

Relating the Knowledge Production Function to Total Factor Productivity: An Endogenous Growth Puzzle

YASSER ABDIH AND FREDERICK JOUTZ*

The knowledge production function is central to research and development–based growth models. This paper empirically investigates the knowledge production function and intertemporal spillover effects using cointegration techniques. Time-series evidence suggests there are two long-run cointegrating relationships. The first captures a long-run knowledge production function; the second captures a long-run positive relationship between total factor productivity (TFP) and the knowledge stock. The results indicate that strong intertemporal knowledge spillovers are present and that the long-run impact of the knowledge stock on TFP is small. This evidence is interpreted in light of existing theoretical and empirical evidence on endogenous growth. [JEL O4, O3, C5]

The most recent advancement of endogenous growth theory has been the emergence of research and development–based (R&D-based) models of growth in the seminal papers of Romer (1990), Grossman and Helpman (1991a and 1991b), and Aghion and Howitt (1992). This class of models aims to explain the role of technological progress in the growth process. R&D-based models view technology as the primary determinant of growth and treat it as an endogenous variable.

*Yasser Abdi is an Economist in the IMF Institute. Frederick Joutz is a Professor of Economics at The George Washington University. The authors would like to thank Ralph Chami, Robert Flood, Tim Fuerst, Arthur Goldsmith, Jim Hirabayashi, Costas Mastrogianis, Stephanie Shipp, Holger Wolf, and an anonymous referee for providing valuable comments. Yasser Abdi also thanks John Kendrick. A preliminary version of this research was awarded the John Kendrick Prize, which supports research in the areas of productivity and growth.

At the heart of R&D-based growth models is a knowledge/technology production function that describes the evolution of knowledge creation. According to that function, the rate of production of new knowledge depends on the amount of labor engaged in R&D and the existing stock of knowledge available to these researchers. A crucial debate framed by the work of Romer (1990) and Jones (1995a) (within the R&D-based growth literature) is centered on the functional form of the knowledge production function. Specifically, the debate is centered on how strongly the flow of new knowledge depends on the existing stock of knowledge. Intuitively, the dependence of new knowledge on the existing stock is intended to capture an “intertemporal spillover of knowledge” to future researchers—that is, knowledge or “ideas” discovered in the past may facilitate the discovery or creation of “ideas” in the present. Hence, the debate is concerned with the magnitude or the strength of these intertemporal knowledge spillovers. As we will discuss, different assumptions on the magnitude of knowledge spillovers generate completely different predictions for long-run growth.

This paper contributes to the empirical understanding of R&D-based growth models in the following ways. We use time-series data for the U.S. economy over the postwar period and directly estimate the parameters of the knowledge production function. This allows us to directly assess the magnitude of knowledge spillovers, the source of the Romer-Jones debate. To achieve this goal, we exploit historical time series of patent filings to construct knowledge flows and stocks. Hence, this paper draws on an extensive body of work that uses patents as measures of innovative output and regards them as useful statistics for measuring economically valuable knowledge—for example, Hausman, Hall, and Griliches (1984); Griliches (1989 and 1990); Joutz and Gardner (1996); and Kortum (1997). We employ Johansen’s (1988 and 1991) maximum-likelihood cointegration procedure to estimate the U.S. knowledge production function. Cointegration techniques are needed because, like most macroeconomic time series, the inputs and output of the knowledge production function can be plausibly characterized as nonstationary and integrated of order one, or I(1) time series. Hence, if estimated using conventional methods like ordinary least squares (OLS), the knowledge production function will suffer from spurious correlations. Johansen’s cointegration procedure corrects for any spurious correlations that may exist in the data and explicitly accounts for the potential endogeneity of the inputs of the knowledge production function.

In his seminal paper, Romer (1990) assumes a knowledge production function in which new knowledge is linear in the existing stock of knowledge, holding the amount of research labor constant. The implication of this strong form of knowledge spillovers is that the growth rate of the stock of knowledge is proportional to the amount of labor engaged in R&D. Hence, policies—such as subsidies to R&D—that increase the amount of labor allocated to research will increase the growth rate of the stock of knowledge. Because the Romer model is one in which long-run per capita growth is driven by technological progress/knowledge growth, such policies will increase long-run per capita growth in the economy.

In an influential paper, Jones (1995b) questions the empirical validity of the Romer model. The Romer model predicts that an increase in the amount of research

labor should increase the growth rate of the stock of knowledge, a prediction that depends critically on strong positive spillovers in knowledge production. Jones tests the validity of this prediction by appealing to data on total factor productivity (TFP) growth (as a proxy for knowledge growth) and R&D scientists and engineers (as a proxy for research labor). He argues that, in the United States, the number of R&D scientists and engineers has increased sharply over the postwar period, while TFP growth has been characterized by relative constancy at best. This weak relationship between the number of R&D scientists and engineers and TFP growth led Jones to conclude that the magnitude of knowledge spillovers assumed by Romer is too large. To be consistent with the empirical evidence, Jones argues that a smaller magnitude of knowledge spillovers needs to be imposed. Imposing a smaller magnitude of knowledge spillovers, however, alters the key implication of Romer's model. Specifically, in the modified model developed by Jones (1995a), long-run growth depends only on exogenously given parameters and, hence, is invariant to policy changes such as subsidies to R&D.

We study the cointegration properties of data on new knowledge (measured by the flow of new patents), the existing knowledge stock (measured by the patent stock), R&D scientists and engineers, and TFP. We include TFP in the empirical model for three reasons. First, Jones (1995b) uses TFP as a measure of knowledge, whereas we use the patent stock. The inclusion of TFP in the empirical model allows us to capture how closely our patent measure relates to Jones's measure. Second, long-run economic growth depends on TFP, which is the application and embodiment of knowledge. Third, it enables the estimated empirical model to shed some light on the observed weak relationship between TFP growth and the number of R&D scientists and engineers.

The paper finds two long-run cointegrating relationships. The first captures a long-run knowledge production function in which the flow of new knowledge depends positively on the existing stock of knowledge and the number of R&D scientists and engineers. The second captures a long-run positive relationship between TFP and the stock of knowledge (patents). The results indicate the presence of strong intertemporal knowledge spillovers, which is consistent with the Romer (1990) model. The long-run elasticity of new knowledge creation, with respect to the existing stock, is at least as large as unity. However, the long-run impact of the knowledge (patent) stock on TFP is small: Doubling the stock of knowledge (patents) is estimated to increase TFP by only 10 percent in the long run. In other words, the results suggest that although R&D scientists and engineers greatly benefit from the knowledge and ideas discovered by prior research, the knowledge they produce seems to have only a modest impact on measured TFP.

These results seem to suggest a new interpretation of the empirical evidence documented by Jones (1995b). The observed weak relationship between the number of R&D scientists and engineers and TFP growth found by Jones is not necessarily an indication of weak intertemporal knowledge spillovers. We feel that knowledge—the output from researchers' effort—is an important intermediate step to TFP. This paper provides some evidence that the rate of diffusion of new knowledge into the productive sector of the U.S. economy has been slow over the past 20 years. The application and embodiment of knowledge into productivity is com-

plex and the diffusion is slow. Our empirical work contributes to understanding and reconciling some of the spillover effects and issues raised by Jones (1995b).

The rest of the paper is organized as follows. Section I presents a simple R&D-based growth model with the focus on the Romer-Jones debate and the knowledge production function. Section II describes the data on the inputs and output of that function. Section III looks at the univariate and multivariate time-series properties of the data and estimates the knowledge production function. Section IV discusses the main results of the paper and their interpretation. Finally, Section V offers some concluding remarks.

I. The Romer-Jones Debate on Knowledge Production

In this section, we present a simplified version of the R&D-based growth models of Romer (1990) and Jones (1995a). We focus on the basic elements and the key macroeconomic implications for long-run growth. As such, we present the model in “reduced form” and, in doing so, suppress the microfoundation and market structure components. This is done purely for ease of exposition.

A Simple R&D-Based Growth Model

The model has four variables: output (Y), capital (K), labor (L), and technology or knowledge (A).¹ There are two sectors: a goods sector that produces output, and an R&D sector that produces new knowledge. Labor can be freely allocated to either of the two sectors, to produce output (L_Y) or to produce new knowledge (L_A). Hence, the economy is subject to the following resource constraint: $L_Y + L_A = L$.

Specifically, output is produced according to the following Cobb-Douglas production function with labor augmenting (Harrod-neutral) technological progress:

$$Y = K^\alpha (AL_Y)^{1-\alpha}, \text{ where } 0 < \alpha < 1. \tag{1}$$

New knowledge or new ideas are generated in the R&D sector. Let A denote the stock of knowledge/technology available in the economy. The knowledge stock can be thought of as the accumulation of all ideas that have been invented or developed. Then, \dot{A} represents the flow of new knowledge or the number of new ideas generated in the economy at a given point in time. New ideas are produced by researchers, L_A , according to the following production function:

$$\dot{A} = \bar{\delta} L_A, \tag{2}$$

where $\bar{\delta}$ denotes (average) research productivity, that is, the number of new ideas generated per researcher. In turn $\bar{\delta}$ is modeled as a function of the existing stock of knowledge/ideas (A) and the number of researchers (L_A) according to the following:

$$\bar{\delta} = \delta A^\phi L_A^{\lambda-1}, \delta > 0, \tag{3}$$

¹In this paper, knowledge, technology, and ideas are used interchangeably.

where δ , ϕ , and λ are constant parameters. The presence of the term A^ϕ in equation (3) is intended to capture the dependence of current research productivity on the stock of ideas already discovered. Ideas formulated in the past may facilitate the discovery or creation of ideas in the present, in which case current research productivity is increasing in the stock of knowledge ($\phi > 0$). Hence, $\phi > 0$ captures a positive “spillover of knowledge” to future researchers and is referred to as the “standing-on-shoulders” effect. Alternatively, it is possible that the most obvious ideas are discovered first and new ideas become increasingly harder to find over time. In this case, current research productivity is decreasing in the stock of ideas already discovered. This corresponds to $\phi < 0$, the “fishing-out” effect.²

The presence of the term $L_A^{\lambda-1}$ in equation (3) captures the dependence of research productivity on the number of people seeking new ideas at a given point in time. For example, it is quite possible that the greater the number of people searching for ideas, the more likely it is that duplication or overlap in research would occur. In that case, if we double the number of researchers (L_A), we may less than double the number of unique ideas or discoveries (\dot{A}). This notion of duplication in research, or the “stepping-on-toes” effect, can be captured mathematically by allowing for $0 < \lambda < 1$, in which case research productivity is decreasing in L_A .

Taken together, equations (2) and (3) suggest the following knowledge production function:

$$\dot{A} = \delta L_A^\lambda A^\phi. \quad (4)$$

That is, the number of new ideas or new knowledge at any given point in time depends on the number of researchers and the existing stock of ideas.

Growth Implications of the Model

Given the above setup, it can be easily shown that a balanced growth path/steady state exists for this economy, which is defined as a situation in which all variables grow at constant (possibly zero) rates. Along this path, output per worker (y) and the capital-labor ratio (k) grow at the same rate as technology (A):

$$g_y = g_k = g_A, \quad (5)$$

where g_y , g_k , and g_A respectively denote the steady state growth rate of y , k , and A . Hence, R&D-based growth models share the prediction of the neoclassical Solow model that technological progress is the source of sustained per capita growth. If technological progress ceases, so will long-run per capita growth. Therefore, to solve for the steady state per capita growth rate in this economy, it suffices to solve for g_A , which is in turn determined by the knowledge production function as shown below. We focus on two versions of that function: Romer (1990) and Jones (1995a). Their versions have completely different implications for long-run growth.

²The case in which $\phi = 0$ allows the fishing-out effect to completely offset the standing-on-shoulders effect. That is, current research productivity is independent of the stock of knowledge.

Those implications depend critically on the magnitude of the knowledge spillover parameter assumed (ϕ in equation (4)).

Romer's Model

Romer (1990) assumes a particular form of the knowledge production function in equation (4). He imposes the restrictions $\phi = 1$ and $\lambda = 1$. The key restriction made by Romer, however, is $\phi = 1$. This makes \dot{A} linear in A , and generates growth in the stock of knowledge (\dot{A}/A) that depends on L_A unit homogeneously:

$$\frac{\dot{A}}{A} = \delta L_A. \quad (6)$$

Equation (6) pins down the steady state growth rate of the stock of knowledge, g_A , as

$$g_A = \delta L_A. \quad (7)$$

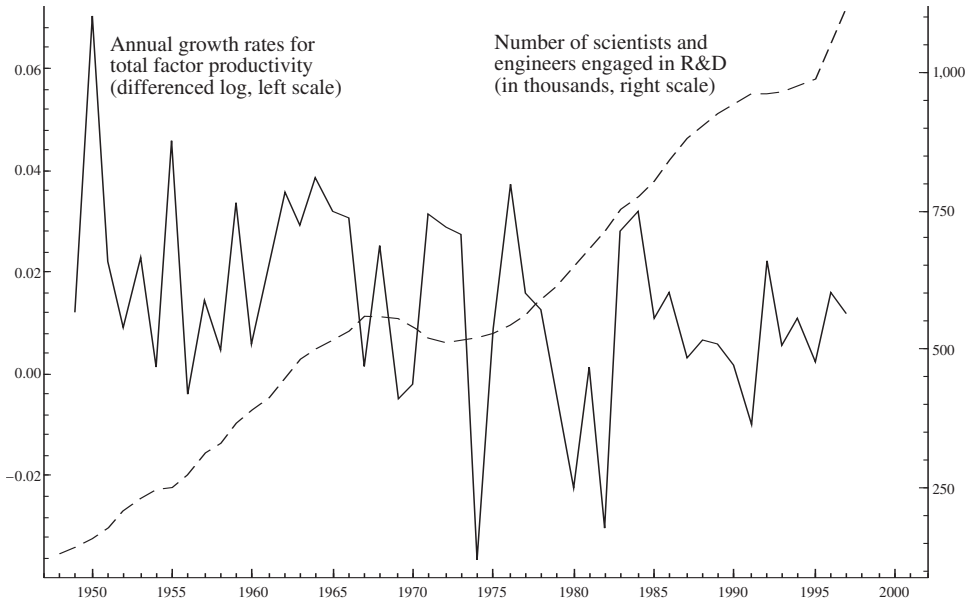
That is, the steady state growth rate of the stock of knowledge and per capita output by equation (5) depend positively on the amount of labor devoted to R&D. This key result has important policy implications: Policies that permanently increase the amount of labor devoted to R&D—a subsidy that encourages research, for example—have a permanent long-run effect on the growth rate of the economy. This “growth effects” result is a hallmark of the Romer (1990) model and many existing R&D-based endogenous growth models, including the important contributions of Grossman and Helpman (1991a and 1991b) and Aghion and Howitt (1992). This result stands in sharp contrast to the neoclassical Solow model, in which changes in variables that are potentially affected by policy have short- and medium-run effects but no long-run growth effects.

Jones's Critique

Equation (7) predicts “scale effects”: An increase in the level of resources devoted to R&D—as measured by L_A —leads to an increase in the growth rate of the economy. In an influential paper, Jones (1995b) presents time-series evidence against scale effects using one measure for L_A and one for \dot{A}/A for the United States over the postwar period. He represents L_A by the number of scientists and engineers engaged in R&D. This is perfectly reasonable because, theoretically, L_A captures the R&D workforce. Jones uses TFP growth as a proxy for \dot{A}/A , which is shown in Figure 1 for the U.S. economy. TFP growth appears to fluctuate around a relatively constant mean of about 1.4 percent per year over the postwar period. Therefore, L_A should, like \dot{A}/A , be relatively constant and exhibit no persistent increase. Otherwise, Romer's knowledge production function and the resulting scale effects are inconsistent with the time-series evidence.

Figure 1 also plots L_A , as measured by the number of scientists and engineers engaged in R&D for the U.S. economy. As Figure 1 reveals, L_A is not relatively

Figure 1. Total Factor Productivity Growth and Research and Development (R&D) Scientists and Engineers in the United States



Sources: National Science Foundation, *Science and Engineering Indicators-2000*; Jones (2002); Machlup (1962); and Bureau of Labor Statistics.

constant over the postwar period. Rather, it exhibits a strong upward trend, rising from about 100,000 in 1950 to about 1 million by 1997. Therefore, the knowledge production function in equation (6), which lies at the heart of the Romer (1990) model, is inconsistent with the time-series data.³

Jones's Alternative

Because the rejection of the scale effects prediction is rooted in the incongruence of the knowledge production function with the time-series data, it seemed sensible for Jones to tackle and modify its functional form to develop an alternative specification that is consistent with the observed time-series pattern of the data. Jones (1995a) actually shows that relaxing the assumption $\phi = 1$ generates a steady state that is consistent with the rising number of research workers observed in the data. To do that, consider once again the knowledge production function in equation (4) and divide both sides of that equation by A :

$$\frac{\dot{A}}{A} = \delta \frac{L_A^\lambda}{A^{1-\phi}}. \tag{8}$$

³The criticism by Jones (1995b) is not exclusive to the Romer (1990) model, but rather it is a criticism against many existing R&D-based endogenous growth models that share Romer's knowledge production function.

In the steady state, the growth rate of A is constant by definition. Therefore, the right-hand side of equation (8) must be constant in the steady state, which means that L_A^λ and $A^{1-\phi}$ must grow at the same rate. That is,

$$\lambda \frac{\dot{L}_A}{L_A} = (1-\phi) \frac{\dot{A}}{A}. \quad (9)$$

Now λ is a positive parameter and \dot{A}/A is always positive and constant in the steady state. Therefore, equation (9) implies that a constant steady state growth of A will be consistent with a rising L_A , that is, $\dot{L}_A/L_A > 0$, provided that ϕ is less than unity. Hence, Jones (1995a) argues that assuming $\phi < 1$ is consistent with the observed relative constancy of TFP growth (the proxy of \dot{A}/A used by Jones) in spite of the rising trend of R&D scientists and engineers. Moreover, with $\phi < 1$ imposed, the scale effects of the Romer (1990) model are removed. This can be seen formally by solving for the steady state growth rate of A from equation (9) as follows:

$$g_A = \frac{\lambda}{1-\phi} \frac{\dot{L}_A}{L_A}. \quad (10)$$

That is, the long-run *growth rate* of the stock of knowledge, which is also the long-run growth rate of per capita output by equation (5), depends on the *growth rate* of L_A rather than its *level*. Note that positive knowledge spillovers are not ruled out. The parameter capturing knowledge spillovers, ϕ , may plausibly be positive and large. What the above discussion does suggest is that the degree of positive knowledge spillovers assumed by Romer is arbitrary and inconsistent with the time-series evidence. A weaker magnitude of such spillovers is needed to achieve congruency with the evidence.

Now, along the balanced growth path/steady state, the growth in the number of research workers will be equal to the growth rate of the labor force/population. If it were greater, then the number of researchers would eventually exceed the labor force, which is not feasible. Let n denote the growth rate of the labor force/population, which Jones (1995a), following the literature, assumes to be exogenously given. Then, in the steady state, $\dot{L}_A/L_A = \dot{L}/L = n$. Substituting this relationship into equation (10) yields the following:

$$g_A = \frac{\lambda}{1-\phi} n. \quad (11)$$

Equation (11) implies that long-run growth depends on ϕ , λ , and n , parameters that usually are assumed to be exogenously given. Hence, long-run growth in the Jones (1995a) model is independent of policy changes such as subsidies to R&D. Because the returns to knowledge accumulation are assumed to be less than unity ($\phi < 1$), such changes will affect the growth of A along the transition path to a new steady state, and these “transitional growth effects” will be translated into long-run level effects. Simply stated, subsidies to R&D will alter the long-run level of the

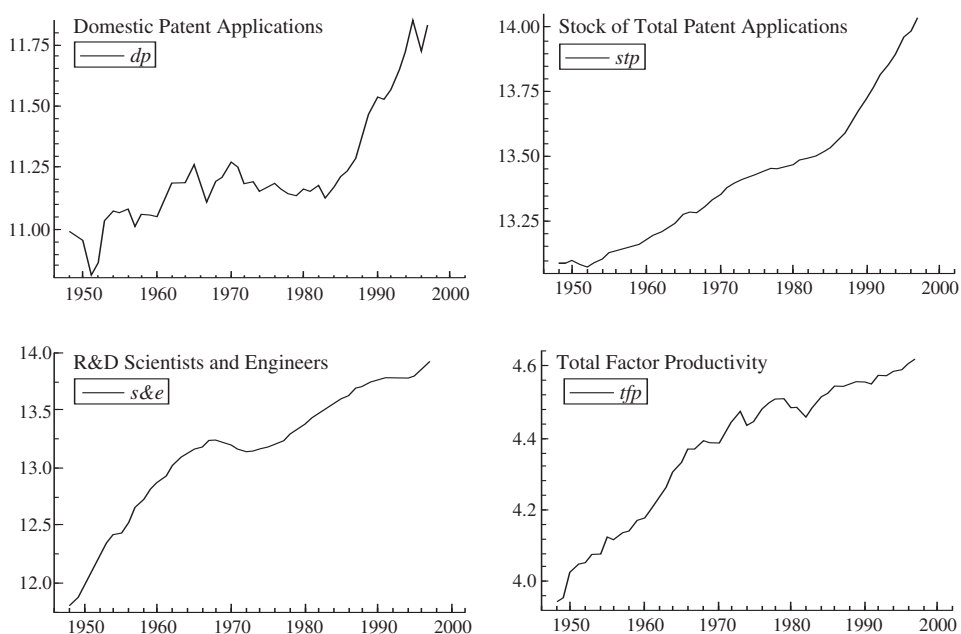
stock of knowledge but not its long-run growth rate. In Jones's (1995a) modified model, long-run growth is invariant to policy.

II. Data

In this section, we describe the variables used to empirically reconsider the theoretical relationships between the knowledge production function and productivity. The four variables include patent applications, the stock of patents, the number of scientists and engineers engaged in R&D, and TFP. The sample frequency is annual and is available from 1948 to 1997. Variables in levels will be transformed into natural logarithms and are shown in Figure 2. Details on the construction and sources of the data are found in the appendix.

Patent applications serve as a valuable resource for measuring innovative activity and have been extensively used in the patent literature as measures of technological change (see, for example, Hausman, Hall, and Griliches, 1984; Griliches, Pakes, and Hall, 1987; and Kortum, 1997). Also, Griliches (1989 and 1990) argues that the aggregate count of patents can serve as a measure of shifts in technology. Joutz and Gardner (1996) argue that patent application trends are a good approximation for technological output over the long run. Firms have invested resources to

Figure 2. Data
(In logs)



Sources: National Science Foundation, *Science and Engineering Indicators-2000*; Jones (2002); Machlup (1962); Bureau of Labor Statistics, and U.S. Patent and Trademark Office.

develop new technologies, which they feel have economic value and for which they are willing to submit applications to capture rents from their initial investments. This paper follows the patent literature and uses patent applications to construct knowledge flows and stocks.⁴ The output of the knowledge production function should reflect new knowledge created by U.S. researchers. As such, we use domestic patent applications (DP) filed at the U.S. Patent and Trademark Office (USPTO) to measure new knowledge.

Figure 2 plots the log of domestic patent applications (dp). There is an overall upward trend in the series, with domestic patent applications growing at an average annual rate of 1.7 percent between 1948 and 1997. The behavior of the series since the mid-1980s is particularly striking: Since about 1985, domestic patent applications have increased dramatically at an average annual rate of 5.1 percent. Jaffe and Lerner (2004) claim that much of this increase in patent applications may be spurious and does not necessarily reflect a true increase in the flow of new knowledge. They argue that regulatory and institutional changes simply made it easier to get a patent on unoriginal ideas, and this encouraged the filing of dubious patent applications.⁵

However, there is sizable evidence that much of the increase in patent applications since the early 1980s *has not* been due to institutional and regulatory changes that made it easier to patent ideas that dubiously constitute an innovation, but rather the increase reflects a true surge in discovery and innovation. First, Greenwood and Yorukoglu (1997) document that the 1980s and 1990s witnessed an explosion of formation of new firms and innovation in the high-tech industries,

⁴Note that we measure knowledge/technology using patent applications rather than patent grants. The lag between application and grants could be quite long, and it varies over time partly because of changes in the availability of resources to the U.S. Patent and Trademark Office. This notion is best articulated by Griliches (1990, p. 1690): “A change in the resources of the patent office or in its efficiency will introduce changes in the lag structure of grants behind applications, and may produce a rather misleading picture of the underlying trends. In particular, the decline in the number of patents granted in the 1970s is almost entirely an artifact, induced by fluctuations in the U.S. Patent Office, culminating in the sharp dip in 1979 due to the absence of budget for printing the approved patents.” This paper views patent applications as a much better measure of knowledge and technology than patent grants. Also, it is widely believed that patent application data are a better measure of new knowledge produced in an economy than R&D expenditures (see, for example, Joutz and Gardner, 1996). R&D expenditures are more properly thought of as inputs to technological change, whereas patents are an output. Hence, patent applications more closely approximate the output of the knowledge production function in R&D-based growth models than R&D expenditures.

⁵Specifically, in 1982, Congress established the Court of Appeals of the Federal Circuit (CAFC), a specialized and centralized appellate court to hear patent cases. (Before 1982, patent appeals cases were heard before various district courts, which differed considerably in their interpretation of the patent law.) Jaffe and Lerner (2004, p. 2) argue that the court “has interpreted patent law to make it easier to get patents, easier to enforce patents against others, easier to get large financial awards from such enforcement, and harder for those accused of infringing patents to challenge the patents’ validity.” Moreover, the court’s rulings regarding the standard of “novelty” and “nonobviousness” may have made it easier for applicants to file and get a patent of dubious validity. In addition, in the early 1990s, Congress changed the patent office from an agency funded by taxpayer money to a self-financed agency, that is, one that relied exclusively on patent application fees to conduct its business. Jaffe and Lerner argue that that might have created a strong incentive for the patent officer to process applications more quickly and at minimum cost. This might have reduced the rigor by which the standards of novelty and nonobviousness are exercised when reviewing patent applications. This, in turn, encouraged the filing of dubious patent applications.

particularly in the information technology, biotechnology, and software industries. Hence, the sharp increase in patenting may indicate a “technological revolution” as emphasized by those authors. Second, it is quite possible that the use of information technology in the discovery of new ideas might have substantially boosted research productivity. Arora and Gambardella (1994) argue that this was an important source of accelerating technological change. A third possibility, emphasized by Kortum and Lerner (1998), is that the sharp increase in patenting since the mid-1980s indicates an increase in innovation driven by improvements in the management of R&D. In particular, there has been a reallocation of resources from basic research toward more applied activities and hence a resulting surge in patentable discoveries. As Kortum and Lerner (1998, p. 287) point out, “Firms are restructuring, redirecting and resizing their research organizations as part of a corporate-wide emphasis on the timely and profitable commercialization of inventions combined with the rapid and continuing improvement of technologies in use.”

In addition, several studies have argued that (the inverse of) the relative price of capital is a good indicator of the quantity of economically useful knowledge—for example, Krusell (1998) and Cummins and Violante (2002).⁶ Samaniego (2005) compares (the inverse of) the relative price of capital with patent applications for the U.S. economy over the postwar period and finds that the two series are highly positively correlated, which is supportive of the use of patent applications as a measure of knowledge. More important, he observes that the growth in both series accelerated starting in the 1980s. This is consistent with the argument that the surge in patenting that started in the 1980s is not spurious but rather reflects an actual increase in the rate of innovation in the U.S. economy.

The stock of knowledge is derived from the cumulated number of total patents applied for by U.S. and foreign inventors. Patent filings are converted into a stock measure (STP) using the perpetual inventory method with a depreciation rate of 15 percent. This is typical in the U.S. patent literature (for example, Griliches (1989) and Joutz and Gardner (1996)). While this approach is ad hoc and not necessarily justified by theory, researchers have typically checked the robustness of their results against changes in the depreciation rate. We experimented with constructing stocks using 0, 5, and 10 percent depreciation rates and found that the precise rate made little difference. Hence, the results presented in this paper are not sensitive to changes in the depreciation rate on the stock.

As shown in Figure 2, the log of the stock of total patent applications (*stp*) follows a strong upward trend, with the stock growing at an average annual rate of 1.9 percent between 1948 and 1997. There appears to be a substantially stronger trend since the mid-1980s, capturing the more rapid increase in (the number of) domestic patent applications that occurred over that period.

Figure 2 also plots the log of the total number of scientists and engineers engaged in R&D activities (*s&e*) in the United States, as compiled by the National

⁶However, although the relative price of capital has merit insofar as it captures embodied knowledge in the capital stock, it does not capture the sources of new knowledge, which are not in physical capital.

Science Foundation. This measure was used by Jones (2002) and represents scientists and engineers employed in industry, the federal government, educational institutions, and nonprofit organizations. It is accepted as the best proxy for the primary input or effort in the knowledge production process. The series exhibits a very strong upward trend over the past 50 years, with the number of R&D scientists and engineers growing at an average annual rate of 4.3 percent over the period 1948–97.⁷

Finally, Figure 2 shows the plot for the log of total factor productivity (*tfp*), as compiled by the Bureau of Labor Statistics. TFP follows an upward trend over the postwar era, growing at an average annual rate of about 1.4 percent between 1948 and 1997. Growth appears to have slowed since 1973—the well-known productivity slowdown. Before 1973, the average annual growth rate of TFP was 2.1 percent. After 1973, the average annual growth rate declined to about 0.7 percent.

III. Estimation of Knowledge Production Functions

We employ the general-to-specific modeling approach advocated by Hendry (1986). This approach attempts to characterize the properties of the sample data in simple parametric relationships that remain reasonably constant over time and are interpretable in an economic sense. Rather than using econometrics to illustrate theory, the goal is to “discover” which alternative theoretical views are tenable and test them scientifically. The approach begins with a general hypothesis about the relevant explanatory variables and dynamic process (that is, the lag structure of the model). The general hypothesis should be considered acceptable to all adversaries. Then the model is narrowed down by testing for simplifications or restrictions on the general model.

The four macroeconomic and innovation variables are linked through two main relationships. The long-run knowledge production function and the long-run relationship between TFP and the stock of total patents (knowledge) can be specified as follows:

$$dp = F(stp, s \& e) \quad (12)$$

$$tfp = G(stp), \quad (13)$$

where lowercase letters denote variables in natural logarithms. That is, *dp* denotes the log of the number of domestic patent applications, *stp* denotes the log of the stock of total patent applications, and *s&e* denotes the log of the number of scientists and engineers engaged in R&D. According to the above production function, U.S. R&D scientists and engineers produce U.S. patents, but

⁷In the late 1960s through the early 1970s, however, employment of R&D scientists and engineers seems to have declined. The National Science Foundation (1998) documents that this is probably due to the substantial decline in federal funding for space-related R&D in the late 1960s and early 1970s after the thrust of funding in the early to mid-1960s, during which time the United States invested substantial resources in the “space race” with the Soviet Union.

they draw upon the “world” stock of knowledge.⁸ Also, the function $F(.)$ is assumed to be linear.⁹

Since Jones (1995b) used TFP as a measure of knowledge, we also include the relation $G(.)$ for the log of TFP (tfp) in the model. This allows us to capture how closely our patent measure and Jones’s measure are related and allows us to interpret the results in terms of Jones’s time-series evidence on TFP and R&D scientists and engineers. We look at the transmission mechanism by separating the R&D effort and output. The total stock of patents represents the cumulative R&D output, which leads to higher productivity. This is consistent with the substantial microproductivity literature (Jaffe, Trajtenberg, and Henderson, 1993; and Thompson and Fox-Kean, 2005) that postulates a positive dependence of TFP on the stock of patents.

The first step in the modeling approach examines the time-series properties of the individual data series. We look at patterns and trends in the data and test for stationarity and the order of integration. Second, we form a vector auto regression (VAR) system. This step involves testing for the appropriate lag-length of the system, including residual diagnostic tests and tests for model/system stability. Third, we test the system for potential cointegration relationship(s). Data series integrated of the same order may be combined to form economically meaningful series that are integrated of lower order. Fourth, we interpret the cointegrating relations and test for weak exogeneity. Based on these results, a conditional error correction model of the endogenous variables may be specified, further reduction tests are performed, and economic hypotheses are tested. This last step will not be performed, because the primary goal is to understand the long-run relationships.

Integration Analysis

Figure 2 shows significant trends in the series and the autocorrelations were quite strong and persistent. Nelson and Plosser (1982) found that many macroeconomic and aggregate level series are shown to be well modeled as stochastic trends, that is, integrated of order one, or $I(1)$. Simple first differencing of the data will remove the nonstationarity problem, but with a loss of generality regarding the long-run “equilibrium” relationships among the variables. We performed the standard augmented Dickey-Fuller (ADF) test in both levels and differences with a constant and trend. Table 1 contains the results in five columns and is divided in two. The top half is for the tests in levels and the bottom half is for the tests in first differences or whether the series in levels are $I(1)$ and $I(2)$, respectively. The first column lists the variables. The Akaike information criterion was used to set the

⁸In section IV, we use the stock of domestic patents as an alternative measure of the stock of knowledge. We compare the results from using such a measure with the results in which the stock of total (domestic and foreign) patents is used.

⁹Recall that the R&D-based growth models of Romer (1990) and Jones (1995a) assume a Cobb-Douglas specification for the knowledge production function expressed in terms of the levels of the variables. Because the function $F(.)$ in the text is expressed in terms of the log levels of the variables, it is assumed to be linear.

Table 1. Augmented Dickey-Fuller (ADF) Test Results for Levels and Differences, 1953–97

Variable	Model with a Constant		Model with a Constant and Trend	
	Lag-length	t-ADF	Lag-length	t-ADF
<i>dp</i>	0	0.22	0	-0.78
<i>tfp</i>	0	-2.25	0	-1.43
<i>s&e</i>	1	0.89	1	-2.05
<i>stp</i>	1	1.7	1	0.44
Δdp	0	-6.51**	0	-6.52**
Δtfp	0	-6.71**	0	-6.59**
$\Delta s\&e$	0	-3.75**	0	-3.36*
Δstp	0	-2.03	0	-2.64
Δstp (Perron)	1	-2.12	1	-2.08

Notes:

For a given variable x , the augmented Dickey-Fuller equation with a constant term included has the following form: $\Delta x_t = \pi x_{t-1} + \sum_{i=1}^p \theta_i \Delta x_{t-i} + a + \varepsilon_t$ where ε_t is a white noise disturbance. The augmented Dickey-Fuller equation with a constant and trend included adds a trend term as a right-hand-side variable to the above specification. For a given variable and specification, the table reports the number of lags on the dependent variable, p , chosen using the Akaike information criterion, and the augmented Dickey-Fuller statistic, t-ADF, which is the t-ratio on π . The statistic tests the null hypothesis of a unit root in x , i.e., $\pi = 0$, against the alternative of stationarity. Critical values at the 5 percent and 1 percent significance levels, respectively, are -2.927 and -3.581.

The symbols * and ** denote rejection of the null hypothesis at the 5 percent and 1 percent critical values, respectively.

The Perron adjusted results report the test for stationarity with a structural shift in the mean with the break point at 1985, approximately 80 percent from the starting observation. The critical values tabulated by Perron (1989) are -3.82 and -4.38 at the 5 percent and 1 percent significance levels, respectively.

appropriate lag-length for the dependent variable in each test, which is provided in the second and fourth columns. The t-ADF statistics are reported in the third and fifth columns.

We cannot reject the null hypothesis of a unit root for all four variables in levels. Domestic patents, TFP, and R&D scientists and engineers reject the null of a unit root in first differences, while the stock of patents does not. However, a recursive analysis of the coefficient estimate and the t-ADF suggest that it is nonconstant with a break right where one might expect it: 1985. In our preliminary look at the data, we saw the acceleration in the propensity to patent and its impact on the stock of patents. We have also inspected the plot of the first difference in the logarithm of the patent stock measure and observed that there appears to be a permanent shift in the mean starting in the mid-1980s. The Perron (1989) structural break procedure was used to test whether there was a mean shift in the first difference process that caused the I(1) findings. We could not reject the (null) hypothesis of stationarity in the first difference process after correcting for the (structural) mean shift. We con-

clude that all of our variables are I(1) in levels, or equivalently stationary in first differences.

Cointegration Analysis

Our analysis of the inputs and output of knowledge production suggests that the processes are nonstationary. This has implications with respect to the appropriate statistical methodology. Although focusing on changes in knowledge production eliminates the problem of spurious regressions, it also results in a potential loss of information on the long-run interaction of variables (for example, Davidson and others, 1978). We examine the hypothesis of whether there exist economically meaningful linear combinations of the I(1) series: (domestic) patent filings, the stock of patents, R&D scientists and engineers, and TFP that are stationary or I(0). The Johansen (1988 and 1991) maximum likelihood procedure is used for the analysis. The procedure begins with specifying a VAR system,

$$Y_t = \pi_0 + \sum_{i=1}^p \pi_i Y_{t-i} + \Psi D_t + e_t, \quad (14)$$

where

$$Y_t = \begin{bmatrix} \textit{Patent Filings}_t \\ \textit{Patent Stock}_t \\ \textit{R \& D Scientists \& Engineers}_t \\ \textit{Total Factor Productivity}_t \end{bmatrix}, e_t ; IN(0, \Omega), \textit{ and } D_t = \begin{bmatrix} \textit{Stepdum86}_t \\ \textit{Impulse9495}_t \\ \textit{Impulse96}_t \\ \textit{Trend}_t \end{bmatrix}.$$

Y_t is (4×1) and the π_i s are (4×4) matrices of coefficients on lags of Y_t . D_t is a vector of deterministic variables that can contain a linear trend, dummy-type variables, or other regressors considered to be fixed and nonstochastic. Finally, e_t is a (4×1) vector of independent and identically distributed errors assumed to be normal with zero mean and covariance matrix Ω —that is, $e_t \sim i.i.d. N(0, \Omega)$. As such, the VAR is composed of a system of four equations, in which the right-hand side of each equation includes a common set of lagged and deterministic regressors.

The VAR includes our four series: the log of domestic patent applications, dp ; the log of TFP, tfp ; the log of the (lagged) stock of total patent applications, $stp1$ ¹⁰; and the log of the number of R&D scientists and engineers, $s\&e$. The VAR also includes a constant, a trend term, and three dummy variables. The first dummy variable is $Stepdum86$, which takes the value of one after 1985 and zero otherwise. The inclusion of this variable is intended to capture the dramatic increase in patenting since the mid-1980s as discussed in detail above. The second dummy variable is $Impulse9495$, which takes the value of one in 1994 and 1995 and zero otherwise, and the third is $Impulse96$, which is zero except for unity in 1996. $Impulse9495$

¹⁰The variable $stp1$ is simply stp lagged one period. Because stp is calculated as end of period stocks, we enter it with a lag in the VAR and cointegration analysis.

captures several institutional changes in the U.S. patent policy: the movement toward the typical international patent system policy of granting 20-year awards instead of 17-year awards, and the fact that 12-year patent renewal fees were collected for the first time in the United States in 1994–95 (Kortum, 1997). *Impulse96* captures the instantaneous negative response by agents facing an increased cost of patent applications (see Figure 2).¹¹

Following Johansen and Juselius (1990), the VAR model provides the basis for cointegration analysis. Adding and subtracting various lags of Y yields an expression for the VAR in first differences. That is,

$$\Delta Y_t = \pi_0 + \pi Y_{t-1} + \sum_{i=1}^{p-1} \Gamma_i \Delta Y_{t-i} + \Psi D_t + e_t, \quad (15)$$

where $\Gamma_i = -(\pi_{i+1} + \dots + \pi_p)$, $i = 1, \dots, p-1$ and $\pi \equiv \left(\sum_{i=1}^p \pi_i \right) - I$.

If π is a zero matrix, then modeling in first differences is appropriate. The matrix π may be of full rank or less than full rank, but of rank greater than zero. When $\text{rank}(\pi) = 4$, then the original series are not $I(1)$, but in fact $I(0)$; modeling in differences is unnecessary. But, if $0 < \text{rank}(\pi) \equiv r < 4$, then the matrix π can be expressed as the outer product of two full column rank ($4 \times r$) matrices α and β where $\pi = \alpha\beta'$. This implies that there are $4 - r$ unit roots in πY . The VAR model can then be expressed in error correction form. That is,

$$\Delta Y_t = \pi_0 + \alpha\beta' Y_{t-1} + \sum_{i=1}^{p-1} \Gamma_i \Delta Y_{t-i} + \Psi D_t + e_t. \quad (16)$$

The matrix β' contains the cointegrating vector(s) and the matrix α has the weighting elements for the r th cointegrating relation in each equation of the VAR. The matrix rows of $\beta' Y_{t-1}$ are normalized on the variable(s) of interest in the cointegrating relation(s) and interpreted as the deviation(s) from the long-run equilibrium condition(s). In this context, the columns of α represent the speed of adjustment coefficients from the long-run or equilibrium deviation in each equation. If the coefficient is zero in a particular equation, that variable is considered to be weakly exogenous and the VAR can be conditioned on that variable.

Unrestricted Model and Testing for Cointegration

Before conducting the cointegration tests, the appropriate lag-length for the VAR must be determined and a constant model found. The lag-length is not known a priori, so some testing of lag order must be done to ensure that the estimated residuals of the VAR are white noise—that is, they do not suffer from autocorrelation,

¹¹Statistically, the inclusion of *Impulse9495*, *Impulse96*, and *Stepdum86* in the VAR results in a substantial improvement in the fit of the model and much better residual diagnostics, and ensures a statistically stable/constant VAR.

non-normality, and so on. We started with a VAR that includes four lags on each variable, denoted VAR(4), then we estimated a VAR with three lags, VAR(3), and tested whether the simplification from VAR(4) to VAR(3) was statistically valid. The process was repeated sequentially down to a VAR with a single lag, VAR(1). Based on sequential F-tests for model reduction, we concluded that the simplification to a VAR with one lag is statistically valid. This result is also supported by the Schwarz criterion and the Hannan-Quinn criterion, which were minimized when the VAR had a single lag. Moreover, the VAR with a single lag produced residuals that are serially uncorrelated, normal, and homoskedastic. We have also estimated VAR(1) recursively to test for model constancy. The recursively estimated chow tests indicated that VAR(1) is statistically stable. Hence, we proceed with the analysis using the VAR(1) model. For a more detailed discussion on model reduction, residual diagnostics, and model stability results, please refer to the working paper version of this paper (Abdih and Joutz, 2005).

The cointegration analysis proceeds in several steps: testing for the existence of cointegration, interpreting and identifying the relationship(s), and conducting inference tests on the coefficients from theory and weak exogeneity. Testing permits reduction of the unrestricted general model to a final restricted model without loss of information.

Table 2 presents the initial test for cointegration and is divided into three panels. Panel A contains results on the possible number of cointegrating relations. There are four columns for the eigen-values, null hypothesis, Trace statistic, and its associated p -value. In the first row, the null hypothesis ($r = 0$) is that there are zero cointegrating vectors, as opposed to the alternative that there are more than zero cointegrating vectors. This hypothesis is soundly rejected with a trace statistic of 184.03 and no measurable p -value. When the possible maximum number of cointegrating relations is one against the alternative hypothesis that there is more than one, the test statistic is 58.87 and the p -value is [0.00]. This suggests that there are at least two cointegrating vectors. We cannot reject the null hypothesis that there are at most two cointegrating relations in the third row.

Panel B presents the two cointegrating vectors normalized on (domestic) patent filings and TFP, respectively. We interpreted the two vectors as a knowledge production function and a function for the determinants of TFP. Panel C reports the feedback coefficients and their standard errors associated with each long-run equation for the variables of the system in first differences.

The cointegrating vectors or relationships as they appear are not uniquely identified and hence the standard errors of these vectors cannot be computed. Any linear combination of the two vectors forms another stationary vector, so the estimates produced by any particular vector are not necessarily unique. Therefore, to achieve identification, it is necessary to impose restrictions on the cointegrating vectors. The restrictions are motivated by economic theory and enable us to test for overidentification and obtain standard errors for the overidentified parameters.

For ease of exposition and to more easily understand the nature of the restrictions, the model can be written in terms of equation (16). The error term and short-run components are omitted to focus on the long-term model. Also, the trend is restricted to lie in the cointegration space.

Table 2. Cointegration Analysis of the Data

(A) Johansen's Cointegration Test

Eigen-Values	Null Hypothesis	Trace Statistic	p-Value
0.937	r = 0	183.03**	[0.000]
0.687	r ≤ 1	58.87 **	[0.000]
0.120	r ≤ 2	6.65	[0.993]
0.019	r ≤ 3	0.88	[0.997]

(B) Estimated Cointegrating Vectors β'

Vector	<i>dp</i>	<i>tfp</i>	<i>stpl1</i>	<i>s&e</i>	Trend
1	1	-0.255	-1.415	-0.195	0.025
2	0.170	1	-0.418	-0.106	-0.002

(C) Feedback Coefficients α and Their Standard Errors SE(α)

	α		SE(α)	
	1	2	1	2
<i>dp</i>	-0.398	-0.342	0.120	0.121
<i>tfp</i>	0.062	-0.109	0.050	0.050
<i>stpl1</i>	0.132	0.022	0.008	0.008
<i>s&e</i>	-0.162	-0.402	0.071	0.072

Notes:

(1) The VAR includes a single lag on each variable (*dp*, *tfp*, *stpl1*, *s&e*), a constant, trend, and three dummy variables: *Stepdum86*, *Impluse9495*, and *Impluse96*. The estimation sample is 1953 to 1997.

(2) * and ** indicate rejection of the null hypothesis at the 5 percent and 1 percent critical values, respectively.

$$\begin{pmatrix} \Delta dp_t \\ \Delta tfp_t \\ \Delta stpl1_t \\ \Delta s \& e_t \end{pmatrix} = \begin{pmatrix} \alpha_{11} & \alpha_{12} \\ \alpha_{21} & \alpha_{22} \\ \alpha_{31} & \alpha_{32} \\ \alpha_{41} & \alpha_{42} \end{pmatrix} \begin{pmatrix} 1 & -\beta_{12} & -\beta_{13} & -\beta_{14} & -\beta_{15} \\ -\beta_{21} & 1 & -\beta_{23} & -\beta_{24} & -\beta_{25} \end{pmatrix} \begin{pmatrix} dp_{t-1} \\ tfp_{t-1} \\ stpl1_{t-1} \\ s \& e_{t-1} \\ trend \end{pmatrix}. \quad (17)$$

The βs and αs are those reported in Panels B and C, respectively, in Table 2. These are the implied unrestricted long-run (cointegrating) solutions. The implied (unrestricted) long-run solution of the model is given by the following:

$$dp = \beta_{12} \, tfp + \beta_{13} \, stpl1 + \beta_{14} \, s \& e + \beta_{15} \, Trend \quad (18)$$

$$tfp = \beta_{21} \, dp + \beta_{23} \, stpl1 + \beta_{24} \, s \& e + \beta_{25} \, Trend. \quad (19)$$

Two restrictions are required to just-identify the model; any additional restrictions are overidentified and thus testable. The first restriction is on the knowledge

production relation and relates the dependence of new knowledge on the stock of knowledge and the R&D scientists and engineers. There does not seem to be a reason to include a direct effect from TFP, in this relation, from the R&D-based growth theory. We impose $\beta_{12} = 0$. Second, the literature does not suggest that current patent applications should determine productivity. This restriction is imposed by setting $\beta_{21} = 0$.

The above two restrictions produce a just-identified model. We consider and test three overidentifying restrictions on the β s and α s. The first type of test is for the specification of the cointegrating relation and the latter are tests for weak exogeneity.

First, scientists and engineers in the R&D sector are unlikely to have a direct long-run impact on TFP. Again, the effect is only indirect: R&D scientists and engineers produce new knowledge. New knowledge ultimately augments the stock of knowledge and the latter has a potential impact theoretically on tfp . Thus, we exclude $s\&e$ from the second cointegrating vector by testing $\beta_{24} = 0$. Statistically, the likelihood ratio test of the restriction $\beta_{24} = 0$ cannot be rejected. The test statistic is $\chi^2(1) = 1.43$ [0.23]; the degrees of freedom are in parentheses and the p -value is in brackets.

Weak exogeneity is an important issue in model reduction. It implies that inference testing can be conducted for the parameters of interest from a conditional density rather than a joint density without loss of information. The modeling effort is simpler yet still efficient. The first hypothesis is that the cointegrating relationship for new patent flows (dp_t) does not explain changes in TFP (Δtfp_t). The restriction is $\alpha_{21} = 0$. The “discovery” of new knowledge is unlikely to explain fluctuations in productivity. Our second hypothesis is that TFP does not provide information about the change in the (lagged) stock of knowledge. The timing issue aside, it suggests that $stp11$ is weakly exogenous with respect to the second cointegrating vector. Because $\alpha_{32} = 0$, the second cointegrating relationship does not enter the equation for $\Delta stp11_t$. These two feedback coefficients appear numerically small in Panel C of Table 2. If all three overidentifying restrictions are imposed, the joint hypothesis cannot be rejected: $\chi^2(3) = 4.05$ [0.26].

Restricted Cointegration Model

Table 3 contains the results from the three restrictions on the just-identified model. The table is divided into two panels. Panel A reports the restricted (and identified) estimates for the cointegrating vectors, the β s, together with their standard errors. Panel B reports the feedback coefficients estimates, α s, and their standard errors. In Panel A, the first cointegrating vector is interpreted as a long-run knowledge or “idea” production function. The implied long-run or cointegrating relationship is given by the following:

$$dp = 1.436^{**} stp11 + 0.208^{**} s \& e - 0.023^{**} Trend. \quad (20)$$

The coefficient of the lagged stock of knowledge, $stp11$, is highly significant and indicates the presence of positive spillovers of knowledge or a standing-on-shoulders

Table 3. Restricted Cointegration Analysis of the Data

(A) Cointegrating Vectors β' and Their Standard Errors (in parentheses)					
Vector	<i>dp</i>	<i>tfp</i>	<i>stp11</i>	<i>s&e</i>	Trend
1	1	0.0	-1.436** (0.062)	-0.208** (0.020)	0.023** (0.001)
2	0.0	1	-0.108 (0.141)	0.0	-0.009** (0.002)
(B) Feedback Coefficients α and Their Standard Errors SE(α)					
	α		SE(α)		
	1	2	1	2	
<i>dp</i>	-0.449**	-0.184	0.121	0.111	
<i>tfp</i>	—	-0.108*	—	0.050	
<i>stp11</i>	0.135**	—	0.008	—	
<i>s&e</i>	-0.190**	-0.405**	0.068	0.053	

Notes:

(1) The VAR includes a single lag on each variable (*dp*, *tfp*, *stp11*, *s&e*), a constant, trend, and three dummy variables: *Stepdum86*, *Impluse9495*, and *Impulse96*. The estimation sample is 1953 to 1997.

(2) * and ** indicate the rejection (at the 5 percent and 1 percent critical values) of the null hypothesis that a particular coefficient is zero. These tests are based on the likelihood ratio statistic, which is distributed under the null hypothesis as χ^2 with 1 degree of freedom.

effect. The sign of the coefficient is consistent with the R&D-based growth models of Romer (1990) and Jones (1995a). However, its magnitude is significantly greater than unity, indicating a stronger degree of spillovers than the theoretical models.¹² After the second cointegrating vector has been examined, this result will be discussed further.

The productivity of researchers, *s & e*, increases with the stock of cumulated knowledge discovered by others in the past. The coefficient of *s & e* is positive, highly significant, and less than unity. Our estimate of 0.21 is within the range 0.1 to 0.6 that Kortum (1993) finds in the microliterature on patents and R&D effort. It supports Jones's (1995a) argument for decreasing returns owing to duplicative research. Duplication itself is not wasteful. Replication is an essential exercise in science and a component of learning by doing.

The negative coefficient for the time trend may at first seem surprising. However, it reflects the fact that the number of R&D scientists and engineers grew at a much higher average annual rate than domestic patent applications over the period 1948–97; the growth rate of the former is 4.3 percent and the latter is 1.7 percent. Their difference roughly matches the coefficient of the trend term (-2.3 percent). The negative coefficient of the trend term is also consistent with

¹²In fact, the likelihood ratio statistic rejects the null hypothesis that that coefficient is unity: $\chi^2(1) = 29.12$ [0.00].

findings by Griliches (1990) based on microdata of patents and R&D. He interprets the negative trend as capturing a decrease in the propensity to patent inventions because of the rising cost of dealing with the patent system.

Now consider the feedback coefficients for the first cointegrating vector (the knowledge production function) in Panel B of Table 3. They are all significant from zero; this means that dp , $stpl1$, and $s\&e$ are *not* weakly exogenous with respect to the parameters of the knowledge production function. That is, in the face of any deviation from long-run equilibrium, dp , $stpl1$, and $s\&e$ jointly respond and move the system back to equilibrium. This finding supports our system approach to estimating the knowledge production function. If a single equation approach had been adopted instead, we would have invalidly conditioned on $stpl1$ and $s\&e$; the result would have produced biased and inconsistent estimates of the knowledge production function.

The feedback coefficient for the Δdp equation—that is, α_{11} in equation (17)—is -0.45 and significant from zero, suggesting stability of the error correction mechanism. The coefficient implies that a positive deviation of dp from its long-run path, given by equation (20), in this period is not permanent, leading to explosive growth. The growth in (domestic) patent filings declines in the next period.

The feedback coefficient for the $\Delta stpl1$ equation—that is, α_{31} in equation (17)—is positive and significant from zero. The positive sign suggests that if dp is above its long-run equilibrium path, then this has a positive effect on the growth of the stock of knowledge in the next period.

Finally, the feedback coefficient for the $\Delta s\&e$ equation—that is, α_{41} in equation (17)—is negative and significant from zero. The negative sign of the coefficient makes sense; it implies that if dp is above its long-run equilibrium path this period, then the growth of the R&D scientists and engineers slows down in the next period to correct for the disequilibrium. This is consistent with the theory that, on a balanced growth path, the proportion of the labor force devoted to knowledge production should remain constant.

Next, consider the second cointegrating vector in Panel A of Table 3. The implied long-run cointegrating relationship is given by the following:

$$tfp = 0.108 \text{ } stpl1 + 0.009^{**} \text{ } Trend. \quad (21)$$

That is, TFP growth depends positively on the lagged stock of knowledge and a time trend. The positive coefficient for the trend term is highly significant and implies a trend growth of TFP of about 1 percent, roughly matching its average annual growth rate over the sample 1948–97. The coefficient for the lagged stock of knowledge is positive but small. It implies that doubling the stock of knowledge will increase TFP by only about 10 percent in the long run. The coefficient actually matches that found by Porter and Stern (2000) for aggregate data on Organization for Economic Cooperation and Development (OECD) countries and falls within the range of estimates of 0.06 to 0.2 found by Griliches (1990) in the micro-productivity literature.


However, the coefficient is not significant. The likelihood ratio statistic under the null hypothesis that the coefficient is zero yields the following:

$\chi^2(1) = 0.62$ [0.43]. This result is consistent with the insignificant estimates of R&D on TFP obtained by Jones and Williams (1998) for a panel of U.S. sectors using a fixed effects model. It is also consistent with the Bureau of Labor Statistics (1989) whose growth accounting computation found little contribution of R&D to TFP growth for the U.S. economy, and with Comin (2004) who calibrates a small effect of R&D on productivity using an endogenous growth model with a free entry condition into the R&D sector.

One possible interpretation from Porter and Stern (2000) is that realizing the full benefits from new knowledge and new technologies depends critically on the diffusion of these technologies into the productive sector of the economy. They provide some evidence that the rate of diffusion of new technologies has been slow and incomplete in the OECD countries over the past 20 years. This may explain the weak relationship between the stock of knowledge and TFP. In fact, this idea is supported by or appears to be evident in the U.S. data depicted in Figure 3. Since the mid-1980s, it does not look like the benefits of knowledge have been realized into measured productivity growth. Kortum (1997) identifies the trend shift in patenting behavior. We suspect that the diffusion process of this trend in knowledge growth did not have its full effect on TFP until the late 1990s, just beyond our sample.

The cointegrating relation for TFP enters into the Δtfp_t equation with the appropriate sign and is significant. The feedback effect is rather quick. Changes in scientists and engineers, $\Delta s\&e_t$, are not weakly exogenous as well. If tfp is above

Figure 3. Total Factor Productivity and (Lagged) Stock of Total Patents
(In logs)



Sources: Bureau of Labor Statistics; U.S. Patent and Trademark Office.

Note: The series are adjusted by their means.

the long-run level, it reduces *s&e* growth over several years. There is less demand for knowledge-producing workers. This has a negative effect on changes in patents, Δdp_t , and is seen in the marginally significant feedback coefficient as well. The stock of knowledge is weakly exogenous with respect to the *tfp* cointegrating relation. We are not as confident in the interpretation of the feedback coefficients from the TFP relation, because the relationships are perhaps more complex and there is not as much direction from economic theory.

IV. Discussion

This section focuses on two important results mentioned above. The first is the estimated long-run impact of the lagged stock of total patents on domestic patents. This is the coefficient of *stp11* in equation (20), which governs the magnitude of knowledge spillovers. The coefficient is 1.4, significantly greater than unity. The second result is the estimated long-run impact of the lagged stock of total patents on TFP. This is the coefficient of *stp11* in equation (21), which is 0.108.

The microliterature on patents and R&D has little to say on knowledge spillovers, because the knowledge production function from that literature does not derive from the R&D-based models of growth and therefore does not include the stock of existing knowledge as a variable explaining new knowledge. Also, simple aggregation from the microlevel may not capture the potential externalities across sectors.

However, we can compare our results with Porter and Stern (2000). They employ aggregate panel data on OECD countries to estimate a knowledge production function, in which domestic patents in each country do depend on that country's existing stock of knowledge. However, they use the stock of domestic patents for a particular country as the proxy for the stock of knowledge in that country, rather than using the stock of total (domestic and foreign) patents, which we use for the U.S. economy. Their results indicate a spillover coefficient of unity supporting the Romer (1990) model.

To check the robustness of our model, we reestimate a knowledge production function, but replace the stock of total (domestic and foreign) patents with the stock of domestic patents. To explain, let *sdpl1* denote the (lagged) *stock* of domestic patents in the U.S. economy. Starting from an unrestricted VAR that includes *dp*, *tfp*, *sdpl1*, and *s&e*, we implement the Johansen (1988 and 1991) cointegration methodology and again estimate two cointegrating vectors. We find that the estimated long-run impact of the (lagged) stock of domestic patents on the flow of new domestic patents (*dp*) is unity and significant, but the long-run impact of the (lagged) stock of domestic patents on TFP still is small and insignificant (the coefficient is 0.1 with a standard error of 0.1). We also find, as in equation (21), the trend term to be highly significant in explaining TFP (the coefficient is 0.009 with a standard error of 0.001). For more detailed results, see Abdih and Joutz (2005).

The result that the impact of the (lagged) stock of domestic patents on the flow of new domestic patents is unity is consistent with Porter and Stern (2000) and Romer (1990). This result supports the coefficient of 1.4 on *stp11* in equation (20): When the stock of total (domestic + foreign) patents is used to measure the stock

of knowledge, researchers have a larger pool of knowledge to draw on, therefore the spillover effect is stronger than the case in which the stock of domestic patents is used.¹³

We still need to reconcile these results with the time-series evidence presented by Jones (1995b). Romer's (1990) model assumes a spillover parameter of unity with the implication that the growth rate of the stock of knowledge is proportional to the R&D scientists and engineers. Jones (1995b) directly measured the stock of knowledge using TFP and rejected Romer's (1990) assumption of a spillover parameter of unity on the basis of the observed weak relationship between TFP growth and the number of R&D scientists and engineers: The number of R&D scientists and engineers has been trending strongly upward over the postwar period with no apparent benefit in terms of faster TFP growth. Jones argued that this weak relationship is therefore an indication of weak knowledge spillovers.

The empirical results presented in this paper suggest that once patents are used to measure knowledge, and hence the knowledge production function is directly estimated, the knowledge spillover parameter is unity or larger. That is, R&D scientists and engineers appear to have greatly benefited from the knowledge and ideas discovered through prior research. Therefore, the observed weak relationship between TFP growth and R&D scientists and engineers documented by Jones is not necessarily inconsistent with the presence of large knowledge spillovers once knowledge is measured using patent statistics.

The results of this paper point to an alternative explanation of this weak relationship: The knowledge that R&D scientists and engineers produce, as proxied by the patent stock, seems to have had only a limited impact on measured TFP. But why is this so? One possible interpretation is that knowledge is what ultimately matters for growth. Patents may be an imperfect proxy for knowledge, and this is why the effect of patents on TFP is insignificant. In other words, measurement error in the patent stock may be driving this result.¹⁴ In a simple OLS-type regression, measurement error could bias the estimated coefficient of the patent stock downward and reduce its precision. However, in a dynamic approach like that of

¹³In Romer (1990), Grossman and Helpman (1991b), and Aghion and Howitt (1992), R&D is the sole source of long-run growth. Although these models are consistent with the large intertemporal knowledge spillovers found in this paper, they are probably inconsistent with the fact that the trend term is highly significant in explaining TFP, in the sense that the trend may be capturing variables or innovations (apart from those associated with R&D) that influence TFP.

¹⁴Measurement error could arise because not all knowledge is patentable or will be patented. For example, firms may use secrecy as an alternative means of appropriation. Measurement error could also arise because our patent measure does not account for changes over time in the distribution of average patent quality. We are aware that patents are heterogeneous in how much they contribute to knowledge. An ideal measure would probably weigh patents by their "importance" or "quality." The work of Jaffe, Trajtenberg, and Henderson (1993) and Jaffe, Trajtenberg, and Fogarty (2000) has attempted to do that at the microlevel, usually for a sample of patents in a cross-sectional framework. Patent citations have been used as a measure of quality or importance in that work. However, to our knowledge, there has been no attempt to address or measure how the distribution of patent quality has evolved at the aggregate level over time for the entire postwar period in the United States. The data are simply not available. This probably explains why virtually all aggregate growth studies that use patents as an indicator of technology/knowledge have relied on patent counts following the tradition of Griliches (1990). Our paper follows this tradition.

Johansen (1988 and 1991)—from which our estimate is derived—this effect is minimized because lagged variables act like instruments. To investigate this issue further, we omitted patents altogether from our empirical model and tested for the existence of a “reduced form” cointegration relationship between log TFP and log R&D scientists and engineers, starting from a bivariate VAR model. We found no evidence of cointegration between log TFP and log R&D scientists and engineers. This is consistent with our more “structural” approach and indirectly suggests that our finding of no statistically significant effect of the patent stock on TFP is not driven by the claim that patents are a poor proxy for knowledge. That is, with or without patents, R&D seems to have an insignificant effect on TFP.

A second and more plausible explanation of the insignificant impact of the stock of knowledge on TFP was mentioned in Section III. We argued that the impact of knowledge on TFP is complex and slowly diffusing even in the long run, and therefore the diffusion process of the sharp increase in knowledge growth since the mid-1980s may not have had its full effect on measured TFP until the late 1990s, just beyond our sample.

Comin and Mulani (2005) provide a third explanation. They developed a model in which productivity growth increases with the development of two types of innovations. The first is standard R&D-driven innovations for which firms can invent a new product or improve the quality of an existing product. These innovations are typically sector-specific and patented. Second, firms can develop what the authors call general innovations, such as the mass production system, mass customization, testing practices and preproduction planning, new personnel and accounting practices, financial innovations, and new credit instruments such as the credit card and so on. This type of innovation is hard to patent and measure.

Comin and Mulani (2005) argue that, in equilibrium, there is a negative relationship between the resources spent on R&D and those spent on the development of general innovations: R&D leads to turnover in market leaders and to a decline in the value of leading firms. This decreases the private return to a general innovation and dampens its development. Therefore, a shock in the economy that increases investment in R&D may induce a decline in the arrival rate of general innovations. Productivity growth increases with the development of both standard R&D-driven and general innovations. Because an increase in R&D intensity leads to a decline in the rate of development of general innovations, the relationship between R&D and productivity growth is ambiguous at the aggregate level.

Our results are consistent with Comin and Mulani’s (2005) prediction of a lack of a clear relationship between R&D and TFP at the aggregate level. Our patent measure could be thought of as capturing standard R&D-driven innovations. However, it does not include general innovations, which are extremely difficult to measure as the authors argue. To the extent that general innovations have important aggregate productivity gains and are negatively correlated with our patent measure, it is not surprising that patents would have a weak and insignificant impact on TFP. That is, the impact of our patent measure on TFP is dampened by the negative impact a decline in general innovations might have on TFP. Comin and Mulani argue that the negative correlation between R&D-driven innovations and general

innovations is extremely hard to test. But their idea is theoretically appealing and could offer a plausible explanation of our empirical result.

V. Reconciliation and Conclusion

This paper began with a theoretical presentation of R&D-based models of economic growth. It showed that this class of models derives long-run growth through the accumulation of knowledge/technology. Knowledge, in turn, is a variable whose evolution is modeled endogenously. At any point in time, the rate of production of new knowledge depends positively on the existing stock of knowledge and the number of research workers.

The dependence of new knowledge on the existing stock captures knowledge spillovers. A large magnitude for knowledge spillovers implies that long-run growth is potentially dependent on policy variables such as subsidies to R&D (Romer, 1990). Conversely, a small magnitude for knowledge spillovers implies that long-run growth is invariant to policy (Jones, 1995a). Hence, empirically assessing the strength of knowledge spillovers is important, at least from a policy perspective.

Extensive empirical microliterature looks at different kinds of spillovers from R&D activities. These include market spillovers, network spillovers, and knowledge spillovers (see, for example, Jaffe, Trajtenberg, and Fogarty, 2000). In this paper, we did not consider market and network spillovers. Rather, the analysis was confined to testing the magnitude of knowledge spillovers at the macrolevel, as postulated by the R&D-based growth literature. The existing empirical macroresearch on the magnitude of knowledge spillovers seems to suggest that these spillovers are small.

In this paper, we focused on estimating the parameters of the knowledge production function. This allowed us to directly assess the magnitude of knowledge spillovers. To achieve this goal, we exploited historical time series of patent applications to construct knowledge flows and stocks. Cointegration techniques were used for analyzing and estimating models that involve nonstationary data and the specified long-run relationships.

We found evidence for two potential long-run relationships. The first is interpreted as a knowledge production function; the second is a relationship that captures a positive dependence of TFP on the stock of knowledge. The results indicate that knowledge spillovers are large. The estimated long-run elasticity of new knowledge with respect to the existing stock of knowledge is unity, when the domestic stock of patents is used as a proxy for the knowledge stock, and greater than unity with a global measure of knowledge based on both domestic and foreign stocks of patents. However, the estimated long-run elasticity of TFP with respect to the stock of knowledge is positive, small, and imprecise.

Our empirical work contributes to understanding and reconciling some of the spillover effects and issues raised by Jones (1995b). These results seem to suggest that while research workers benefit greatly from “standing on the shoulders” of prior researchers, the knowledge that they produce seems to have complex and slowly diffusing impacts on TFP. Therefore, the results suggest a new interpretation of the empirical evidence. The observed weak relationship between TFP growth and the number of research workers that Jones found does not necessarily indicate that

knowledge spillovers are small. Rather, it could be due to the small impact knowledge has on TFP. This small impact may be only very partially attributed to measurement error in the patent stock, our proxy for knowledge. The application and embodiment of knowledge into productivity is complex and diffuses slowly (even in the long run). Alternatively, the positive impact of the patent stock (knowledge stock) on TFP may be dampened by the negative impact that a potential decline in the rate of development of general innovations might have had on productivity, similar to Comin and Mulani (2005). Further work can examine the transmission mechanism(s) between the knowledge production function and productivity.

APPENDIX

Data Construction and Sources

Stock of Total (Domestic and Foreign) Patent Applications

The U.S. Patent and Trademark Office provides information on the number of patent applications filed from 1840 to the present. These include patents for inventions, designs, and plants. The data are available online from the U.S. Patent and Trademark Office website at http://www.uspto.gov/web/offices/ac/ido/oeip/taf/h_counts.htm.

Patents for inventions, designs, and plants were added to get the total number of patent applications. The total number of patent applications is used to construct a patent stock measure following the methodology proposed by Joutz and Gardner (1996) whereby

$$Stock_t = (\text{total number of patent applications})_t + Stock_{t-1}(1-d), \quad (\text{A.1})$$

where d denotes the depreciation rate. To make the above formula operational, an initial stock needs to be estimated. We calculate the initial stock of patents using two alternative methodologies. In the first, we follow the procedure suggested by Coe and Helpman (1995) and calculate the initial stock as

$$Stock_0 = (\text{number of patent applications})_0 / (g + d), \quad (\text{A.2})$$

where g is the average annual growth rate of the number of patent applications over the period for which data are available (1840–1999). The $(\text{number of patent applications})_0$ denotes the number of patent applications in the first year for which data are available (1840), and $Stock_0$ is the stock for that year (1840). Alternatively, we simply assume that the initial stock in 1840 is equal to the number of patent applications in 1840. That is, we assume that prior to 1840 there was no “knowledge.”

With an assumed depreciation rate, usually 15 percent (see Griliches, 1989 and 1990), and an estimate for the initial stock in 1840, equation (A.1) is used to construct the subsequent patent stocks. Both procedures for calculating the initial stock yield virtually identical patent stock series over the postwar period, and, therefore, it does not matter which one is used.

Domestic Patent Applications

For 1940 to 1999, data are obtained from the following sources:

- Tabulations of the U.S. Patent and Trademark Office available online at www.uspto.gov/web/offices/ac/ido/oeip/taf/us_stat.pdf;

- Unpublished memorandum of P.J. Federico, the U.S. Patent and Trademark Office, January 18, 1961;
- *Journal of the Patent Office Society*, Vol. 44, No. 2 February 1964, page 168; and
- Commissioner of Patents and Trademarks, Annual Report, 1966, page 26, the U.S. Patent and Trademark Office.

For 1840 to 1939, data on domestic patent applications are not available. We follow Kortum and Lerner (1998) and proxy for the number of domestic patent applications by multiplying the total number of patent applications by the fraction of total patent grants issued to U.S. inventors.

Stock of Domestic Patent Applications

We cumulate the number of domestic patent applications into a stock measure using the perpetual inventory method with an assumed depreciation rate of 15 percent.

Scientists and Engineers Engaged in R&D

The data for the period 1979–97 are obtained from the National Science Foundation, *Science and Engineering Indicators–2000*. This source is available online at <http://www.nsf.gov/statistics/seind00/>. For years before 1979, the data are taken from Jones (2002) and Machlup (1962), who in turn obtain their data from the National Science Foundation.

Total Factor Productivity (TFP)

The data are for the private business sector of the U.S. economy. They were obtained from Larry Rosenblum at the Office of Productivity and Technology, the Bureau of Labor Statistics, the U.S. Department of Labor. TFP measures output per combined unit of labor and capital inputs. The capital input measures the service flows from the level of the physical capital stock. The latter includes equipment, structures, inventories, and land. The labor input is based on hours at work data, adjusted to allow for changes in the skill composition (as measured by education and experience) of the labor force.

REFERENCES

- Abdih, Yasser, and Frederick Joutz, 2005, “Relating the Knowledge Production Function to Total Factor Productivity: An Endogenous Growth Puzzle,” IMF Working Paper 05/74 (Washington: International Monetary Fund).
- Aghion, Philippe, and Peter Howitt, 1992, “A Model of Growth Through Creative Destruction,” *Econometrica*, Vol. 60 (March), pp. 323–51.
- Arora, A., and A. Gambardella, 1994, “The Changing Technology of Technological Change: General and Abstract Knowledge and the Division of Innovative Labour,” *Research Policy*, Vol. 23 (September), pp. 523–32.
- Bureau of Labor Statistics, 1989, “The Impact of Research and Development on Productivity Growth,” Bulletin No. 2331 (Washington).
- Coe, David T., and Elhanan Helpman, 1995, “International R&D Spillovers,” *European Economic Review*, Vol. 39 (May), pp. 859–87.
- Comin, Diego, 2004, “R&D: A Small Contribution to Productivity Growth,” *Journal of Economic Growth*, Vol. 9 (December), pp. 391–421.

- , and Sunil Mulani, 2005, “A Theory of Growth and Volatility at the Aggregate and Firm Level,” NBER Working Paper No. 11503 (Cambridge, Massachusetts: National Bureau of Economic Research).
- Cummins, Jason, and Giovanni Violante, 2002, “Investment-Specific Technical Change in the United States (1947–2000): Measurement and Macroeconomic Consequences,” *Review of Economic Dynamics*, Vol. 5 (April), pp. 243–84.
- Davidson, J. E. H., David F. Hendry, F. Srba, and S. Yao, 1978, “Econometric Modeling of the Aggregate Time Series Relationship Between Consumer’s Expenditure and Income in the United Kingdom,” *Econometric Journal*, Vol. 88, pp. 661–92.
- Greenwood, J., and M. Yorukoglu, 1997, “1974,” *Carnegie-Rochester Conference Series on Public Policy*, Vol. 46, pp. 49–95.
- Griliches, Zvi, 1989, “Patents: Recent Trends and Puzzles,” NBER Working Paper No. 2922 (Cambridge, Massachusetts: National Bureau of Economic Research).
- , 1990, “Patent Statistics as Economic Indicators: A Survey,” *Journal of Economic Literature*, Vol. 18, No. 4, pp. 1661–707.
- , Ariel Pakes, and Bronwyn H. Hall, 1987, “The Value of Patents as Indicators of Inventive Activity,” in *Economic Policy and Technological Performance*, ed. by P. Dasgupta and P. Stoneman (Cambridge, United Kingdom: Cambridge University Press).
- Grossman, Gene, and Elhanan Helpman, 1991a, *Innovation and Growth in the Global Economy* (Cambridge, Massachusetts: MIT Press).
- , 1991b, “Quality Ladders in the Theory of Growth,” *Review of Economic Studies*, Vol. 58 (January), pp. 43–61.
- Hausman, J., B. H. Hall, and Z. Griliches, 1984, “Econometric Models for Count Data with an Application to the Patents-R&D Relationship,” *Econometrica*, Vol. 52 (July), pp. 909–38.
- Hendry, David F., 1986, “Econometric Modeling with Cointegrated Variables: An Overview,” *Oxford Bulletin of Economics and Statistics*, Vol. 48, No. 3, pp. 201–12.
- Jaffe, Adam, and Josh Lerner, 2004, *Innovation and Its Discontents: How Our Broken Patent System is Endangering Innovation and Progress, and What to Do About It* (Princeton, New Jersey: Princeton University Press).
- Jaffe, Adam, M. Trajtenberg, and M. Fogarty, 2000, “Knowledge Spillovers and Patent Citations: Evidence From a Survey of Inventors,” *American Economic Review, Papers and Proceedings*, Vol. 90 (May), pp. 215–18.
- Jaffe, Adam, Manuel Trajtenberg, and Rebecca Henderson, 1993, “Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations,” *Quarterly Journal of Economics*, Vol. 108 (August), pp. 577–98.
- Johansen, Soren, 1988, “Statistical Analysis of Cointegrating Vectors,” *Journal of Economic Dynamics and Control*, Vol. 12, No. 2–3, pp. 231–54.
- , 1991, “Estimation and Hypothesis Testing of Cointegrating Vectors in Gaussian Vector Autoregressive Models,” *Econometrica*, Vol. 59, No. 6, pp. 1551–80.
- , and Katarina Juselius, 1990, “Maximum Likelihood Estimation and Inference on Cointegration, with Applications to the Demand for Money” *Oxford Bulletin of Economics and Statistics*, Vol. 52, pp. 169–210.
- Jones, Charles, 1995a, “R&D-Based Models of Economic Growth,” *Journal of Political Economy*, Vol. 103, No. 4, pp. 759–84.
- , 1995b, “Time Series Tests of Endogenous Growth Models,” *Quarterly Journal of Economics*, Vol. 110 (May), pp. 494–525.
- , 2002, “Sources of U.S. Economic Growth in a World of Ideas,” *American Economic Review*, Vol. 92 (March), pp. 220–39.

- , and J. Williams, 1998, “Measuring the Social Return to R&D,” *Quarterly Journal of Economics*, Vol. 113, pp. 1119–35.
- Joutz, Frederick L., and Thomas A. Gardner, 1996, “Economic Growth, Energy Prices and Technological Innovation,” *Southern Economic Journal*, Vol. 62, No. 3, pp. 653–66.
- Kortum, Samuel S., 1993, “Equilibrium R&D and the Patent-R&D Ratio: U.S. Evidence,” *American Economic Review*, Vol. 83, No. 2, pp. 450–57.
- , 1997, “Research, Patenting, and Technological Change,” *Econometrica*, Vol. 65 (November), pp. 1389–419.
- , and Josh Lerner, 1998, “Stronger Protection or Technological Revolution: What Is Behind the Recent Surge in Patenting?” *Carnegie-Rochester Conference Series on Public Policy*, Vol. 48, pp. 247–304.
- Krusell, Per, 1998, “Investment-Specific R and D and the Decline in the Relative Price of Capital,” *Journal of Economic Growth*, Vol. 3 (June), pp. 131–41.
- Machlup, Fritz, 1962, *The Production and Distribution of Knowledge in the United States* (Princeton, New Jersey: Princeton University Press).
- National Science Foundation, 1998, *National Patterns of R&D Resources: 1998* (Washington).
- Nelson, C. R., and C. R. Plosser, 1982, “Trends and Random Walks in Macroeconomic Time Series: Some Evidence and Implications,” *Journal of Monetary Economics*, Vol. 10, No. 2, pp. 139–62.
- Perron, Pierre, 1989, “The Great Crash, The Oil Price Shock, and The Unit Root Hypothesis,” *Econometrica*, Vol. 57 (November), pp. 1361–401.
- Porter, Michael E., and Scott Stern, 2000, “Measuring the ‘Ideas’ Production Function: Evidence from International Patent Output,” NBER Working Paper No. 7891 (Cambridge, Massachusetts: National Bureau of Economic Research).
- Romer, Paul M., 1990, “Endogenous Technological Change,” *Journal of Political Economy*, Vol. 98 (October), pp. S71–S102.
- Samaniego, Roberto, 2005, “R&D Growth: The Missing Link?” Department of Economics Working Paper (Washington: The George Washington University).
- Thompson, Peter, and Melanie Fox-Kean, 2005, “Patent Citations and the Geography of Knowledge Spillovers: A Reassessment,” *American Economic Review*, Vol. 95 (March), pp. 450–60.