J. Krizan. 2006. "Market Selection, Restructuring and Reallocation in the Retail Trade Sector in the 1990s." *Review of Economics and Statistics*, 88(4): 748-58.

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

Foster, Lucia, John C. Haltiwanger, and Chad Syverson. 2009. "The Slow Growth of New Plants: Learning about Demand?" Working Paper.

Fox, Jeremy and Valerie Smeets. 2009. "Does Input Quality Drive Measured Differences in Firm Productivity?" Working Paper.

Greenstone, Michael, John A. List, and Chad Syverson. 2009. "The Effects of Environmental Regulation on the Competitiveness of U.S. Manufacturing." Working Paper.

Griffith, Rachel, Rupert Harrison and John Van Reenen. 2007. "How Special is the Special Relationship? Using the Impact of U.S. R&D Spillovers on U.K. Firms as a Test of Technology Sourcing." *American Economic Review*, 96(5): 1859-75.

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

Haltiwanger, John, Stefano Scarpetta, and Helena Schweiger. 2008. "Assessing Job Flows Across Countries: The Role of Industry, Firm Size and Regulations." NBER Working Paper 13920.

Hamilton, Barton H., Jack A. Nickerson and Hideo Owan. 2003. "Team Incentives and Worker Heterogeneity: An Empirical Analysis of the Impact of Teams on Productivity and Participation." *Journal of Policital Economy, 111*(3): 465-497.

Holmes, Thomas J., David K. Levine, and James A Schmitz Jr. 2008. "Monopoly and the Incentive to Innovate When Adoption Involves Switchover Disruptions." NBER Working Paper 13864.

Hortacsu, Ali and Chad Syverson. 2007. "Cementing Relationships: Vertical Integration, Foreclosure, Productivity, and Prices." *Journal of Political Economy*, *115*(2): 250-301.

Hortaçsu, Ali and Chad Syverson. 2009. "Why Do Firms Own Production Chains?" Working Paper.

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

Ichniowski, Casey and Kathryn Shaw. 2003. "Beyond Incentive Pay: Insiders' Estimates of the Value of Complementary Human Resource Management Practices." *Journal of Economic Perspectives*, *17*(1): 155-78.

Ichniowski, Casey, Kathryn Shaw and Giovanna Prennushi. 1997. "The Effects of Human Resource Management Practices on Productivity: A Study of Steel Finishing Lines." *American Economic Review*, *87*(3): 291-313.

Ilmakunnas, Pekka, Mika Maliranta, and Jari Vainiomäki. 2004. "The Roles of Employer and Employee Characteristics for Plant Productivity." *Journal of Productivity Analysis*, 21(3): 249-276.

Jorgenson, Dale W., Mun S. Ho and Kevin J. Stiroh. 2005. *Information Technology and the American Growth Resurgence*. Cambridge: MIT Press.

Jorgenson, Dale W., Mun S. Ho and Kevin J. Stiroh. 2008. "A Retrospective Look at the U.S. Productivity Growth Resurgence." *Journal of Economic Perspectives*, *22*(1): 3-24.

Keller, Wolfgang and Stephen R. Yeaple. Forthcoming. "Multinational Enterprises, International Trade, and Productivity Growth: Firm-Level Evidence from the United States." *Review of Economics and Statistics*.

Kellogg, Ryan. 2009. "Learning by Drilling: Inter-Firm Learning and Relationship Persistence in the Texas Oilpatch." NBER Working Paper 15060.

Klette, Tor Jakob and Samuel Kortum. 2004. "Innovating Firms and Aggregate Innovation." *Journal of Political Economy*, *112*(5): 986–1018.

Knittel, Chirstopher R. 2002. "Alternative Regulatory Methods and Firm Efficiency: Stochastic Frontier Evidence the US Electricity Industry." *Review of Economics and Statistics*, 84(3): 530-540.

Lazear, Edward P. 2000. "Performance Pay and Productivity." *American Economic Review, 90*(5): 1346-61.

Lentz, Rasmus and Dale T. Mortensen. 2009. "An Empirical Model of Growth through Product Innovation." *Econometrica*, *76*(6): 1317–73.

Levinsohn, James and Amil Petrin. 2003. "Estimating Production Functions Using Inputs to Control for Unobservables." *Review of Economic Studies*, *70*(2): 317-41.

Maksimovic, Vojislav and Gordon Phillips. 2002. "Do Conglomerate Firms Allocate Resources Inefficiently across Industries? Theory and Evidence." *Journal of Finance*, 57(2), 721-767.

Maksimovic, Vojislav and Gordon Phillips. 2001. "The Market for Corporate Assets: Who Engages in Mergers and Asset Sales and Are There Efficiency Gains?" *Journal of Finance*, 56(6), 2019-2065

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

Mas, Alexandre. 2008. "Labour Unrest and the Quality of Production: Evidence from the Construction Equipment Resale Market." *Review of Economic Studies*, 75(1): 229-58.

Melitz, Marc J. 2003. "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity." *Econometrica*, 71(6): 1695-1725.

Moretti, Enrico. 2004. "Workers' Education, Spillovers, and Productivity: Evidence from Plant-Level Production Functions." *American Economic Review*, *94*(3): 656-690.

Muendler, Marc A. 2004. "Trade, Technology, and Productivity: A Study of Brazilian Manufacturers, 1986-1998." CESifo Working Paper, 1148.

Nicoletti, Giuseppe and Stefano Scarpetta. 2005. "Regulation and Economic Performance: Product Market Reforms and Productivity in the OECD." OECD Economics Department Working Paper No. 460.

Oliner, Stephen D., Daniel Sichel, and Kevin Stiroh. 2007. "Explaining a Productive Decade." *Brookings Papers on Economic Activity*, 2007 Issue 1: 81-137.

Olley, G. Steven and Ariel Pakes. 1996. "The Dynamics of Productivity in the Telecommunications Equipment Industry." *Econometrica*, 64(6), 1263-97.

Pavcnik, Nina. 2002. "Trade Liberalization, Exit, and Productivity Improvement: Evidence from Chilean Plants." *Review of Economic Studies*, 69(1): 245-76.

Petrin, Amil and James Levinsohn. 2008. "Measuring Aggregate Productivity Growth Using Plant-Level Data." Working Paper no. 11887.

Petrin, Amil and Jagadeesh Sivadasan. 2006. "Job Security Does Affect Economic Efficiency: Theory, a New Statistic, and Evidence from Chile." Working paper.

Sakellaris, Plutarchos and Daniel J. Wilson. 2004. "Quantifying Embodied Technological Change." *Review of Economic Dynamics*, *7*(1): 1-26.

Schmitz, James A., Jr. 2005. "What Determines Productivity? Lessons from the Dramatic Recovery of the U.S. and Canadian Iron Ore Industries following Their Early 1980s Crisis." *Journal of Political Economy*, 113(3), 582-625.

Schoar, Antoinette. 2002. "The Effect of Diversification on Firm Productivity." *Journal of Finance*, 62(6): 2379-2403.

Syverson, Chad. 2004a. "Product Substitutability and Productivity Dispersion." *Review of Economics and Statistics*, 86(2): 534-550.

Syverson, Chad. 2004b. "Market Structure and Productivity: A Concrete Example." *Journal of Political Economy*, 112(6): 1181-1222.

Syverson, Chad. 2007. "Prices, Spatial Competition, and Heterogeneous Producers: An Empirical Test." *Journal of Industrial Economics*, 55(2): 197-222.

Thornton, Rebecca Achee and Peter Thompson. 2001. "Learning from Experience and Learning from Others: An Exploration of Learning and Spillovers in Wartime Shipbuilding." *American Economic Review*, 91(5): 1350-68.

UK Office of Fair Trading. 2007. "Productivity and Competition: An OFT Perspective on the Productivity Debate."

van Ark, Bart, Mary O'Mahony, and Marcel P. Timmer. 2008. "The Productivity Gap between Europe and the United States: Trends and Causes." Journal of Economic Perspectives, *22*(1): 25-44.

Van Biesebroeck, Johannes. 2003. "Productivity Dynamics with Technology Choice: An Application to Automobile Assembly." *Review of Economic Studies*, 70(1): 167-98.

Van Biesebroeck, Johannes. 2005. "Exporting Raises Productivity in sub-Saharan African Manufacturing Firms." *Journal of International Economics*, 67(2): 373-91.

Van Biesebroeck, Johannes. 2008. "The Sensitivity of Productivity Estimates: Revisiting Three Important Productivity Debates." *Journal of Business and Economic Statistics*, 26(3): 321-38.

Verhoogen, Eric A. 2008. “Trade, Quality Upgrading and Wage Inequality in the Mexican Manufacturing Sector.” *Quarterly Journal of Economics*, 123(2): 489-530.

Walker, Francis A. 1887. “The Source of Business Profits.” *Quarterly Journal of Economics*, *1*(3): 265-288.