

R&D AND GROWTH: THE MISSING LINK?

ROBERTO M. SAMANIEGO

George Washington University

The presumption that R&D is a key driver of economic growth is difficult to reconcile with empirical evidence. For example, in most studies, which identify technical change with total factor productivity (TFP), the link between TFP and measures of knowledge is found to be weak.

This paper shows that a reconciliation may be possible in a model where R&D contributes to growth through *investment-specific technical change*. Such a model predicts that the empirical link between knowledge and productivity would be weak even if the generation of knowledge is the predominant factor of economic growth. The paper also shows that estimates of the production function for knowledge using patent data may be biased.

Keywords: R&D, Ideas' Production, Endogenous Growth, Patent Data, Investment-Specific Technical Change

1. INTRODUCTION

An extensive literature links economic growth to increases in technical knowledge. Although such models of *ideas-based growth* are theoretically appealing, empirical work tends not to support a significant role for R&D in the growth process. First, the link between measures of knowledge and productivity is weak. Second, estimates point to the presence of constant or even *increasing* returns to the production of ideas, which has the counterfactual implication that rates of economic growth should increase with the population size.¹

This paper proposes that a quantitatively important role for ideas in the process of economic growth is easily reconciled with these empirical “puzzles” if ideas are *investment-specific*. In this case, ideas contribute to economic growth through the factor accumulation process, and do not enter TFP directly. Measured TFP

I would like to thank the Editor, the Associate Editor, and an anonymous referee for suggestions that substantially improved the paper. I am indebted to Gianluca Violante for providing me with data on the relative price of capital. I also am grateful to Anna Ilyina, Fred Joutz, and to participants at the 2004 Midwest Macro Meetings at Iowa State University; the 2004 Annual Meetings of the Society for Economic Dynamics at New York University, La Pietra; the 2004 Macro International Seminar at George Washington University; the 2004 Macroeconomics Seminar at Indiana University—Bloomington; and the 2005 Meetings of the Society for Computational Economics at George Washington University and the Board of Governors of the Federal Reserve, for suggestions and comments. All errors are the author's. Address correspondence to: Roberto M. Samaniego, George Washington University, Department of Economics, 1922 F St. NW #208, Washington, DC 20052, USA; e-mail: roberto@gwu.edu.

may then display only a weak relationship to indicators of new knowledge, even if technical knowledge were the predominant factor of economic growth.

Several authors identify investment-specific technical change using the relative price of capital, which has declined steeply since 1947.² If ideas lead to growth through investment-specific technical change, this price is itself an indicator of the quantity of economically relevant knowledge in use. I use price data to construct an implicit series for the knowledge stock, and compare it with a measure derived from a more traditional indicator: patent activity. The two series do indeed co-move.³ However, there is also evidence of a change over time in the relationship between ideas and patents, suggesting that patent “quality” has varied—or, more broadly, that patent data may not adequately proxy for the stock of ideas. Once this is taken into account, the paper shows that mismeasurement can lead to the appearance of increasing returns in the production function for ideas even if they are not present. Indeed, the results are supportive of the existence of *decreasing* returns, with an elasticity possibly as low as [0.2, 0.6]. In this way, the empirical evidence is in fact consistent with a simple aggregate framework in which economic growth is driven by the production and implementation of new knowledge.

Section 2 lays out the empirical findings that the model will address. Section 3 describes the model framework, and Section 4 studies the relationship between patent data and the relative price of capital. Section 5 discusses the empirical “puzzles,” and Section 6 concludes with a discussion of the robustness of the results.

2. EMPIRICAL CONTEXT

Empirical work to identify the macroeconomic relationship between research and growth has focused on linking measures of research activity either to total factor productivity or to long-run patterns of economic growth. This work typically begins by identifying a measure of technical knowledge. Let q_t be the measure of new knowledge in period t . The stock of knowledge T_t then evolves according to the equation

$$T_{t+1} = (1 - \delta_T)T_t + q_t, \quad (1)$$

where δ_T is the rate at which ideas “depreciate.”⁴

The key to ideas-based growth models is that T_t is an important determinant of macroeconomic variables, and that it changes over time as the result of volitional research activity by agents. Thus, the aggregate behavior of the economy is given by *two* production functions. First, the production function of aggregate output y_t depends upon T_t , in addition to aggregate capital k_t and labor n_t

$$y_t = z_t T_t^\theta k_t^{\alpha_k} n_t^{1-\alpha_k}, \quad (2)$$

where T_t is a factor of TFP, and z_t is a productivity residual. Second, the quantity of new knowledge is given by an *ideas production function*

$$q_t = s_t T_t^\phi x_t^\psi, \quad (3)$$

where x_t indicates the input of resources into research, and s_t is a residual analogous to z_t . Empirically, the variable s_t captures the fact that measures of the input into research display far less low-frequency variation than measures of the output. It grows on average by a factor γ_s .

Estimates of the parameters of equations (2) and (3) tend to yield results that are challenging for ideas-based growth models. First, using patent applications as an indicator of knowledge T_t , Porter and Stern (2000) and Abdi and Joutz (2006) find that the contribution of ideas to total factor productivity is small. Point estimates of θ lie in the range [0.05, 0.2], and often lack statistical significance.

The second challenge concerns the parameter ϕ . A value of $\phi > 0$ implies that past ideas are useful for the production of new ones, whereas a value of $\phi < 0$ suggests that research uncovers the ideas that are easiest to find, so that discovery becomes progressively more difficult. These effects are known respectively as “standing on shoulders” and “fishing-out.” Endogenous growth models of this class are typically constructed so that $\phi = 1$: for example, this is true of the model in Romer (1990), in which increases in capital variety drive growth, and of Krusell (1998), in which ideas lead to growth through investment-specific technical change as here. However, if $\phi = 1$, then the growth rate depends on the population size (or, more generally, the size of the economy). Jones (1995) strongly rejects any empirical growth-population link in postwar U.S. data: population rises monotonically, whereas growth rates display no upward trend, implying that $\phi < 1$. Greenwood and Jovanovic (2001) argue that this constitutes a key empirical shortcoming of ideas-based growth models.⁵ Direct empirical estimates of the ideas production function have recently become available—see Porter and Stern (2000) and Abdi and Joutz (2006). The estimates suggest that returns are in fact close to constant or even *increasing* over the postwar period—in other words, $\phi \gtrsim 1$. It is unclear how to square this result with the absence of accelerating growth.

A third empirical puzzle also regards equation (3). Estimates of the ideas’ production function generally detect a downward trend in s_t . This decline is robust to a diversity of approaches. For example, Caballero and Jaffe (1993) use a selection of citation-weighted patent grants. Porter and Stern (2000) use aggregate patent application data, also identifying the effects of international patenting. Abdi and Joutz (2006) also use patent application data, simultaneously estimating the aggregate production and ideas’ production functions. The downward trend is perplexing, as it lacks a theoretical basis. For example, if it is the case that “easy” ideas are discovered first, so that ideas become progressively harder to uncover, this is precisely the “fishing out” hypothesis, which should be reflected in a negative value of ϕ : it should depend on the number of ideas that have already been uncovered, not on the *date*.

To sum up, the empirical evidence on ideas-based growth models is puzzling, and has difficulty preserving a central role for the production of knowledge in the process of economic growth. The remainder of this paper asks whether these

findings can in fact be reconciled with such a role within the context of a simple aggregate model.⁶

3. THEORETICAL FRAMEWORK

Consider now an aggregate framework with the same structure as (1), (2), and (3). The difference is that technical knowledge T_t does not enter the production function directly. Hence, the production function for output is

$$y_t = z_t k_t^{\alpha_k} n_t^{1-\alpha_k}, \tag{4}$$

where output may be used for consumption c_t , investment⁷ i_t or research input x_t :

$$y_t \geq c_t + i_t + x_t. \tag{5}$$

Instead, technical knowledge leads to economic growth through *investment-specific technical change*. Thus, it appears in the *capital accumulation equation*:

$$k_{t+1} = (1 - \delta_k)k_t + T_t i_t, \tag{6}$$

where δ_k is the rate of physical depreciation.

Investment-specific technical change is a feature of many macroeconomic models—for example, Greenwood et al. (1997). The hallmark of such models is that the price of capital relative to final output (denoted p_t) is the inverse of the extent of technical knowledge dedicated solely to the production of capital:

$$T_t = p_t^{-1}. \tag{7}$$

Krusell (1998) develops a model in which investment-specific technical change is the endogenous result of R&D activity, so that it is the accumulation of new knowledge that is responsible for the observed decline in p_t in postwar U.S. data.

In the remainder of the paper, I refer to the class of ideas-based growth models in which new knowledge leads to growth through equation (6) rather than through equation (2) as the *embodied knowledge* growth framework, or the EKG framework.⁸

If, as suggested by equation (7), investment-specific technical change is the endogenous result of R&D activity, the dynamics of the relative price of capital p_t and the stock of economically useful knowledge T_t should be related.

The long-run properties of the EKG framework are as follows. Let γ_χ be the growth factor for any variable χ . Equation (7) implies that, on a balanced growth path,

$$\gamma_T = \gamma_p^{-1}. \tag{8}$$

Combining (4) and (6), the relationship between ideas, TFP, population and output growth is given by

$$\gamma_y = \gamma_z^{\frac{1}{1-\alpha_k}} \gamma_T^{\frac{\alpha_k}{1-\alpha_k}} \gamma_n, \tag{9}$$

as, from (5), investment and R&D input must grow at the same rate as output along a balanced growth path.

Independently, (3) yields an expression for the growth rate of ideas itself, in terms of other factors.

$$\gamma_T = \gamma_s^{\frac{1}{1-\phi}} \gamma_y^{\frac{\psi}{1-\phi}}. \tag{10}$$

Eliminating γ_T using (9) and (10), the growth rate of final output in terms of the exogenous variables is

$$\gamma_y = (\gamma_z^{1-\phi} \gamma_s^{\alpha_k} \gamma_n^{(1-\alpha_k)(1-\phi)})^{\frac{1}{(1-\alpha_k)(1-\phi)-\alpha_k\psi}}. \tag{11}$$

Although (11) may appear complicated, it captures the notion that the production of ideas is an important channel of growth—although ultimately, as in Jones (1995), long-run growth is driven by exogenous factors such as population change.

A necessary condition for balanced growth assuming that $\gamma_z \geq 1$, $\gamma_s \geq 1$ and $\gamma_n \geq 1$ is that

$$\psi < \frac{(1 - \alpha_k)(1 - \phi)}{\alpha_k}, \tag{12}$$

which will be satisfied if either of α_k , ψ or ϕ is sufficiently small. In this event, there exists a sufficient degree of concavity that the feedback from ideas production back into output production is limited. Notably, it does *not* hold as $\phi \rightarrow 1$, which is the Romer (1990) model. This result is analogous to that of Jones (1995): if there is population growth then no balanced growth path exists in an ideas-based growth model without decreasing returns. Nonetheless, large values of ϕ may still be consistent with balanced growth, provided that $\gamma_z < 1$ and/or $\gamma_s < 1$. Greenwood et al. (1997) do in fact argue that, at least since the 1970s, $\gamma_z < 1$, whereas Griliches (1990), Porter and Stern (2000), and Abdih and Joutz (2006) find that $\gamma_s < 1$. Hence, the structure of the model is agnostic as to the value of ϕ .

4. QUANTITATIVE ANALYSIS

The key trait of the EKG framework is that changes in the relative price of capital p_t reflect the implementation of economically useful ideas. In this case, equation (7) suggests that there should be an observable empirical relationship between the relative price of capital p_t and indicators of the quantity of economically useful knowledge.

To put it another way, the model suggests two ways to construct an index of the quantity of economically useful ideas:

1. using a direct measure of knowledge, such as patent applications,⁹ as in the related literature;
2. using the inverse of the relative price of capital, p_t^{-1} .

Constructing the patent-based series for the stock of knowledge uses equation (3), which requires a value for the depreciation rate for ideas δ_T . We will examine




FIGURE 1. Correlation between the growth rates of the two constructed measures of the knowledge stock, for different values of δ_T . Data are from the USPTO and Cummins and Violante (2002).

different values of this parameter, and ask whether the patent-based series and the price-based series co-move appreciably. If they do, this supports the EKG framework—provided they do so for reasonable parameter values. Wilson (2002) finds support in industry cross-section for the identification of capital-embodied R&D with investment-specific technical change: finding such a link in aggregate time-series would be significant.

There are two ways in which the data lend support to the EKG framework. First, equation (8) suggests that the values of γ_T and γ_p^{-1} should be identical. This is indeed the case. As measured by patent applications at the USPTO (1947–2000) $\gamma_T = 1.0261$. By contrast, as measured by the relative price of capital in Cummins and Violante (2002) over the same time period, $\gamma_p^{-1} = 1.0262$. Second, the two time series are positively correlated regardless of δ_T —see Figure 1. Notably, this co-movement is the strongest at empirically reasonable values of δ_T —Nadiri and Prucha (1996) find that $\delta_T = 12\%$, and Pakes and Schankerman (1984) find values up to $\delta_T = 26\%$. The correlation is maximized at $\delta_T = 14.5\%$, which is in the middle of this range. In what follows, we shall assume that $\delta_T = 12\%$: results are robust to reasonable variations in this value.

Figure 2 displays the growth rates of the two series. The price-based series is more volatile: this is at least partly because of known sources of measurement error,¹⁰ but also may reflect variations over time in the knowledge content of




FIGURE 2. Patent stock growth and growth in the (inverse) relative price of capital. Correlation: 35%. Data are from the USPTO and Cummins and Violante (2002).

patents, an issue that has been raised in the literature¹¹ and which we will address in detail later. Figures 3 and 4 display the growth rates of the two series once more, this time Hodrick-Prescott filtered with smoothing parameters that are “small” to remove only very high-frequency movements. Notably, very little smoothing is necessary to raise the correlations between the two series even further. Observe that both measures of the knowledge stock accelerate in the 1980s, beginning somewhat earlier in the price-based series. The rise in patent growth is documented by Kortum and Lerner (1998); that it coincides with a rise in the growth of the relative price of capital is consistent with their conclusion that the surge in patenting is not spurious but rather reflects an actual increase in the rate of innovation in the U.S. economy.

Figure 5 displays the cross-correlogram of the growth rate of the two knowledge stocks. There is some co-movement between the patent-based series and the price-based series, contemporaneously and with a two-year lag. This suggests that innovation leads to implementation fairly quickly, as found by Caballero and Jaffe (1993).

5. ON THE EMPIRICS OF IDEAS-BASED GROWTH

This section addresses the empirical “puzzles” described at the beginning of this paper, and their relationship to the EKG framework.

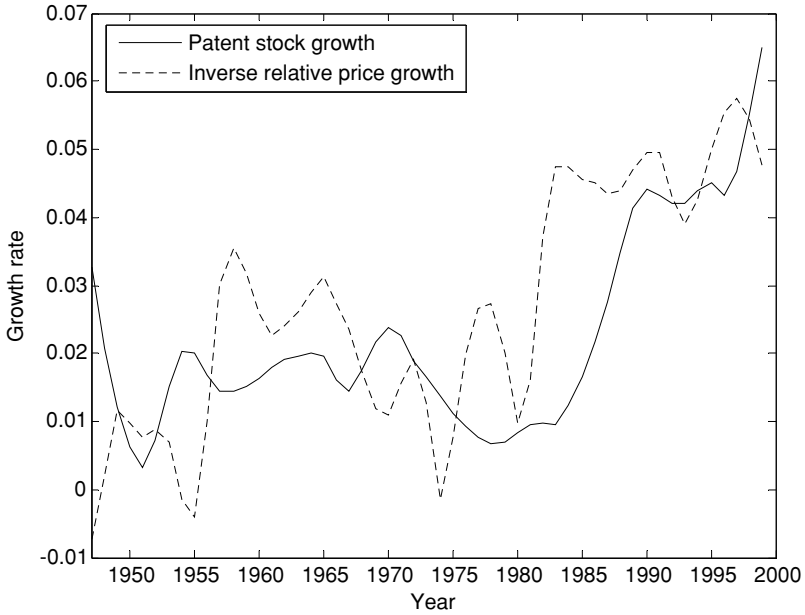


FIGURE 3. Patent stock growth and growth in the (inverse) relative price of capital, Hodrick-Prescott filtered growth rates with smoothing parameter $\lambda = 1$. Correlation: 57%. Data are from the USPTO and Cummins and Violante (2002).

5.1. Knowledge and Productivity

The first challenge confronted by ideas-based models is that the empirical literature generally finds a weak link between the knowledge stock and TFP. In the EKG framework, however, this may not be surprising—even if the entirety of economic growth is driven by the knowledge channel.

Let T_t be a measure of the knowledge stock. Several authors estimate the following specification:

$$\log y_t = \theta \log T_t + \alpha_k \log k_t + (1 - \alpha_k) \log n_t + \eta t + \varepsilon_t, \quad E_{t-1}[\varepsilon_t] = 0, \quad (13)$$

which, defining $z_t \equiv e^{\eta t + \varepsilon_t}$, is equivalent to the production function (2). Typically, estimates of θ are very low, ranging from 0.05 to 0.2 and sometimes lacking statistical significance.¹²

Assume that the EKG framework is a correct representation of the world. Consider an economist who ignores the presence of investment-specific technical change, and who wishes to identify TFP using a standard aggregate model. The aggregate production function and law of motion for capital in such a model




FIGURE 4. Patent stock growth and growth in the (inverse) relative price of capital, Hodrick-Prescott filtered growth rates with smoothing parameter $\lambda = 10$. Correlation: 71%. Data are from the USPTO and Cummins and Violante (2002).

will be:

$$y_t = \hat{z}_t \hat{k}_t^{\alpha_k} n_t^{1-\alpha_k}, \tag{14}$$

$$\hat{k}_{t+1} = (1 - \hat{\delta}_t) \hat{k}_t + i_t, \tag{15}$$

where \hat{k}_t is the capital stock derived according to (15), $\hat{\delta}_t$ equals *economic* depreciation and \hat{z}_t is a residual. If economic depreciation is correctly measured, then the following are the relationships between the aggregates of the standard model and those of the present framework:¹³

$$\begin{aligned} \hat{k}_t &= k_t p_{t-1} \\ \hat{\delta}_t &= 1 - (1 - \delta) \left(\frac{p_t}{p_{t-1}} \right), \\ \hat{z}_t &= z_t \left(\frac{1}{p_{t-1}} \right)^{\alpha_k}. \end{aligned} \tag{16}$$

Let us refer to \hat{z}_t as “measured TFP.” These changes of variables are an equivalent manner in which to write down the present model, so long as the relationships between \hat{z}_t , z_t and p_t are kept in mind. By contrast, (16) is the accounting link that misattributes investment-specific to neutral technical change if




FIGURE 5. Cross-correlogram for growth in the patent-based knowledge stock T_t and in the price-based stock p_t^{-1} . Horizontal lines represent 95% confidence bounds under the hypothesis that the two series are uncorrelated. Data are from the USPTO and Cummins and Violante (2002).

investment-specificity is ignored. To see this, using equation (7), the relationship between knowledge T_t and measured TFP \hat{z}_t becomes

$$\hat{z}_t = z_t T_{t-1}^{\alpha_k}. \tag{17}$$

Thus, in the current framework, there are several reasons why the empirical link between ideas and aggregate TFP should be weak. First, there could be lags between inspiration and implementation, aside from the one-period lag¹⁴ in equation (17). Second, there may be measurement error—something we will address in more detail below. Most importantly, however, if it is true that the stock of ideas affects growth through investment-specific technical change, then (17) states that the “true” value of θ equals α_k : estimates of the contribution of research to productivity will be bounded by the capital share. Measured capital shares tend to be around 0.3 and, depending on the exact methodology, they can be as low as 0.2 (see Maddison, 1987). To conclude, if research leads to growth through investment-specific technical change, low estimates of θ are to be *expected*—even if the knowledge channel may account for a significant proportion of economic growth.

In fact, the contribution of research to economic growth through investment-specific technical change can be assessed using standard growth-accounting

TABLE 1. Statistics used in Growth Accounting

Variable	U.S. Data
γ_p	1.0262 ⁻¹
γ_T	1.0261
γ_y	1.0333
γ_n	1.0119
δ_T	12%
α_k	0.3

Sources: USPTO, Cummins and Violante (2002) and NIPA. See Appendix B for further details.

techniques. Equation (9) and Table 1 imply that, through investment-specific technical change, the production of knowledge can account for about 55% of U.S. economic growth over the postwar era. This is very close to the contribution of investment-specific technical change to growth found by Greenwood et al. (1997) via growth accounting using only equipment price data, and by Cummins and Violante (2002) using structures also but with a different methodology.¹⁵ That this number is so robust to different concepts of embodiment is remarkable.

5.2. The Production of Knowledge and Measurement Bias

The second challenge confronted by ideas-based growth models is that the empirical literature generally estimates large values of ϕ , something that contradicts the absence of observable “scale effects.” This puzzle can be resolved on the basis that these estimates are likely to be biased as a result of the use of patent data.

Let q_t be the true flow of new ideas as before, and let b_t equal the flow of new patent applications. Then, q_t can be decomposed according to

$$q_t = m_t b_t, \quad (18)$$

where m_t is a factor that relates the number of patents to the number of ideas, and captures the extent of measurement error inherent in using patent data as an indicator of new knowledge. Denote m_t as average “patent quality.” It cannot be distinguished whether a given change in m_t is due to patents being of higher or lower informativeness, or due to changes in the fraction of ideas that are in fact patented. However, in either case, any variation in m_t implies that knowledge has grown by more (or by less) than is indicated by patent data.

Let γ_m be the trend in m_t . If $\gamma_m = 1$, then patents should serve as a reasonable index of ideas and m_t merely represents noise. There is, however, evidence that m_t may have increased over time. Recall that papers that estimate the ideas’ production function typically use patent applications as an indicator of new ideas,

including a time trend among their regressors. They estimate variants of the following equation:

$$\log b_t = \mu t + \phi \log T_t + \psi \log x_t + \varepsilon_t, \tag{19}$$

Griliches (1990) and Abdih and Joutz (2006) both find a significant negative time trend, with μ ranging between -1% and -2.3% . Porter and Stern (2000) also find a negative time trend for many of their specifications, averaging about -3% .¹⁶

Equation (19) is equivalent to the present setup, net of a relabeling of the variables. In the ideas production function (3), set $q_t = m_t b_t$. Taking logarithms, this becomes simply

$$\log b_t + \log m_t = \log s_t + \phi \log T_t + \psi \log x_t \tag{20}$$

Combining (19) and (20) yields $m_t = s_t e^{-(\mu t + \varepsilon_t)}$ so that

$$\mu = \log \gamma_s - \log \gamma_m. \tag{21}$$

Thus, the measured time trend μ cannot distinguish between two factors: changes in the effectiveness of research over time γ_s , and changes in patent quality γ_m . In particular, if $\gamma_s \approx 1$, then the time trend reflects not a downward trend in the ideas’ production function but, instead, an *upward* tendency in patent quality.

If it is true that the patent stock underestimates growth in the stock of ideas ($\gamma_m > 1$), then estimates of ϕ using patent data will be *biased upward*. To see this, recall that the ideas’ production function and the ideas’ accumulation equation can be written:

$$\log b_t + \log m_t = \log s_t + \phi \log T_t + \psi \log x_t, \tag{22}$$

$$T_{t+1} = b_t m_t + (1 - \delta_T) T_t. \tag{23}$$

Observe that m_t appears in (22) as a time trend, and also in (23) as a “correction” to measuring the quantity of new ideas using patents.

By contrast, the empirical implementation of these equations is

$$\log b_t = \mu t + \hat{\phi}_{\text{measured}} \log B_t + \psi \log x_t + \varepsilon_t, \tag{24}$$

$$B_{t+1} = b_t + (1 - \delta_T) B_t. \tag{25}$$

Here B_t is the patent stock, and $\hat{\phi}_{\text{measured}}$ is the estimate of ϕ obtained when using patent data as an indicator of ideas.

Writing the long-run forms of equations (22) and (24) and equating their right-hand sides yields

$$\phi \log \gamma_T = \mu + \hat{\phi}_{\text{measured}} \log \gamma_b - \log \gamma_s + \log \gamma_m,$$

$$\phi = \hat{\phi}_{\text{measured}} \frac{\log \gamma_b}{\log \gamma_T},$$

TABLE 2A. Patent and Trend Data

Source	Period	γ_b	μ	$\hat{\phi}_{\text{measured}}$	Underlying ϕ
Porter and Stern	73–93	1.0273	−0.03	0.84–1.2	0.40–0.57
Abdih and Joutz	48–97	1.0212	−0.23	1.4	0.67

γ_b is the growth rate of patent applications over the relevant sample period.
 μ is drawn from the results of each paper. Because γ_b is somewhat sensitive to the period of measurement, values are based on data for the sample period.

TABLE 2B. Price and Trend Data

Source	Period	γ_p^{-1}	μ	$\hat{\phi}_{\text{measured}}$	Underlying ϕ
Porter and Stern	73–93	1.0316	−0.03	0.84–1.2	0.42–0.61
Abdih and Joutz	48–97	1.0247	−0.23	1.4	0.73

γ_p is the growth rate of the relative price of capital over the relevant sample period.
 μ is drawn from the results of each paper.

since $\mu = \log\gamma_s - \log\gamma_m$. Finally, $\gamma_T = \gamma_b\gamma_m$. Hence, the “true” value of ϕ underlying their estimates is:

$$\phi = \hat{\phi}_{\text{measured}} \left(\frac{\log\gamma_b}{\log\gamma_b + \log\gamma_m} \right). \tag{26}$$

If $\gamma_m > 1$, then $\hat{\phi}_{\text{measured}}$ is biased upward.

Table 2A displays patent growth and the measured time trend over the periods for which each paper uses data, as well as the value of ϕ that is consistent with their estimates after being “corrected” with equation (26). The point estimates turn out to be consistent with values of ϕ that are well below unity.

Table 2B repeats the same exercise, this time replacing the patent stock measure with the (inverse) relative price of capital. This implements the premise that underlies the EKG framework: that the relative price of capital is an indicator of the quantity of economically useful ideas. Remarkably, the “corrected” values of ϕ turn out to be very close to those obtained using patent data in Table 2A—and all are below unity.

Another way to assess the sensitivity of estimates of the ideas production function to the measurement issues just raised is to use the growth accounting relationships (10) and (11) to obtain bounds on ϕ that are consistent with the estimates. Equation (10) does not impose the identification of ideas with investment-specific technical change, whereas equation (11) does.

First, on a balanced growth path $\gamma_q = \gamma_T$, so the long-run version of the ideas production function (10) can be rewritten:

$$\phi = 1 - \frac{\psi \log\gamma_y + \log\gamma_s}{\log\gamma_T}. \tag{27}$$

TABLE 3A. Predicted and Measured ϕ

ψ	0.3		0.6		$\hat{\phi}_{\text{measured}}$
	ϕ_{max}	ϕ_{min}	ϕ_{max}	ϕ_{min}	
Porter and Stern	1.7	0.83	1.3	0.65	0.84–1.2
Abdih and Joutz	1.6	0.78	1.1	0.55	1.4

γ_T is measured using the patent stock.

TABLE 3B. Predicted and Measured ϕ

ψ	0.3		0.6		$\hat{\phi}_{\text{measured}}$
	ϕ_{max}	ϕ_{min}	ϕ_{max}	ϕ_{min}	
Porter and Stern	1.7	0.68	1.3	0.37	0.84–1.2
Abdih and Joutz	1.6	0.60	1.1	0.19	1.4

γ_T is measured using the relative price of capital.

Given ψ , suppose that the patent stock is indeed an accurate index for the ideas stock—as does the empirical literature. In this case, over the long run, $\gamma_m = 1$, $\gamma_T = \gamma_b$ and $\gamma_s = e^\mu < 1$. Applying equation (27) delivers a value for ϕ that assumes that the measured trend represents decreases in s_T . Denote this value ϕ_{max} .

On the other extreme, suppose that the entire time trend is attributable to measurement error, so that $\gamma_T = \gamma_b \gamma_m$, $\gamma_s = 1$ and $\gamma_m = \gamma^{-\mu}$. This delivers a lower bound on ϕ that is consistent with the estimates, denoted ϕ_{min} .

The microeconomic estimates of ψ surveyed by Griliches (1990) range from 0.3 to 0.6, which turns out to be consistent with the estimates of other authors based on aggregate data.¹⁷ We will focus on values within this range.

Results are reported in Table 3A. Again, the column for ϕ_{max} lists the values that the model predicts the authors will find by taking the negative time trend at face value. These values are all larger than one. By contrast, if patent data systematically underestimate growth in the stock of ideas, the resulting values are all below unity.¹⁸

Table 3B repeats the same exercise, this time replacing the patent stock measure with the relative price of capital. There is a striking similarity between the results of Table 3A and Table 3B, the difference being that when the relative price of capital is used the estimates are consistent with values of ϕ that are even lower than when patent data are used.

Second, the same procedure can be applied to equation (11), the long-run relationship between parameters and aggregates under the EKG framework. It can be rearranged as follows:

$$\phi = 1 - \frac{\alpha_k \psi \log \gamma_y + \alpha_k \log \gamma_s}{(1 - \alpha_k) \log \gamma_y - \log \gamma_z - (1 - \alpha_k) \log \gamma_n} \quad (28)$$




FIGURE 6. Values of ϕ consistent with the estimates, given assumptions upon ψ and γ_s . Based on equation (28) and on data from the NIPA, the USPTO, and Cummins and Violante (2002), and on parameter values from Porter and Stern (2000) and Abdih and Joutz (2006).

Values of ϕ that are consistent with this equation are displayed graphically in Figure 6. Once more, when the measured trend is attributed to decreases in s_t , the model predicts that measures of ϕ will be large and exceed unity—exactly as in the empirical literature. By contrast, when the trend is attributed to measurement error, the value of ϕ falls in the range [0.23, 0.62]. Interestingly, this is close to the range of “corrected” estimates in Tables 2A and 2B.

Aside from the values themselves, the following results should be emphasized. First, both the estimates and growth accounting point to values of $\phi > 0$. This is consistent with the “standing on shoulders” effect that prior art is useful for the generation of current art, whereas the “fishing out” hypothesis is not substantiated. Second, growth accounting suggests that there are likely to be *decreasing returns* in the ideas’ production function: $\phi < 1$. The framework is in fact consistent with values of ϕ that exceed unity: however, this hinges on the presence of a decreasing time trend in the ideas production function. Although such a trend has been detected, there are several reasons to interpret it as an indicator of bias. Third, the empirical estimates themselves are also consistent with decreasing returns once allowance is made for this bias. Clearly, further quantitative work is desirable to obtain more reliable point estimates. However, “quality bias” appears to be as important an issue for the measurement of *ideas* as it is for the measurement

of *capital*. This paper shows how taking this into consideration points to values of ϕ that tally much better with theory, and that growth accounting within the EKG framework is consistent with this interpretation.

6. DISCUSSION

6.1. Trend or Mismeasurement?

Does the literature offer any guidance as to whether the measured time trends μ are best attributed to measurement error or to changes in the ideas' production function? To begin, the concern that aggregate patent data may not adequately proxy for the quantity of ideas has been raised in the literature previously. For example, Griliches (1990) points out that a potential problem with the use of patent data is that patents may vary over time in terms of *quality*. The results of Hall, Jaffe, and Trajtenberg (2000) suggest that patent quality varies substantially in cross-section, in that citation-weighted patents are more closely related to the market value of firms than are "raw" patents. Empirical work using the sum of large numbers of patents as an indicator assumes implicitly that aggregation will eliminate the influence of such heterogeneity: however, the distribution of patent quality (as measured by future citations or other measures of value) is so highly skewed that this aggregation result may not always hold—see Griliches (1990), Harhoff et al. (2003), and Leiva (2004). As a result, several authors conclude that there may be discrepancies between the stock of economically useful knowledge and the patent stock *over time* as well as in cross-section.

Second, it is difficult to see a theoretical basis for the existence of a downward time trend in the ideas production function. One might propose that a downward trend reflects the fact that R&D digs up ideas that are progressively more difficult to find. However, in this event, the volume of new ideas should depend on the quantity of ideas that have already been discovered, not on the date: this is precisely the "fishing out" hypothesis, and should be reflected in a negative estimate of ϕ , not in a time trend.

Third, empirically, a downward trend in the ideas production function seems at odds with the evidence of Kortum and Lerner (1998), who argue that innovativeness has *increased* over time, particularly since the 1980s. It is also at odds with Griliches (1986) and (1990), who finds no evidence of a decline in the returns to R&D. Cohen et al. (2000) find that there has been a tendency away from patenting and toward secrecy as a means of protecting intellectual property among U.S. manufacturing firms. Lanjouw and Schankerman (1999) find that adjusting patent data for forward citations and other measures of "quality" eliminates the apparent downward trend in the productivity of research spending. All of this points to the interpretation of μ as reflecting not changes in the productivity of ideas as such but, rather, changes in the empirical link between patents and ideas.




FIGURE 7. Patent quality m_t , as derived from equation (29). $\delta_T = 0.12$. Data are from the USPTO and Cummins and Violante (2002).

Finally, the EKG framework can be used to derive a series for the discrepancy between patents and new ideas m_t . From equations (1) and (18), it follows that

$$m_t = \frac{p_{t+1}^{-1} - p_t^{-1}(1 - \delta_T)}{b_t}. \quad (29)$$

Iterating on (29) with b_t as new patent applications yields the desired series, which is displayed in Figure 7. Patent quality is highly variable and does indeed display an upward trend, increasing on average by 1.0% each year. This provides further support for the interpretation of a time trends μ as the extent of measurement error underlying the use of patent data.

6.2. Ideas and Sector-Neutral Technical Change

The paper assumes that all ideas are investment-specific, and that none result in *sector-neutral* technical change. This contrasts with the empirical literature that assumes that *all* ideas result in neutral technical change. Nonetheless, it is reasonable to ask, what are the implications of assuming that only a *fraction* of knowledge is investment-specific?

Suppose in what follows that a fraction ζ of knowledge T_t is investment-specific, and that a fraction $(1 - \zeta)$ leads to sector neutral technical change:

1. Regarding the quantitative results of section 4, this simply multiplies the knowledge series by a constant, so that the results of the correlation analysis of growth rates that follows are not affected.
2. The results concerning measurement bias in estimates of ϕ are not affected either, for the same reason and because most of those results are independent of whether T_t and p_t^{-1} are identified. Consequently, none of the above results in fact depend on whether *all* knowledge is investment-specific: they all continue to hold so long as *part* of knowledge is investment-specific.
3. What of the puzzle regarding the link between TFP and ideas? Suppose there is a one-period lag between innovation and implementation as far as neutral technical change is concerned.¹⁹ Then, equation (17) becomes

$$\hat{z}_t = z_t[(1 - \zeta)T_{t-1}]^\theta [\zeta T_{t-1}]^{\alpha_k} \tag{30}$$

$$\Rightarrow \log \hat{z}_t = Z + \log z_t + (\theta + \alpha_k) \log T_t \tag{31}$$

where Z is a constant. In this case, a log regression of the measured residual \hat{z}_t on a lagged measure of knowledge T_{t-1} will yield a coefficient $\theta + \alpha_k$. If $\alpha_k = 0.3$, the low empirical estimates of the relationship between ideas and TFP point to values of θ close to zero.

In any case, if neutral technical change is a channel through which knowledge affects growth, one would expect a correlation between the residual z_t in equation (17) and the patent series T_t . Figure 8 does not display any significant




FIGURE 8. Cross-correlogram for growth in the patent-based knowledge stock T_t and in the residual z_{t+k} . 95% confidence bounds under the hypothesis that the two series are uncorrelated. Data are from the USPTO, the NIPA, and Cummins and Violante (2002).

correlation. This contrasts with Figure 5, which displays a clear correlation between T_t and p_t^{-1} , as the EKG framework would suggest.

To sum up, the results are robust to allowing only *part* of the ideas stock to lead to growth through investment-specific technical change. In addition, the sector-neutral channel is not supported by the aggregate data presented herein.

Perhaps this is to be expected. If firms have a choice between performing neutral or investment-specific R&D, they might prefer to conduct the latter if it is more easily appropriable. Investment-specific knowledge should be embodied in the capital goods that it is used to produce, and this detectability may make patent enforcement easier. The decision of whether or not to embody knowledge could be an interesting avenue for future research.

6.3. Concluding Remarks

This paper proposes a view of economic growth that is consistent with several empirical paradoxes regarding the link between knowledge and TFP, while preserving the centrality of the accumulation of knowledge as a factor of economic growth. The key is to identify the implementation of economically useful knowledge with investment-specific technical change. In this case, the absence of a strong relationship between measures of knowledge and TFP is to be expected: instead, it is co-movement among measures of knowledge and measures of investment-specific technical change that provides evidence of a link between ideas and growth. This identification, in combination with the issues raised involving the measurement of ideas, suggest that there may be positive but decreasing returns to old knowledge in the production of new ideas.

NOTES

1. See Romer (1990), Jones (1995), Porter and Stern (2000), and Abdih and Joutz (2006).
2. See Greenwood et al. (1997), Krusell (1998), and Cummins and Violante (2002).
3. Also in support of this account, Wilson (2002) finds cross-sectional evidence linking measured industry-level capital price declines to accumulated R&D in upstream capital goods.
4. Ideas may be superseded by others; certain avenues of research may be exhausted so that they cease to be important for the production of new ideas and goods; or they may simply be forgotten. That this occurs empirically is seen in that patent maintenance fees are not always paid—see Griliches (1990).
5. The implication was that ideas-based growth models should have weaker policy implications, because taxation or subsidy schemes could not affect growth rates, although in a richer framework Howitt (1999) demonstrates that this need not be the case. Independently of this debate, Jones (1995) argues that the extent of returns remains an important determinant of other model properties, such as transition dynamics.
6. Of course, one interpretation of the results is that patent data do not accurately reflect the aggregate quantity of technical knowledge. I defer a discussion of such measurement questions until Section 5.
7. The term “investment” here refers to foregone consumption goods that are used to make new capital goods.

8. Although the Krusell (1998) model articulates the key EKG relationships (7) and (6), it does not aggregate to the empirical equations (1) and (3). The Appendix describes a generalization of the Krusell (1998) model, which satisfies this property.

9. See Appendix B for further comments on the use of patent data.

10. In particular, the sharp dip in 1974 is related to price controls at the time—see Cummins and Violante (2002).

11. See, for example, Griliches (1990), Hall, Jaffe, and Trajtenberg (2000), Harhoff et al. (2003), and Leiva (2004).

12. See Griliches (1990) for a survey of microeconomic estimates, and Porter and Stern (2000) and Abdi and Joutz (2006) for macroeconomic estimates.

13. See Greenwood et al. (1997) and Krusell (1998).

14. Aware of this, empirical work often assumes such a lag. For example, Porter and Stern (2000) assume a three-year lag between patent applications and TFP. As for the empirical consistency of the assumed lag, in related models that also satisfy equation (16), Greenwood et al. (2000) and Fisher (2006) find that 30% of business cycle fluctuations and 50% of time-series variation in labor productivity can be attributed to variations in p_t , respectively.

15. Interestingly, in a calibrated model of the establishment life-cycle, Samaniego (2006) finds a broadly similar value for *plant-embodied* technical change.

16. Porter and Stern (2000) find annual trends between +3% to -10%, depending on the exact method and time period of analysis. I adopt a value of -3%, which is around the middle of the range and that is not very different from the values found by other authors.

17. It is simple to show that the model of ideas' production is equivalent to one in which each sector uses labor and capital rather than the numeraire. The corresponding specification is:

$$Q(T_{jt}, \bar{T}_t, k_t, n_t) = \left(T_{jt}^{1-\sigma} \bar{T}_t^\sigma \right)^\phi k_{jt}^\eta n_{jt}^\lambda$$

for the special case in which $\eta = \alpha_k \psi$ and $\lambda = (1 - \alpha_k) \psi$, this restriction being consistent with the absence of a trend in factor shares. Abdi and Joutz (2006) estimate an equation of the form $q_t = e^{\lambda t} T_t^\phi k_t^\eta n_t^\lambda$, estimating $\lambda = 0.21$. Setting $\psi = \frac{\lambda}{1-\alpha_k}$ implies that $\psi = 0.3$. Porter and Stern (2000) obtain values of λ between 0.21 and 0.45 when there is a year control present, which corresponds to the range $\psi \in [0.3, 0.64]$.

18. Growth in reported R&D spending has amounted to over 4% per year: if this is replaced for γ_y in equation (27) then the values of ϕ obtained are even lower.

19. If there is no lag, results are essentially the same but more complicated to derive.

20. Henceforth, let γ_χ denote the average growth factor for any variable χ_t .

21. Alternatively, define the opportunity cost of new capital goods in terms of the numeraire to be p_{jt} . Then, (7) emerges from this definition.

22. Positive population growth is necessary in order to address the question of scale-effects. the Krusell (1998) assumes that $\phi = 1$ and does not have population growth: if there were, that model would not have a balanced growth path.

REFERENCES

- Abdi, Yasser and Frederick Joutz (2006) Relating the knowledge production function to total factor productivity. *International Monetary Fund Staff Papers* 53, 242–271.
- Caballero, Ricardo and Adam B. Jaffe (1993) How high are the giant's shoulders: An empirical assessment of knowledge spillovers and creative destruction in a model of economic growth. *NBER Macroeconomics Annual* 8, 15–74. Cambridge, Mass.: MIT Press.
- Cohen, Wesley, Richard Nelson and John Walsh (2000) Protecting Their Intellectual Assets: Appropriability Conditions and Why US Manufacturing Firms Patent (or Not). NBER working paper w7552.
- Cooley, Thomas F. and Edward C. Prescott (1995) Economic growth and business cycles. In Thomas F. Cooley (ed.), *Frontiers of Business Cycle Research*, 1–30. Princeton: Princeton University Press.

- Cummins, Jason and Giovanni Violante (2002) Investment-specific technical change in the US (1947–2000): Measurement and macroeconomic consequences. *Review of Economic Dynamics* 5, 243–284.
- Fisher, Jonas D.W. (2006) The dynamic effects of neutral and investment-specific technology shocks. *Journal of Political Economy* 114, 413–451.
- Greenwood, Jeremy, Zvi Hercowitz and Per Krusell (1997) Long-run implications of investment-specific technological change. *American Economic Review* 87, 342–362.
- Greenwood, Jeremy, Zvi Hercowitz and Per Krusell (2000) The role of investment-specific technological change in the business cycle. *European Economic Review* 44, 91–115.
- Greenwood, Jeremy and Boyan Jovanovic (2001) Accounting for growth. In *New Developments in Productivity Analysis*, NBER Studies in Income and Wealth 63, 179–222. Chicago and London: University of Chicago Press.
- Griliches, Zvi (1986) Productivity, r&d, and basic research at the firm level in the 1970's. *American Economic Review* 76, 141–154.
- Griliches, Zvi (1990) Patent statistics as economic indicators: A survey. *Journal of Economic Literature* 28, 1661–1707.
- Hall, Bronwyn H., Adam B. Jaffe and Manuel Trajtenberg (2000) Market Value and Patent Citations: A First Look. NBER working paper w7741.
- Harhoff, Dietmar, Frederic M. Scherer and Katrin Vopel (2003) Citations, family size, opposition and the value of patent rights. *Research Policy* 32, 1343–1363.
- Howitt, Peter (1999) Steady endogenous growth with population and r&d inputs growing. *Journal of Political Economy* 107, 715–730.
- Jones, Charles (1995) R&d-based models of economic growth. *Journal of Political Economy* 103, 759–784.
- Kortum, Samuel and Josh Lerner (1998) Stronger protection or technological revolution: What is behind the recent surge in patenting? *Carnegie-Rochester Conference Series on Public Policy* 48, 247–304.
- Krusell, Per (1998) Investment-specific r&d and the decline in the relative price of capital. *Journal of Economic Growth* 3, 131–141.
- Lanjouw, Jean and Mark Schankerman (1999) The Quality of Ideas: Measuring Innovation with Multiple Indicators. NBER working paper w7345.
- Leiva, Fernando (2004) Pricing Patents Through Citations. Mimeo: University of Iowa.
- Maddison, Angus (1987) Growth and slowdown in advanced capitalist economies: Techniques of quantitative assessment. *Journal of Economic Literature* 25, 649–698.
- Nadiri, M. Ishaq and Ingmar R. Prucha (1996) Estimation of the depreciation rate of physical and r&d capital in the U.S. total manufacturing sector. *Economic Inquiry* 34, 43–56.
- Pakes, Ariel and Mark Schankerman (1984) The rate of obsolescence of patents, research gestation lags, and the private rate of return to research resources. In *R&D, Patents, and Productivity*, NBER conference report, 73–88. Chicago and London: University of Chicago Press.
- Porter, Michael E. and Scott Stern (2000) Measuring the “Ideas” Production Function: Evidence from International Patent Output.” NBER working paper w7891.
- Romer, Paul (1990) Endogenous technological change. *Journal of Political Economy* 98, S71–S101.
- Samaniego, Roberto (2006) Organizational capital, technology adoption and the productivity slowdown. *Journal of Monetary Economics* 53, 1555–1569.
- Wilson, Daniel J. (2002) Is embodied technology the result of upstream R&D? Industry-level evidence. *Review of Economic Dynamics* 5, 285–317.

APPENDIX A: DECENTRALIZED MODEL

This section describes a decentralized economy that fits within the EKG framework. As in Krusell (1998), in this framework the production of ideas generates economic growth

through investment-specific technical change and, with some minor modifications, the Krusell (1998) model aggregates to a production function and law of motion for ideas that is the same as those estimated in the empirical literature—equations (1) and (3). The interested reader may refer to that paper for further theoretical analysis.

A.1. OUTPUT

Time is discrete. There are two sectors, one that produces the numeraire (final) good and another that produces new capital goods. Output y_t is produced according to a Cobb-Douglas technology that uses labor n_t and a continuum of different types of capital k_{jt} , $j \in [0, 1]$.

$$y_t = z_t \left(\int k_{jt}^{\alpha_k} dj \right) n_t^{1-\alpha_k}, \tag{A.1}$$

where sector-neutral technical change z_t is stochastic, and grows by a factor γ_z on average each period.²⁰ Capital of type j commands a rental rate r_{jt} , and labor a wage w_t . Define aggregate capital $k_t = \int k_{jt} dj$.

Output has three uses in this model. It may be used for household consumption (c_t), as investment (i_{jt}) for transformation into capital goods of any type j , or used as an input into R&D (x_{jt}). The feasibility constraint is

$$y_t \geq c_t + \int i_{jt} dj + \int x_{jt} dj. \tag{A.2}$$

A.2. RESEARCH AND INVESTMENT-SPECIFICITY

Each type of capital j is produced by a monopolist, who also may perform R&D in order to increase her productivity. Let T_{jt} denote the quantity of investment-specific ideas relevant for the production of capital type j . Capital is accumulated according to the equation

$$k_{j,t+1} = (1 - \delta_k)k_{jt} + T_{jt}i_{jt}. \tag{A.3}$$

where δ_k is the rate of physical depreciation. Ideas evolve according to

$$T_{j,t+1} = (1 - \delta_T)T_{jt} + q_{jt}. \tag{A.4}$$

Define $T_t = \int T_{jt} dj$ to be the average level of knowledge across sectors, and let q_{jt} be the quantity of new ideas generated in sector j in period t , which is given by a production function

$$q_{jt} = s_t Q(T_{jt}, T_t, x_{jt}).$$

Q depends on the quantity of ideas in sector j , as well as T_t and physical input x_{jt} . Thus, there may be cross-industry spillovers. Krusell (1998) adopts the following functional form for Q :

$$Q(T_{jt}, \bar{T}_t, x_t) = (T_{jt}^{1-\sigma} T_t^\sigma)^\phi x_{jt}^\psi. \tag{A.5}$$

The parameter σ indicates the extent of cross-sectoral spillovers, whereas ϕ and ψ govern the return to old ideas and to research input in the production of new knowledge. This formulation is a generalization of the function in Krusell (1998), which assumes that $\phi = 1$. As mentioned earlier, the value of ϕ has been the subject of much debate, and as a

result the current model does not restrict it. Krusell (1998) also differs in that it sets $\delta_T = 1$. To strengthen the connection of the model to the related empirical literature, the paper does not restrict δ_T either. The value of σ is interesting for theoretical reasons, but does not affect the long run growth properties of the model. A third generalization of the Krusell (1998) model concerns population growth: see later.

The monopolists earn profits from renting their capital to the final goods' sector. Hence, the first-order condition of the final goods firms with respect to capital $r_{jt}(k_{jt}) = \alpha_k k_{jt}^{\alpha_k - 1} n_t^{1 - \alpha_k}$ represents an inverse demand function for each investment firm. Let V be the value function of a monopolist, which depends on the idiosyncratic and aggregate state variables. Then,

$$V(k_{jt}, T_{jt}, \mathbf{k}_t, \mathbf{T}_t) = \max_{i_t, x_t} \left\{ r_{jt}(k_{jt})k_{jt} - i_{jt} - x_{jt} + \frac{1}{1 + t} E_t V(k_{j,t+1}, T_{j,t+1}, \mathbf{k}_{t+1}, \mathbf{T}_{t+1}) \right\}, \tag{A.6}$$

subject to the production functions and laws of motion of capital and ideas.

The first-order condition of this problem with respect to investment is

$$1 = T_{jt} \frac{1}{1 + t} E_t V_1(k_{j,t+1}, T_{j,t+1}, \mathbf{k}_{t+1}, \mathbf{T}_{t+1}).$$

Define the price of capital p_{jt} as the marginal benefit of creating a new unit of capital, so that²¹

$$p_{jt} = \frac{1}{1 + t} E_t V_1(k_{j,t+1}, T_{j,t+1}, \mathbf{k}_{t+1}, \mathbf{T}_{t+1}).$$

Then the first-order conditions can be reformulated in terms of p_{jt} to yield the decentralized equivalent of (7):

$$T_{jt} = p_{jt}^{-1}. \tag{A.7}$$

A.3. HOUSEHOLDS

Closing the model requires a formulation of the household sector.

The population at date²² t is $N_t = \gamma_n^t$. There is a unit continuum of households characterized by the following dynastic utility function:

$$E_0 \sum_{t=0}^{\infty} \beta^t N_t \{\log c_t + \eta \log l_t\}, 0 \leq l_t \leq \Lambda, \tag{A.8}$$

where l_t is leisure and Λ is their time endowment. Consumption and labor are in per capita terms. Households own the capital stock and all firms. Their budget constraint is

$$c_t \leq w(\Lambda - l_t) + \Pi_t, \tag{A.9}$$

where Π_t represents any dividends they earn from firms that they own.

A.4. SYMMETRIC EQUILIBRIUM

Let us focus on an equilibrium in which all capital types are treated equally, so that $T_{jt} = T_t$ and $k_{jt} = k_t \forall j$. From inspection of (A.3) and (A.5), it is easy to show that the real sector and the knowledge sector aggregate to yield equations (2) and (3). Simply put, a recursive competitive equilibrium for this economy is a set of prices, allocations and decision rules that jointly satisfy the optimization, market-clearing, feasibility and rational expectations conditions at every date. The equilibrium definition is standard and is the same as in Krusell (1998), *mutatis mutandis*.

APPENDIX B: NOTES ON THE DATA

Patent data are from the United States Patent and Trademark Office (USPTO). Price p_t is the quality-adjusted price of capital relative to consumption of Cummins and Violante (2002), which encompasses both equipment and structures. Output and investment data are from the National Income and Product Accounts. Labor hours are from the Bureau of Labor Statistics, and population data are from the United States Census Bureau. The years covered are 1947–2000: all series are annual.

There are two alternative measures of the output of research activity: patent grants, and patent applications. The paper uses patent *applications* for the following reasons. The “application-to-grant” lag is about two years but varies widely over time, sometimes taking considerably longer. In fact, grants are strongly correlated with the number of patent examiners at the USPTO—see Griliches (1990). As a result, they are inadequate as contemporaneous indicators of ideas’ production activity. Moreover, at the point in time in which the patent is applied for, the applicant must have conceived of an innovation that, although possibly in need of further development and marketing, is close to the point of implementability. The model allows for a one-year lag at this stage.

Another option is R&D spending. The drawback is that it is a measure of the *input* into the ideas production function, not the output, and as such displays very little short-term variation. R&D spending may not constitute an accurate measure even of the input into the production of ideas: managers and other workers whose positions are not nominally connected with R&D may have patentable or otherwise economically useful ideas. Kortum and Lerner (1998) argue that there has been an acceleration of innovative activity in the United States in recent decades, attributed to increasingly active management of knowledge at the establishment level. A likely effect of such a change would be an increase in the proportion of business expenditures that are *classified* as R&D, as managers distinguish and target R&D activities. Hence, measured aggregate R&D growth may overestimate ideas growth.

Another alternative is the share of employment made up by scientists and engineers in R&D. Again, this is a measure of the input into R&D—not the output—and will suffer from some of the same problems as R&D spending, such as the fact that managers and other nonscientists may have economically useful ideas. It, too, displays very little high-frequency variation.

To initialize the two knowledge stocks (T_t and p_t^{-1}), I assume that the economy was on the balanced growth path in 1947: small deviations in this assumption do not affect results. Patent data are available as far back as the late 18th century: constructing patent data based on the entire series does not change results.