

International Technology Diffusion

WOLFGANG KELLER¹

1. Introduction

Productivity differences explain a large part of the variation in incomes across countries, and technology plays a key role in determining productivity (William Easterly and Ross Levine 2001; Robert Hall and Charles Jones 1999; Edward Prescott 1998).² For most countries, foreign sources of technology account for 90 percent or more of domestic productivity growth. At present, only a handful of rich countries account for most of the world's creation of new technology.³ The pattern of worldwide technical change is thus determined in large part by international technology diffusion. What determines the extent of technology flows between countries and the means by which

technology diffuses? This paper surveys what is known about the extent of international technology diffusion and the channels through which technology is spread.

International technology diffusion is important because it determines the pace at which the world's technology frontier may expand in the future. Do countries' incomes converge over time or not? The answer turns on whether technology diffusion is global or local, as several widely used models show (Gene Grossman and Elhanan Helpman 1991; Peter Howitt 2000; Luis Rivera-Batiz and Paul Romer 1991).⁴ A better understanding of technical change therefore provides insights on the likelihood that certain less-developed countries will catch-up to rich countries.

Some of the major channels for technology diffusion across countries may include international trade and foreign direct investment (FDI), both of which are discussed in this survey. There is evidence that imports are a significant channel of technology diffusion. At the same time, the evidence for benefits associated with exporting is weaker. The importance of FDI has long been emphasized in the case-study literature, and recently that

¹ University of Texas, Austin; NBER and CEPR. I have benefited from discussions related to this paper with Bob Baldwin, Elhanan Helpman, Danny Quah, Andres Rodriguez-Claré, Dani Rodrik, Jim Tybout, and Stephen Yeaple. Oded Galor, Peter Howitt, Sam Kortum, Pete Klenow, Pierre Mohnen, Jim Rauch, Carol Shiue, Dan Treffer, and seminar participants at Industry Canada, three anonymous referees and John McMillan provided helpful comments. Support by NSF grant SES 9818902 is gratefully acknowledged.

² In addition, institutions and policy have been emphasized as well; see, e.g., Daron Acemoglu, Simon Johnson, and James Robinson (2001); Abhijit Banerjee and Lakshmi Iyer (2003); and Carol Shiue and Wolfgang Keller (2004).

³ The G-7 countries (the largest seven industrialized countries) accounted for about 84 percent of the world's research and development (R&D) spending in 1995, for example; their share in world GDP was only 64 percent. The world's distribution of production is much less skewed.

⁴ The models differ in that some predict global diffusion would lead to convergence in productivity levels while others predict convergence in productivity growth rates; see Philippe Aghion and Peter Howitt (1998) and Charles Jones (1995), as well as the framework in section 2.

evidence has been complemented by some micro-econometric findings.

Despite the global reach of computer programs, there is no indication that a global pool of technology yet exists. This localized character of technology suggests that an important component of it is tacit in nature. Although the relative importance of international technology diffusion appears to be increasing along with higher levels of economic integration, international diffusion of technology is neither inevitable nor automatic. Domestic technology investments are necessary.

Before discussing some of the evidence in more detail (sections 6–8), section 2 presents a conceptual framework that will guide this discussion. The following sections review various measures of technology (section 3) and technology diffusion (section 4) that are available, while section 5 discusses the pros and cons of a number of empirical approaches that have been employed. Section 9 summarizes the overall evidence on international technology diffusion, and section 10 contains a synthesis and some suggestions for future work.

2. International Technology Diffusion: A Framework for Analysis

Theories of endogenous technical change of the early 1990s (Aghion and Howitt 1992; Grossman and Helpman 1991; Romer 1990; and Paul Segerstrom, T.C.A. Anant, and Elias Dinopoulos 1990) emphasize two aspects of technology:^{5, 6}

⁵ See Grossman and Helpman (1995) and Aghion and Howitt (1998) for broader overviews. I will also discuss some related work on learning-by-doing and human-capital accumulation that falls into the broad category of models of knowledge accumulation.

⁶ Many of these ideas have been discussed in the literature before; important contributors include Paul David, Giovanni Dosi, Robert Evenson, Jan Fagerberg, Richard Nelson, Keith Pavitt, Nathan Rosenberg, Luc Soete, Sidney Winter, Larry Westphal, and others (see Fagerberg 1994, and Robert Evenson and Larry Westphal 1995 for overviews). What distinguishes the recent work is that it includes fully specified general equilibrium models. This means that these technology effects can in principle be estimated in a well-defined framework.

1. Technology is non-rival in the sense that the marginal costs for an additional agent to use the technology are negligible.⁷
2. The return to technological investments is partly private and partly public.

Point 1 distinguishes technology from rival factor inputs such as human and physical capital; the latter can only be used by one firm at a time, or put differently, the marginal costs of using the same factor somewhere else are infinite. Point 2 highlights that while private returns must be strong enough to keep innovation ongoing, technological investments often create benefits to individuals other than the inventor.⁸ These external effects are called technology, or knowledge, spillovers. As an example, the introduction of one product might speed up the invention of a competing product, because the second inventor can learn from the first by carefully studying the product or its product design (the “blueprint”).

The focus of this paper is the empirics of international technology diffusion and its effect on productivity.⁹ Before discussing the empirical results, I will lay out a model of international technology diffusion, based on Jonathan Eaton and Samuel Kortum (1999) and Kortum (1997), to help clarify the issues.

Consider a model with $n=1, \dots, N$ countries. Output in n at time t , Y_{nt} , is produced by combining intermediate inputs subject to a constant-returns-to-scale Cobb-Douglas production function:

$$\ln(Y_{nt} / J) = J^{-1} \int \ln[Z_{nt}(j)X_{nt}(j)]dj, \quad (1)$$

where $X_{nt}(j)$ is the quantity of intermediate

⁷ Other authors have coined the terms “perfectly expandable” (David 1992) and “infinitely expandable” (Danny Quah 2001a,b) to positively define this characteristic.

⁸ The focus of this work is on technical change driven by private, not public, research. The majority of all research and development (R&D) is privately funded, although public R&D is substantial in many countries.

⁹ International technology diffusion in this paper refers to both technology spillovers as well as non-spillovers. As will be discussed below, the two are often difficult to separate.

input j at time t in country n and $Z_{nt}(j)$ is the quality of that input. This quality is increased over time through new ideas, or technologies. The range of inputs is the same across countries and constant over time. Output is homogeneous and tradable, while intermediates are nontraded.

Each input j is produced anywhere by a Cobb-Douglas combination of capital $K(j)$ and labor $L(j)$, $X(j) = K(j)^\phi L(j)^{1-\phi}$, with $\phi \in [0, 1]$. New technologies are the result of research effort. Consider country i at time t with L_{it} workers. If a fraction s_{it} of them are engaged in research, they create new technologies at rate $\alpha_{it} s_{it}^\beta L_{it}$, where α_{it} is the research productivity, and $\beta > 0$ characterizes the R&D talent distribution.

Each technology has three dimensions: (i) quality, (ii) use, and (iii) diffusion time lag. First, the quality of a technology is a random variable drawn from the cumulative distribution $F(q)$, $F(q) = 1 - q^\theta$, $\theta > 0$. This quality is common to all countries to which the technology diffuses. Second, the technology can be used only in one intermediate sector, which is determined randomly.

Third, technologies become productive only once they have diffused. Diffusion is a stochastic process, with a mean diffusion lag between country i and country n of ε_{ni}^{-1} , or, ε_{ni} measures the speed of technology diffusion between the countries. The rate at which technologies diffuse to country n at time t is given by

$$\dot{\mu}_{nt} = J^{-1} \sum_{i=1}^N \underbrace{\varepsilon_{ni}}_A \int_{-\infty}^t \underbrace{e^{-\varepsilon_{ni}(t-s)} \alpha_{is} s_{is}^\beta L_{is}}_B ds, \quad (2)$$

where A in equation (2) is the bilateral speed of diffusion, and B is country i 's cumulative output of technologies up to time t .¹⁰

Even though every technology will eventually be in every country (assuming $\varepsilon_{ni} > 0$), a given technology may not be employed in

production, because by the time it has diffused it is dominated by a higher quality. The stock of technologies in country n at time t is $\mu_{nt} = \int_{-\infty}^t \dot{\mu}_{ns} ds$. The fraction of intermediates in country n at time t that are below some quality threshold \bar{q} is decreasing in μ_{nt} , because the higher is the number of diffused technologies, the greater is the likelihood that the quality in country n 's sectors is relatively high.

Let A_{nt} be the geometric mean of qualities across sectors. Final output (1) is maximized when factors are evenly divided across all sectors, and it is given by $A_{nt} K_{nt}^\phi [L_{nt}(1-s_{nt})]^{1-\phi}$. This means that A_{nt} is equal to output divided by a factor-share weighted sum of inputs, or total factor productivity (TFP).¹¹

To obtain a steady-state equilibrium where μ_{nt} grows at the same constant rate in all countries in the long run, we assume that the research productivity α_{it} is given by $\alpha_{it} = \alpha (\mu_{it}/\bar{\mu}_i) \bar{\mu}_i^\gamma$; $\mu_t = \sum_{i=1}^N \mu_{it}$, $\gamma \leq 1$, and $\alpha > 0$. The higher is γ , the greater is the strength of the international research spillovers.¹² Finally, steady-state relative TFP levels can be computed as

$$\frac{A_{nt}}{A_{Nt}} = \left(\frac{\mu_{nt}}{\mu_{Nt}} \right)^{1/\theta}, \quad n = 1, \dots, N-1, \quad (3)$$

and world TFP growth is proportional to

$$g = \frac{\dot{\mu}_n}{\mu_n} = \frac{\alpha}{J} \sum_{i=1}^N \frac{\varepsilon_{ni}}{\varepsilon_{ni} + g} \frac{\mu_i}{\mu_n} s_i^\beta L_i, \quad n = 1, \dots, N, \quad (4)$$

where (μ_i/μ_n) and the share of labor in research, $s_i L_i$, are constant in steady-state. Equation (4) says that world growth is proportional to a weighted sum of research efforts in all countries.

¹¹ Eaton and Kortum (1999) show that due to imperfect competition, TFP is not identical to A_{nt} , but it is proportional to it; also see their paper for the exact definition of A_{nt} .

¹² For simplicity, we set γ to equal one. The parameter γ is related to the question of whether the model has "scale effects"; see, e.g., Jones (1995), and Aghion and Howitt (1998) ch. 12, for more details.

¹⁰ Past R&D is discounted because the technologies might have been superseded by higher-quality, more recent technologies.

The model has a number of implications:

- 1) Foreign R&D raises domestic TFP: An increase in foreign research leads to a greater inflow of technologies and higher TFP. This holds in both the short and the long-runs (equations 2, 3 respectively).¹³
- 2) Receiving technology from abroad: A given research effort abroad has a greater effect on domestic TFP, the faster foreign technologies diffuse to the domestic economy (ε_{ni} in equation 2).
- 3) Global sources of technology: A country is important in determining the world's rate of growth if it has (i) a relatively high share of the world's research labor and technologies ($s_i L_i$ and μ_i in equation 4), and/or (ii) a relatively high rate of technology diffusion to other countries (the weight $\varepsilon_{ni}/(\varepsilon_{ni}+g)$ in equation 4 is increasing in ε_{ni}).

International technology diffusion in this framework captures the notion that product quality innovations of country i become available in country n at rate ε_{ni} . Moreover, this is a spillover, because although research is costly in country i (it withdraws labor from final output production), conditional on the technology having diffused, it can be used to produce final output there without incurring additional costs. The diffusion of technology from country i also raises the productivity of research in country n due to the international research spillover (as long as $\gamma > 0$). An alternative, partial-spillover interpretation may be that the diffusion lag ε_{ni}^{-1} is indicative of the strength of costly country- n technology investment to acquire the technology from country i .

In either case, it becomes clear that our ability to give the diffusion rates ε_{ni} a structural interpretation that can be empirically tested is crucial for making further progress. Which are the major channels of international technology diffusion, or, what is behind the rates ε_{ni} ? This is the question that

is addressed by the major portion of the literature to date, and will be reviewed in sections 6–9. A number of issues are important when thinking about the strength of different diffusion channels. I will briefly turn to them now.

How do technologies move from one country to another? One possibility is through international trade in intermediate goods (Rivera-Batiz and Romer 1991; Grossman and Helpman 1991; and Eaton and Kortum 2002).¹⁴ Employing a foreign intermediate good in final-output production involves the implicit usage of the technology in embodied form. There is a spillover in this process of international technology diffusion to the extent that the intermediate good costs less than its opportunity costs—which include the R&D costs of product development. If trade is an important diffusion channel, one should expect that international technology diffusion is geographically localized, because of the well-documented fact that trade falls with geographic distance (e.g., Edward Leamer and James Levinsohn 1995). In any case, this is a relatively weak form of technology diffusion, because the technology as such is not available domestically—only the manufactured outcome of it is.

Of course, international technology diffusion is not limited to the channel of trade. In principle, just as researchers today “stand on the shoulders” of researchers of the past, one might expect researchers in one country to directly benefit from research conducted in other countries.¹⁵ The model above captures this in that research productivity α_{it} is proportional to the global stock of technologies ($\bar{\mu}_t$), which is increasing in world R&D. This is an international R&D spillover (i.e., no domestic opportunity

¹³ The same equations indicate that domestic R&D raises TFP as well.

¹⁴ In the model above, this is not explicitly modeled, but it would suggest that ε_{ni} is related to intermediate inputs trade; see the analysis in Eaton and Kortum (1996).

¹⁵ See Ricardo Caballero and Adam Jaffe (1993) for empirical results in a closed economy.

costs).¹⁶ Moreover, this seems to be a stronger form of diffusion than the intermediate goods trade discussed above: if the foreign technology raises the productivity of domestic researchers, it suggests full mastery of the technology, as opposed to only the ability to use the technology that is embodied in the intermediate good.

Which patterns of technology diffusion should thus be expected? The first type of diffusion above suggests that it should follow the pattern of intermediate goods trade. International R&D spillovers appear to be more difficult to pin down. They are the result of acquiring technology that is not tied to any particular form. In addition to libraries, conferences and other sources, technology may be stored in as little as a couple of bytes, and thus sent over the internet to practically every country in the world at close to zero marginal cost. At the same time, this does not mean that this technology diffusion in fact creates equal levels of technological knowledge in all countries, for the following reasons.

First, irrespective of feasibility, it is not in the interest of the original inventor—who has incurred the R&D cost—to send it to others at no charge. On the contrary, the inventor may decide to spend additional resources to keep the technology secret. Second, even if some domestic technology becomes available abroad, it may be possible to preclude others from using it by patenting the technology. Third, even if the technology is non-rival and can be moved from one country to another at zero marginal costs, operating the technology efficiently often requires making costly investments in terms of complementary skills (see section 8).

Another issue is that technology may in fact not be transferable at essentially zero cost even if all parties involved desire this, such as multinational parents providing

technology to their subsidiaries. David Teece (1977) estimates that the costs of such within-firm transfers were on average almost 20 percent of the total project costs in the cases he analyzed.¹⁷ One reason for that is that only the broad outlines of technology are codified—the remainder remains “tacit” (Michael Polanyi 1958).¹⁸ In this view, knowledge is to some extent tacit because the person who is actively engaged in a problem-solving activity cannot necessarily define (and hence prescribe) what exactly she is doing. Technology is only partially codified because it is impossible or at least very costly to fully codify it.

What does this imply for the channels of international technology diffusion? Polanyi (1958, p. 53) argues that tacit knowledge can be passed on only “by example from master to apprentice.” A broader view is that non-codified knowledge is often transferred through person-to-person demonstrations and instructions (David 1992). The most effective way of doing this, despite telephone, video, and other remote methods, is face-to-face interaction. This, however, means that technology diffusion implies incurring the costs of moving teacher and student into the same geographic location.

We can summarize the implications of this for the pattern of international technology diffusion as follows.

- 1) The partial codified nature of technology means that technology diffusion will be incomplete, and technology stocks in different countries will vary. Diffusion will tend to be more geographically localized the higher is the non-codified share in total technology.¹⁹

¹⁷ For transferring machinery equipment technology, this share was 36 percent (Teece 1977, p. 248); see also Eric von Hippel (1994).

¹⁸ See also Kenneth Arrow (1969), David (1992), and Evenson and Westphal (1995), and references given there.

¹⁹ This follows because the costs of moving people in space are typically increasing in geographic distance; see von Hippel (1994) and Maryann Feldman and Frank Lichtenberg (1997) for empirical support.

¹⁶ Recall that we assume γ is equal to 1. More generally, there is an international R&D spillover as long as $\gamma > 0$.

2) Because international economic activities (trade, FDI, etc.) lead to additional contacts with foreign persons who may possess advanced technological knowledge (exporter, importer, engineers, researchers), this may stimulate the diffusion of (non-codified) foreign technology.

Trade and interaction with multinationals may thus lead not only to technology diffusion of the limited kind (technology embodied in intermediate goods), but it may also raise the probability of international R&D spillovers.²⁰

I now turn to how recent empirical work has studied these issues. The following two sections discuss the data that is available on technology (section 3) and technology diffusion (section 4).

3. *Measures of Technology*

Technology is an intangible that is difficult to measure directly. Three widely used indirect approaches are to measure (1) inputs (R&D), (2) outputs (patents), and (3) the effect of technology (higher productivity).

First, internationally comparable data on R&D expenditures have been published by the Organization of Economic Cooperation and Development (OECD) since about 1965. According to the OECD's definition (OECD 2002), only about two dozen relatively rich countries report substantial amounts of R&D, because the definition captures primarily resources spent towards innovation, and not those spent on imitation and technology adoption. Technology investments of middle and poor countries can therefore typically not be analyzed using R&D data.²¹

²⁰ This process could also in part be sequential: (1) using foreign technology in embodied form, and gradually (2) learning of the technology per se, (3) imitation, and (4) (original) innovation.

²¹ However, R&D data becomes more widely available as countries' incomes are rising. There is also increasingly information on poorer countries because surveys encompass the R&D conducted by affiliates of multinational companies located abroad; see e.g. NSF (2004) for R&D expenditures of U.S. firms in China. The main OECD R&D statistics are on a geographic, not ownership, basis.

A drawback of R&D as a measure of technology is that it ignores the stochastic nature of the process of innovation. The current flow of R&D expenditures is thus a noisy measure of technology improvements in that period. Many authors construct R&D stocks from the flows using the perpetual inventory method.²² Beyond year-to-year noise, the return to R&D expenditures may vary substantially across agents and over time, which limits comparability. One important aspect of this is that the return to publicly funded R&D is lower than the return to privately funded R&D (Frank Lichtenberg 1993), so that many studies focus on business research and development spending.

Second, a patent gives its holder a temporary legal monopoly to use an innovation in a specific market at the price of public disclosure of technical information in the patent description.²³ An innovation must be sufficiently important to be worthy of a patent, which is judged by a trained official (called patent examiner). Relative to R&D, patents have the advantage that patent data has been collected for a longer time (more than 150 years for some countries), and also poorer countries have a substantial number of patents (see WIPO 2003).

There are some issues with using patent data as well. First, a small number of patents accounts for most of the value of all patents. This means that simple patent counts may not measure technology output well. Recent work has addressed this issue in part by using citation-weighted patent data (see Adam Jaffe and Manuel Trajtenberg 2002).²⁴ Second, the patent decision is an act of choice on the part of the firm, and a large set of innovations is not

²² This requires an assumption on the rate of R&D depreciation; standard values range between 0 and 10 percent; see Zvi Griliches (1995).

²³ See Griliches (1990) for more discussion.

²⁴ Typically, a patent contains one or more references to other patents that indicate which earlier knowledge was used to come up with the technology underlying the present patent application; these references are called patent citations.

ever patented.²⁵ And third, if technology is in part non-codifiable, as argued above, patent statistics will necessarily miss that part.

The third measure of technology discussed here is total factor productivity (TFP). The idea, well-known since the 1950s, is that if one subtracts from output the contribution of inputs such as labor and capital, the remainder is due to the factor “technology.”²⁶ A simple example is the term A in the Cobb-Douglas production function $Y = AK^\alpha L^{1-\alpha}$. Other TFP measures are more general and have certain desirable properties that are important for comparability (e.g. the “superlative” index proposed by Richard Caves, Laurits Christensen, and Erwin Diewert 1982).

In contrast to R&D and patents, TFP is a derived measure of technology, as it is computed from data on inputs and output. This introduces measurement error and perhaps biases, because the appropriate data on inputs and outputs is rarely, if ever, available. Hajime Katayama, Shihua Lu, and James Tybout (2003), for example, show that the use of (1) real sales revenues, (2) depreciated capital spending, and (3) real input expenditures; instead of (unavailable) data on the physical quantities (1) of output, (2) of capital, and (3) of intermediate inputs, as is frequently done, will often confound higher productivity with higher mark-ups.²⁷ Other factors might thus contaminate the use of TFP as a measure of technological efficiency, which ultimately goes back to the concern that TFP is constructed as a residual, and may potentially capture a host of spurious influences (Moses Abramovitz 1956).

²⁵ The fact that there is variation in the propensity to patent in several dimensions—for instance, it is profitable to patent only in a large-enough market—could also provide important information.

²⁶ See Charles Hulten (2000) for more details on TFP.

²⁷ This could have major social welfare implications as consumers benefit from higher productivity but not from higher mark-ups. See also Tor Klette and Griliches (1996), as well as Eric Bartelsman and Mark Doms (2000) for a recent analysis of micro productivity data.

Because of these difficulties in computing TFP, researchers have pursued a number of strategies. One is to consider changes in TFP as opposed to TFP levels. This will help in identifying technology (or rather, technical change) if spurious factors do not change over time, or more generally, if they change less than technology. For example, in Katayama, Lu, and Tybout’s (2003) case from above, if a firm faces higher adjustment costs to changing its mark-up, this will reduce the mark-up variability in equilibrium. A second strategy has been to employ TFP measures in studies of technical change together with data on R&D (e.g., Griliches 1984). By establishing a relationship between TFP changes and its presumed major cause, R&D spending, the likelihood of measuring changes in technology inappropriately is substantially reduced.

I now turn to measures of technology diffusion and spillovers.²⁸

4. *Measures of International Technology Diffusion*

The diffusion of technology involves both market transactions and externalities (see above), and obtaining data on the former is fairly straightforward. For instance, firms make royalty payments for their use of patents, licenses, and copyrights, and this data is available for major countries in the international services balance (e.g., OECD 2003). Many economists believe, though, that most international technology diffusion occurs not through market transactions but instead through externalities (spillovers).²⁹

²⁸ Sometimes firm-level datasets provide information on the fraction of computer-controlled machinery, or whether a firm satisfies certain technological standard; either of these would also provide information about the technological sophistication of the firm.

²⁹ One reason is that technology is not fully codifiable, so that complete contracts cannot be written. Another reason is market failure in the market for technology due to asymmetric information: the buyer does not know (the productivity of) the technology, while the seller cannot commit to truthful claims about it. This is considered to be a major reason why international technology transfers are often internalized within firms (between multinational parent and subsidiary); see, e.g., William Ethier (1986).

Naturally, data on spillovers does not exist. Measures that are related to it do exist, but typically they capture spillovers only partially, because the measures do not account for costs of acquisition (learning). For instance, if one patent application cites an earlier patent, this generally indicates that the applicant has benefited from the earlier patent.³⁰ At the same time, it is impossible to know how large these benefits are net of the learning costs that the patent applicant had to incur.

Among the different methods that try to measure international spillovers, the largest set of papers employs international R&D spillover regressions. In one set of papers, if R&D of firm j is positively correlated with TFP in firm i , all else equal, this is consistent with international technology spillovers from firm j to firm i (e.g., Keller 2002a).³¹ A variant of this approach replaces TFP by the number of patents (e.g. Lee Branstetter 2001b), and Giovanni Peri (2002) presents a hybrid approach by relating patents in region i to patents in other regions, where the latter is instrumented by R&D expenditures.

There are two alternatives to this basic approach, a generalization and a simplification. In the former, a particular channel of technology diffusion is added to the analysis. David Coe and Elhanan Helpman (1995), e.g., analyze the relationship between productivity and foreign R&D conditional on imports from that foreign country. The other alternative to the international R&D spillover regressions is to relate productivity not to foreign R&D, but to other measures of foreign activity; Brian Aitken and Ann Harrison (1999), e.g., study the correlation of inward FDI and domestic firm productivity (so-called FDI spillover regressions).

A common concern in these studies is that, for various reasons, the estimate might give

us correlations but not causal effects. These, as well as issues in other approaches, such as estimates of spillovers as the parameter ε_{ni} in the model of section 2, are discussed in the following section.

5. Empirical Methods

5.1 Case Studies

Case studies can offer a rich description of the setting and the major factors that determine international technology diffusion. For instance, the paper by Felipe Larrain, Luis Lopez-Calva, and Andres Rodriguez-Claré (2000) analyzes whether Intel's foreign direct investment into Costa Rica in the 1990s generated any technology spillovers. The authors have spoken to the main actors, including Intel's top management and the government of Costa Rica, about their concerns and motivations. This is informative. The major limitation of case studies is that one does not know how general one particular case is. Nevertheless, there seems to be a role for case study research, especially if the study is—as in the Intel case—backed up by quantitative evidence.³²

5.2 Econometric Studies

5.2.1 Association Studies

In association studies, authors ask whether a specific foreign activity (FA) leads to a particular domestic technology outcome (DTO):

$$DTO = f(X, FA) + u \quad (5)$$

In equation (5), X is a vector of control variables, and u is a regression error. For instance, Aitken and Harrison (1999) examine the importance of foreign direct investment (FDI) as a source of international technology spillovers; here, FA is the industry share of employment in foreign-owned firms, and DTO is increases of domestic firm productivity.

³⁰ It is not a perfect indicator because sometimes citations are added by the patent examiner, not the applicant.

³¹ This approach goes back to the closed-economy work by Griliches (1995) and Frederic Scherer (1984).

³² See section 6.2 for more on this particular case study.

This approach is based on economic theory in the following sense. Often, there are several models that have been proposed to explain, in this case, FDI spillovers, and model-specific evidence often does not yet exist. Association studies try to shed light on the most interesting models by proposing what might be the common reduced-form equation of all these FDI spillover models. In order to accommodate several models, the framework cannot be very specific. This explains why authors of FDI spillover studies employ a simple measure such as the foreign employment share. If FDI has a major effect on domestic firm productivity, it will be picked up by foreign employment no matter what the particular mechanism is.

The generality is attractive, but it also precludes a precise interpretation of the results. Moreover, by using a foreign activity (FA) instead of a foreign technology variable, this approach is prone to estimating technology diffusion only with some error.³³ This matters because calculating a proper causal effect (instead of a correlation) becomes more difficult as the possible causes of spurious correlation multiply. This could be due to business cycle effects or to unobserved heterogeneity. If productivity were exogenous and FDI would vary inversely with trade costs across industries, a positive coefficient of FDI on productivity would suggest spurious FDI spillovers if high-trade cost industries would exhibit on average higher productivity, for example.³⁴ Endogeneity is another problem; here, it could take the form of foreign investors choosing to move FDI into industries that have high rates of productivity growth (those industries are typically also relatively profitable).

How can these issues be addressed? First, spurious business-cycle effects can be avoided

³³ Here, using the foreign employment share as the measure of FDI may be justified because multinational enterprises tend to be R&D-intensive and share their technology between parent and subsidiaries.

³⁴ Lael Brainard (1997) finds that FDI is relatively high in high-trade cost industries.

by including time fixed effects, as long as it is the case that the business cycle is common to the entire sample. More difficult is the question whether trend variables are appropriate, given the time-series properties of the data, in particular whether it is stationary. This is an important issue because including trend variables can have a major impact on the results in this area.³⁵ Panel unit root tests and cointegration analysis can be applied, but that approach is still relatively new, and results vary.³⁶ This suggests trying out alternative approaches; if results change dramatically depending on the assumption on the data generation process, this warrants at least some discussion.

Second, a standard approach in the presence of unobserved heterogeneity is to use fixed effects.³⁷ The consequence of course is to reduce the variation in the dependent variable that is left to explain. As such, fixed effects tend to be of little interest, as they rarely capture the economics that is involved. Moreover, the differencing that is implied by the fixed effects can exacerbate measurement error problems (Zvi Griliches and Jerry Hausman 1986). These considerations suggest keeping the number of fixed effects low. At the same time, there is often some uncertainty about the 'true' model, and if a low number of fixed effects means omitted variables bias, the results can be seriously misleading. In general, it will be worth trying alternative specifications.

Fixed effects will of course not work as a control for unobserved heterogeneity if the

³⁵ For instance, Coe and Helpman (1995) report that including a generalized time trend in their regression reduces the international spillover coefficient by about 60 percent, and using conventional standard errors, the parameter might be insignificantly different from zero (table B.1 and table 3, coefficient on m^*S).

³⁶ For instance, critical values are dependent on the exact specification of the intercept and trend components, as well as how much cross-sectional heterogeneity is allowed for; see G.S. Maddala and In-Moo Kim (1998) for an introduction, and Chris Edmonds (2000) for an application to spillover regressions.

³⁷ Time-differencing is often used for the same purpose; see e.g. Keller (2000).

heterogeneous effects are actually not time-invariant. Not much is known on the extent to which this has led to biases in the existing studies. At any rate, the dynamic industry framework of Steven Olley and Ariel Pakes (1996) does not rely on this assumption—it allows for time-varying heterogeneity—and it has been employed in a number of micro studies.³⁸

The third point, endogeneity, has been recognized in the literature, but it is rarely fully addressed. For instance, Keller (2002a) splits his sample into two, the high- and low-R&D expenditure sectors, based on the fact that four out of his twelve industries account for more than 80 percent of all R&D. For a variety of reasons, endogeneity problems are much more likely to arise in the former than in the latter sectors. Keller's result for the low-R&D sample turns out to be similar to that for the whole sample. In general, association studies conduct often a wide number of robustness analyses. These analyses are hardly definitive in dealing with endogeneity: they are sequential in nature (looking at one suspected problem at a time), they proceed under assumptions that might not hold (e.g., lagging a regressor while assuming that there is no serial correlation), and for other reasons. Nevertheless, such robustness analysis can be very helpful in revealing whether endogeneity might have a first-order effect in a particular context.

Endogeneity issues can in principle be fully addressed by using instrumental variable (IV) techniques. So far these techniques have not been used widely in this literature, although that seems to be changing now. In general, the application of IV techniques has been limited because finding good instruments for technology variables such as R&D stocks is difficult. These are

estimated to begin with, and moreover, most variables that are uncorrelated with the error term tend to be not particularly correlated with R&D, especially at the industry- or micro-level.

The lack-of-instruments problem should become smaller because some recent IV techniques use primarily lags of right hand side variables as instruments. The instruments are valid as long as there is not too much persistence—which can be tested for—and the robustness of the approach exceeds earlier techniques because it uses both the variables' levels and their changes (see Richard Blundell and Stephen Bond 1999).

Two papers that have used IV estimation recently are Rachel Griffith, Stephen Redding, and Helen Simpson (2003) and Keller and Yeaple (2003). Both of these papers study the importance of technology spillovers associated with FDI.³⁹ The Griffith et al. (2003) paper studies the influence of foreign-owned subsidiaries in the United Kingdom. Their instruments, the economic conditions in France and the United States, are quite plausible, because these are the countries from which many of the U.K. subsidiaries come. Using over-identification tests, the authors find that endogeneity is not an issue in their sample. This may be somewhat surprising at first, perhaps suggesting that the instruments are not as good as one might hope for. However, Keller and Yeaple's (2003) IV analysis does not suggest an endogeneity bias either. These authors use industry level transport costs and tariffs as instruments for FDI industry variation in the United States. Keller and Yeaple's IV estimates suggest a higher effect through FDI spillovers than is obtained with OLS. This is the opposite of the suspected endogeneity bias, and could be the result from a FDI variable that does not measure too well the source of technology spillovers.

³⁸ See, for instance, Sourafel Girma and Katharina Wakelin (2001), Garrick Blalock and Paul Gertler (2002), and Wolfgang Keller and Stephen Yeaple (2003). On Olley and Pakes' (1996) method, see also Zvi Griliches and Jacques Mairesse (1998).

³⁹ See section 8 for a broader discussion of this.

Despite these results, it is too early to argue that endogeneity does not play a major role in many studies. More research is clearly needed. Other approaches that should be used more broadly are difference-in-difference and other control-group estimations, especially when a clear regime shift (“natural experiment”) can be identified.

I now turn to a second class of empirical papers that has been employed in the econometric literature.

5.2.2 Structure Studies

These studies incorporate more of the structure of the underlying model than association studies. I distinguish between two types of structure studies. Generally, the first set is given by

$$DTO=f(X,M,FT)+u \quad (6)$$

where the foreign technology variable, FT , replaces the foreign activity variable in equation (5), and the specification adds a specific channel, or mechanism of diffusion (denoted by M).

An influential example is the study by Coe and Helpman (1995). These authors test the prediction of the trade and growth models of Grossman and Helpman (1991), and Rivera-Batiz and Romer (1991), in which foreign R&D creates new intermediate inputs and perhaps spillovers that the home country can access through imports. Assume that output is produced according to

$$z=Al^\alpha d^{1-\alpha}, 0<\alpha<1, \quad (7)$$

where A is a constant, l are labor services, and d is a CES aggregator of differentiated intermediate inputs x of variety s

$$d=\left(\int_0^{n^e} x(s)^{1-\alpha} ds\right)^{\frac{1}{1-\alpha}}, \quad (8)$$

where n^e is the range of intermediate goods that are employed in this country.⁴⁰ It can be

distinct from n , the range of intermediate goods produced in the country. The latter is augmented through R&D (denoted χ); if intermediate goods do not become obsolete, the range of intermediate goods at time T is given by the cumulative resources devoted to R&D, $n_T=S_T\equiv\int_0^T\chi(\tau)d\tau$. The goods $x(s)$ are best thought of as differentiated capital goods, produced with foregone consumption. If the $x(s)$ are symmetric and linearly produced with foregone consumption, one can write the stock of capital as $k=\int_0^n x(s)ds=nx$, which can be used to obtain a reduced-form expression for output as

$$z=A(n^e)^\alpha l^\alpha k^{1-\alpha}. \quad (9)$$

Defining TFP as $f\equiv\frac{z}{l^\alpha k^{1-\alpha}}$, one obtains

$$\ln f=\ln A+\alpha\ln n^e. \quad (10)$$

Equation (10) shows that TFP is positively related to the range of intermediate products employed in this country (Ethier 1982).

In a symmetric model with C countries and no trade barriers, in equilibrium all countries will employ all intermediates from all countries. Given that the intermediate designs are produced through national R&D, the range n^e will be the same in all countries, equal to the world's cumulative R&D spending:

$$n_t^e=\sum_{c=1}^C n_{ct}=\sum_{c=1}^C S_{ct}. \quad (11)$$

Coe and Helpman's (1995) test of the model recognizes that there is both country heterogeneity as well as trade barriers by separating domestic and foreign R&D:

$$\ln f_c=\alpha_c+\beta^d\ln S_c+\beta^f\ln S_c^f+\varepsilon_c, \quad (12)$$

where country c 's so-called foreign knowledge stock S_c^f is defined as the bilateral import-share weighted R&D stocks of its trade partners:

$$S_c^f=\sum_{c'\neq c} m_{cc'}S_{c'}, \quad (13)$$

This captures the prediction of Grossman and Helpman's (1991) model that if a country

⁴⁰ This follows Romer (1990) and Grossman and Helpman (1991); see also Keller (2000); time subscripts are omitted when no ambiguity arises.

imports primarily from high-R&D partner countries, it is likely to receive relatively much technology embodied in intermediate goods, which should be reflected in a higher productivity level; and vice versa.

This approach has been influential because of its plausibility, simplicity, and versatility, and has been widely used to examine alternative channels of international diffusion; for instance, Frank Lichtenberg and Bruno van Pottelsberghe de la Potterie (2000) have examined FDI by substituting bilateral measures of FDI instead of imports.⁴¹ The approach is more specific about the mechanism, which helps in interpreting the results. At the same time, such regressions estimate different models only partially; what determines R&D, for instance, is typically not modeled, which means that endogeneity could be an issue. Given the partial equilibrium nature of the estimation, authors should and generally do consider alternative specifications.

The second set of studies in this section puts relatively less emphasis on a particular model of growth and diffusion, and more emphasis on econometric specification and estimation. A good example is the paper by Sofronis Clerides, Saul Lach, and James Tybout (1998), where the authors provide evidence on learning externalities from exporting using micro data from Columbia, Morocco, and Mexico. Clerides, Lach, and Tybout (1998) extend the sunk cost model of exporting due to Richard Baldwin (1989), Avinash Dixit (1989), and Paul Krugman (1989) to include the possibility that exporting experience lowers the costs of production. Clerides et al. take account of the important point that it is on average the already-productive firms that self-select into the export market. They do this by estimating simultaneously a dynamic discrete choice equation that determines export market participation

$$y_{it} = \begin{cases} 1 & \text{if } 0 \leq \beta^x X_{it} + \beta^e e_{it} + \sum_{j=1}^J \beta_j^c \ln(AVC_{it-j}) \\ & + \sum_{j=1}^J (F^0 - F^j) y_{it-j} + \eta_{it} \\ 0 & \text{otherwise,} \end{cases} \quad (14)$$

and an autoregressive cost function,

$$\ln(AVC_{it}) = \gamma_0 + \sum_{j=1}^J \gamma_j^k \ln(K_{it-j}) + \gamma^e \ln(e_t) + \sum_{j=1}^J \gamma_j^c \ln(AVC_{it-j}) + \sum_{j=1}^J \gamma_j^y y_{it-j} + v_{it} \quad (15)$$

where y_{it} is the export indicator of plant i in period t , X_{it} is a vector of exogenous plant characteristics, e_t is the exchange rate, AVC_{it} are average costs, K_{it} is capital, and F^0 and F^j are sunk costs terms (of export market participation).

Equation (14) states that one only sees a plant exporting if the profits from doing so are greater than from not exporting (the latent threshold is expressed in terms of observables). Equation (15) estimates whether past exporting experience reduces current cost (captured by the parameters γ_j^y), conditional on past costs and size (proxied by capital). These learning-by-exporting parameters are not constrained in a significant way—recall that y_{it-j} are 0/1 indicator variables—which is probably a good starting point given that it is unknown how learning from exporting works (if indeed it exists). Clerides, Lach, and Tybout (1998) employ alternative estimation methods in their analysis.⁴² Moreover, the econometric evidence is combined with descriptive evidence in form of over-time plots of average costs for different types of plants (exporters, non-exporters, etc). The emphasis throughout is to distill the results that emerge consistently from all empirical approaches (they are discussed in section 6.1.2 below). I would

⁴¹ The approach of computing a weighted sum of R&D goes back to Griliches (1979).

⁴² This includes full information maximum likelihood and generalized methods of moments, the latter with a generalized cost function.

expect more work along these lines to be very fruitful.

5.2.3 *Empirical Analysis with General Equilibrium Models*

Eaton and Kortum (1996, 1997, 1999) have studied international technology diffusion using general equilibrium models. In their models productivity growth is related to increases in the quality of intermediate goods, which is key to Aghion and Howitt's (1992) quality ladder model of growth. Eaton and Kortum add a process that governs technology diffusion between countries (see section 2 above). They then explore the quantitative relationships between R&D, technology diffusion, and domestic productivity.

Eaton and Kortum's work is important because instead of focusing on the reduced-form relationship between a subset of variables, they use their full model for making predictions on all (endogenous) variables: endogeneity as in the single-equation studies above is therefore no longer an issue. Eaton and Kortum study both the long-run equilibrium and the models' predictions for the transitional dynamics (Eaton and Kortum 1999 and 1997, respectively). Eaton and Kortum (1996) shed additional light on the rates of diffusion in their model by using reduced-form techniques. Several of their papers study comparative-static changes in the long-run equilibrium, which is especially useful for economic policy analysis (Eaton, Eva Gutierrez, and Kortum 1998).⁴³

There are costs to this approach as well though. First, in order to arrive at a framework that can be analyzed, typically some strong assumptions have to be made. In Eaton and Kortum (1999), for instance, the quality of technology that is discovered in a country is a random variable with a Pareto distribution, while the distribution of the

diffusion lag to other countries is exponential. Within the context of a given model, such assumptions are very difficult to test. It means that this empirical work does not so much test or select among several models, as it estimates one particular model. Second, the models are usually too complicated for estimation to identify all model parameters. If, as is common practice, a number of parameters are fixed based on values from earlier studies, this means that the results are partly simulated, and not estimated.⁴⁴

Third, it is often difficult to judge how empirically successful a model is. One measure of success is the difference between actual versus predicted values for the endogenous variables. If a model predicts productivity levels that lie within an error margin of 15 percent of the actual levels, for instance, how much of this is due to "free" model parameters (i.e., those that are set ex-ante)? With those parameters, how much worse would a different, perhaps simpler model fare? And what if the set parameters are adjusted to maximize the fit in the alternative model? Peri (2002), for instance, uses a partial equilibrium approach to estimate technology diffusion in a model similar to Eaton and Kortum (1996) who work in a general equilibrium framework. With N countries, $N \times N$ equations pin down bilateral technology diffusion, and the implied technology stocks determine $(N-1)$ relative productivity levels, as in equation (3) above. According to Peri (2002), the technology stocks do not have a significant effect on productivity. In Eaton and Kortum (1996), that relationship is one element of their model's general equilibrium structure, which means that it could be that the structural assumptions have too strong an influence on the empirical results.

⁴³ See also Eaton and Kortum (2002); Andrew Bernard et al. (2003); and the discussion in 6.2 below.

⁴⁴ In addition, it can be difficult to find other studies that have estimated a particular model parameter, for instance θ , which governs technology discovery in Eaton and Kortum (1999).

It is important to ask which part of the variation in the data, or which assumptions in the model, are primarily responsible for the results. Despite the sensitivity analysis that may be presented, the informed reader will typically find this question impossible to answer, and this reduces the usefulness of the general equilibrium results.⁴⁵

To summarize, the different empirical approaches each have their advantages and disadvantages. There is a case for adopting a broad estimation framework, as in association studies, if there is model uncertainty, and this case is stronger if endogeneity issues can be addressed. It is useful to impose more structure as more information becomes available, because the structure aides interpretation. General equilibrium models allow us to conduct counterfactuals, which is crucial for policy and welfare analysis, and they will be most influential if the sources of variation that underlie the results are clear.

The following section discusses the importance of different potential channels of international technology diffusion.

6. *Channels of International Technology Diffusion*

6.1 *Trade*

6.1.1 *Imports and Technology Diffusion*

First, I turn to the evidence on diffusion through intermediate input imports. In Eaton and Kortum (2002, 2001), the authors have combined the structure of technology diffusion in Eaton and Kortum (1999)—see

section 2—with the Ricardian model of trade due to Rüdiger Dornbusch, Stanley Fischer, and Paul Samuelson (1977).⁴⁶ In Eaton and Kortum's model, trade augments a country's production possibilities for the classic Ricardian reason: trade gives access to foreign goods or, implicitly, technologies. By specializing in their respective comparative advantage goods, countries can gain from trade in the sense that given a country's resources, the efficient level of output with trade is higher than without trade. At the same time, there are no spillovers in this model, in the sense that importers pay the competitive price and importing has no effect on innovation.⁴⁷

Eaton and Kortum assume that unit transport costs are increasing in geographic distance. This implies that the price of intermediate (or equipment) goods in remote countries is relatively high, or, equivalently, that productivity in these countries is relatively low. These effects are shown to be quantitatively important, as differences in the relative price of equipment account for 25 percent of the cross-country productivity differences in a sample of 34 countries (Eaton and Kortum 2001).

However, it is not clear at this point whether this provides strong support for imports as a major channel for technology diffusion. This is because the equipment goods prices predicted by Eaton and Kortum's (2001) model are inversely related to those reported in Robert Summers and Alan Heston's International Comparison Program of Prices data (CIC 2003): while according to Summers and Heston's data, rich countries have higher equipment prices than poorer countries, Eaton and Kortum's model predicts the opposite (Eaton and Kortum 2001, figure 7).

⁴⁶ See also the extension by Bernard et al. (2003) to plant-level analysis.

⁴⁷ As far as I know, the literature to date lacks a model in which imports have a direct effect on technology diffusion, at least for the multi-country case with sectoral differences.

⁴⁵ In comparison, it is straightforward to assess partial-equilibrium work. For instance, according to Coe and Helpman (1995), bilateral import patterns are important for explaining variation in productivity; see equation (13). A simple test of that is to drop the import shares in (13), and see whether it leads to a major reduction in explained variation of productivity (Keller 1998; see section 6.1). Empirical work on technology diffusion in general equilibrium has not been subjected to the same level of scrutiny, presumably because the nature of the framework has more to do with simulation, and less with testing to begin with.

Summers and Heston's price data might be imperfect, but it is unlikely that revisions of this data will reverse how equipment prices correlate with income. At this point, Eaton and Kortum's model, although appealing because it suggests plausibly large productivity effects come from importing foreign technologies, seems to require a major modification in order to be consistent with international equipment goods price data. The latter is not the only core prediction of the model, but it appears to be important enough to put at least a question mark behind the estimated quantitative importance of trade for technology diffusion.

Second, there is evidence on the importance of imports that comes from international R&D spillover regressions. Based on the framework laid out in section 5.2.2 above, Coe and Helpman (1995) relate TFP to domestic and foreign R&D,

$$\ln f_{ct} = \alpha_c + \beta^d \ln S_{ct} + \beta^f \ln S_{ct}^f + \varepsilon_{ct}, \quad (12')$$

where S_{ct}^f is defined as the bilateral import-share weighted R&D stocks of its trade partners, $S_{ct}^f = \sum_{c' \neq c} m_{cc'} S_{c't}$. A positive effect from the foreign R&D variable S_{ct}^f would imply that a country's productivity is increasing in the extent to which it imports from high- as opposed to low-R&D countries, supporting the hypothesis that imports are a channel of technology diffusion along the lines of the models discussed in Grossman and Helpman (1991). In a sample with 22 OECD countries, Coe and Helpman (1995) estimate a positive and quantitatively large effect from import-weighted foreign R&D. Similar effects are found for technology diffusion from highly industrialized to 77 less developed countries (Coe, Helpman, and Alexander Hoffmaister 1997).

There are some reasons to remain skeptical here as well. First, the analysis of Keller (1998) has shown that the import shares in the construction of the foreign R&D variable S_{ct}^f are not, in fact, essential to obtain Coe and Helpman's (1995) results. Specifically, Keller (1998) uses randomly created shares,

denoted by $\mu_{c'c}$, in place of the actual bilateral import shares to create the counterfactual foreign knowledge stock $\tilde{S}_{ct}^f = \sum_{c' \neq c} \mu_{c'c} S_{c't}$. Using this alternative foreign R&D variable yields similarly high coefficients and levels of explained variation as the regressions using S_{ct}^f (that is, imports data).⁴⁸ Given that import shares are not essential for Coe and Helpman's (1995) results, their analysis does not allow us to draw strong conclusions regarding the importance of imports as a vehicle for diffusion.⁴⁹

A number of authors have made progress by examining the international R&D spillover regressions further. Bin Xu and Jianmao Wang (1999) emphasize that technology diffusion in recent trade and growth models is associated specifically with differentiated capital goods trade. This is in contrast to the trade data Coe and Helpman (1995) use to construct their import shares (from overall trade). Xu and Wang (1999) show that this distinction matters: the capital goods-foreign R&D variable accounts for about 10 percent more of the variation in productivity than does Coe and Helpman's analysis, and it also performs better than Keller's (1998) counterfactual variable.⁵⁰

It has also been noted that the foreign R&D variable S_{ct}^f captures only current-period bilateral trade; it is clearly possible though that country A benefits from country C's technology without importing from this source, if country C exports to country B, which in turn exports to A. Olivier Lumenga-Neso, Marcelo Olarreaga, and Maurice Schiff (2001) use a specification

⁴⁸ Alternatively, Keller (1998) sets all μ_{cc} equal to 1, which produces similar results. This confirms that the import shares are inessential for the results, whether or not they are truly random (see Keller 1997, 2000; and Coe and Hoffmaister 1999).

⁴⁹ Keller's (1998) analysis has the same implication for Coe, Helpman, and Hoffmaister's (1997) paper.

⁵⁰ Keller (2000) uses specialized machinery imports data at the industry level in a related analysis. His results show that technology diffusion through imports does affect productivity growth for high skewed trade patterns (e.g. for Canada, which imports to about 70 percent from the United States), suggesting that the effect is nonlinear.

that captures such indirect R&D spillovers, and show that it performs better than Coe and Helpman's (1995) and Keller's (1998) models. These results are consistent with the importance of dynamic effects from imports, but more research in an explicitly dynamic framework is needed to learn more about this.

The importance of imports for technology diffusion has also been assessed with patent citation data. Frederic Sjöholm (1996) studies citations in patent applications of Swedish firms to patents owned by foreign inventors. Controlling for a number of other correlates and also conducting an extreme-bounds analysis, Sjöholm finds a positive correlation between Swedish patent citations and bilateral imports. This is consistent with imports contributing to international knowledge spillovers.

Overall, the evidence points to a significant role for imports in international technology diffusion. However, the various strands of the literature leave still some questions open, and we do not have yet a firm estimate of the quantitative importance of imports for international technology diffusion.

6.1.2 *Exports and Technology Diffusion:*

Is There Learning-By-Exporting?

A major question is whether firms learn about foreign technology through exporting experience. There is anecdotal evidence claiming that firms do benefit from interacting with foreign customer, for instance because the latter impose higher product quality standards than domestic customer, while at the same time providing information on how to meet the higher standards. Case studies of the export success of a number of East Asian countries starting in the 1960s are particularly strong in their emphasis on learning-by-exporting effects (Yung-Whee Rhee, Bruce Ross-Larson, and Garry Pursell 1984). The question is whether this evidence can be supported with econometric evidence.

There is abundant evidence that in a given cross-section, exporters are on average more productive than nonexporters (e.g., Bernard and Jensen 1999; Clerides, Lach, and Tybout 1998; Mary Hallward-Driemeier, Giuseppe Iarossi, and Kenneth Sokoloff 2002).⁵¹ That does not settle the issue of causality, however: are exporting firms more productive because of learning effects associated with exporting, or is it rather the case that firms that are more productive to begin with start exporting? The conventional wisdom today is that learning-by-exporting effects are non-existent. This position is consistent with current evidence, but at the same time there are a number of issues which in my view are worth further research before the question should be considered settled.

An important paper is by Clerides, Lach, and Tybout (1998). These authors study manufacturing plants in Columbia, Morocco, and Mexico during the 1980s to find evidence on learning-by-exporting effects. Their framework is based on a model with sunk costs of (export) market entry, leading to a dynamic discrete choice equation for participation and an autoregressive cost function as a performance measure (see section 5 above).

Clerides, Lach, and Tybout show results for each country separately, and also by major industry. In general, they tend to show no significant effects from past exporting experience on current performance (this is parameter γ_i^y in equation 15). Clearly, this does not lend support for the existence of strong positive learning-by-exporting effects. In fact, to the extent that Clerides et al.'s estimates are significant, they go into the wrong direction (exporting raising costs). It would be surprising if indeed there would be negative learning effects, and the authors give a number of

⁵¹ Note that the evidence on learning-by-exporting effects discussed in the following comes from micro studies, while the evidence on diffusion associated with imports discussed above is based on more aggregate studies.

plausible reasons of why this finding may have to be discounted. Another interpretation of the generally insignificant estimates may be that the estimation framework is demanding too much of the data. However, Clerides et al.'s plots of average cost before and after export market entry seem to support their main estimation result of no evidence for learning-by-exporting effects.

While learning-by-exporting has been emphasized to be important primarily for low- and middle-income countries' firms, there is in principle no reason why it is limited to these countries, especially given the firms' heterogeneity in terms of productivity in any given country. Bernard and Jensen (1999) study the learning-by-exporting question using data on U.S. firms. This has the advantage that the sample is relatively large and there is comparatively much experience with data collection and preparation, which may result in lower measurement error.

Unlike Clerides, Lach, and Tybout (1998), Bernard and Jensen (1999) do not model export market participation explicitly. Instead, they study the performance of the different sets of firms separately.⁵² Bernard and Jensen estimate that labor productivity growth is about 0.8 percent higher among exporters than nonexporters.⁵³ This estimate is fairly small, and it becomes even smaller (and insignificant) for longer time horizons.

The estimate of 0.8 percent, however, appears to be a downward-biased estimate of the learning-by-exporting effect because it comes from an analysis conditional on plant survival. In fact, Bernard and Jensen show that conditional on size, exporters are 10 percent more likely to survive than non-

exporters.⁵⁴ It is plausible that this 10 percent survival probability difference is indicative of higher productivity growth for exporters than nonexporters, because plants tend to fail because their productivity growth is low. This suggests that the overall difference in productivity growth between exporters and nonexporters may be larger than 0.8 percent, although at this point it is not clear how much larger.

Some of these estimates come from a relatively small number of years, and there may be an argument that this time horizon is too short to see major learning-by-exporting effects. Alternatively, Hallward-Driemeier, Iarossi, and Sokoloff (2002) focus on the time before entering the export market. These authors use data from five Southeast Asian countries to show that firms which eventually become exporters make more investments to raise productivity and the quality of their goods than firms that plan to stay out of the export market. This is plausible, but if these investments require—which is likely—real resources, those need to be subtracted from any learning effects the firms receive after they have entered the export market.⁵⁵ Moreover, given that the productivity increases predate the firm's entry into the export market, at best these are indirect learning-by-exporting effects.

The analysis has shown that there is no econometric evidence for a strong learning-from-exporting effect. At the same time, there are a number of issues that still need to be addressed. First, it is puzzling that the

⁵⁴ Size is the main predictor of survival in recent industry-equilibrium models (e.g. Olley and Pakes 1996), because a small firm might have to exit after only one bad shock (or, a small number of successive bad shocks), whereas a large firm has substance enough to weather a longer succession of bad shocks.

⁵⁵ This point is related to the fact that none of the estimates I discuss are claimed to be spillover effects; rather, they are "learning" effects, which could be costly to acquire. Clerides, Lach, and Tybout (1998) do estimate spillovers, those that might accrue to *other* plants; this evidence is mixed.

⁵² Four types of firms can be distinguished: exporters, nonexporters, starters (plants that start exporting), and quitters (plants that stop exporting).

⁵³ Bernard and Jensen's estimates using TFP instead of labor productivity are lower, but the labor productivity figures are preferred in this case. The TFP measure is a simple regression residual that is fallible to a number of problems (e.g. Griliches and Mairesse 1998).

econometric evidence is so strongly at odds with the case-study evidence. Second, there could be heterogeneity across industries masking strong but industry-specific learning effects, especially with respect to high-technology products.⁵⁶ Third, the analysis could be improved if we knew more on both the export destination and the exporter, instead of simply an indicator variable (exporting yes/no). For instance, to which firms, in which countries, do the exports go? What are the characteristics of these foreign firms—are they engaged in R&D, for example?

6.2 Foreign Direct Investment

Foreign direct investment (FDI) has long been considered as an important channel for technology diffusion. It is a plausible channel from a theoretical standpoint, because an influential theory says that firm-specific technology is transferred across international borders by sharing technology among multinational parents and subsidiaries (see, e.g., James Markusen 2002). There are a number of models showing how multinational enterprises (MNEs) might generate technological learning externalities for domestic firms, for instance through labor training and turnover (Andrea Fosfuri, Massimo Motta, and Thomas Rønde 2001) or through the provision of high-quality intermediate inputs (Rodríguez-Claré 1996).⁵⁷

Whether FDI indeed generates substantial technological externalities for domestic firms is also a major policy issue, because governments all over the world spend large amounts of resources to attract subsidiaries of MNEs to their jurisdiction. For instance, in 1994, the U.S. state of Alabama spent \$230 million, or \$150,000 per newly created job, to attract a new plant of Mercedes-Benz

(see Jonathan Haskel, Sonia Pereira, and Matthew Slaughter 2001).

What is the evidence on FDI spillovers? A number of surveys have recently concluded that there is *no* evidence for substantial FDI spillovers (Gordon Hanson 2001; Holger Görg and David Greenaway 2002).⁵⁸ However, this view, based primarily on a number of micro-level productivity studies, seems to be overly pessimistic. First of all, some more recent evidence for the micro productivity approach suggests that, on the contrary, FDI spillovers are large and economically important. Second, there are some results outside the micro productivity literature that do provide evidence for FDI spillovers.

A case study of Intel's FDI into Costa Rica, for example, already mentioned in section 5 above, provides interesting information on how widespread the changes can be that FDI by a major high-technology company can trigger in a relatively small country (Larrain, Lopez-Calva, and Rodríguez-Claré 2000).⁵⁹

Other authors have provided econometric evidence on whether multinational enterprises (MNEs) raise the rate of international technology transfer as measured by patent citations (Steven Globermann, Kokko, and Sjöholm 2000; Branstetter 2001a; and Jasjit Singh 2003). Here, the results are less clear. First of all, MNE subsidiaries could either disseminate technology to domestic firms of their host country (inward FDI technology transfer), or MNE subsidiaries might pick up new technologies from the firms in the host country (outward FDI technology sourcing). It turns out that the relative strength of these effects is estimated differently, although the technology sourcing effect appears to be

⁵⁶ Industry heterogeneity has recently been emphasized in the literature on FDI (see 6.2); Clerides, Lach, and Tybout (1998), however, do generally not estimate major differences across industries.

⁵⁷ See Magnus Blomström and Ari Kokko (1998) and Kamal Saggi (2000) for a more detailed discussion of possible spillover channels associated with FDI.

⁵⁸ In the words of another author, "today's policy literature is filled with extravagant claims about positive spillovers from FDI, [but] the hard evidence is sobering" (Dani Rodrik 1999, p. 37).

⁵⁹ It is an open question whether, or to what extent these are spillover effects; see also the discussion at the end of this section.

stronger.⁶⁰ This result—MNE subsidiaries learn more from the firms in their host country than vice versa—might be indicative of a number of problems, however.

The first is firm heterogeneity: MNE subsidiaries are larger and more technologically intensive than the average firm in the host country, and this might be the reason why they are good at sourcing technology. This interpretation seems to be confirmed by Singh's (2003) finding that patent citations between two MNE subsidiaries are stronger than either from MNE subsidiary to a domestic firm or the reverse. The second issue is endogeneity. It could be that one finds MNE subsidiaries to be sourcing more technology than they provide because the MNE parent set up the subsidiary with the explicit goal of technology sourcing, while the average host-country firm, in contrast, has not made a comparable location decision. This suggests that the estimates are not fully comparable, and future research is needed to settle this issue.

The value of a patent is difficult to estimate (see, e.g., Pakes 1986). An important issue in the patent citation studies is therefore how the economic significance of technology diffusion measured in this way should be assessed. A large literature has thus tried to estimate directly the extent to which FDI leads to productivity increases for domestic firms. Xu (2000), for instance, uses the U.S. Bureau of Economic Analysis' comparable data on U.S. outward FDI into forty countries over almost thirty years (between 1966 and 1994). He finds generally a positive relation between FDI and productivity growth, which is stronger in the richer than in the poorer countries.

Xu's (2000) analysis is at the manufacturing level, which may cause aggregation bias

because of heterogeneity across sectors and across firms. For this reason, the literature on FDI spillovers has moved towards using micro (firm or plant) level data (e.g., Aitken and Harrison 1999; Girma and Wakelin 2001; Haskel, Pereira, and Slaughter 2001; Griffith, Redding, and Simpson 2003; Keller and Yeaple 2003).⁶¹ These papers present regressions of productivity on FDI and a number of control variables (they are association studies in the sense of section 5.2 above). If productivity growth of domestic firms is systematically higher in industries with more FDI than in industries with less FDI—and endogeneity and other problems can be ruled out, see section 5.1 above—this provides support for the FDI spillovers hypothesis. Irrespective of the precise spillover mechanism, a greater presence of foreign-owned subsidiaries in a particular industry is plausibly conducive to domestic firms' technological learning if they belong to this industry.

What is the evidence on this? In Aitken and Harrison (1999), the authors estimate a *negative* relationship between FDI and productivity for a sample of Venezuelan plants. The finding of a negative relationship points to a specification error: while FDI spillovers may or may not exist, there is no reason to believe that they would be negative. One possibility, raised by Aitken and Harrison (1999), is that their results capture the increased competition through FDI in addition to spillovers.⁶² Haskel, Pereira, and Slaughter (2001), among others, control for such competition effects. These authors, as well as Girma and Wakelin (2001) and Griffith, Redding, and Simpson (2003) study inward FDI for the United Kingdom. While Girma and Wakelin's results are somewhat mixed, Haskel, Pereira, and Slaughter (2001) and Griffith, Redding, and Simpson

⁶⁰ Branstetter (2001a) finds evidence for spillovers of both types for Japanese FDI in the U.S., Singh (2003) finds the sourcing effect to be stronger than the transfer effect in a sample of 10 OECD countries, and Gliberman, Kokko, and Sjöholm (2000) find only evidence for the outward FDI effect.

⁶¹ Görg and Greenaway (2002) is a recent survey.

⁶² Another possibility is endogeneity: FDI into Venezuela targets sectors in which Venezuelan firms are relatively weak. Aitken and Harrison do not treat endogeneity explicitly.

(2003) present evidence in support of a significant positive FDI spillover effect. The implied economic magnitudes are fairly small, however, certainly relative to the subsidies that have been paid to attract FDI.⁶³

In contrast, Keller and Yeaple's (2003) study of recent FDI activity in the United States suggests that FDI spillovers are positive and large. They estimate that FDI in the United States accounts for about 11 percent of U.S. manufacturing productivity growth in that period. This estimate is much larger than Haskel, Pereira, and Slaughter's (2001) as well as other estimates, and the question is why. There seem to be two primary reasons.

The first reason is that FDI spillovers appear to be much stronger in relatively high-technology than in relatively low-technology sectors.⁶⁴ This explains in part Keller and Yeaple's larger FDI spillover estimate, because in their *Compustat* sample, the share of high-technology firms is higher than in the broader sample of Haskel, Pereira, and Slaughter (2001), for example. The strong difference between sectors, which is implicitly present also in Griffith, Redding, and Simpson (2003), table 5—suggests that FDI might be an important conduit for international technology diffusion, but only some forms of FDI that occur in a limited set of sectors. Other FDI, for instance that primarily seeking lower factor costs, is not likely to generate major spillovers. Moreover, the practice of pooled FDI spillover estimation—across all manufacturing industries—is likely to lead to misleading results.

The second finding of Keller and Yeaple underscores the importance of good measures of foreign economic activity. They use detailed U.S. statistics to allocate the employees of a given MNE subsidiary to potentially *different* industries. In contrast,

in most previous studies, all employees are allocated to one, the so-called major industry. To the extent that MNE subsidiaries are diversified and their industry mix changes over time, the major-industry method of measuring FDI is prone to year-to-year volatility that is essentially noise. Keller and Yeaple (2003) show that for their sample, the preferred FDI data yields estimates that are seven times as large as those with the inferior by major-industry FDI data.

We can summarize the literature on FDI spillovers as follows. In contrast to the earlier literature, recent micro productivity studies tend to estimate positive, and in some cases also economically large spillovers associated with FDI. This difference does not appear to be primarily due to endogeneity (or other problems).⁶⁵ Moreover, although the current evidence from micro productivity studies comes from the United Kingdom and the United States, there are reasons to believe that the findings might apply in other countries as well. If these micro productivity FDI spillovers hold up in the future, it would also provide support for the FDI learning effects that are found in some of the case studies.

There are other important questions. First, are there effects of FDI outside the industry in which MNE subsidiaries are active? It is plausible to assume that to the extent that MNE subsidiaries buy local intermediate inputs, they have an incentive to provide technological knowledge to their suppliers; a stronger incentive at any rate than to local competitors that produce the same product. Maurice Kugler (2002) as well as Blalock and Gertler (2002) have found some evidence for such vertical FDI spillovers.⁶⁶

Second, could there be more general externalities? Larrain, Lopez-Calva, and

⁶³ Haskel, Pereira, and Slaughter (2001) present a calculation of costs and benefits.

⁶⁴ The two types of sectors are distinguished by R&D intensity; see Keller and Yeaple (2003).

⁶⁵ As noted in section 5, instrumental variable estimations are presented in Griffith, Redding, and Simpson (2003) and Keller and Yeaple (2003). In neither case are the least-squares estimated FDI spillovers larger than the corresponding IV-estimates.

⁶⁶ See also Beata Smarzynska Javorcik (2004).

Rodriguez-Claré (2000) argue that among the most important consequences of the chipmaker Intel's recent FDI into Costa Rica was that (1) Intel funded schools that taught local workers certain skills (but by no means Intel specific skills), and (2) Intel's FDI served as a signal for other potential foreign investors that it might be time to invest into Costa Rica now. Estimating such signaling effects (if they exist) is a challenge, because the more diffuse and indirect the alleged externalities are, the more difficult it becomes to identify clean causal effects.

7. *Geographic Effects on International Technology Diffusion*

Global technology spillovers favor income convergence, and local spillovers tend to lead to divergence, no matter through which channel technology diffuses. In addition to research on channels of diffusion, another strand of the literature has examined international technology diffusion in its geographic dimension (Jaffe, Trajtenberg, and Rebecca Henderson 1993; Douglas Irwin and Peter Klenow 1994; Eaton and Kortum 1999; Branstetter 2001b; Laura Bottazzi and Peri 2000; and Keller 2002a). An advantage of this is that geography is arguably exogenous in this process.

One question is whether technology diffusion within countries is stronger than across countries. The evidence generally supports this hypothesis, although there are exceptions. In particular, Jaffe, Trajtenberg, and Henderson (1993) compare the geographic location of patent citations with that of the cited patents in the United States. They find that U.S. patents are significantly more often cited by other U.S. patents than they are cited by foreign patents. This is confirmed by Branstetter (2001b), who uses R&D and patenting data on U.S. and Japanese firms to compute weighted R&D spillover stocks analogous to Coe and Helpman's (1995)

import share weights (see section 5).⁶⁷ He finds that within-country spillovers are much stronger than between-country spillovers. More evidence for stronger diffusion within than across countries comes from Eaton and Kortum's (1999) study: these authors estimate that for the G-5 countries, the rate of domestic technology diffusion is about 200 times the size of the average rate of international technology diffusion between the G-5 countries.⁶⁸

In contrast, Irwin and Klenow (1994) do not find stronger within country spillover compared to across country spillovers. Irwin and Klenow estimate that for eight vintages of semiconductors introduced between 1974 and 1992, the spillovers from one U.S. firm to another U.S. firm are *not* significantly stronger than those between an U.S. firm and a foreign firm. The different results might be obtained because Irwin and Klenow's spillovers, which are identified from the effects of cumulative production on market shares, are different from knowledge spillovers as measured in the other studies.⁶⁹ It could also have to do with the particulars of the semiconductor industry at the time. Most of the relatively small number of firms were located in the United States and in Japan, which means that the scope for identifying the within versus between country difference is limited.

⁶⁷ Patenting data by technological field allows Branstetter to compute weights that are increasing in the similarity of two firms' patenting activities; this captures the idea that R&D expenditures of another firm are more likely to generate spillovers, the closer the two firms are in technology space.

⁶⁸ The average international diffusion rate, ε_{ni} , for $n \neq i$, in Eaton and Kortum (1999) is about 0.088, while the domestic rate of diffusion (ε_{ni} , for $n=i$) is fixed at 17.7; see their table 2.

⁶⁹ Irwin and Klenow (1994) estimate learning-by-doing spillovers, which are more closely related to human-capital models (e.g. Robert Lucas 1988) than to models of technological change. They might be different in that they measure improvements in the efficiency of an already adopted technology, whereas international technology spillovers may capture to a greater extent learning about new and not-yet-adopted technologies. Empirical analysis has not been able to clearly separate one from the other so far.

The analysis has therefore been extended beyond the national versus international distinction by estimating spillovers conditional on geographic distance and the countries' locations relative to each other (Keller 2001, 2002a). Keller (2002a) uses industry-level productivity in nine mostly smaller OECD countries to R&D in the G-5 countries (France, Germany, Japan, the United Kingdom, and the United States) using a simple exponential decay function in distance:

$$\ln TFP_{cit} = \beta \ln \left(S_{cit} + \gamma \sum_{c' \in G5} S_{c'it} \times e^{-\delta D_{c'c}} \right) + \alpha' X + \varepsilon_{cit}. \quad (16)$$

Here, $D_{c'c}$ is the geographic distance between countries c' and c , and X is a vector of control variables. If δ is estimated to be greater than zero, variation in productivity is best accounted for by giving a lower weight to R&D conducted in countries that are located relatively far away, whereas if $\delta=0$, geographic distance and relative location do not matter. Keller (2002a) finds that δ is positive, and moreover, the decay of technology diffusion implied by the estimate is substantial: with every additional 1,200 kilometers distance there is a 50-percent drop in technology diffusion. Applying this estimate to Australia, for example, with its remote geographic location relative to the G-5 countries, would suggest that Australia benefits extremely little from technology created in the G-5 countries. Along the same lines, Bottazzi and Peri (2003) find a strong geographic decay in their analysis of technology diffusion between European regions using a similar framework. These studies suggest that technology is highly geographically localized in particular regions and countries.

A related question is whether the degree of localization has fallen in recent years. This may be expected as a consequence of transport cost improvements, information and communication technology innovations, increased multinational activity, and

other changes. Keller (2002a) examines this by estimating different decay parameters (δ in equation 16) for the late 1970s and the early 1990s. The estimates indicate that δ has shrunk substantially over time in absolute value, suggesting that the degree of localization has become smaller. This is consistent with the changes mentioned above leading to a greater international diffusion of technology.

A major concern is that the estimated geography effect may be spurious, perhaps due to unobserved heterogeneity across locations. This issue is addressed in a number of studies, and while the proposed solutions are clearly imperfect, overall the results suggest that geography in fact is an important determinant of technology diffusion.⁷⁰ Another concern seems to be more important. This is that we still need to find out exactly what the geography effect means in terms of economic analysis: does this capture trade costs, for instance? We know that trade volumes are strongly declining with distance (Leamer and Levinsohn 1995), so this seems clearly a possibility. Keller (2001) is a first attempt to explain the geographic effects in international technology diffusion in terms of imports, foreign direct investment patterns, and person-to-person communication, but more work is needed that reveals what geography stands for in terms of economic models and behavior.

8. Human Capital and R&D as Determinants of International Technology Diffusion

As international technology diffusion becomes stronger over time, this favors income convergence, and vice versa, weaker

⁷⁰ See Jaffe, Trajtenberg, and Henderson (1993), e.g., who adopt a control group approach: they examine cited and not cited patents that are matched in terms of otherwise similar characteristics. See also Keller's (2002a) discussion of whether his results may be explained by geographic distance being correlated with technological distance (i.e., that geographically proximate countries produce relatively similar products).

diffusion makes divergence more likely. However, this result applies only if the trend towards (or away from) greater diffusion is global, affecting all countries to the same extent. In fact, there seem to be big differences in how effective countries are in adopting foreign technology. Given that the bulk of new technology is created in a handful of the world's richest countries, greater technology diffusion could in fact lead to divergence in the world income distribution: this will happen if today's richer countries are on average better at adopting foreign technology than today's poorer countries. Alternatively, poorer countries could gain more from better technology diffusion than richer countries. After all, poorer countries have more to gain from it, or, in the words of Alexander Gerschenkron (1962), latecomers may benefit from their relative backwardness.

This underlines the importance of identifying the major determinants of successful technology diffusion. Among the many that have been proposed, two stand out: human capital and R&D expenditures. Both are associated with the notion of absorptive capacity, the idea that a firm or country needs to have a certain type of skill in order to be able to successfully adopt foreign technology; see Keller (1996) for a formalization of this idea.⁷¹ These skills can come, first of all, in the form of human capital (Richard Nelson and Edmund Phelps 1966). They can also be in the form of R&D spending, a notion that was first emphasized by Wesley Cohen and Daniel Levinthal (1989). These authors argue that R&D investments are necessary for a firm to acquire outside technology. This is because R&D is critical to enabling the firm to understand and evaluate new technological trends and innovations.

⁷¹ Others distinguish specific and currently embodied from more general skills, and argue that technology diffusion destroys the former skills in the short-run, while in the long-run there is a complementary relationship between technical change and skills (Oded Galor and Daniel Tsiddon 1997).

There is evidence that human capital facilitates the adoption of new technology in a closed economy setting (e.g., Ann Bartel and Lichtenberg 1987). Does the same also hold true for international technology adoption? The evidence presented in a number of papers suggests the affirmative. Eaton and Kortum (1996) find that inward technology diffusion is increasing in the level of a country's human capital. In Francesco Caselli and Wilbur Coleman's (2001) study, inward technology diffusion is captured by imports of office, computing, and accounting machinery, on the grounds that many countries do not have a domestic computer industry, so that computer technology comes necessarily from abroad. They find that computer imports are positively correlated with human capital. The results presented by Xu (2000) are consistent with the importance of human capital for technology diffusion as well. He finds that the reason relatively rich countries benefit from hosting U.S. multinational subsidiaries while poorer countries do not as much has to do with a threshold level of human capital in the host country.

Does absorptive capacity in the form of R&D have a similar effect on technology diffusion? To date there exist only a few studies. In one of them, Griffith, Redding, and John Van Rens (2000) use industry-level data from twelve OECD countries for the years 1974 to 1990 to study the main determinants of productivity dynamics. They show that conditional on a certain productivity gap to the leader country, subsequent productivity growth in an industry is higher, the higher are its R&D expenditures. This is consistent with R&D playing a role similar to that of human capital in providing the necessary skills for technology adoption.

The second question is whether conditional on human capital and R&D speeding up the diffusion of technology, there is more likely to be convergence or divergence. As discussed above, this will depend on the distribution of human capital and R&D across countries, and the evidence on this is mixed

so far. First, the human-capital threshold identified by Xu (2000) might not have been surpassed by poor countries, which suggests that differences in human capital could lead to divergence. Second, a number of FDI spillover studies have found that it is the less productive firms that benefit most from foreign technology (Haskel, Pereira, and Slaughter 2001; Griffith, Redding and Simpson 2003; and Keller and Yeaple 2003), and that favors convergence.

One difference between Xu (2000) on the one side, and the FDI studies on the other side, is that the latter do not specifically focus on the role that a firm's own R&D (or skills) might have for convergence. A possible explanation for these seemingly contradictory results could be that while the distribution of R&D and skills in favor of the relatively rich is as such a force towards divergence, this is more than outweighed by the stronger benefits that the relatively poor derive from foreign technology that are not conditional on skills and R&D. It is of course also possible that convergence of countries and convergence of firms in a particular country are not comparable in this way. More work in this area is clearly needed.

9. Quantifying Domestic and Foreign Sources of Productivity Growth

Building on the analysis in section 5, the following equation shows the relation between a country's productivity (f_c) and domestic (S_c) as well as foreign R&D (S_c^f):

$$\ln f_c = \alpha_c + \beta^d \ln S_c + \beta^f \ln S_c^f + \varepsilon_c, \quad (12'')$$

Two hypotheses regarding extreme cases of international technology diffusion can be stated easily. If β^f is equal to zero, there is no international technology diffusion at all, whereas if β^d is equal to zero, then there is perfect international technology diffusion, or, a global pool of technological knowledge. The discussion so far has indicated that neither of these two cases finds empirical support— β^f is clearly greater than zero, but at

the same time geographic localization means that β^d is not equal to zero either. Going beyond these results, this section asks what we know about the relative magnitude of domestic and foreign technology sources.

Most of the estimates we have come from analyzing samples of relatively rich countries. In poor countries, there is very little comparable data on domestic technology investments. The data on the most important foreign technology sources that is available—R&D in the G-7 countries—by itself is thus not sufficient to assess the *relative* foreign contribution for poorer countries. In order to get some idea nevertheless, I will rely on the differences in relative contribution of foreign technology among OECD countries.

There are several different ways of assessing the relative importance of foreign technology diffusion, roughly corresponding to the different empirical methods described in section 5. One approach is to compare the TFP elasticities of domestic and foreign R&D. Coe and Helpman (1995) estimate a domestic R&D elasticity of about 23 percent for the relatively large G-7 countries, whereas the corresponding foreign R&D elasticity is only about 6 percent (these are approximately equal to β^d and β^f in equation 12''). Thus, for the G-7 countries the relative contribution from foreign R&D is about 21 percent (0.06 divided by 0.06 plus 0.23). This is almost the same as in Keller's (2002b) analysis of the G-7 countries plus Sweden using data at the industry level, where he arrives at a share of about 20 percent for foreign R&D in the total effect on productivity.

For the smaller OECD countries, Coe and Helpman estimate a domestic R&D elasticity of about 8 percent and a foreign R&D elasticity of about 12 percent. Thus, the ratio of foreign to domestic effect is estimated to be 3:2 for smaller and 1:4 for larger OECD countries. What lies behind the result that foreign R&D is relatively much more important for smaller countries? Here, this follows because smaller countries are more open than larger countries—recall

from section 6.1 that the foreign R&D effect is proportional to a country's overall import share in this specification.

This openness effect points to country size as an important determinant of the relative importance of foreign versus domestic technology sources. Is the relative contribution of foreign R&D indeed so much larger for small compared with large countries? A first look at other evidence seems to suggest yes. For instance, Eaton and Kortum (1999) estimate that the part of productivity growth that is due to domestic as opposed to foreign R&D is between 11 percent and 16 percent in Germany, France, and the United Kingdom, around 35 percent for the larger Japan, and about 60 percent for the United States. Keller (2002a) computes the domestic and foreign share from his distance-adjusted estimate of effective R&D (equation 16). For nine countries that are smaller than the United Kingdom, Keller estimates that the domestic source share is about 10 percent, even smaller than Eaton and Kortum's estimate of 11 percent for the United Kingdom.

That size matters in some way is clear, however, because in the limit, as a country's share in the world approaches one, the share due to foreign technology goes to zero. To control for size differences, one possibility is to use GDP data. Doing this gives for the United States a ratio of domestic technology share to GDP share of 60 percent to 48 percent, or 1.24, whereas the corresponding value for the Netherlands is 2.18.⁷² This suggests that relative to GDP, productivity growth in the United States relies *less* on domestic sources of technology than in the Netherlands.

⁷²These figures are based on results in Eaton and Kortum (1999) and Keller (2002a); GDP data from Heston, Summers, and Bettina Aten (2002). In the year 1990, the United States contributed about 48 percent to the combined size of the G-5 countries plus Sweden, Denmark, and the Netherlands, while the Netherlands accounted for 2.2 percent. Eaton and Kortum estimate that about 60 percent of U.S. productivity growth was due to U.S. R&D, while Keller finds that about 4.7 percent of Dutch productivity growth can be attributed to Dutch R&D.

This may suggest that more generally, the smaller a country, the more intensively it relies on domestic technology sources. If so, it would be interesting to find out whether this is only a country size, or also a distance-to-the-frontier phenomenon. The latter would mean that even though poor countries receive almost all of their technology from abroad—there is a vast pool of technology out there—it is domestic technology efforts that matter most for them. This would be consistent with theories where technology that is invented in frontier countries is less appropriate for poorer countries (e.g., Susanto Basu and David Weil 1998).

10. *Concluding Discussion*

What have we learned from this literature so far, and what do we still need to know? I will begin by giving a snapshot summary of the evidence in some major areas. The empirical literature is still quite fragmented, but there are signs of consolidation which I will highlight. I will turn then to an outlook on what research might have a particularly high payoff in the future.

For most countries, foreign sources of technology are of dominant importance (90 percent or more) for productivity growth (section 9). This fact underlines the significance of international technology diffusion and the importance of finding out which activities determine its success, and whether substantial technology spillovers are associated with them. Foreign sources of technology are more important for small (and relatively poor) than for bigger countries, which is what one expects given the size difference in domestic R&D investments.

There is no indication, however, that international learning is inevitable or automatic. Neither does it seem to be easier for countries that are relatively backward, as Gerschenkron (1962) has called it. It could well be that controlling for size, the poorer a country is, the greater is the importance of domestic technology investments. Instead of

simply being far from the frontier, the success of countries is in part explained by how their firms engage in international economic activities. What is the evidence on this so far?

Early on, international trade has been suggested as a major channel for technology diffusion. With regard to imports, I think that overall the evidence now supports the notion that importing is associated with technology spillovers. At the same time, there is still an abundance of disparate results that need to be brought together and quantified. Moreover, we do not know yet how strong diffusion is through embodied technology in intermediate goods versus other technology diffusion associated with imports. Learning effects from exporting have been found in the case study literature, whereas authors of econometric studies take a much more skeptical view. While the econometric evidence on learning-by-exporting effects is not as clear cut and negative as it is often stated, right now I do not see evidence in favor of strong important learning-by-exporting effects either.

What about learning effects for domestic firms from FDI? The FDI literature seems to be closer to a consensus than the trade literature is right now. For instance, both micro-econometric studies and case studies point in the same direction. The evidence suggests that there can be FDI spillovers, but they do not occur everywhere to the same degree. The remaining questions include the following: how large are vertical compared to horizontal spillovers, and which firms—stronger or weaker—benefit most from spillovers. In addition, we need to know whether FDI spillovers in richer and poorer countries are equally strong, and last not least whether they are quantitatively large enough to justify the large subsidies that governments provide to attract multinationals.

There is a thin line between recognizing the importance of international technology diffusion and getting carried away in the belief that in today's world technology is

global. In fact, technology is not global today, as the studies on the localization of technology diffusion have shown. Yes, we can zap technology in the form of computer programs or designs easily around the globe, but no, this does not mean that there is effectively a global pool of technology. Perhaps it is because technology is in part tacit, requiring the sender at times to actually go there and implement the technology, that there is still a geographic pattern to technology diffusion—although the grip of geography has become weaker recently. We still need to know more about why geography matters—explaining changes in the geographic pattern of technology diffusion will tell us much, I think, about the key mechanisms, especially in analyses that include both rich and poor countries.

In order to decide what kind of research would be most useful at this point, it is worth recalling what the options are. Broadly speaking, those are case studies, econometric analyses, and simulation studies. As long as their respective strengths and weaknesses are recognized, all three can potentially prove to be very useful. Of great importance is whether the data that is employed is closely related to issues of technology and technology diffusion. For instance, FDI spillovers estimated from data on foreign subsidiaries' (and their parents') R&D should tell us much more on technology transfer than a variable like the foreign share of employment. The general principle is: if you know something about actual technological activity, try to use it.

An important question is the appropriate level of data aggregation. This issue was moot for a long time in international economics because there was very little micro data. Recently however, micro data sets have started to become available. It is thus important to know the differences between micro and more aggregate data for studying international technology diffusion, and which of the two, if any, is preferred. Not only can one use micro data today, but moreover it seems to matter a great deal for the results:

generally, the higher is the level of aggregation, the stronger tends to be the evidence for externalities and learning effects. What can explain this difference between micro and more aggregate evidence?

Micro data can capture heterogeneity across firms, something that features prominently even in narrowly defined industries, whereas industry- and aggregate-level studies cannot control for this and may suffer from composition and aggregation biases that tend to lead to inflated spillover estimates. I think that this is indeed sometimes the case. At the same time, firm heterogeneity is less important if one is primarily interested in industry-level outcomes. For example, it may not matter too much for economic policy whether industry productivity rises because all firms benefit from spillovers or because only the more productive firms benefit and become larger while the least productive firms exit.

Simultaneity and endogeneity seem to be more important issues than aggregation, and in this respect there is little difference between micro and more aggregate studies: for instance, interpreting a cross-sectional correlation of foreign ownership and productivity as evidence for FDI spillovers would be just as inappropriate at the firm level as it is at the aggregate level. Instead, the work in question is more or less important depending on how thorough the authors are in trying to identify a truly causal effect.

Most strategies for doing that rely on comparing sets of firms. This comparison, or, taking differences, may allow us to see whether the treatment (say, inward FDI) has an effect. There is no doubt that differencing is crucial for capturing a causal effect. At the same time, the assumptions on the nature of the differences are often quite simple (e.g., firm-specific time-invariant heterogeneity), and it is not clear whether the results would remain the same if firms are actually allowed to react to a major change in their environment.

More generally, I expect the biggest contribution of micro data sets to come from a

better estimation of micro behavior, as the data is recorded right at the decision-taking level. And to reach their full potential, empirical micro studies of technology diffusion will probably involve modeling micro structure as well, because it is the structure that shapes behavior. As of now, there is little of that. For instance, currently, empirical studies of technology diffusion are not based on models whose structure explicitly allows for learning and spillover externalities between firms. That is, so far the micro evidence is based on “no-externality” models. This suggests that the micro evidence may become stronger once externalities are included in the models as a matter of theory (some work in this direction is Clerides, Keller, and Tybout 2003).

If the idea of differencing is to control for the averse effects of heterogeneity (across firms or over time), where does one stop? Of course, solving the heterogeneity issue fundamentally means conducting a case study, with a sample of one observation—but it is not necessarily advisable to go that far, at least not as the sole empirical methodology. More generally, can there be a problem of overdifferencing? An important issue is that technology data tends to be not of very high quality, and this limits the amount of differencing that one can do. Once the estimated coefficients look like a collection of random numbers scattered around zero with typically large standard errors, the researcher has probably asked too much of the data. Similarly, is it possible to probe so deeply into the structure of micro behavior that one fails to notice more diffuse, more aggregate, and slower-acting effects? As noted above, one benefit of attracting a major foreign investor, according to Larrain, Lopez-Calva, and Rodriguez-Claré (2000), is that it sends a signal to other potential foreign investors. It is difficult to see how a deeper analysis of micro structure would help to detect these types of effects.

Summarizing, I think that we can expect to obtain additional major results on interna-

tional technology diffusion in the future, and this will be in part because of the increased usage of micro data and estimation. At the same time, it will be important to keep an open mind and not expect too much from any one empirical methodology.

What are the policy implications of this literature? The evidence is not strong enough yet to provide support for specific policy measures, such as a particular subsidy to a multinational enterprise for locating in a country. This would require more agreement on quantitative effects than there is right now. At this point, the results suggest that the international dimension of technological change is of key importance for most countries. In this situation, a closed-off international economic regime must have detrimental welfare consequences for a country.

There is also evidence that the relative importance of international technology diffusion has been increasing with the level of economic integration in the world. This suggests that the performance advantage of outward-oriented economies over inward-oriented economies will be higher in the future than it is today. What if future research were to establish that international technology diffusion raises productivity more in advanced than in less-developed countries? The implication for less-developed countries would *not* be that a turn towards autarky is beneficial—on the contrary, given the evidence on positive productivity effects from international technology diffusion. Rather, if the benefits from operating in an international economic environment differ across countries, this suggests we should investigate further the major reasons for this—what works, and why?

While there is no consensus yet on the exact magnitude of spillover benefits, it is clear that well-functioning markets and an undistorted trade and foreign-investment regime are conducive to these learning effects. However, the evidence suggests that the latter are quite difficult to pin down. For one, it does not seem to be primarily a matter of simply specializing in the production

of high-tech goods. Rather, the goods characteristics are only part of what is important. India, for example, aimed at producing relatively advanced products during its era of import-substitution policy, but their quality was low compared to international standards. In contrast, Southeast Asian countries such as Hong Kong and South Korea were initially specializing in relatively low-tech products, and moved slowly but successfully into the range of higher-tech products.

Instead, technological knowledge spillovers appear to be resulting from a deliberate commitment to learning and matching international performance standards through ongoing interaction with foreigners. Local efforts are clearly necessary for successful technology adoption. At the same time, the ongoing interaction with foreign firms and consumers seems to be a process of knowledge discovery for firms that cannot be had from interacting only with other domestic firms. To model as well as empirically capture these in a convincing way continues to be a challenge.

REFERENCES

- Abramovitz, Moses. 1956. "Resource and Output Trends in the United States since 1870," *Amer. Econ. Rev.* 46:2, pp. 5–23.
- Acemoglu, Daron; Simon Johnson and James Robinson. 2001. "The Colonial Origins of Comparative Development," *Amer. Econ. Rev.* 91:5, pp. 1369–401.
- Aghion, Philippe and Peter Howitt. 1992. "A Model of Growth through Creative Destruction," *Econometrica* 60, pp. 323–51.
- . 1998. *Endogenous Growth Theory*. Cambridge, MA: MIT Press.
- Aitken, Brian and Ann Harrison. 1999. "Do Domestic Firms Benefit from Foreign Direct Investment? Evidence from Venezuela," *Amer. Econ. Rev.* 89:3, pp. 605–18.
- Arrow, Kenneth. 1969. "Classificatory Notes on the Production and Transmission of Technological Knowledge," *Amer. Econ. Rev. Pap. Proceed.* 59:2, pp. 29–35.
- Baldwin, Richard. 1989. "Sunk Costs Hysteresis," NBER work. pap. 2911, Cambridge, MA.
- Banerjee, Abhijit and Lakshmi Iyer. 2003. "History, Institutions, and Economic Performance: The Legacy of Colonial Land Tenure Systems in India," BREAD work. pap.
- Bartel, Ann and Frank Lichtenberg. 1987. "The Comparative Advantage of Educated Workers in

- Implementing Technology," *Rev. Econ. Statist.* 69:1, pp. 1–11.
- Bartelsman, Eric and Mark Doms. 2000. "Understanding Productivity: Lessons from Longitudinal Microdata," *J. Econ. Lit.* 38:3, pp. 569–94.
- Basu, Susanto and David Weil. 1998. "Appropriate Technology and Growth," *Quart. J. Econ.* 113:4, pp. 1025–54.
- Bernard, Andrew and Bradford Jensen. 1999. "Exceptional Exporter Performance: Cause, Effect, or Both?" *J. Int. Econ.* 47:1, pp. 1–25.
- Bernard, Andrew; Jonathan Eaton, Bradford Jensen and Samuel Kortum. 2003. "Plants and Productivity in International Trade," *Amer. Econ. Rev.* 93:4, pp. 1268–90.
- Blalock, Garrick and Paul Gertler. 2002. "Technology Diffusion from Foreign Direct Investment through Supply Chains," work pap., Haas School Bus. UC Berkeley.
- Blomström, Magnus and Ari Kokko. 1998. "Multinational Corporations and Spillovers," *J. Econ. Surveys* 12:3, pp. 247–77.
- Blundell, Richard, and Stephen Bond. 1999. "GMM Estimation with Persistent Panel Data: An Application to Production Functions," work. pap. 99/4, Inst. Fiscal Studies, London.
- Bottazzi, Laura and Giovanni Peri. 2003. "Innovation, Demand, and Knowledge Spillovers: Evidence from European Patent Data," *Europ. Econ. Rev.* 47, pp. 687–710.
- Branard, Lael. 1997. "An Empirical Assessment of the Proximity-Concentration Trade-off between Multinational Sales and Trade," *Amer. Econ. Rev.* 87:4, pp. 520–44.
- Branstetter, Lee. 2001a. "Is Foreign Direct Investment a Channel of Knowledge Spillovers? Evidence from Japan's FDI in the United States," NBER work. pap. 8015.
- . 2001b. "Are Knowledge Spillovers International or Intranational in Scope? Microeconomic Evidence from the U.S. and Japan," *J. Int. Econ.* 53, pp. 53–79.
- Caballero, Ricardo and Adam Jaffe. 1993. "How High Are the Giant's Shoulders? An Empirical Assessment of Knowledge Spillovers and Creative Destruction in a Model of Growth," NBER *Macroeconomic Annual*, vol. 8. Olivier Blanchard and Stanley Fischer, eds.
- Caselli, Francesco and Wilbur Coleman II. 2001. "Cross-Country Technology Diffusion: The Case of Computers," *Amer. Econ. Rev. Pap. Proceed.* 91:2, pp. 328–35.
- Caves, Richard; Laurits Christensen and Erwin Diewert. 1982. "Multilateral Comparisons of Output, Input, and Productivity Using Superlative Index Numbers," *Econ. J.* 92:365, pp. 73–86.
- CIC. 2003. The International Comparison of Prices Program. Center Int. Comparisons, U. Penn.
- Clerides, Sofronis; Wolfgang Keller and James Tybout. 2003. "Foreign Direct Investment, Exports, and Spillovers," in progress, Brown U.
- Clerides, Sofronis; Saul Lach and James Tybout. 1998. "Is Learning by Exporting Important? Micro-dynamic Evidence from Colombia, Mexico, and Morocco," *Quart. J. Econ.* 113, pp. 903–48.
- Coe, David and Elhanan Helpman. 1995. "International R&D Spillovers," *Europ. Econ. Rev.* 39:5, pp. 859–87.
- Coe, David; Elhanan Helpman and Alexander Hoffmaister. 1997. "North-South Spillovers," *Econ. J.* 107, pp. 134–49.
- Coe, David and Alexander Hoffmaister. 1999. "Are There International R&D Spillovers among Randomly Matched Trade Partners? A Response to Keller," IMF work. pap. 99/18.
- Cohen, Wesley and Daniel Levinthal. 1989. "Innovation and Learning: The Two Faces of R&D," *Econ. J.* 99:397, pp. 569–96.
- David, Paul. 1992. "Knowledge, Property, and the Systems Dynamics of Technological Change," in *Proceedings of the World Bank Annual Conference on Development Economics 1992*, Larry Summers and Anwar Shah, eds. pp. 215–48.
- Dixit, Avinash. 1989. "Entry and Exit Decisions under Uncertainty," *J. Polit. Econ.* 97:3, pp. 620–38.
- Dornbusch, Rüdiger; Stanley Fischer and Paul Samuelson. 1977. "Comparative Advantage, Trade, and Payments in a Ricardian Model with a Continuum of Goods," *Amer. Econ. Rev.* 67:5, pp. 823–39.
- Easterly, William and Ross Levine. 2001. "It's Not Factor Accumulation: Stylized Facts and Growth Models," *World Bank Econ. Rev.* 15:2, pp. 177–219.
- Eaton, Jonathan; Eva Gutierrez and Samuel Kortum. 1998. "European Technology Policy," *Econ. Pol.* 27, pp. 405–38.
- Eaton, Jonathan and Samuel Kortum 1996. "Trade in Ideas: Patenting and Productivity in the OECD," *J. Int. Econ.* 40:3-4, pp. 251–78.
- . 1997. "Engines of Growth: Domestic and Foreign Sources of Innovation," *Japan World Econ.* 9, pp. 235–59.
- . 1999. "International Patenting and Technology Diffusion: Theory and Measurement," *Int. Econ. Rev.* 40, pp. 537–70.
- . 2001. "Trade in Capital Goods," *Europ. Econ. Rev.* 45:7, pp. 1195–235.
- . 2002. "Technology, Geography, and Trade," *Econometrica* 70:5, pp. 1741–79.
- Edmonds, Chris. 2001. "Some Panel Cointegration Models of International R&D Spillovers," *J. Macroecon.* 23, pp. 241–60.
- Ethier, William. 1986. "The Multinational Firm," *Quart. J. Econ.* 101, pp. 805–33.
- . 1982. "National and International Returns to Scale in the Modern Theory of International Trade," *Amer. Econ. Rev.* 72, pp. 389–405.
- Evenson, Robert and Larry Westphal. 1995. "Technology Change and Technology Strategy," in *Handbook of Development Economics*, vol. 3A. Jere Behrman and T.N. Srinivasan, eds. North-Holland.
- Fagerberg, Jan. 1994. "Technology and International Differences in Growth Rates," *J. Econ. Lit.* 32:3, pp. 1147–75.
- Feldman, Maryann and Frank Lichtenberg. 1997. "The

- Impact and Organization of Publicly-Funded Research and Development in the European Community," NBER work. pap. 6040.
- Fosfuri, Andrea; Massimo Motta and Thomas Rønnde. 2001. "Foreign Direct Investment and Spillovers through Workers' Mobility," *J. Int. Econ.* 53, pp. 205–22.
- Galor, Oded and Daniel Tsiddon. 1997. "Technological Progress, Mobility, and Economic Growth," *Amer. Econ. Rev.* 87, pp. 363–82.
- Gerschenkron, Alexander. 1962. *Economic Backwardness in Historical Perspective*. Cambridge, MA: Harvard U. Press.
- Girma, Sourafel and Katharina Wakelin. 2001. "Regional Underdevelopment: Is FDI the Solution? A Semi-Parametric Analysis," GEP Research Pap. 2001/11, U. Nottingham.
- Görg, Holger, and David Greenaway. 2002. "Much Ado About Nothing? Do Domestic Firms Really Benefit from Foreign Direct Investment?" work pap. U. Nottingham.
- Globerman, Steven; Ari Kokko and Frederic Sjöholm. 2000. "International Technology Diffusion: Evidence from Swedish Patent Data," *Kyklos* 53, pp. 17–38.
- Griffith, Rachel; Stephen Redding and Helen Simpson. 2003. "Productivity Convergence and Foreign Ownership at the Establishment Level," CEPR work. pap. 3765.
- Griffith, Rachel; Stephen Redding, and John Van Reenen. 2000. "Mapping Two Faces of R&D: Productivity Growth in a Panel of OECD Industries," forthcoming, *Rev. Econ. Statist.*
- Griliches, Zvi. 1979. "Issues in Assessing the Contribution of Research and Development to Productivity Growth," *Bell J. Econ.* 10, pp. 92–116.
- . ed. 1984. *R&D, Patents, and Productivity*. U. Chicago for NBER.
- . 1990. "Patent Statistics as Economic Indicators: A Survey," *J. Econ. Lit.* 28:3, pp. 1661–707.
- . 1995. "R&D and Productivity: Econometric Results and Measurement Issues," in *Handbook of the Economics of Innovation and Technological Change*. Paul Stoneman ed., Oxford: Blackwell, pp. 52–89.
- Griliches, Zvi, and Jerry Hausman. 1986. "Errors in Variables in Panel Data," *J. Econometrics* 31, pp. 93–118.
- Griliches, Zvi, and Jacques Mairesse. 1998. "Production Functions: The Search for Identification," in *Econometrics and Economic Theory in the 20th Century: Ragnar Frisch Centennial Symposium*. Steinar Strøm, ed. Cambridge: Cambridge U. Press, pp. 169–203.
- Grossman, Gene and Elhanan Helpman. 1991. *Innovation and Growth in the World Economy*, Cambridge, MA: MIT Press.
- . 1995. "Technology and Trade," in *Handbook of International Economics* vol.3. Gene Grossman and Kenneth Rogoff, eds. North-Holland.
- Hall, Robert and Charles Jones. 1999. "Why Do Some Countries Produce So Much More Output Per Worker Than Others?" *Quart. J. Econ.* 114:1, pp. 83–116.
- Hallward-Driemeier, Mary; Giuseppe Iarossi and Kenneth Sokoloff. 2002. "Exports and Manufacturing Productivity in East Asia: A Comparative Analysis with Firm-Level Data," work. pap., UCLA.
- Hanson, Gordon. 2001. "Should Countries Promote Foreign Direct Investment?" G-24, UN Discus. Pap. Series, NY and Geneva.
- Haskel, Jonathan; Sonia Pereira and Matthew Slaughter. 2001. "Does Inward Foreign Direct Investment Boost the Productivity of Domestic Firms?" pap. for NBER Summer Institute.
- Heston, Alan; Robert Summers and Bettina Aten. 2002. *Penn World Tables Version 6.1*, Center Int. Comparisons, U. Penn.
- Howitt, Peter. 2000. "Endogenous Growth and Cross-Country Income Differences," *Amer. Econ. Rev.* 90:4, pp. 829–46.
- Hulten, Charles. 2000. "Total Factor Productivity: A Short Biography," NBER work. pap. 7471.
- Irwin, Douglas and Peter Klenow. 1994. "Learning Spillovers in the Semi-Conductor Industry," *J. Polit. Econ.* 102, pp. 1200–27.
- Jaffe, Adam and Manuel Trajtenberg. 2002. *Patents, Citations and Innovations: A Window on the Knowledge Economy*. Cambridge: MIT Press.
- Jaffe, Adam; Manuel Trajtenberg and Rebecca Henderson. 1993. "Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations," *Quart. J. Econ.* 108:3, pp. 577–98.
- Jones, Charles. 1995. "R&D-Based Models of Growth," *J. Polit. Econ.* 103:4, pp. 759–84.
- Katayama, Hajime; Shihua Lu and James Tybout. 2003. "Why Plant-Level Productivity Studies Are Often Misleading, and an Alternative Approach to Inference," work pap. Penn State U.
- Keller, Wolfgang. 1996. "Absorptive Capacity: On the Creation and Acquisition of Technology in Development," *J. Devel. Econ.* 49, pp. 199–227.
- . 1997. "How Trade Patterns and Technology Flows Affect Productivity Growth," World Bank Policy Research work. pap. 1831.
- . 1998. "Are International R&D Spillovers Trade Related? Analyzing Spillovers among Randomly Matched Trade Partners," *Europ. Econ. Rev.* 42, pp. 1469–81.
- . 2000. "Do Trade Patterns and Technology Flows Affect Productivity Growth?" *World Bank Econ. Rev.* 14, pp. 17–47.
- . 2001. "Knowledge Spillovers at the World's Technology Frontier," CEPR work. pap. 2815.
- . 2002a. "Geographic Localization of International Technology Diffusion," *Amer. Econ. Rev.* 92, pp. 120–42.
- . 2002b. "Trade and the Transmission of Technology," *J. Econ. Growth*: 7, pp. 5–24.
- Keller, Wolfgang and Stephen Yeaple. 2003. "Multinational Enterprises, International Trade, and Productivity Growth: Firm Level Evidence from the United States," IMF work. pap. 248.
- Klette, Tor and Zvi Griliches. 1996. "The Inconsistency of Common Scale Estimators When Output Prices Are Unobserved and Endogenous," *J. Appl. Econometrics* 11, pp. 343–61.

- Krugman, Paul. 1989. *Exchange Rate Instability*. Cambridge, MA: MIT Press.
- Kugler, Maurice. 2002. "The Diffusion of Externalities from Foreign Direct Investment: Theory Ahead of Measurement," work pap. U. Southampton.
- Larrain, B. Felipe; Luis Lopez-Calva, and Andres Rodriguez-Claré. 2000. "Intel: A Case Study of Foreign Direct Investment in Central America," work. pap. 58, Center Int. Develop., Harvard U.
- Leamer, Edward and James Levinsohn. 1995. "International Trade theory: The Evidence," in *Handbook of International Economics*, vol. 3. G. Grossman and K. Rogoff, eds. Elsevier: Amsterdam.
- Lichtenberg, Frank. 1993. "R&D Investment and International Productivity Differences," in *Economic Growth in the World Economy*. Symposium 1992. Horst Siebert, ed. Tübingen: Mohr.
- Lichtenberg, Frank and Bruno van Pottelsberghe de la Potterie. 2000. "Does Foreign Direct Investment Transfer Technology Across Borders?" forthcoming, *Rev. Econ. Statist.*
- Lucas, Robert. 1988. "On the Mechanics of Economic Development," *J. Monet. Econ.* 22, pp. 3–42.
- Lumenga-Neso, Olivier; Marcelo Olarreaga and Maurice Schiff. 2001. "On 'Indirect' Trade-Related R&D spillovers," mimeo, World Bank.
- Maddala, G.S. and In-Moo Kim. 1998. *Unit Roots, Cointegration, and Structural Change*. Cambridge: Cambridge U. Press.
- Markusen, James. 2002. *Multinational Firms and the Theory of International Trade*. Cambridge: MIT Press.
- Nelson, Richard and Edmund Phelps. 1966. "Investment in Humans, Technological Diffusion and Economic Growth," *Amer. Econ. Rev.* 56:2, pp. 69–75.
- NSF. 2004. "US–China R&D Linkages: Direct Investment and Industrial Alliances in the 1990s," NSF Info. Brief NSF 04-236.
- OECD. 2002. *Frascati Manual 2002: Proposed Standard Practice for Surveys on Research and Development*. Paris: OECD.
- . 2003. *Main Science and Technology Indicators*. Paris: OECD.
- Olley, Steven and Ariel Pakes. 1996. "The Dynamics of Productivity in the Telecommunications Equipment Industry," *Econometrica* 64:6, pp. 1263–97.
- Pakes, Ariel. 1986. "Patents as Options: Some Estimates of the Value of Holding European Patent Stocks," *Econometrica* 54:4, pp. 755–84.
- Peri, Giovanni. 2002. "Knowledge Flows and Innovation," work. pap., UC Davis.
- Polanyi, Michael. 1958. *Personal Knowledge: Towards a Post-Critical Philosophy*. Chicago: U. Chicago Press.
- Prescott, Edward. 1998. "Needed: A Theory of Total Factor Productivity," *Int. Econ. Rev.* 39:3, pp. 525–51.
- Quah, Danny. 2001a. "Technology Dissemination and Economic Growth," work. pap., London School Econ.
- . 2001b. "Demand-Driven Knowledge Clusters in a Weightless Economy," work pap., London School Econ.
- Rhee, Yong-Whee; Bruce Ross-Larson and Garry Pursell. 1984. *Korea's Competitive Edge: Managing Entry into World Markets*. Baltimore: Johns Hopkins U. Press for World Bank.
- Rivera-Batiz, Luis and Paul Romer. 1991. "Economic Integration and Endogenous Growth," *Quart. J. Econ.* 106:2, pp. 531–55.
- Rodriguez-Claré, Andres. 1996. "Multinationals, Linkages, and Economic Development," *Amer. Econ. Rev.* 86, pp. 852–73.
- Rodrik, Dani. 1999. "The New Global Economy and Developing Countries: Making Openness Work," Policy Essay 24, Overseas Develop. Council, Washington DC.
- Romer, Paul. 1990. "Endogenous Technological Change," *J. Polit. Econ.* 98:5, pp. S71–S102.
- Saggi, Kamal. 2000. "Trade, Foreign Direct Investment, and International Technology Transfer: A Survey," forthcoming in *World Bank Res. Observer*.
- Scherer, Frederic. 1984. "Using Linked Patent and R&D Data to Measure Interindustry Technology Flows," in *R&D, Patents, and Productivity*. Zvi Griliches ed. U. Chicago for NBER, pp. 417–61.
- Segerstrom, Paul; T.C.A. Anant and Elias Dinopoulos. 1990. "A Schumpeterian Model of the Product Life Cycle," *Amer. Econ. Rev.* 80:4, pp. 1077–92.
- Shiue, Carol H. and Wolfgang Keller. 2004. "Markets in China and Europe on the Eve of the Industrial Revolution," CEPR discus. pap. 4420.
- Singh, Jasjit. 2003. "Knowledge Diffusion and Multinational Firms: Evidence using Patent Citation Data," work. pap., Grad. School Bus. Admin. and Econ. Dept., Harvard U.
- Sjöholm, Frederic. 1996. "International Transfer of Knowledge: The Role of International Trade and Geographic Proximity," *Weltwirtschaftliches Archiv* 132, pp. 97–115.
- Teece, David. 1977. "Technology Transfer by Multinational Firms: The Resource Cost of Transferring Technological Know-How," *Econ. J.* 87, pp. 242–61.
- von Hippel, Eric. 1994. "'Sticky Information' and the Locus of Problem Solving: Implications for Innovation," *Manage. Sci.* 40:4, pp. 429–39.
- Xu, Bin. 2000. "Multinational Enterprises, Technology Diffusion, and Host Country Productivity Growth," *J. Devel. Econ.*, 62:2, pp. 477–93.
- Xu, Bin and Jianmao Wang. 1999. "Capital Goods Trade and R&D Spillovers in the OECD," *Can. J. Econ.* 32:5, pp. 1258–74.
- WIPO. 2003. World Intellectual Property Organization, www.wipo.org.
- World Bank. 1999. *Knowledge for Development*. World Development Report 1998/99, Washington DC: World Bank.