

NBER WORKING PAPER SERIES

MEASURING THE “IDEAS” PRODUCTION FUNCTION:
EVIDENCE FROM INTERNATIONAL PATENT OUTPUT

Michael E. Porter
Scott Stern

Working Paper 7891
<http://www.nber.org/papers/w7891>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2000

We thank Joshua Gans, Chad Jones, Sam Kortum, Paul Romer, and seminar participants at the NBER Summer Institute for helpful comments. Jeff Furman, Irineu Carvalho, and Greg Bond provided exceptional research assistance. The authors acknowledge the financial support of the Council on Competitiveness. The views expressed herein are those of the authors and not necessarily those of the National Bureau of Economic Research.

© 2000 by Michael E. Porter and Scott Stern. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Measuring the “Ideas” Production Function: Evidence from International Patent Output
Michael E. Porter and Scott Stern
NBER Working Paper No. 7891
September 2000
JEL No. O3, O47, O57

ABSTRACT

This paper estimates the parameters of the “ideas” production function central to recent models of economic growth. We do so by evaluating the determinants of “international” patenting rates across the OECD, where an international patent is one granted by the U.S. patent office to a foreign establishment. Taking advantage of variation in the flow of ideas produced by different countries over time, we provide evidence for three main findings. First, at the level of the production of international patents, country-level R&D productivity increases proportionally with the stock of ideas already discovered, a key parametric restriction associated with the Romer model of ideas-driven growth (Romer, 1990; Jones, 1995). Second, we find that ideas productivity in a given country is constant or declining in the *worldwide* stock of ideas. Ideas production by other countries raises the bar for producing new-to-the-world technology domestically, outweighing the positive effects of international knowledge spillovers. Finally, ideas productivity is concave in the size of the R&D workforce and the linkage between ideas production and overall productivity growth is small. These results suggest that while the parametric restrictions required to generate endogenous technological change may be satisfied for individual economies, the growth rate associated with such effects may be modest. There seems to be a gap between the the sustained production of ideas by advanced economies and the ability to translate ideas into measured productivity growth.

Michael E. Porter
Harvard Business School
Aldrich Hall 200
Soldiers Field Road
Boston, MA 02163
mporter@hbs.edu

Scott Stern
MIT Sloan School
E52-554
50 Memorial Drive
Cambridge, MA 02142
and NBER
sstern@mit.edu

I. Introduction

A central debate within new growth theory is centered around the role of the “ideas” sector in sustaining equilibrium productivity growth. In Romer’s seminal model of endogenous technological change, productivity growth is driven by a constant allocation of resources to an ideas-producing sector (Romer, 1990), a result which depends critically on strong positive intertemporal spillovers in ideas production. Specifically, to generate ideas-driven growth, ideas sector productivity must increase proportionally with the stock of ideas already discovered.

The significance of ideas-driven growth therefore depends on whether the production function for ideas satisfies this critical property. To evaluate this claim, several authors have examined the relationship between the TFP growth rate and the size of the workforce devoted to the production of ideas (Jones, 1995; Coe and Helpman, 1995). The Romer model predicts that expansion in the number of ideas workers leads to a permanent increase in the TFP growth rate. In contrast, the empirical evidence suggests that most OECD economies have increased the size of their R&D workforce while experiencing (at best) constant TFP growth rates. This weak relationship between the size of the ideas workforce and the TFP growth rate has led some to question the viability of ideas-driven growth for the long run (Jones, 1999).¹

This paper contributes to the empirical understanding of economic growth by estimating the shape of the ideas production function and the strength of the intertemporal and international spillovers in ideas. We examine the time-series pattern of international patenting (patents granted in the United States to inventors from other OECD countries) to evaluate the determinants of the flow of “new-to-the-world” ideas directly.² By using a panel dataset of patents, we achieve several goals which cannot be addressed in the TFP growth rate literature. First, by examining patenting (an observable manifestation of ideas production), we can separate

¹ Several authors instead argue that productivity growth rates can be explained by factor accumulation (as first suggested in Solow (1956)) including the accumulation of human capital (Barro, 1991; Mankiw, Romer, and Weil, 1992). Jones’ synthesis model incorporates these insights and includes an ideas sector with weaker intertemporal spillovers, and whose predictions and empirical motivation are similar to factor accumulation theories (Jones, 1995).

² In its use of international patenting data to evaluate specific theories of economic growth, this paper builds on Eaton and Kortum (1996; 1998). However, we are able to substantially extend these prior analyses by exploiting the panel nature of the data and so go beyond the “snapshot” methodology employed in this prior work.

ideas production from the more general relationship between the ideas sector and overall productivity growth. Accordingly, we distinguish the properties of the ideas production function from the sensitivity of TFP growth to ideas production. Second, by exploiting the accumulation of patents over time, we can explicitly estimate the strength of the spillover from ideas-to-ideas. In other words, we assess the strength of the linkage between patenting in the past and current ideas sector productivity. Finally, by evaluating patenting patterns using a panel dataset of OECD countries, we contribute to an emerging literature on the differences between domestic and international knowledge spillovers (Porter, 1990; Nelson, 1993; Coe and Helpman, 1995; Park, 1995; Eaton and Kortum, 1998). Specifically, we differentiate the shape of the national ideas production function (which includes country-specific spillovers) from the more general world ideas production function.

Our analysis employs a parameterized version of the ideas production function, $\dot{A}_t = \delta H_{A,t}^\lambda A_t^\phi$, where A is a measure of the total stock of knowledge/ideas and H_A is the quantity of human capital devoted to the ideas-producing sector. The relationship between the stock of ideas, A , and the flow, \dot{A}_t , represents the key intertemporal spillover – the relationship between ideas productivity today and the number of ideas discovered in the past. As a result, the empirical debate centers around the value of ϕ and the relevant measure of A , the stock of knowledge upon which current ideas productivity draws. Romer (1990) assumes strong positive spillovers (setting $\phi = 1$) and that all actors at a given point in time draw upon the same stock of knowledge. Under these assumptions, a given percentage increase in the stock of ideas in a national economy results in a proportional increase in the productivity of the ideas-producing sector, and the growth rate of ideas is a function of the level of human capital devoted to ideas

$$\left(\frac{\dot{A}}{A} = \delta H_A^\lambda\right).$$

Our analysis amends the Romer ideas production function in two important ways. First, following Jones (1995), we allow for concavity in ideas production ($0 < \phi < 1$, $0 < \lambda < 1$).³ Second, we differentiate the stock of domestic knowledge from the world knowledge pool. In

³ The principal consequences of such concavity is that a constant growth rate for ideas depends on constant *growth* in the resources devoted to ideas production so that policies which induce an increase in ideas production will shift

evaluating the drivers of ideas sector productivity, we allow for separate contributions from the country-specific and rest-of-world knowledge stocks, respectively. By separately considering multiple knowledge sources for each country's national ideas production function, we evaluate the relative importance played by international versus domestic knowledge spillovers in fostering the production of new ideas.⁴

We explore the properties of the national ideas production function by examining the determinants of the flow of international patents. For each of the 16 OECD countries in our dataset, we define an “international” patent to be a patent granted by the U.S. patent office to foreign establishments. We exploit variation in the rate of international patenting – a (mostly) comparable measure of the flow of commercially relevant ideas across time and countries – to identify the key parameters of the ideas production function.⁵ Of course, this measure is a subset of the total flow of new ideas produced by an economy in a given year, as it excludes those ideas which are (a) novel to a country but have been already discovered elsewhere and (b) those ideas which are not worthwhile patenting internationally. To overcome these two potential sources of measurement bias, we derive the expected flow of international patents (i.e., valuable “new-to-the-world” technologies) from the more primitive national ideas production function. Specifically, we propose a structural model of international patenting which depends on a year-specific effect, a country-specific effect, and the national ideas production function.⁶ In other words, we are able to test for key parametric restrictions in ideas-based growth by taking advantage of the differences across advanced economies in their knowledge stocks and their allocation of resources to invention.

the *level* of productivity but will have no long-term impact on the growth *rate*.

⁴ We also consider extensions involving “intermediate” spillover regimes (e.g., specific geographic regions, countries with a common language base, or spillovers proportional to trading shares.)

⁵ The use of patent data to evaluate spillovers and economic growth goes back at least to Schmookler (1963). Over the last decade (beginning with the work in Griliches (1984)), there has been increased attention to both the potential uses (and misuses) of this type of data, both in the domestic and international context (Jaffe, 1986; Griliches, 1990; Trajtenberg, 1990; Dosi, Pavitt and Soete, 1990).

⁶ This derived relationship is in the spirit of Eaton and Kortum (1996, 1998) and Jones (1998a). However, Eaton and Kortum focus on patenting and technology transfer while Jones focuses on the world ideas production function.

With this theoretical structure in mind, we employ two distinct estimation strategies. In the first, hereafter referred to as the *direct* approach, we exploit the time-series nature of our data to construct a “knowledge stock” ($A_{jt} = \sum_{s=0}^{t-1} \dot{A}_{js}$), where A_{jt} is the sum of international patents granted to country j up until year t .⁷ We then use this measure to evaluate the sensitivity of the flow of international patenting to the stock. In the second approach, referred to as the *inversion* approach, we invert the production function ($Y = A^\sigma K^\alpha H_L^{1-\alpha} \Rightarrow A = \left[\frac{Y}{K^\alpha H_L^{1-\alpha}} \right]^{\frac{1}{\sigma}}$) substituting aggregate output and factor inputs for the unmeasured ideas stock. The key difference between the direct and inversion approaches is in the measurement of the knowledge stock that each country draws upon in discovering new ideas. Finally, rather than imposing complicated cross-equation restrictions, our empirical methodology seeks to let the data speak for itself as much as possible. We employ relatively simple regressions and the empirical argument proceeds from the simplest procedure to more sophisticated treatments, highlighting the main robust empirical relationships.

We provide evidence for three main findings. First, there is compelling empirical evidence that the crucial parametric restriction of the Romer model, $\phi = 1$, operates at the country level. Whether the ideas stock is calculated under the direct or the inversion approaches, patenting productivity is proportional (or nearly proportional) to our estimate of the domestic knowledge stock. Second, we find a negative relationship between domestic R&D productivity and various measures of the world knowledge stock. Ideas production by other countries raises the bar for producing new-to-the-world technology domestically, and this effect outweighs the positive effects of international knowledge spillovers. While cross-border knowledge spillovers are stronger for groups of countries which share a common language base or are geographically proximate to each other, their strength is small compared to the relationship found between the domestic knowledge stock and current patenting productivity. Our third finding results from our attempt to reconcile these first two findings with the poor performance of the Romer model in

⁷ Adams (1990) provides a useful discussion of the issues involved in constructing stocks of scientific and engineering knowledge.

explaining aggregate productivity growth. We explore how TFP is affected by changes in the domestic and international stock of patents and find a small and positive effect (i.e., a doubling of the patent stock is predicted to have only a 10% increase on the level of TFP). This suggests that while endogenous growth may be feasible, the size of the ideas-driven growth effects may be quite modest. Concavity in the returns to additional researchers and the modest relationship between increments to the measured stock of ideas and overall productivity growth dampen the size of endogenous growth effects.

These findings suggest that while R&D workers may “stand on the shoulders” of prior (domestic) research, these technical achievements have had only a limited impact on measured aggregate productivity. Indeed, there seems to be a gap between the increasing level of resources devoted to R&D and the ability of advanced economies to translate these new technologies into higher TFP. Our evidence suggests that, relative to modern ideas-driven growth models, there is a weaker (but still positive) link between ideas production and realized productivity growth. Identifying the factors which enhance this linkage is a promising area for further research.

The remainder of the paper proceeds as follows. In Section II, we review the parameterization and functional form of the ideas production function, focusing on the elements which distinguish competing models of the process of economic growth. In Section III, we extend the underlying model to incorporate national ideas production in an international context. In Sections IV and V, we turn to the development of an empirical model based on the use of international patenting data; to do so, we derive the relationship between national ideas production and the production of “new-to-the-world” technologies and describe how we measure the knowledge stock of each country. After a short description of the data (Section VI), the principal empirical results are reviewed in Section VII. Section VIII is devoted to reconciling the evidence on patenting with data on TFP. A final section concludes.

II. Parameterizing the Ideas Production Function

Recent models of technology-driven growth extend the neoclassical growth model by endogenizing the rate of technological progress (\dot{A}) through the introduction an ideas sector for

the economy with an associated production function, $\dot{A}_t = \delta H_{A,t}^\lambda A_t^\phi$.⁸ According to this production structure, the rate of new ideas production is simply a function of the number of ideas workers (H_A) and the stock of ideas available to these researchers (A_t). This function makes the rate of technological change endogenous in two distinct ways. First, the share of the economy devoted to the ideas sector is a function of the R&D labor market; allocation to the ideas sector depends on R&D productivity and the private economic return to new ideas. Second, the productivity of new ideas production is sensitive to the stock of ideas discovered in the past. The sign and size of this intertemporal externality is the critical issue for the feasibility of long-term ideas-based growth.⁹ When $\phi > 0$, prior research increases current R&D productivity (the so-called “standing on shoulders” effect); when $\phi < 0$, prior research has discovered the ideas which are easiest to find, making new ideas discovery more difficult (the “fishing out” hypothesis).

The magnitude of ϕ is crucial to the debate about whether equilibrium ideas-driven growth is feasible. To see this, consider the differences between the initial specification suggested by Romer (1990) and the synthesis model proposed by Jones (1995).¹⁰ In the Romer model, $\phi = \lambda = 1$. This implies that a given percentage increase in the stock of ideas results in a proportional increase in the productivity of the ideas-producing sector. Under this assumption,

⁸ This specific structure was pioneered in Romer (1990) and is employed (though occasionally in modified form) in several modern growth models (Grossman and Helpman, 1991; Jones 1995; Eaton and Kortum, 1998; and Jones (1998) for an introduction).

⁹ We therefore are focusing on the second aspect of endogeneity in the ideas production framework (sufficient positive feedbacks to ideas production itself). The empirical analysis does not address directly the viability of preserving economic incentives for ideas production except insofar as we separately evaluate the salience of ideas production for generating TFP growth (see Section VIII). However, given that the ideas sector in many country is subject to large public subsidies and that it is difficult to measure the contribution of new technologies to TFP, we believe that evaluating the viability of preserving innovation incentives over the long run awaits further research.

¹⁰ While we focus on the specifics of these two alternatives, each reflects a distinct stream of research which offers dramatically different assessments of the sources of long-term economic growth. Jones, in particular, proposes a model consistent with the capital investment oriented interpretations of economic growth proposed by, among others, Mankiw, Romer and Weil (1992) and Barro (1991). In contrast, Romer emphasizes the endogenous production of economically useful new knowledge rather than simple capital accumulation.

the growth rate in ideas is a function of the level of effort devoted to ideas production

($\frac{\dot{A}}{A} = \delta H_A$), ensuring a positive growth rate even in the absence of population growth. Under

the Romer restrictions, policies which permanently shift the level of H_A have permanent effects on the growth rate.

Jones questions the empirical relevance of this functional form and suggests that the strength of intertemporal spillovers may be less than proportional ($\phi < 1$) and there may be concavity in the returns to research effort at a point in time ($\lambda < 1$). Under this logic, the constant

growth rate in ideas is $\frac{\dot{A}}{A} = \frac{\delta L_A^\lambda}{A^{1-\phi}}$; taking logs and the time derivative, the ideas growth rate is a

function of the *growth rate* in the effort devoted to ideas production, $\frac{\dot{A}}{A} = \frac{\lambda (\frac{\dot{L}_A}{L_A})}{1-\phi}$. In other

words, when ideas spillovers are weaker than assumed by Romer, the steady-state growth rate in ideas will be zero in the absence of a sustained rate of increase in the size of the labor force devoted to ideas production (see Jones, 1999, for further exposition on this point).

The relationship between these two functional forms can be understood in terms of their implications for transition dynamics. Consider a small shift in the size of the workforce devoted to ideas production. For values of λ and ϕ approaching one, a large increase in the level of ideas productivity would result (according to $\frac{\lambda}{1-\phi}$); achieving this uptick in productivity would result from a long period of transition dynamics in which the economy would mimic the behavior of the Romer model. In other words, while steady-state predictions of the model are starkly different, the models behave quite similarly in the “medium term” as ϕ approaches one.

It is important to note that, in either model, equilibrium aggregate productivity growth depends on the rate of growth in ideas. In particular, both models share an aggregate production function, $Y_{it} = A_{it}^\sigma K_{it}^\alpha L_{it}^{1-\alpha}$, which depends directly on the growth in the stock of ideas.¹¹ A separate empirical issue arises about the relationship between ideas production and aggregate

¹¹ For expositional purposes, we abstract away from the intermediate goods sector emphasized in the Romer model and impose the Solow production function (Solow, 1956, 1957). Replacing the capital stock with the output of the

productivity growth. Models of endogenous technological change do not distinguish explicitly between ϕ , the return of ideas-to-ideas, and σ , the return of output-to-ideas. These parameters need not be the same: a high value of ϕ can coexist with a low value of σ . Moreover, endogenous ideas-driven growth requires only that ϕ be equal to one (yielding $\frac{\dot{A}}{A} = \delta H_A$). Realized productivity growth, in contrast, is a function of *all* of the parameters of both the ideas and output production functions ($\frac{\dot{Y}}{Y} = \sigma \delta H_A^\lambda$). While the feasibility of endogenous growth depends only on $\phi = 1$, its empirical relevance depends on the degree of concavity of ideas production in ideas workers and in the sensitivity of output growth to ideas growth.¹²

III. The National Ideas Production Function in an International Context

The ideas production function can be enriched by distinguishing between the domestic and international stock of ideas. Most earlier research which has attempted to derive the implications of growth theory for international data has focused on how the stock of ideas (or technology) in one country influences productivity growth in other countries through imitation and trade.¹³ Our more narrow emphasis is on how the stock of international ideas impacts

intermediate goods sector has no impact on our main conclusions.

¹² As well, equilibrium ideas-driven growth depends on preserving the incentives for innovation. Under the Romer restriction, these incentives are preserved to the extent that the marginal productivity of the ideas sector increases proportionally with TFP (and ideas producers are assumed to be able to extract the marginal product of their ideas). To the extent that $\phi < 1$, the MRP of the ideas sector must decline over time. When $\phi = 1$ but $\sigma < 1$, the viability of the ideas sector depends either on public subsidies for knowledge production or increases over time in the level of rent extraction by the ideas sector. While both of these policies may not be viable for the very long-run (see Jones (1999)), both are relevant concerns in understanding patterns of R&D employment and investment in advanced economies over the past century (Romer, 1990; Nelson, 1993; Stern, Furman and Porter, 1999).

¹³ Indeed, the implications of international knowledge spillovers has been an active area of both theoretical and empirical research within growth theory (Grossman and Helpman, 1991; Rivera-Batiz and Romer, 1991; Coe and Helpman, 1995; Eaton and Kortum, 1996, 1998; Keller, 1997). This literature emphasizes that ideas produced abroad contribute to aggregate productivity through licensing, imitation, and the supply of intermediate capital goods (Romer, 1996). A principal contention of this literature is that under equal factor prices, and as a consequence of knowledge spillovers, the growth rate of laggards will be higher than leading nations until convergence is achieved (Abromowitz, 1986; Barro and Sala-i-Martin, 1992).

domestic ideas production. The main reason for this narrow focus is that we are attempting to disentangle the specific properties of ideas production from the more general relationship between ideas production and overall productivity growth.

Our proposed functional form incorporates both domestic and international sources of knowledge. Consider the *national* ideas production function for country j ,

$$\dot{A}_{j,t} = \delta H_A^\lambda A_j^\phi A_{-j}^\psi \quad (1)$$

where A_j is the stock of ideas native to country j , and A_{-j} is the stock of ideas which have been discovered elsewhere but have not yet diffused into country j . As the flow of “new-to-the-country” ideas, $\dot{A}_{j,t}$ may include ideas which have been discovered and utilized elsewhere in the world but not yet used in country j . In this sense, a positive value of ψ captures the possibility of international imitation -- creating new-to-the-country technologies is easier as the foreign stock of ideas increases.¹⁴ This functional form draws upon a microeconomic literature which emphasizes the potential for complementarity between domestic and international knowledge (Jaffe and Trajtenberg, 1998): ψ determines the strength of the complementarity between the foreign and domestic ideas stock.¹⁵

IV. A Structural Model of International Patenting

In this section, we derive the link between the national ideas production function and the production of international patents (i.e., patents granted in the United States to foreign inventors). In so doing, we develop a specific econometric model which we will use to estimate the parameters of (1). Several distinct issues arise: (a) the relationship between the production of new-to-the-country and new-to-the-world ideas; (b) concavity in global ideas production; (c) variation across countries in terms of their propensity-to-patent; (d) the lag between invention and patent granting; and (e) identifying the parameters of the national ideas production function. The remainder of this section derives an empirical model which accounts for each of these

¹⁴ However, diffusion is not free -- acquiring ideas from abroad requires resources (H_A) and sophistication (A_j).

¹⁵ Of course, alternative formulations could be explored (e.g., the CES ($\dot{A}_{j,t} = \delta H_A^\lambda (A_j^{\rho/\phi} + A_{-j}^{\rho/\phi})^\phi$). Preliminary empirical research in this direction suggests that, while computationally more complex, our qualitative results are robust to the CES specification.

issues.

For a country at a point in time, only a subset of the ideas contributing to its own economic growth will also be globally novel (the requirement for international patent protection). In other words, (1) represents the production of new-to-the-country ideas, not new-to-the-world ideas. These two concepts can be linked by making assumptions about how ideas differ in terms of their novelty. While some ideas are relatively “close” to the current knowledge frontier, others are more advanced. Our empirical work depends on two concrete assumptions. First, we assume that ideas which are new-to-the-country but *not* new-to-the-world are drawn from the portion of the novelty distribution which is “closest” to the current domestic frontier. Second, we assume that the novelty of ideas follows a Pareto distribution; letting z note the relative novelty of an idea compared to the country’s knowledge baseline,

$$f(z_{j,t}) = \beta \left[\frac{\tau}{z_{j,t}} \right]^\beta z_{j,t}^{-1} \quad (2)$$

where τ is a constant which determines the lower-range of the ideas distribution and β determines the slope of the ideas distribution (see Figure 1). While stylized, the Pareto distribution captures an important feature of the distribution of novelty, namely that the density of novelty is decreasing in the distance from the current frontier.¹⁶ Under these assumptions, the flow of new-to-the-world ideas, $\dot{A}_{j,t}^*$, is the proportion of the total ideas flow which exceeds the world knowledge “threshold” A_{-j} , or

$$\begin{aligned} \dot{A}_{j,t}^* &= (1 - F(A_{-j,t})) \dot{A}_{j,t} = \left[\frac{\tau}{A_{-j,t}} \right]^\beta \delta H_A^\lambda A_j^\phi A_{-j}^\psi \\ &= \theta H_{A,t}^\lambda A_{j,t}^\phi A_{-j,t}^{\psi-\beta} \end{aligned} \quad (3)$$

Equation (3) suggests that the stock of ideas discovered in the world but not yet discovered in an individual country plays two distinct roles in the flow of new-to-the-world ideas. On the one hand, ideas productivity benefits from cross-border knowledge spillovers (ψ); on the other hand, there is a “raising the bar” effect determined by the number of ideas which are new-to-the-

¹⁶ Kortum (1997), Eaton and Kortum (1998), and Bental and Peled (1996) all use the Pareto distribution to characterize the R&D novelty distribution. As discussed in Kortum (1997), a key advantage of the Pareto distribution is that the distribution function for new inventions does not depend on the current state of knowledge.

country but *not* new-to-the-world.

In addition, we account for potential concavity in ideas production at the global level. In a given year, (3) describes the “pool” of new-to-the-world ideas produced by a country inclusive of those ideas which have been duplicated contemporaneously elsewhere in the world.

International patenting patterns, however, will be net of such duplication. In other words, the global ideas production function may be concave in the ideas flow of individual countries:

$$\dot{A}_{w,t} = \left[\sum_{j \in J} \dot{A}_{j,t}^* \right]^\rho \quad (4)$$

As a result, equal levels of international patenting in different years may reflect different levels of overall ideas production, since the degree of competition for patents may vary by year. Further, countries may differ in terms of their propensity-to-patent (perhaps because of legal, corporate or institutional norms) with the consequence that countries with the same flow of new-to-the-world ideas may be observed with different international patenting levels. Finally, patents granted in any year reflect inventive effort undertaken in prior years, and so we impose a lag of three years between observed patenting and the variables associated with national ideas production. Taken together, our empirical framework must therefore account for heterogeneity among countries, concavity in global ideas production, and the lag between invention and patent grants. To do so, we assume that beyond a country-specific propensity effect, the expected levels of international patents in year $t+3$ by a given country will be proportional to its contribution to the “pool” of innovations described in (4) in year t :

$$E \left(\frac{PATS_{j,t+3}}{PATS_{k,t+3}} \right) = \frac{\dot{A}_{j,t}^* \gamma_j}{\dot{A}_{k,t}^* \gamma_k} = \frac{\gamma_j}{\gamma_k} \frac{H_{A,j,t}^\lambda A_{j,t}^\phi A_{-j,t}^{\psi-\beta}}{H_{A,k,t}^\lambda A_{k,t}^\phi A_{-k,t}^{\psi-\beta}} \quad (5)$$

(5) suggests that, in order to evaluate the relationship between observed levels of international patenting and the national ideas production function, it will be crucial to account for both year and country effects. In other words, after accounting for each of the above issues, the structural relationship between international patenting and national ideas production is given by

$$E(PATS_{j,t+3}) = \gamma_j \theta_t H_{A,t}^\lambda A_{j,t}^\phi A_{-j,t}^{\psi-\beta} \quad (6)$$

Linking (6) to an empirical model requires specifying two additional features of the model. First, the realized level of patenting will differ from the predicted value given by (6); the

source of this disturbance is assumed to arise from an idiosyncratic country-specific technology shock each year unrelated to either the level of research effort or the measured stock of ideas.¹⁷ Second, (6) highlights the fact that there is a tradeoff between carefully controlling for the year-by-year competition for patents at the world level (the year effect, θ_t) and measuring the impact of international spillovers. Specifically, the coefficient on $A_{-j,t}$ is not separately identified from individual time effects in a regression specification (note that $A_j + A_{-j}$ is constant across countries in a given year). However, the realized level of patenting by all countries grows slowly and at a relatively constant rate across years; therefore, it is not unreasonable to assume that the variance due to concavity in the global ideas production function can be captured through the inclusion of an overall time trend rather than individual year effects. Under such an assumption, one can separately identify the impact of international knowledge spillovers, A_{-j} .

This discussion suggests two alternative regression specifications:

$$\ln \text{PATS}_{j,t+3} = \delta_t^Y + \gamma_j^C + \lambda \ln H_{A_{j,t}} + \phi \ln A_{j,t} + \eta_{j,t} \quad (7A)$$

$$\ln \text{PATS}_{j,t+3} = \delta_{\text{YEAR}} Y_t + \gamma_j^C + \lambda \ln H_{A_{j,t}} + \phi \ln A_{j,t} + (\psi - \beta) A_{-j,t} + \eta_{j,t} \quad (7B)$$

In (7A), we include specific year effects; without further assumptions, this specification yields consistent estimates for ϕ and λ . In (7b), we replace the individual year effects with the single time trend coefficient, allowing for separate identification of $(\psi - \beta)$. Consistency of OLS for (7B) requires us to additionally assume that the country-specific error is not autocorrelated and that the error is independent of both the country and world stock of ideas. In our empirical work, we explore both specifications (and variants), being careful to compare the results and highlight the differences resulting from different approaches.

V. Measuring the Domestic and “Rest of World” Ideas Stock

To estimate (7), we must measure both the domestic and international stock of ideas, A_j and A_{-j} , respectively. We do so by using a dataset composed of international patenting patterns

¹⁷ After we include a country and year fixed effect, the main sources of endogeneity and serial correlation should be accounted for. However, in the empirical work, we specifically check for various forms of autocorrelation and present several different specifications which allow us to evaluate the potential for certain sources of endogeneity. As well, all of the standard errors are adjusted for arbitrary heteroskedasticity.

by 17 OECD countries between 1973-1993. Two distinct alternatives are available. In one approach, we exploit the panel nature of the data and construct relative knowledge stocks by equating the knowledge stock with the relative contribution of a given country to the worldwide patent stock (thus we call this the direct approach since the knowledge stock is directly measured from international patenting patterns). In our second approach, we invert the aggregate production function by exploiting the functional relationship between the level of realized total factor productivity and the knowledge stock of a given country (thus we call this the inversion approach).

A. The Direct Approach

In the direct approach, the country-specific technological stock in year t is equal to a country-specific effect and the observed accumulation of international patents until $t-1$

($A_{j,t} = A_{j,0} \cdot \sum_{s=0}^{t+2} \text{PATS}_{j,s}$), resulting in the following OLS implementation of (7A):

$$\ln \text{PATS}_{j,t+3} = \delta_t^Y + \gamma_j^C + \lambda \ln H_{A_{j,t}} + \phi \ln \cdot \sum_{s=0}^{t+2} \text{PATS}_{j,s} + \eta_{j,t} \quad (8)$$

The underlying assumption of this model is that the specific knowledge stock upon which different countries draw when contributing to worldwide ideas is closely connected to their prior innovative output. In the current context, after controlling for both the year of observation (thus accounting for all “common” elements of the knowledge stock across different countries) and including a country-specific effect (accounting for the baseline knowledge stock prior to the beginning of the observed sample), this assumption may be a reasonable approximation of the degree to which countries differ in terms of the knowledge stock they can draw upon,

particularly in the medium-run. Moreover, this approach can be easily generalized to measure

the “rest-of-world” knowledge stock in (7B) ($A_{-j,t} = \sum_{k \neq j} \sum_{s=0}^{t+2} \text{PATS}_{k,s}$); as well, the empirical

work explores variants involving regional ideas stocks ($A_{R,t} = \sum_{k \in R, k \neq j} \sum_{s=0}^{t+2} \text{PATS}_{k,s}$), and models

which account for depreciation in the value of ideas over time ($A_{j,t} = A_{j,0} \cdot \sum_{s=0}^{t+2} e^{\phi(t-s+2)} \text{PATS}_{j,s}$).

B. The Inversion Approach

The inversion approach is a second, distinct methodology which exploits the functional relationship between the level of total factor productivity and the implied level of a country's

knowledge stock ($Y = A^\sigma K^\alpha H_L^{1-\alpha} \Rightarrow A = \left[\frac{Y}{K^\alpha H_L^{1-\alpha}} \right]^{\frac{1}{\sigma}}$). In other words, according to the Romer

model, there is an equivalence in the level of knowledge upon which the country draws upon in generating ideas and final output; we use this restriction to back out A_j from the variables which determine total factor productivity.¹⁸ Under such an approach, we are able to derive the following OLS implementation of (7A):

$$\ln \text{PATS}_{j,t+3} = \delta_t^Y + \gamma_j^C + \lambda \ln H_{A,j,t} + (\phi / \sigma) \ln \left[\frac{Y}{K^{.33} H_L^{.67}} \right] + \eta_{j,t} \quad (9)$$

It is important to note that, in contrast to the direct approach, (9) cannot separately identify ϕ and σ ; only their ratio is identified. However, this is of some independent interest. For example, if $\sigma < 1$, then the estimate provides an upper bound estimate for ϕ . In conjunction with an analysis of how different measures of A affect the level of productivity itself, we are also able to highlight the potential differences between ϕ and σ . Further, note that it is possible to unpack the elements of TFP (Y , K , and L) and separately analyze each of their contributions to international patenting productivity (we make a first pass at exploring these possibilities in the empirical work). Finally, note that we do not choose in the current analysis to impose the cross-equation restrictions implied by joint estimation of the productivity equation and the ideas equation together (as is done in an alternative model pursued in Eaton and Kortum (1998)). Instead, we focus on reporting the simplest relationships associated with international patenting, guided by the Romer model of economic growth which allows us to isolate the underlying production relationship.

IV. THE DATA

The data consists of a novel dataset of patenting activity and its determinants from 1973

¹⁸ Following many earlier authors, we impose $\hat{\alpha} = .33$ and so we do not compute a “first-stage” regression to determine the relative weights on labor and capital in determining final output.

through 1993 (Table 1 provides definition and sources, and Table 2 presents summary statistics). To allow us to estimate the ideas production function in several different ways, the data include (a) the flow of international patents for each country in each year, (b) measures of the factor inputs into ideas production, and (c) aggregate production function variables. We review the definition and summary statistics associated with each of the measures now in turn.

Patents. The principal dependent variable, $PATENTS_{j,t+3}$, is the number of U.S. patents granted to establishments in country j in year $(t+3)$. In other words, we assume that ideas production in a given year is reflected in the patents which are granted 3 years in the future.¹⁹ We employ U.S. patents to provide a comparable measure of innovations with substantial commercial importance across the 17 OECD countries in our sample. International establishments obtaining U.S. patent protection for their ideas incur a substantial cost (recent estimates hover at around \$100,000), and this cost is approximately equal for foreign inventors from different countries (Eaton and Kortum, 1996).²⁰

The use of patenting data to measure national ideas production requires some explanation. First, the use of $PATENTS$ in no way implies that we believe that international patents constitute the only relevant innovative outputs of national economies nor an ideal absolute measure of such outputs; instead, we are assuming that the international patents measure provides a useful *index* of overall innovative activity.²¹ Second, there are potential differences

¹⁹ We experimented with other lag structures and found robustness for our principal results for any lag between 2-4 years. An obvious (but time-consuming) remedy to the (strong) assumption about the lag structure is to gather patent application data from worldwide priority date (a measure recorded by the Derwent patent statistics) which would allow us to more specifically match resource expenditures and the timing of invention.

²⁰ The costs of filing a U.S. patent are lower for U.S. inventors perhaps by about half (and the incentives to patent may be higher); consequently, for most of our analysis, we exclude observations for U.S. inventors. However, we do show robustness to the inclusion of the U.S. inventor data and use U.S. inventor data for the purposes of comparison. When we do, we use a measure of U.S. patenting which is equal to the number of issued non-individual U.S. patents (i.e., patents issued to an inventor associated with a corporation, governmental body, educational or other non-profit institution).

²¹ As suggested in Eaton and Kortum (1996, 1998), the key assumption is that the value of innovations is distributed according to a fixed distribution across economies and a constant *fraction* of innovative output turns out to be valuable enough to justify an international patent. To the extent that this fractional value varies across countries or across time, the framework we developed in Section III derives how to overcome these issues through the use of

among countries in terms of their industrial composition and their underlying “propensity to patent.” Such heterogeneity introduces variance into the regression and potential mismeasurement of the *stock* of knowledge.²² Finally, across the sample, there has been an increasing number of international patents granted over the last quarter century. Beyond reflecting true increases in new-to-the-world ideas production, this increase may reflect increasing levels of international trade, the increasing strength of intellectual property protection, or cyclical waves in the realization of important ideas (the so-called “recharge” phenomena). However, the use of year and trend effects provides a control for these common shifters.

With these caveats in mind, the average number of international patents produced by a country in a given year in our sample is 2022 with a standard deviation of 3876. As can be seen in Figure 1, there has been a temporal increase in the number of patents issued to foreign inventors over the past 20 years, reaching a peak in the last year of the sample (which corresponds to patents issued in 1996). Figure 2 shows “per-capita” patenting rates (i.e., PATENTS / POPULATION) for different countries over the sample. Three facts are of immediate interest. First, there is substantial heterogeneity among countries in terms of their realized patenting intensity. Second, at the beginning of the sample, the only country with a per-capita patenting rate similar to the United States’ own rate is Switzerland.²³ Third, while the most dramatic increase in per capita patenting is associated with Japan, *convergence* towards the U.S. level is experienced by a substantial number of OECD countries (though countries such as Italy and the U.K. do not share in this phenomena). The empirical work explains this heterogeneity in terms of more structural factors such as the resources devoted to R&D and the stock of ideas of individual countries.

Factor Inputs for the Ideas Production Function. We estimate the determinants of PATENTS using both the direct and inversion approaches. In both approaches, we estimate the sensitivity of ideas production to human capital resources devoted to the ideas-producing sector

fixed year and country effect in the empirical model.

²² In terms of differences in patenting propensity, the most likely candidate for a large difference is Japan. We have confirmed that our results are robust to the exclusion of Japan.

²³ Recall that the U.S. patenting level is determined by the number of patents issued to U.S. inventors associated with an institution such as a company, governmental body, or university.

(i.e., the λ parameter in the Jones' variant model). We measure this input with the number of full-time equivalent scientists and engineers working in R&D in an economy (FTE S&E),²⁴ FTE S&E averages over 136,000 in this sample (compare this to a 1996 figure for the United States of over 3 million), and can be seen to exhibit an upward trend for most countries in the sample over the twenty-year period under study (Figure 3). As can be seen in Figure 4, there exists substantial heterogeneity among countries and over time in terms of underlying international patenting productivity (PATENTS / FTE S&E), ranging from over .05 patents per R&D worker for Switzerland during the 1970s to less than .001 patents per R&D workers for Spain. However, there is a strong correlation between PATENT and FTE S&E ($\rho = .94$).

In the direct approach, we exploit the panel aspect of our data and measure the stock of knowledge as each individual country's cumulative patent stock (PATENT STOCK) which is simply the sum over patents granted for all prior years since 1970 (Mean = 18704). Controlling for country fixed effects (accounting for each country's baseline stock of knowledge) and annual fixed effects (accounting for the diffusion of knowledge across international borders), changes in PATENT STOCK reflect the accretion of country-specific knowledge. It is useful to note that while the level of PATENT STOCK is increasing by construction for all countries, there is still a close correlation between PATENT STOCK and PATENT ($\rho = .91$).

We extend the ideas production function to separate out the effect of the within-country patent stock from international spillovers effects. WORLD PATENT STOCK is the sum of PATENT STOCK over all countries in each year. However, as mentioned earlier, the relevant concept from the view of country j is that part of the STOCK which has been invented elsewhere. As such, our dataset includes the "REST OF WORLD" PATENT STOCK (WORLD PATENT STOCK – PATENT STOCK).²⁵ Further, we also calculate more nuanced measures

²⁴ Of course, just as PATENT is only a noisy measure of innovative output, FTE S&E represents only an index of the number of workers in an economy devoted to ideas production. This measure tends to exclude technical investment by small firms (venture start-ups, etc.) as well as technical problem-solving which occurs on the "shop room" floor, yielding key process and product improvements for individual firms as well as entire industries (REF). In related work, we confront these issues to some extent by breaking down this measure into private, public, and university sector employment, by using expenditure-based rather than employment-based measures, and including alternative measures of the size of the ideas sector (Stern, Porter and Furman, 1999; Porter and Stern, 1999).

²⁵ Note that the WORLD PATENT STOCK includes patents granted to U.S. establishment (though we explore

which might capture the impact of international spillovers, including the REGIONAL PATENT STOCK (where only countries in close geographic proximity are included) and the SHARED LANGUAGE PATENT STOCK (dividing the relevant “pools” into English-speaking, German-speaking, and other groups). All of these measures are used to evaluate the role of the absolute versus relative patent stock position on R&D productivity. Finally, following the suggestion of Coe and Helpman (1995), among others, we explore whether the salience of the trading relationship with the U.S. has an effect on the spillovers associated with international patenting productivity by including both the level of imports to and exports from the United States.

“Output” Production Function Variables. To proxy the stock of knowledge with the determinants of total factor productivity in the inversion approach, we use aggregate output and factor inputs in our analysis. These measures include GDP (measured in PPP-adjusted 1985 US\$), POPULATION, the labor force employed in the non-ideas producing sector (LABOR), and the non-residential fixed capital stock as calculated from the Penn World Tables (CAPITAL).²⁶

Year and Country Dummies and Controls. We introduce year and country controls to account for potential concavity in global ideas production, differences in patenting propensity, and to control for potential differences in industrial composition. Indeed, it is important that we establish the robustness of our results to the inclusion (or exclusion) of these different controls to evaluate the potential importance of such heterogeneity. The empirical analysis includes both fixed effects (Year and/or Country dummies (Y_t and C_j , respectively) as well as coarser controls, including a YEAR trend variable (ranging from 0 to 20) and a “baseline” per capita productivity measure, real 1967 GDP (GDP67).

VI. EMPIRICAL RESULTS: THE IDEAS PRODUCTION FUNCTION

The empirical argument proceeds in several distinct steps, so it is useful to preview our overall approach and findings. We first review the estimates resulting from the direct approach (Tables 3-6). This analysis begins with a baseline model which includes year and country-

variants where we calculate this concept excluding the U.S. data).

²⁶ Except for the adjustment to LABOR to account for FTE S&E, these measures are traditional aggregate productivity variables and have been thoroughly examined in many prior studies (Mankiw, Romer and Weil (1991)).

specific effects, following (7A); our main finding is that there exists a strong, proportional relationship between PATENTS and PATENT STOCK. In Tables 4 & 5, we expand the analysis to include international knowledge pools (following (7B)); our main result is that while the restriction on the domestic stock is robust, the net impact of international knowledge is negative in terms of the production of new-to-the-world ideas. Table 6 explores the possibility of serial correlation and potential misspecification and endogeneity of the patent stock itself. In each of these robustness exercises, we reconfirm our general finding in favor of the Romer assumption ($\hat{\phi} \approx 1$) and the negative net impact of international spillovers.

We then turn to the Inversion approach (Table 7), where we find a similarly strong relationship between ideas sector productivity and our estimate of the knowledge stock. Finally, in Table 8, we briefly explore how our measures of the knowledge stock contribute to the level of TFP, providing evidence about the sensitivity of economywide productivity to the flow of knowledge from the ideas sector of the economy. We find a small but significant impact, suggesting that ideas-driven growth may be a feasible but empirically modest effect.

A. Evidence from the Direct Approach

Table 3 presents several models closely connected to the single-country models of Romer (1990) and Jones (1995), namely, a simple log-log regression of the flow of patents on the amount of labor employed in the ideas-producing sector, a measure of the stock of knowledge, and controls for year and country. The results are striking. The baseline specification (3-1) includes fixed effects both for each country in the sample as well as for each year.²⁷ While the returns to research effort are declining in its level ($\hat{\lambda} = .48$), the estimated sensitivity to increases in the patent stock is greater than 1 ($\hat{\phi} = 1.19$ in the baseline). Indeed, we can reject $H_0: \phi < 1$.²⁸ These parameters cannot reject the presence of an important precondition for ideas-driven growth, namely, the existence of a positive, proportional feedback between the flow and stock of ideas production. In model (3-3), we go to the other extreme in terms of controlling for heterogeneity across years and countries by including only a baseline

²⁷ See Section IV for further discussion of the development of this specification.

²⁸ Similar results are found in (3-2), which includes the data for patents granted to U.S. establishments.

year dummy (Y_{73}), a YEAR trend and a baseline country proxy, L GDP67. While there is a substantial decline in the estimated sensitivity to PATENT STOCK, the parameter is still greater than .8 and precisely estimated ($\hat{\lambda}$ declines (but not significantly) to .39). The reduction in the estimate of ϕ depends on the inclusion of country-specific effects. In other words, evidence in favor of the Romer hypothesis is even stronger when one only relies on the time-series dimension in this data. Rather than simply being the result of differences across countries in terms of their “propensity-to-patent,” international patenting productivity is driven by changes over time in the stock of patents accumulated within individual countries. A further striking result is that, even with only five regressors, nearly all of the variance in PATENTS is explained ($R^2 = .98$).

In (3-3), we see that there is a strong downward evolution of patenting productivity (estimated at the rate of 8% per year). Along with the strong positive impact of the PATENT STOCK (which increases over time), this result suggests that time-series variation in R&D productivity is a “tug-of-war” between the positive impact of each country’s PATENT STOCK and an overall negative effect resulting from the passage of time. Table 4 unpacks this dilemma more precisely and, in so doing, highlights the role of international knowledge spillovers. In (4-1), we exclude all controls for year (excluding both the year effects and the time trend) and find no dependence between PATENT and PATENT STOCK.²⁹ This negative result confirms that the productivity of a given stock of ideas in producing new-to-the-world ideas depends critically on the height of the global innovation “bar.” This intuition is strengthened in the final two columns of Table 4 which follow the specification suggested in (7B) and include the REST OF WORLD PATENT STOCK in the regression. In both specifications, $\hat{\phi} > .9$ and, perhaps as importantly, our estimate of $[\psi - \beta] < -1$. This latter estimate suggests that either (a) the international knowledge spillovers have either a modest positive or even negative effect and/or (b) there exists a substantial “raising the bar” effect by which increases in the number of ideas discovered worldwide substantially decreases the ability of a country to produce new-to-the-world technologies. Indeed, after including REST OF WORLD PATENT STOCK, the measured effect of the YEAR trend is in fact positive and the level of international patenting is

²⁹ As well, there is a substantial increase in the size of the coefficient on FT S&E.

substantially concave in the size of the R&D workforce.

We extend this analysis of spillovers in Table 5. In each of the first two columns, we group all of the countries according to a regional or language affiliation (Table 1 lists these groupings). For each country, we calculate the PATENT STOCK for the remaining countries in the group. We then examine the sensitivity of the flow of international patents to each of these measures using (4-3) as a baseline regression. There is no measured dependence between the REGIONAL or SHARED LANGUAGE PATENT STOCK and the flow of international patents from any given country. However, even including each of these measures, the sensitivity to the domestic PATENT STOCK remains high (estimated at .85 and .94 respectively). Finally, in (5-3), we examine the potential spillovers between each country and the United States (as the United States is the destination patenting country). While our core results remain similar ($\hat{\phi} = .89$), we additionally find that countries with high levels of exports to the United States tend to register a higher flow of international patents and that countries who import heavily from the United States tend to have lower rates of international patenting. This result provides support for prior research which has highlighted international trade as the key mechanism through which international knowledge spillovers occur (Coe and Helpman, 1995; Eaton and Kortum, 1996; Keller, 1997).

The results from Tables 4 and 5 potentially illuminate an important difference between domestic and international spillovers. Specifically, accumulation of knowledge at the world level may not have the same effect on the marginal productivity of R&D as knowledge accumulation at the national level. According to these regressions, higher rates of ideas production in other countries imposes a net *negative* externality on a given country's ideas sector in terms of the production of new-to-the-world technologies. As other countries extend the frontier of realized technologies, it becomes more costly for any given country to further extend that frontier (the "raising the bar" effect). Moreover, this raising the bar effect is not swamped by an international "standing on shoulders" effect. Indeed, even when controlling for this overall effect, there are only modest measured "regional" or "language-based" spillovers. Despite the low levels of these international spillovers, individual countries seem to be able to benefit from their own form of standing on shoulders, as suggested by the consistently high parameters associated with ϕ . Another way of interpreting these results is that as other countries make

new discoveries, individual countries must spend an increasing share of their R&D resources on absorbing and imitating these ideas. The consequence of substitution towards imitative activity is that the production of new-to-the-world ideas by the R&D sector is declining in the number of ideas produced externally.

The analysis so far has focused on characterizing, in a simple way, the drivers of ideas sector productivity as reflected in the production of international patents. As discussed in Section IV, the inclusion of fixed year and country effects (or other appropriate controls) yields a specification which is consistent under OLS (though we do account for heteroskedasticity when computing the standard errors). Of course, it is possible that additional serial correlation (above and beyond that which is controlled for by the fixed effects) is spuriously driving the above results. Consequently, Table 6A specifically presents our results correcting for the simplest form of serial correlation (AR(1)).³⁰ The results are remarkably similar to our earlier estimates; while there is a slight increase in the variance, all of the coefficient are essentially the same or shift in a direction which provides further support for our earlier results.

As well, in Table 6B, we attempt to address several issues which arise in the construction of the PATENT STOCK itself. First, we check whether our results are sensitive to spurious patterns based on a few years of data early in the sample. In (6B-1), we show that the results are essentially unchanged even if observations are drawn only from years after 1980 (of course, the PATENT STOCK variable is still calculated from the 1970 data onwards). In (6B-2) and (6B-3), we extend the analysis to account for the depreciation in PATENT STOCK. In (6B-2), we impose a 20% depreciation rate on the PATENT STOCK. Perhaps surprisingly, the results

³⁰ We only present results for the AR(1) correction; however, we experimented extensively with alternative error structures (specifically higher-order AR) and found no change in our results. Further, we examined whether the PATENT STOCK and the year and country fixed effects are cointegrated. While we cannot reject panel cointegration for our sample (Pedroni, 1999), preliminary analysis using *existing* panel cointegration methods did not substantially impact our results (the standard errors increased modestly but in no way impacted our hypothesis tests). However, because PATENT STOCK is calculated as the *summation* of prior realizations of the dependent variable (and so is not simply the lagged immediate value of the dependent variable), existing panel cointegration techniques (such as Kao and Chiang (1997) and Pedroni (1996)) do not provide the precise correction necessary for our application. As such, we report OLS (and AR1) standard errors, leaving the derivation of the appropriate cointegration adjustment for an “increasing returns” heterogeneous panel data model for future research.

continue to be robust, though the absolute value of the coefficient on REST OF WORLD PATENT STOCK substantially decreases. In (6B-3), we specifically compare the sensitivity to newer and older patenting by specifically including lagged PATENTS. To guard against pure autocorrelation, we instrument lagged PATENTS with the lagged PATENT STOCK and FTE S&E. We find that both PATENT STOCK and PATENTS(T-1) are significant, of roughly equal size, and the sum of these coefficients is over .9, providing additional support for our claim that the estimated parameter on intertemporal spillovers is approximately unity.

Taken together, this evidence points to a novel and intriguing possibility, namely, that while there is strong evidence for positive *local* knowledge spillovers, advances in the international knowledge frontier may *reduce* the flow of frontier ideas by a given country. It is difficult for technological laggards to “leapfrog” over the technological stock of other nations and produce at the world frontier.³¹ For all countries, a substantial amount of technical effort is devoted towards imitative or absorptive activity rather than the production of new-to-the-world technologies.

B. Evidence from the Inversion Approach

We now explore a second avenue for evidence through the inversion approach. As mentioned earlier, exploiting inversion requires making a (relatively) strong assumption about the dependency between current output and patent flows; the errors in the output equation cannot be correlated with the error in the patent equation. While the validity of this assumption is debatable (though including country and year-specific effects may make this assumption more plausible), the inversion method does allow for a simple evaluation of the relationship between productivity in the output and the ideas sectors. The first model presented in Table 7 employs a population-based specification for aggregate output ($GDP_{j,t} = A_{j,t}^{\sigma} POP_{j,t}^{\beta}$),³² yielding the simple empirical specification,

$$L PATENTS_{j,t} = Y_t + C_j + \lambda L FTE S\&E_{j,t} + \phi / \sigma L GDP_{j,t} - \beta \phi / \sigma L POP_{j,t} \quad (10)$$

³¹ We recognize that such effects are predicted by certain growth models but not others (compare the quality ladder model of Grossman and Helpman (1991) which explicitly includes a raising the bar effect) with Rivera-Batiz and Romer (1991) which essentially assumes that integration involves strong positive R&D productivity spillovers.

³² While it will be separately estimated, one could imagine setting $\beta=1$.

As noted earlier, we cannot separately identify the structural parameters ϕ or σ but only their ratio (i.e., the returns of ideas-to-ideas versus the return of ideas-to-output). In (7-1), we examine this simple model, and find that both λ and ϕ/σ are estimated to be relatively high. Indeed, one cannot reject $H_0: \phi / \sigma < 1$. In (7-2), we examine this relationship in a more structural fashion in (7-2), replacing the population-based measure with the level of TFP (setting $\hat{\alpha} = .33$). While λ remains just above .5, the TFP parameter is estimated to be quite low ($\phi/\sigma = .13$). However, we are able to demonstrate (in (7-3)) that this finding results from the separate positive contribution of the level of the CAPITAL STOCK to ideas production at a given point in time (recall that we impose country-specific dummies). While the role of capital in the productivity of ideas production is not discussed in the theoretical literature (at least to our knowledge), comparing (7-2) and (7-3) suggests its importance: R&D productivity at the country level is positively (and separately) related to both TFP and the economywide capital stock. Indeed, this should not be so surprising: the ideas sector is in fact heavily capital-intensive and relies on the availability of advanced capital equipment and scientific instruments (Rosenberg, 1994).

Together, both the inversion and direct approaches provide consistent evidence. First, the domestic ideas stock is a key determinant of a country's production of new-to-the-world technologies. Moreover, this positive domestic spillover seems relatively independent of the impact of knowledge from abroad. The evidence suggests instead that international knowledge has a net negative impact on the flow of new-to-the-world technologies. Moreover, this effect is strongest for "distant" international knowledge. This suggests that rather than assuming that all countries somehow share a common knowledge stock (or that more advanced countries' knowledge pools are strict supersets of poorer countries knowledge stocks), countries may have very heterogeneous and only partially matched knowledge stocks. The specific industrial and technological focus of individual countries (or regions) yields domestic knowledge pools which are substantially differentiated from the rest of the world (Porter, 1990; Nelson, 1993; Stern, Furman, and Porter, 1999).

VII. EMPIRICAL RESULTS: IDEAS PRODUCTION & AGGREGATE OUTPUT/GROWTH

Finally, in Table 8, we attempt to reconcile the evidence for Romer's model for ideas with the poor predictive power of that model in explaining aggregate productivity growth (this evidence has been reviewed by, among other, Jones (1995)). To do so, we explore how the level of GDP or TFP is affected by the PATENT STOCK measure. Our main conclusion from Table 8 is that there seems to be a small (but significant) effect of the PATENT STOCK on the level of TFP. Our analysis begins in (8-1) with a regression which includes both year and country effects; just relying on the relative growth of the PATENT STOCK, we find that TFP is predicted to grow only 22% in response to a doubling of the PATENT STOCK. In the remaining regressions, an even weaker link is found. In (8-2) and (8-3), we explore the productivity boost from the PATENT STOCK given the level of LABOR and CAPITAL. While the coefficient on PATENT STOCK remains significant, its size is substantially reduced; in (8-3), a doubling of the patent stock is associated with only a 5% increase in the level of TFP. Finally, we are not able to separately identify the impact of the domestic and international patent stock on TFP (the results are noisy though consistent with our earlier findings).

Together with the previous section, these results provide some guidance about the feasibility and significance of ideas-driven growth. While the critical production parameter for ideas-driven growth is the proportional feedback between current and historical ideas production ($\phi = 1$), the impact on productivity growth is a function of all the parameters of the model

$(\frac{\dot{Y}}{Y} = \sigma \delta H_A^\lambda)$. Our results suggest that ideas-driven growth may be feasible; however, the size of such effects may be modest. Several effects intervene between ideas production and realized productivity growth: the concavity in research effort and the weak sensitivity of productivity to ideas production.

VIII. CONCLUSIONS

This paper provides evidence for three principal findings. First, as we have emphasized repeatedly, the evidence on international patenting provides support for the strong positive

feedbacks in ideas production at the national level proposed by Romer. However, this support for the Romer model is balanced by our additional findings which shed light on the difficulties in using an ideas-driven growth model to explain realized productivity growth. Specifically, the local nature of knowledge spillovers (at least in the medium run) along with the inability to translate new-to-the-world ideas production into TFP weakens the *measured* empirical significance of the Romer model.

One of the most surprising findings from our analysis is the sharp disjunction between the strong positive relationship between ideas productivity and the national ideas stock and the strong negative relationship with foreign knowledge sources. This finding provides support for one of two complementary interpretations: either foreign patenting substantially “raises the bar” for future international patenting or international knowledge spillovers for *ideas production* are relatively weak. Both of these interpretations are consistent with recent work, at the industry or national level, which has emphasized the important role played by the national (or even regional) environment in understanding the dynamics of the innovative process (Porter, 1990; Nelson, 1993). While our results do not contradict earlier evidence that the aggregate economy benefits from the absorption and imitation of foreign technologies (Grossman and Helpman, 1991; Barro and Sala-i-Martin, 1992; Coe and Helpman, 1995; Park, 1995), the evidence highlights that the capacity for new-to-the-world innovation seems to be a country-specific asset.

It is possible that more far-reaching factors are at work in explaining the weak relationship between ideas sector productivity and TFP. The results highlight that realizing the full social benefits from new technologies depends on the diffusion of these new technologies into the productive sector of the economy and the ability to the productive sector to take best advantage of technical advance. However, there is good reason to believe that diffusion may be both slow and incomplete, even for technologies which are capable of being patented and presumably have at least some commercial value. In most OECD countries, the scientific and engineering communities represent a relatively small part of the overall economy and the linkages among members of technical community tend to be strong relative to linkages between this community and the productive side of the economy. As a result, it should not be very surprising if the productive sector of the economy does not exploit the full range of innovations being developed by the ideas sector. Indeed, the incentives provided to scientists and engineers

may in fact explicitly discourage the development of linkages with the full range of downstream commercial application sectors (Stern,1999).

Moreover, the full translation of ideas and new technologies into productive output depends on many intervening factors including the technical and human capital of the workforce, whether the economy's physical and information infrastructure are positioned to take advantage of new technologies, and whether the industrial organization of adopting sectors is conducive to taking advantage of technology development (e.g., consumers who are willing to experiment with new technologies, the existence of micro-market incentives to commercialize new technologies, and supply relationships which complement and reinforce the social value of novel technologies (Porter, 1990)). While most earlier work has simply assumed the existence of a downstream production sector which is both perfectly competitive and whose physical and human capital is compatible with new technologies, it may be the case that the substantial investments and adjustments which are involved in applying new technologies substantially undercuts the ability of the ideas sector and new technologies to contribute directly to TFP growth in the medium-term (Rosenberg, 1976; David, 1990). In related work, we specifically explore the relationship between the microeconomic foundations of R&D investment and the national capacity for innovation more fully (Stern, Porter, and Furman (1999)). In the spirit of this work, identifying the specific mechanisms by which linkages among the ideas and production sectors of the economy lead to higher rates of diffusion and what factors limit the applicability of new technologies to generating productivity gains seems an extremely promising area for further theoretical and empirical research.

REFERENCES

- Adams, J. (1990). "Fundamental Stocks of Knowledge and Productivity Growth," *Journal of Political Economy* 98: 673-702.
- Ambramowitz, M. (1986). "Catching Up, Forging Ahead and Falling Behind," *Journal of Economic History* 46: 385-406..
- Barro, R. (1991). "Economic Growth in a Cross-Section of Countries," *Quarterly Journal of Economics* 106: 407-43.
- Barro, R. and X. Sala-i-Martin (1992). "Convergence" *Journal of Political Economy* 100:223-51.
- Bental, B. and D. Peled (1996). "The Accumulation of Wealth and the Cyclical Generation of New Technologies: A Search Theoretic Approach," *International Economic Review* 37(3): 687-718.
- Coe, D. and E. Helpman (1995). "International R&D Spillovers," *European Economic Review* 39(859-887).
- Dasgupta, P. and P. David (1994). "Towards a New Economics of Science," *Research Policy*, 23: 487-521.
- David, P. (1990). "The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox," *American Economic Review* 80(2): 355-61.
- Dosi, G., K. Pavitt and L. Soete (1990). *The Economics of Technical Change and International Trade* New York (NY): Columbia University Press.
- Eaton, J. and S. Kortum (1996). "Trade in Ideas: Patenting & Productivity in the OECD," *Journal of International Economics* 40(3-4):251-278.
- Eaton, J. and S. Kortum (1998). "International Technology Diffusion: Theory and Measurement," *International Economic Review*, forthcoming.
- Griliches, Z., ed (1984). *R&D, Patents and Productivity*. Chicago (IL): Chicago University Press.
- Griliches, Z. (1990). "Patent Statistics as Economic Indicators: A Survey," *Journal of Economic Literature*, 92:630-53.
- Grossman, G. and E. Helpman (1991). *Innovation and Growth in the World Economy*. Cambridge (MA): Harvard University Press.

- Jaffe, A. (1986). "Technological Opportunity and Spillovers from R&D: Evidence from Firms Patents, Profits and Market Value," *American Economic Review*.
- Jaffe, A. and M. Trajtenberg (1998). "International Knowledge Flows: Evidence from Patent Citations," *NBER Working Paper 6507*.
- Jones, C. (1995). "R&D Based Models of Economic Growth," *Journal of Political Economy*, 103: 739-784.
- Jones, C. (1998). *Introduction to Economic Growth*. New York: W.W. Norton & Company.
- Jones, C. (1999). "Sources of U.S. Economic Growth in a World of Ideas," mimeo, Stanford University.
- Kao, C. and M.H. Chiang (1997). "On the Estimation and Inference of a Cointegrated Regression in Panel Data," mimeo, Syracuse University.
- Keller, W. (1997). "Trade and the Transmission of Technology," NBER Working Paper #6113.
- Kortum, S. (1997). "Research, Patenting, and Technological Change," *Econometrica*, 65(6): 1389-1419.
- Mankiw, N.G., D. Romer, and D.N. Weil (1992). "A Contribution to the Empirics of Economic Growth," *Quarterly Journal of Economics*, 107:407-437.
- Nelson, R., ed. (1993). *National Innovation Systems: A Comparative Analysis* New York (NY): Oxford University Press.
- Park, W. (1995). "International R&D Spillovers and OECD Economic Growth," *Economic Inquiry* 33:571-591.
- Pedroni (1996). "Fully Modified OLS for Heterogenous Cointegrated Panels and the Case of Purchasing Power Parity," mimeo, Indiana University.
- Pedroni (1999). "Critical Values for Cointegration Tests in Heterogeneous Panels with Multiple Regressors," mimeo, Indiana University.
- Porter, M.E. (1990). *The Competitive Advantage of Nations* New York (NY): Free Press.
- Porter, M.E., and S. Stern (1999). *The New Challenge to America's Prosperity: Findings from the Innovation Index*. Washington (DC): Council on Competitiveness.
- Rivera-Batiz, L., and P. Romer (1991). "Economic Integration and Endogenous Growth," *Quarterly Journal of Economics* 106(2): 531-555.

- Romer, P. (1990). "Endogenous Technological Change," *Journal of Political Economy*, 98: S71-S102.
- Romer, P. (1994). "New Goods, Old Theory, and the Welfare Costs of Trade Restrictions," *Journal of Development Economics* 43:5-38.
- Romer, P. (1996). "Why, Indeed, in America? Theory, History, and the Origins of Modern Economic Growth," *American Economic Review* 86(2): 202-206.
- Rosenberg, N. (1976). *Perspectives on Technology*. New York: Cambridge University Press.
- Rosenberg, N. (1994). *Exploring the Black Box: Technology, Economics, and History*. New York: Cambridge University Press.
- Schmookler, J. (1966). *Invention and Economic Growth*. Cambridge (MA): Harvard University Press.
- Solow, R.M. (1956). "A Contribution to the Theory of Economic Growth," *Quarterly Journal of Economics*, 70:65-94.
- Solow, R.M. (1957). "Technical Change and the Aggregate Production Function," *Review of Economics and Statistics*, 39:312-20.
- Stern, S., M.E. Porter, and J.L. Furman (1999). "The Determinants of National Innovative Capacity," HBS Working Paper, October 1999.
- Trajtenberg, M. (1990). "Patents as Indicators of Innovation," *Economic Analysis of Product Innovation* Cambridge (MA): Harvard University Press.

Primary Data Sources

OECD, Main Science and Technology Indicators, Paris, France. 1998 vols 1 &2, 1997 vols 1 & 2, 1996 vol 2, 1995 vol 2, 1993 vol 2, 1992 vol 2, 1990 vol 2, 1989 vol 2, 1988 (1982-1988) vol 2, 1987 (1981-1987)

National Science Board, Science Indicators, Washington, DC, 1974, 1976, 1978, 1980, 1982, 1985

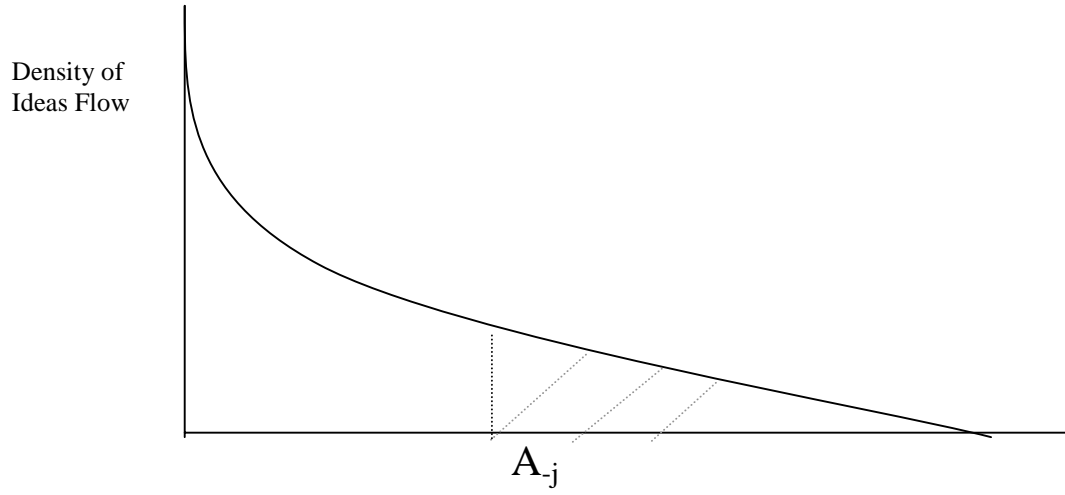
National Science Board, Science & Engineering Indicators, Washington, DC, 1987, 1991, 1993, 1996, 1998

International Monetary Fund, Balance of Payments Statistics Yearbook, 1979-1996

United Nations, National Accounts Statistics: Main Aggregates and Detailed Tables, 1993-Part I.

National Bureau of Economic Research, Penn World Tables, www.nber.org

FIGURE 1



N.B. Dashed area is the portion of national ideas production which is also "new-to-the-world"

FIGURE 2

Patenting Per Capita

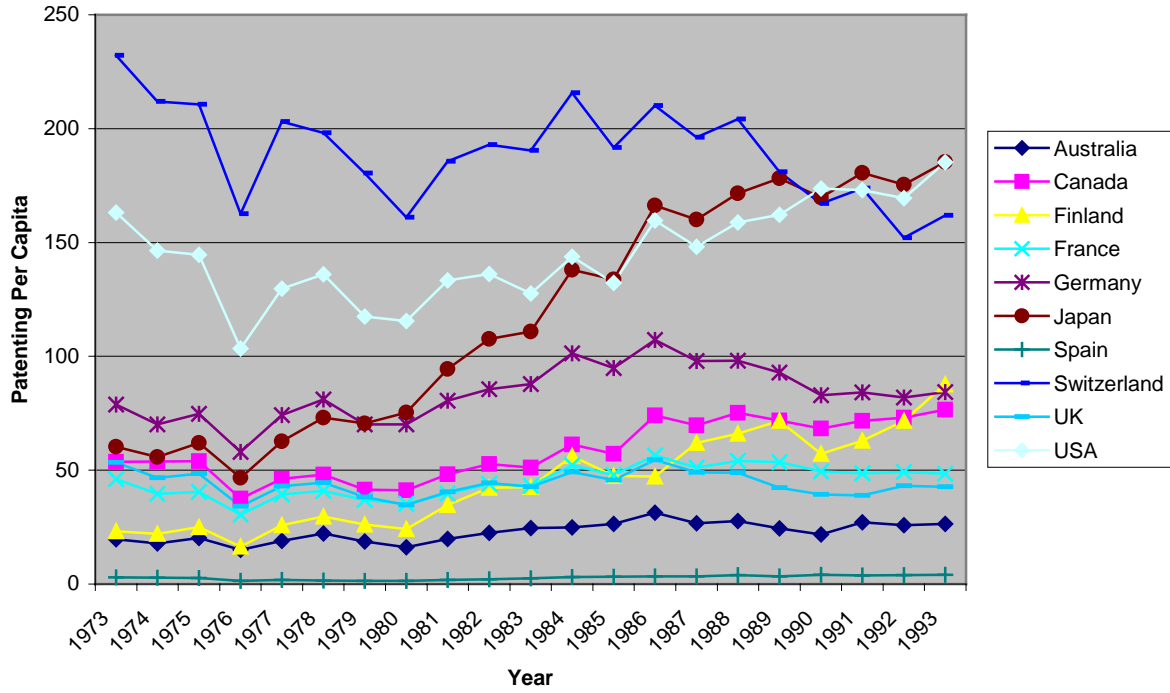


Figure 3

FTE S&E by Country

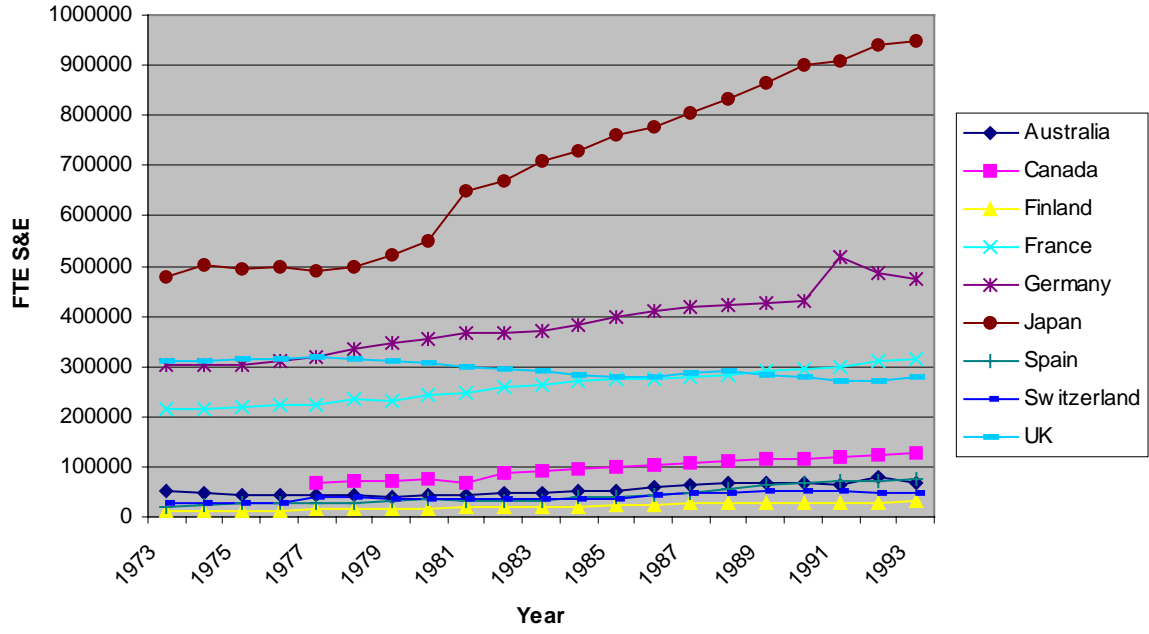


Figure 4

R&D Productivity by Country

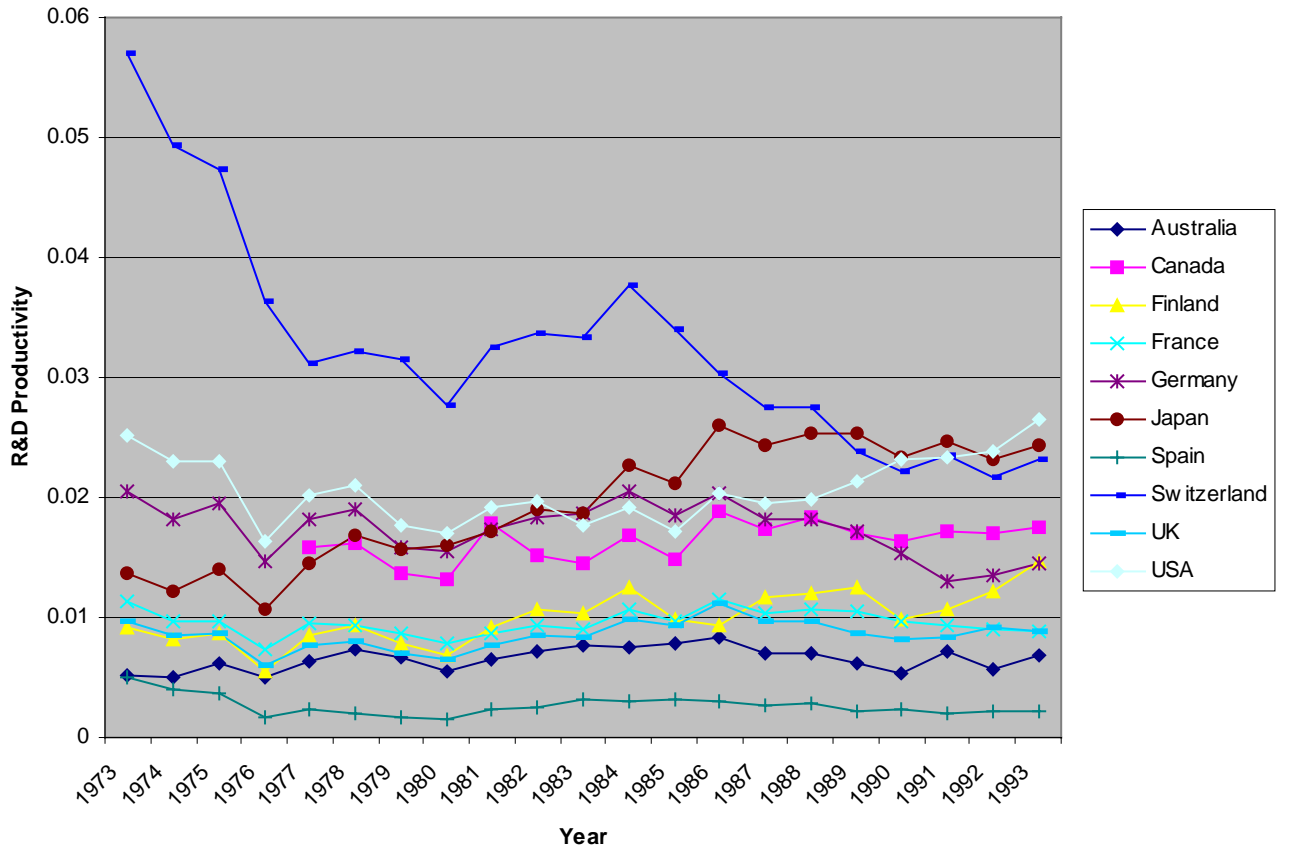


TABLE 1
VARIABLES* & DEFINITIONS

VARIABLE	DEFINITION	SOURCE
INTERNATIONAL PATENTING		
PATENTS _{j,t}	Patents granted in the US to establishments in country y in year (t+3)	CHI US patent database
IDEAS PRODUCTION FUNCTION INPUTS		
FT S&E _{j,t}	Full Time Equivalent Scientists & Engineers in all sectors in country j in year t	OECD Main Science & Technology Indicators
PATENT STOCK _{j,t}	$\sum_{s=1973}^{t+2} \text{PATENTS}_{j,s}$	
SPILLOVER MEASURES		
WORLD PATENT STOCK _t	$\sum_k \text{PATENT STOCK}_{k,t}$	
“REST OF WORLD” PATENT STOCK	WORLD PATENT STOCK _t - PATENT STOCK	
“REGIONAL” PATENT STOCK	Analogous to “REST OF WORLD” PATENT STOCK summed over regions: * Australia, New Zealand * Austria, France, Germany, Italy, Netherlands, Spain, Switzerland * Denmark, Finland, Norway, Sweden * Canada, United States * Japan and UK are ungrouped	
“SHARED LANGUAGE” PATENT STOCK	Analogous to “REST OF WORLD” PATENT STOCK over common language base: * Australia, Canada, New Zealand, UK, United States * Austria, Germany, Switzerland * All other countries are ungrouped	
IMPORTS _{j,t}	Real PPP-Adjusted Imports from country j to US in year t in 1985 US B \$	OECD
EXPORTS _{j,t}	Real PPP-Adjusted Exports from US to country j in year t in 1985 US B \$	OECD
AGGREGATE PRODUCTION FUNCTION VARIABLES		
GDP _{j,t}	Real PPP-Adjusted GDP in 1985 US B \$ in country j in year t	Penn World Tables
POPULATION _{i,t}	Population in millions in country j in year t	Penn World Tables
LABOR _{j,t}	FULL-TIME EMPLOYED LABOR FORCE - FTE S&E	Penn World Tables;
CAPITAL _{i,t}	Non-Residential Capital Stock in 1985 B US\$	Penn World Tables
TFP _{j,t}	$\frac{\text{GDP}_{j,t}}{(\text{LABOR}_{j,t}^{.67} * \text{CAPITAL}_{j,t}^{.33})}$	
TIME & COUNTRY VARIABLES		
YEAR (1972 = Year 0)	(Year - 1972)	

* The natural logarithm of a variable, X, will be denoted L X.

TABLE 2
MEANS & STANDARD DEVIATIONS

OBSERVATIONS = 330	MEAN	STANDARD DEVIATION
“INTERNATIONAL” PATENTING		
PATENTS	2022.24	3876.35
PATENTS PER CAPITA	55.45	48.23
IDEAS PRODUCTION FUNCTION INPUTS		
FT S&E	136715.30	188095.30
PATENT STOCK	18704.41	37909.60
PATENT STOCK PER CAPITA	553.37	643.08
SPILLOVER MEASURES		
WORLD PATENT STOCK	646282.00	392388.00
REGIONAL PATENT STOCK	91516.47	116833.00
SHARED LANGUAGE PATENT STOCK	124790.60	197195.80
IMPORTS	11313.04	21444.67
EXPORTS	8477.76	14405.80
AGGREGATE PRODUCTION FUNCTION VARIABLES		
GDP	538.02	608.36
POPULATION	31.12	32.56
LABOR	12.80	14.30
CAPITAL	388.00	445.00
TFP	1.42	0.22

* Summary statistics exclude US observations.

TABLE 3
IDEAS PRODUCTION FUNCTION:
DIRECT APPROACH

	Dependent Variable = Ln(Patents)		
	(3-1)	(3-2)	(3-3)
	Baseline with Fixed Effects	Baseline (Including US observations)	Baseline including Year & Country Controls
Ideas Inputs			
L FT S&E (λ)	0.48 (0.06)	0.44 (0.06)	0.39 (0.05)
L PATENT STOCK (ϕ)	1.19 (0.07)	1.17 (0.07)	0.84 (0.02)
Year & Country Controls			
Y_0			1.12 (0.07)
YEAR			-0.08 (0.003)
L GDP67			-0.24 (0.05)
Year Dummies (20)	Sig.	Sig.	
Country Dummies (16)	Sig.	Sig.	
Regression Statistics			
R-squared	0.99	0.99	0.98
# of Observations	330.00	351.00	330.00

TABLE 4
INTERNATIONAL SPILLOVERS &
R&D PRODUCTIVITY OVER TIME

	Dependent Variable = Ln(Patents)		
	(4-1)	(4-2)	(4-3)
	(3-1) with no Year control	(4-1) including L "Rest of World" Patent Stock	(3-3) including L "Rest of World" Patent Stock
Ideas Inputs			
L FT S&E (λ)	0.85 (0.09)	0.81 (0.08)	0.30 (0.04)
L PATENT STOCK (ϕ)	0.02 (0.02)	1.03 (0.10)	0.90 (0.02)
L "REST OF WORLD" PATENT STOCK (ψ - β)		-1.05 (0.10)	-1.01 (0.08)
Year & Country Effects			
Y ₀			0.17 (0.10)
YEAR			0.03 (0.01)
L GDP67			-0.21 (0.04)
Year Dummies (20)			
Country Dummies (16)	Sig.	Sig.	
Regression Statistics			
R-squared	0.98	0.99	0.98
# of Observations	330.00	330.00	330.00

TABLE 5
REGIONAL, LANGUAGE &
TRADE-RELATED SPILLOVERS

	Dependent Variable = Ln(Patents)		
	(5-1)	(5-2)	(5-3)
	(4-3) including L “Rest of Region” Patent Stock	(4-3) including L “Rest of Shared Language” Patent Stock	(4-3) including L Imports & L Exports
Ideas Inputs			
L FT S&E (λ)	0.36 (0.06)	0.20 (0.07)	0.30 (0.04)
L PATENT STOCK (ϕ)	0.85 (0.02)	0.94 (0.03)	0.89 (0.02)
L “REST OF REGION” PATENT STOCK	-0.01 (0.004)		
L “REST OF SHARED LANGUAGE” PATENT STOCK		-0.02 (0.01)	
L EXPORTS TO U.S.			0.11 (0.03)
L IMPORTS FROM U.S.			-0.11 (0.02)
L “REST OF WORLD” PATENT STOCK	-0.58 (0.06)	-0.28 (0.07)	-1.05 (0.08)
Year & Country Effects			
Y ₀	0.56 (0.09)	0.94 (0.09)	0.09 (0.09)
YEAR	-0.01 (0.01)	-0.05 (0.01)	0.03 (0.01)
L GDP67	-0.26 (0.05)	-0.15 (0.06)	-0.21 (0.04)
Regression Statistics			
R-squared	0.98	0.98	0.99
# of Observations	330.00	330.00	330.00

TABLE 6A
EXPLORING SERIAL CORRELATION

	Dependent Variable = Ln(Patents)		
	(6-1)	(6-2)	(6-3)
	(3-1) with AR(1) correction	(4-2) with AR(1) correction	(4-3) with AR(1) correction
Ideas Inputs			
L FT S&E (λ)	0.36 (0.15)	0.45 (0.13)	0.22 (0.08)
L PATENT STOCK (ϕ)	1.13 (0.21)	1.07 (0.15)	0.90 (0.02)
L "REST OF WORLD" PATENT STOCK (ψ - β)		-1.30 (0.16)	-1.22 (0.07)
Year & Country Effects			
Y_0		0.04 (0.04)	-0.05 (0.03)
YEAR		0.04 (0.01)	0.06 (0.01)
L GDP67			-0.16 (0.09)
Year Dummies (20)	Sig.		
Country Dummies (16)	Sig.	Sig.	
Regression Statistics			
# of Observations	330.00	330.00	330.00

TABLE 6B
EXPLORING THE CALCULATION OF THE PATENT STOCK

	Dependent Variable = Ln(Patents)		
	(6-1)	(6-2)	(6-3)
	(4-3) using only post-1980 observations	(4-3) using “depreciated” L Patent Stock	(4-3) including lagged L Patents (instrumented)*
Ideas Inputs			
L FT S&E (λ)	0.40 (0.05)	0.40 (0.05)	0.21 (0.04)
L PATENTS (T-1)			0.45 (0.06)
L PATENT STOCK (ϕ)	0.85 (0.02)		0.48 (0.05)
“DEPRECIATED” L PATENT STOCK (ϕ)		0.86 (0.02)	
L “REST OF WORLD” PATENT STOCK (ψ - β)	0.36 (0.45)	-0.47 (0.22)	0.07 (0.13)
Year & Country Effects			
Y ₀	-0.20 (0.06)	0.37 (0.07)	0.29 (0.06)
YEAR	-0.10 (0.04)	0.01 (0.02)	-0.04 (0.01)
L GDP67	-0.26 (0.05)	-0.27 (0.05)	-0.15 (0.04)
Regression Statistics			
R-squared	0.99	0.98	
# of Observations	224.00	287.00	300.00

- Instrumental Variables for L Patents(t-1) = L FTE S&E(t-1), L PATENT STOCK(t-1)

TABLE 7
IDEAS PRODUCTION FUNCTION:
INVERSION APPROACH

	(7-1)	(7-2)	(7-3)
Ideas Inputs			
L FT S&E (λ)	1.33 (0.07)	0.94 (0.09)	0.67 (0.09)
L GDP (ϕ/σ)	1.09 (0.20)		
L POP ($-\phi/\sigma$)	-1.65 (0.09)		
L TFP (ϕ/σ)		0.13 (0.05)	1.10 (0.15)
L CAPITAL			0.99 (0.14)
Year & Country Effects			
Y_D	0.12 (0.12)		
YEAR	-0.03 (0.01)		
L GDP67	0.39 (0.18)		
Year Dummies (20)		Sig.	Sig.
Country Dummies (16)		Sig.	Sig.
Regressions Statistics			
R-squared	0.93	0.99	0.99
# of Observations	330.00	271.00	271.00

TABLE 8
SENSITIVITY OF AGGREGATE PRODUCTIVITY
TO IDEAS PRODUCTION

	(8-1)	(8-2)	(8-3)	(8-4)
DEPENDENT VARIABLE	L GDP	L GDP	L TFP	L TFP
L PATENT STOCK (σ)	0.22 (0.02)	0.11 (0.02)	0.05 (0.01)	0.06 (0.06)
L "REST OF WORLD" PATENT STOCK			-0.31 (0.24)	-0.32 (0.23)
L PATENTS				-0.02 (0.06)
Aggregate Production Function Inputs				
L POPULATION	0.38 (0.10)			
L LABOR ($1 - \alpha$)		0.59 (0.04)		
L CAPITAL (α)		0.28 (0.05)		
Year & Country Effects				
Y_D		0.04 (0.06)	-0.03 (0.08)	-0.03 (0.08)
YEAR		-0.001 (0.002)	0.04 (0.02)	0.03 (0.02)
Year Dummies (20)		Sig.		
Country Dummies (16)		Sig.		
Regressions Statistics				
R-squared	0.97	0.97	0.13	0.13
# of Observations	288.00	271.00	271.00	271.00