Investment-Specific Technical Change and the Production of Ideas

October, 2004.

Abstract

I argue that an aggregate model in which the generation of knowledge is an important factor of economic growth can be reconciled with several otherwise puzzling empirical findings on this link if knowledge affects output through *investment-specific technical change*. In the model, there may be a weak empirical relationship between measures of knowledge and total factor productivity even when the generation of knowledge is the predominant channel through which economic growth takes place. The results also suggest that intertemporal spillovers in the production of knowledge are likely to be small.

JEL Codes: E300, O300, O400, O470.

Keywords: ideas' production, quasi-endogenous growth, patent stock, investment specific technical change, price of capital.

1 Introduction

Are aggregate models empirically consistent with an important role for the generation of knowledge in the process of economic growth?

There is an extensive literature that links economic growth to increases in technical knowledge – known as the "ideas-based" growth literature. Although theoretically appealing, such models have suffered empirically. First, the link between measures of knowledge and productivity is weak. Second, several models in this class rely upon the existence of constant returns in the production of ideas, which has the counterfactual implication that rates of economic growth increase with the population size. Puzzlingly, direct estimates of the ideas production function point to the presence of constant or even *increasing* returns.[1]

In this paper, I argue that a simple macroeconomic model can be formulated in which 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 may then display only a weak relationship to measures of ideas, even if ideas are an important factor of economic growth.

Conversely, several authors identify investment-specific technical change (ISTC) using the relative price of capital, which has declined steeply since 1947. They argue that the contribution of ISTC to growth in the US has been significant, particularly since the 1970s. This begs the question: what has led to this precipitous price decline? The implementation of new economically useful ideas presents itself as a candidate explanation. Krusell (1998) explores some theoretical properties of a model of ideas-based growth via ISTC and, in support of this account, Wilson (2002) finds cross-sectional evidence linking measured industry-level capital price declines to accumulated research and development in upstream capital goods.

If ideas lead to growth through ISTC as reflected in the relative price of capital, then 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, which I then compare with a measure derived from a more traditional indicator: patent activity. I find that the two series do indeed co-move. However, I also find 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, I show that mismeasurement can lead to the appearance of increasing returns in the production function for ideas when in fact there are none. 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 outlines the relationship between ideas-based models and recent empirical findings. Section 3 lays out the model economy, and Section 4 studies its equilibrium properties. Section 5 calibrates the model to post-war US data. In Section 6 I discuss the implications of the model for the structure of the ideas' production function.

[1] See Romer (1990), Jones (1995), Porter and Stern (1998) and Abdih and Joutz (2003).

2 Empirical Context

This section lays out in brief the empirical findings that the model will attempt to reconcile. Empirical work to identify the macroeconomic relationship between research and growth has focussed upon linking measures of research activity either to total factor productivity or to long run patterns of economic growth. In both cases, results have proven challenging for ideas-based models of growth.

First, using patents applications as an indicator of new knowledge, Porter and Stern (2000) and Abdih and Joutz (2003) find that the contribution of ideas to total factor productivity is small. To put it another way, suppose that at date t output y_t depends upon a vector of inputs v_t with I elements; upon the quantity of accumulated technical knowledge T_t; and on a productivity residual z_t, through the aggregate production function

$$y_t = z_t T_t^{\theta} \prod_{i=1}^{I} v_{it}^{\alpha_i}. \tag{1}$$

When patent data is taken as an indicator of T_t, point estimates of θ lie in the range $[0.05, 0.2]$, and often lack statistical significance.

Second, ideas-based growth models imply that growth rates increase with the population size – unless there are decreasing returns to scale in the production of ideas. To be precise, let T_t be the stock of ideas at date t. Let q_t be the quantity of new ideas, and let x_t indicate the input of resources into research, with s_t as a residual analogous to z_t. Then, the *ideas production function* is

$$q_t = s_t T_t^{\phi} x_t^{\psi} \tag{2}$$

Parameter ϕ is widely referred to as the "intertemporal spillover." Theoretically, a value of $\phi > 0$ implies that past ideas are useful for the production of new ones, whereas it is also possible that $\phi < 0$ if past research has uncovered the ideas that were easiest to find and discovery becomes progressively more difficult – these effects are known respectively as "standing on shoulders" and "fishing-out." Endogenous growth models 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. This value of ϕ is important, as it implies that public policy towards private research could have growth effects. However, Jones (1995) strongly rejects any empirical growth-population link in post-War US 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.[2]

Given the interest in the value of ϕ, recently some direct empirical estimates of the ideas production function have become available – see Porter and Stern (1998) and Abdih and Joutz (2003). Their estimates suggest that returns are in fact close to constant or even *increasing* over the post-War period – in other words, $\phi \gtrapprox 1$. It is unclear how to square this result with the absence of accelerating growth.

[2] 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, as noted by Jones (1995), the extent of returns remains an important determinant of other model properties, such as transition dynamics.

A third puzzle also regards equation (2). Estimates of the ideas' production function generally detect a downward trend: s_t appears to decline over time. This finding 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. Abdih and Joutz (2003) also use patent application data, simultaneously estimating the aggregate production and ideas' production functions. The downward trend is puzzling, 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 over time, 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 has difficulty preserving a central role for the production of knowledge in the process of economic growth. In what follows, I ask whether these findings can in fact be reconciled with such a role within the context of a simple aggregate model.[3]

3 Theoretical Model

The objective is to develop a model which is as simple as possible an extension of a standard ideas-based growth model that allows knowledge to affect aggregates through ISTC, and to show that this innovation is sufficient to allow a reinterpretation of the reported empirical findings.

The model economy contains two sectors: a final goods sector and an investment sector. There is a continuum of different types of investment goods, and each type is produced by a monopolist who may also attempt to improve her own productivity through R&D. The monopoly assumption provides the rents that generate provide a positive return to research, as in Romer (1990). The real side of the economy aggregates as in Greenwood et al (1997), which is a standard general equilibrium macroeconomic framework except for the presence of ISTC. The aggregate dynamics of knowledge will also be standard, as in Jones (1995). Thus, the economy aggregates to a production function for output and a production function for ideas that are the same as those estimated in the empirical literature, such as equations (1) and (2). The difference is that the two are linked through ISTC.[4]

3.1 Production of Output

Time is discrete. 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}, $j \in [0, 1]$.

$$y_t = z_t \left(\int k_{jt}^{\alpha_k} dj \right) n_t^{1-\alpha_k}, \tag{3}$$

[3]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 6.

[4]This link is also central in Krusell (1998). The theoretical difference between that paper and this one lies is the way ideas evolve over time, as we shall see below. Krusell (1998) does not attempt any sort of quantitative evaluation.

where sector-neutral technical change z_t is stochastic, and grows by a factor γ_z on average each period.[5] Capital of type j commands a rental rate r_{jt}, and labor a wage w_t. Define aggregate capital $\mathbf{k}_t = \int k_{jt}dj$.

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

$$y_t \geq c_t + \int i_{jt}dj + \int x_{jt}dj. \tag{4}$$

The term "investment" here refers to forgone consumption goods that are used to make new capital goods. Let the *number* of new capital goods produced in each period be u_{jt}, so that the capital stock for each type evolves according to:

$$k_{j,t+1} = (1 - \delta_k) k_{jt} + u_{jt}, \tag{5}$$

where δ_k is the rate of physical depreciation.

3.2 Capital and Ideas

Each type of capital j is produced by a monopolist, who may also perform R&D activities in order to increase her productivity.

Let T_{jt} denote the quantity of investment-specific ideas relevant for the production of capital type j. New capital is produced according to the function $u_{jt} = T_{jt} i_{jt}^{\alpha_i}$, so that (5) becomes:

$$k_{j,t+1} = (1 - \delta_k) k_{jt} + T_{jt} i_{jt}^{\alpha_i}. \tag{6}$$

The literature on investment-specific technical change implicitly sets $\alpha_i = 1$, so that T_{jt} is interpreted directly as the efficiency of investment. I allow $\alpha_i \leq 1$, so the production possibilities set for consumption and capital may be strictly concave. Parameter α_i will play a role in calibration later on. Moreover it has economic content: if $\alpha_i < 1$, then there is an additional source of rents to R&D beyond monopoly rents.

Like capital, ideas "depreciate." They 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.[6] For any given idea, this occurs with probability δ_T each period. Ideas then evolve according to the equation

$$T_{j,t+1} = (1 - \delta_T) T_{jt} + q_{jt}. \tag{7}$$

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

$$q_{jt} = s_t Q (T_{jt}, \mathbf{T}_t, x_{jt}).$$

Q depends on the quantity of ideas in sector j, as well as $\mathbf{T}_t$ and physical input x_{jt}. Thus, there may be cross-industry spillovers. s_t is a random variable that allows the effectiveness of R&D

[5] Let γ_χ denote the average growth factor for any variable χ_t.

[6] That this occurs empirically is seen in that patent maintenance fees are not always paid – see Griliches (1990).

to vary over time. 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.

In what follows, I use the following functional form for Q :

$$Q\left(T_{jt}, \bar{T}_t, x_t\right) = \left(T_{jt}^{1-\sigma} \mathbf{T}_t^{\sigma}\right)^{\phi} x_{jt}^{\psi}. \tag{8}$$

This formulation encompasses a number of other models as special cases. First, parameter σ indicates the extent of spillovers. For instance, as $\sigma \to 1$, $Q\left(T_{jt}, \mathbf{T}_t, x_t\right) \to \mathbf{T}_t^{\phi} x_{jt}^{\psi}$: this is the Jones (1995) model of ideas production, in which knowledge applies equally well across sectors. Second – as mentioned above – the value of ϕ has been the subject of much debate. At this point, the current model places no restriction on the value of ϕ.

Third, the related model of Krusell (1998) lacks the notion of a persistent knowledge stock. This is equivalent to setting $\delta_T = 1$, so that it is the *flow* – not the *stock* – of knowledge that matters for current output. Allowing $\phi \neq 1$ and $\delta_T < 1$ strengthens the connection of the model to the empirical literature.

Parameter	Interpretation
α_k	Capital share
α_i	Concavity of the Production Possibilities' Frontier
ϕ	Concavity of the ideas' production function to past ideas
ψ	Concavity of the ideas' production function to real input
δ_k	Physical depreciation rate of capital
δ_T	Depreciation rate of ideas

Table 1 - Parameters

3.3 Households

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

$$E_0 \sum_{t=0}^{\infty} \beta^t N_t \left\{\log c_t + \eta \log l_t\right\}, 0 \leq l_t \leq \Lambda,$$

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

$$c_t \leq w\left(\Lambda - l_t\right) + \Pi_t,$$

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

4 Equilibrium

Research firms earn profits from renting their capital to the final goods' sector. Hence, the rental rate $r_{jt}\left(k_{jt}\right) = \alpha_k k_{jt}^{\alpha_k - 1} n_t^{1-\alpha_k}$ represents an inverse demand function for each investment firm.

Let V be the value function of such a monopolist, which depends on idiosyncratic and aggregate state variables. Then,

$$\begin{aligned} V\left(k_{jt}, T_{jt}, \mathbf{k}_t, \mathbf{T}_t\right) &= \max_{i_t, x_t} \left\{ r_{jt}\left(k_{jt}\right) k_{jt} - i_{jt} - x_{jt} \right. \\ &\left. + \frac{1}{1+\iota} E_t V\left(k_{j,t+1}, T_{j,t+1}, \mathbf{k}_{t+1}, \mathbf{T}_{t+1}\right) \right\}, \end{aligned} \tag{9}$$

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

$$1 = \alpha_i T_{jt} i_{jt}^{\alpha_i - 1} \frac{1}{1+\iota} E_t V_1\left(k_{j,t+1}, T_{j,t+1}, \mathbf{k}_{t+1}, \mathbf{T}_{t+1}\right)$$

or

$$i_{jt} = \alpha_i u_{jt} \frac{1}{1+\iota} E_t V_1\left(k_{j,t+1}, T_{j,t+1}, k_{t+1}, \mathbf{T}_{t+1}\right)$$

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

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

Then the first order conditions can be reformulated in terms of p_{jt} :

$$1 = \alpha_i p_{jt} T_{jt} i_{jt}^{\alpha_i - 1}, \tag{10}$$

$$i_{jt} = \alpha_i p_{jt} u_{jt}. \tag{11}$$

These equations turn out to be useful.[7] First, (10) can be used to induce an ideas' stock from price and investment data. Second, equation (11) implies that the interpretation ikn the related literature of $\frac{1}{p_{jt}}$ as a "quality adjustment" to new capital goes through, even when the production possibilities set is strictly convex. To see this, observe that the capital accumulation equation can be rewritten

$$k_{j,t+1} = \frac{i_{jt}}{\alpha_i p_{jt}} + \left(1 - \delta_k\right) k_{jt}, \tag{12}$$

which is the same as in Greenwood et al (1997) – net of the constant α_i, which amounts to a change of the units in which quality is measured.

4.1 Symmetric Equilibrium

I now focus on an equilibrium in which all capital types are treated equally, so that $T_{jt} = T_t$ and $k_{jt} = k_t \ \forall j$, suppressing industry subscripts and bold fonts henceforth. 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. As noted, it is easy to show that the real sector and the knowledge sector aggregate to yield equations (1) and (2).

[7] Alternatively, define the opportunity cost of new capital goods in terms of the numeraire to be p_{jt}. Then, (10) emerges from this definition.

The Euler equations reveal the economics that underlie a model of ideas-driven ISTC. In (13), the marginal cost of investment – which decreases in T_t – equals the marginal expected return to new capital in the future:

$$\frac{i_t^{1-\alpha_i}}{\alpha_i T_t c_t} = \gamma_n \beta E_t \left\{ \frac{\alpha_k^2 z_t k_t^{\alpha_k - 1} n_t^{1-\alpha_k} - (1-\delta_k) \frac{i_t^{1-\alpha_i}}{\alpha_i T_{t+1}}}{c_{t+1}} \right\}, \tag{13}$$

or, in terms of prices,

$$\frac{p_t}{c_t} = \gamma_n \beta E_t \left\{ \frac{r_t - (1-\delta_k) p_{t+1}}{c_{t+1}} \right\}.$$

The double exponent on α_k reflects the inefficiency due to the non-competitive market structure of the investment sector, which leads to the underproduction of new capital in equilibrium and increased rents for researchers: the exponent would equal one in the planner's problem, although this does not matter for purposes of growth accounting. Observe that the dynamic return to investment is increasing in the future stock of ideas, as new ideas lower the opportunity cost of physical depreciation. In the limit, as $T_{t+1} \to \infty$, the economic cost of physical depreciation becomes negligible.

The following is the Euler equation for ideas:

$$\begin{aligned}
\frac{x_t^{1-\psi}}{\bar{T}_t c_t} &= \gamma_n \beta E_t \left[\frac{1}{T_{t+1} c_{t+1}} \left(\xi_{1,t+1} + \xi_{2,t+1} + \xi_{3,t+1} + \xi_{4,t+1} \right) \right] \qquad (14) \\
\xi_{1,t} &= \frac{\Psi i_t^{1-\alpha_i} T^{\frac{2\alpha_i - 1}{1-\alpha_i}}}{(1-\alpha_i)} \\
\xi_{2,t} &= \frac{x_t^{1-\psi}}{\psi s_t} (1-\sigma) \phi \frac{T_{t+1}}{T_t^{1+\phi}} \\
\xi_{3,t} &= \frac{x_t^{1-\psi}}{\psi s_t} (1 - \phi(1-\sigma)) (1-\delta_T) T_t^{-\phi} \\
\xi_{4,t} &= - \left[\frac{(k_{t+1} - (1-\delta_k) k_t) i_t^{\alpha_i - 1}}{\alpha_i} \right].
\end{aligned}$$

The left hand side of (14) is the marginal cost of R&D in terms of the numeraire. The marginal benefit, on the right hand side, has several components. First, $\xi_{1,t+1}$ represents a reward that accrues to the researcher-investor due to the concavity of the PPF, where $\Psi \equiv \alpha_i^{\frac{\alpha_i}{1-\alpha_i}} - \alpha_i^{\frac{1}{1-\alpha_i}} \geq 0$. Note that, as $\alpha_i \to 1$, $\xi_{1,t} \to 0$, so that this disappears as an element of the reward for research.[8] Next, $\xi_{2,t+1}$ captures the positive influence that new ideas have on the production of ideas in the future, representing what Krusell (1998) terms the dynamic return to ideas' production. A further component, $\xi_{3,t+1}$, reflects the fact that ideas depreciate only slowly, so that their dynamic influence lasts beyond the subsequent period. Finally, new ideas make capital cheaper to produce in the future. This represents a capital loss $\xi_{4,t+1}$ on current investment which Krusell (1998) terms "planned obsolescence."

[8] Interestingly, this suggests that researchers and investors need not be identified: so long as $\alpha_i < 1$ then there will be rewards to both. However, if $\alpha_i = 1$, then the monopolistic assumption is necessary.

Krusell (1998) also provides an extended discussion of the efficiency properties of a model that is equivalent to the above for the case that $\phi = 1$, $\delta_T = 1$ and $\alpha_i = 1$. Firms will tend to under-invest in research because of the cross-sectoral externality σ. In the present framework, the fact that $\delta_T < 1$ means that this under-investment may be more severe, as ideas and hence spillovers persist over time.

4.2 Balanced Growth

The long-run properties of the model economy will play an important role in growth accounting. Equation (10) implies that, on a balanced growth path (BGP),

$$\gamma_T = \frac{\gamma_y^{1-\alpha_i}}{\gamma_p}. \tag{15}$$

The relationship between ideas, TFP, population and output growth is given by

$$\gamma_y = \gamma_z^{\frac{1}{1-\alpha_i\alpha_k}} \gamma_T^{\frac{\alpha_k}{1-\alpha_i\alpha_k}} \gamma_n^{\frac{1-\alpha_k}{1-\alpha_i\alpha_k}}. \tag{16}$$

Finally, (8) yields an expression for the growth rate of ideas itself, in terms of other factors.

$$\gamma_T = \gamma_s^{\frac{1}{1-\phi}} \gamma_y^{\frac{\psi}{1-\phi}}. \tag{17}$$

Eliminating γ_T, the growth rate of final output is

$$\gamma_y = \left(\gamma_z^{1-\phi} \gamma_s^{\alpha_k} \gamma_n^{(1-\alpha_k)(1-\phi)}\right)^{\frac{1}{(1-\alpha_i\alpha_k)(1-\phi)-\alpha_k\psi}}. \tag{18}$$

While (18) may appear complicated, it captures the notion that the production of ideas is an important channel of growth – although ultimately, as in Jones (1995), growth is ultimately 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

$$\psi < \frac{(1-\alpha_i\alpha_k)(1-\phi)}{\alpha_k}, \tag{19}$$

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 not excessive. Notably, it does *not* hold as $\phi \to 1$, 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. Still, 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, $\gamma_z < 1$, while the results of Griliches (1990), Porter and Stern (2000) and Abdih and Joutz (2003) seem to indicate that $\gamma_s < 1$. Hence, the structure of the model is agnostic as to the value of ϕ. Observe that the spillover σ does not affect the long run growth properties of the economy: I suppress σ henceforth.

Over the short run, the model has the interesting property that it possesses a powerful endogenous propagation mechanism for shocks, in the form of the ideas production function itself. As short run considerations are not the focus of the paper, these results are discussed in Appendix *A*.

4.3 Knowledge and Productivity

As noted, the empirical literature generally finds a weak link between the knowledge stock and TFP. If ideas lead to growth through ISTC, however, this is not 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:

$$\log y_t = \theta \log T_t + \alpha_k \log k_t + (1-\alpha_k)\log n_t + \eta t + \zeta_t,\ \ E_{t-1}[\zeta_t] = 0, \tag{20}$$

which, defining $z_t \equiv e^{\eta t + \zeta_t}$, is equivalent to a production function of the form $y_t = z_t T_t^{\theta} k_t^{\alpha_k} n_t^{1-a_k}$. Typically, estimates of θ are very low, ranging from 0.05 to 0.2 and sometimes lacking statistical significance.[9] This does not imply that ideas do not matter for economic growth, however, if investment-specific technical change is the channel that links them.

Assume that the model is a correct representation of the world, and that $\alpha_i = 1$. 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 will be:

$$y_t = \hat{z}_t \hat{k}_t^{\alpha_k} n_t^{1-\alpha_k}, \tag{21}$$

$$\hat{k}_{t+1} = \left(1-\hat{\delta}_t\right)\hat{k}_t + i_t, \tag{22}$$

where $\hat{k}_t$ is the capital stock derived according to (22), $\hat{\delta}_t$ 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 model with ISTC:[10]

$$\begin{aligned}
\hat{k}_t &= k_t p_{t-1} \\
\hat{\delta}_t &= 1-(1-\delta)\left(\frac{p_t}{p_{t-1}}\right), \\
\hat{z}_t &= z_t\left(\frac{1}{p_{t-1}}\right)^{\alpha_k}.
\end{aligned} \tag{23}$$

I henceforth 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. On the other hand, (23) is the accounting link that mis-attributes investment-specific to neutral technical change if the existence of ISTC is ignored. To be precise, using equation (10), the relationship between knowledge T_t and $\hat{z}_t$ becomes

$$\hat{z}_t = z_t T_{t-1}^{\alpha_k}. \tag{24}$$

In this way, 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, on top of the one-period lag in equation (24). Second, there may be measurement error – for instance, Harhoff et al (2003) and Leiva (2004) argue that patent quality is sufficiently

[9] See Griliches (1990) for a survey of microeconomic estimates, and Porter and Stern (2000) and Abdih and Joutz (2003) for macroeconomic estimates.

[10] See Greenwood et al (1997).

heterogeneous that aggregate patent data is not a useful indicator of aggregate knowledge. Most importantly, however, if it is true that the stock of ideas affects growth through investment-specific technical change, then (24) states that the "true" value of θ equals α_k. Thus, 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 when, depending on the value of γ_z, the knowledge channel may in principle account for the entirety of economic growth.

5 Quantitative Analysis

I calibrate the model in two stages. First, I concentrate on the real side of the economy. Then, in Section 6, I calibrate the production function for ideas.

5.1 Calibration

To calibrate the model to US post war growth, I follow the procedure put forward by Kydland and Prescott (1982). Assuming a balanced growth path, I use (15) and (16) to solve for α_i and γ_z respectively.

Key to this task will be identifying an empirical counterpart for T_t. A premise of the paper is that the relative price of capital p_t is an indicator of the economically useful stock of knowledge and, for most of the paper, I will use p_t to construct T_t. However, there are other observable indicators of research activity. In particular, patents have been widely used as an indicator of the output of research, and estimates of the ideas production function typically use patent data. Hence, I begin by identifying T_t with the patent stock.[11]

Patent data is 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 is from the United States Census Bureau. The years covered are 1947-2000:

[11] See Appendix B for further comments on the use of patent data.

all series are annual.

Variable	Value
γ_p	1.0262^{-1}
γ_T	1.0261
γ_y	1.0333
γ_n	1.0119
α_k	0.3
ι	0.07
γ_z	1.0069
α_i	1
δ_T	12%
δ_k	0.039

Table 2 - Statistics and parameters used in calibration

I calibrate the rate of physical depreciation using the following long-run relationship:

$$\gamma_k = (1 - \delta_k) + \frac{I}{K}.$$

where I/K is the investment-to-capital ratio. In the NIPA data, the investment-to-GDP ratio $\frac{I}{Y}$ is 0.19. Given $\alpha_k = 0.3$, the following represents the Euler equation along a BGP:

$$\frac{\gamma_k}{\beta\gamma_n} = \left(\alpha_k \frac{Y}{K} + 1 - \delta_k\right).$$

Imposing the condition that $\beta = \frac{\gamma_y}{1+\iota}$, and setting $\iota = 7\%$, physical depreciation δ_k equals 3.9%. per annum.[12]

Some authors argue that the depreciation rate of ideas δ_T is low – for example, Wilson (2001) uses a value of 2%. The empirical literature tends to find much larger numbers, however – for instance, Nadiri and Prucha (1996) find a value of 12%, and Pakes and Schankerman (1984) find numbers up to 26%, which seems large. I assume that $\delta_T = 0.12$, and later assess the sensitivity of results to this parameter.

Table 2 lists the parameter values and growth rates that I use to calibrate γ_z and α_i. The values that are consistent with the above are $\alpha_i = 1$ and $\gamma_z = 1.0069$.

5.2 Constructing the knowledge stock

Because the growth rate of the relative price of capital is quite variable, the value of γ_p used in calibration is sensitive to the time period used. This is important for the choice of α_i and γ_z. Hence, I adopt a second strategy for calibrating these variables, that will also turn out to provide

[12] While this number may appear low, the reader is reminded that this figure accounts for physical, not economic, depreciation, and that it includes both equipment and structures. A value of $\delta_k = 0.07$ was also used, without affecting conclusions.

some support for the relationship at the core of the model: the link between the accumulated stock of ideas and the price of capital.

To do this, I construct two independent measures of the stock of ideas. In the model, changes in the relative price of capital reflect the implementation of economically useful ideas. Hence, one approach is to use equation (10) in conjunction with the series for the price of capital and for investment to build an implicit ideas' stock T_t^{price}. A second approach, following the empirical literature, is to use the stock of patent applications $T_t^{patents}$. More concretely, if q_t^{patent} is new patent applications, the two stocks are constructed as follows:

$$T_t^{price} = \frac{i_t^{1-\alpha_i}}{\alpha_i p_t}, \; T_{t+1}^{patent} = (1-\delta_T)\, T_t^{patent} + q_t^{patent} \tag{25}$$

While the construction of the price-based knowledge series requires a value of α_i, the patent-based series requires a value for the depreciation rate for ideas δ_T. I examine different combinations of these parameters, and ask whether the two stocks co-move appreciably. If they do, this supports the approach taken here – provided they do so for reasonable parameter values. Recall that Wilson (2001) 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.

Figure (1) displays the contemporaneous correlation between the growth rate of T_t^{price} and that of T_t^{patent}, for all possible (α_i, δ_T) pairs. An interesting picture emerges. First, there is a positive correlation between the two series almost regardless of parameter values. However, for any given δ_T, this correlation is generally increasing in the value of α_i. In other words, it is strongest when price data is the most important factor in the construction of T_t. That α_i equals or is close to one is informative. It implies that the main factor in changes to the relative price of capital is the use of ideas, rather than the quantity of investment. Moreover, it implies that the return to R&D is primarily the effect that it has upon productivity in the investment sector, rather than rents from the shape of the PPF.

Second, the correlation is also the strongest for values of δ_T below 25% – which is the empirically relevant zone. This suggests that indeed it is the *stock* – not the $flow$ – of ideas that matters for growth. The maximum correlation is 37.5% where $\alpha_i = 1$ and $\delta_T = 14.6\%$, and is statistically significant. That the two measures should co-move in the time-series at all is remarkable, as is the fact that this co-movement is substantial only for parameter values that are plausible and consistent with those suggested by the earlier calibration.[13]

Figure (2) displays the growth rates of the two series for $(\alpha_i, \delta_T) = (1, 12\%)$. The price-based series is more volatile: this is at least partly due to known sources of measurement error,[14] but may be due to variations in patent quality over time, an issue that has been raised in the literature[15] and which I discuss in detail later. Figures (3) and (4) display the growth rates of

[13]The growth rates of T_t^{price} and T_t^{patent}strongly reject a unit-root test. Hence, these correlations are not spurious.

An augmented Dickey-Fuller test cannot reject the hypothesis of a unit root for either knowledge stock series (as distinct from their growth rates), and the Johansen (1991) procedure suggests that they share a single cointegating vector. Again, this is consistent with the two series being driven by the same underlying factors.

[14]In particular, the sharp dip in 1974 is related to price controls at the time.

[15]See for example Griliches (1990), Hall, Jaffe and Trajtenberg (2001), Harhoff et al (2003) and Leiva (2004).

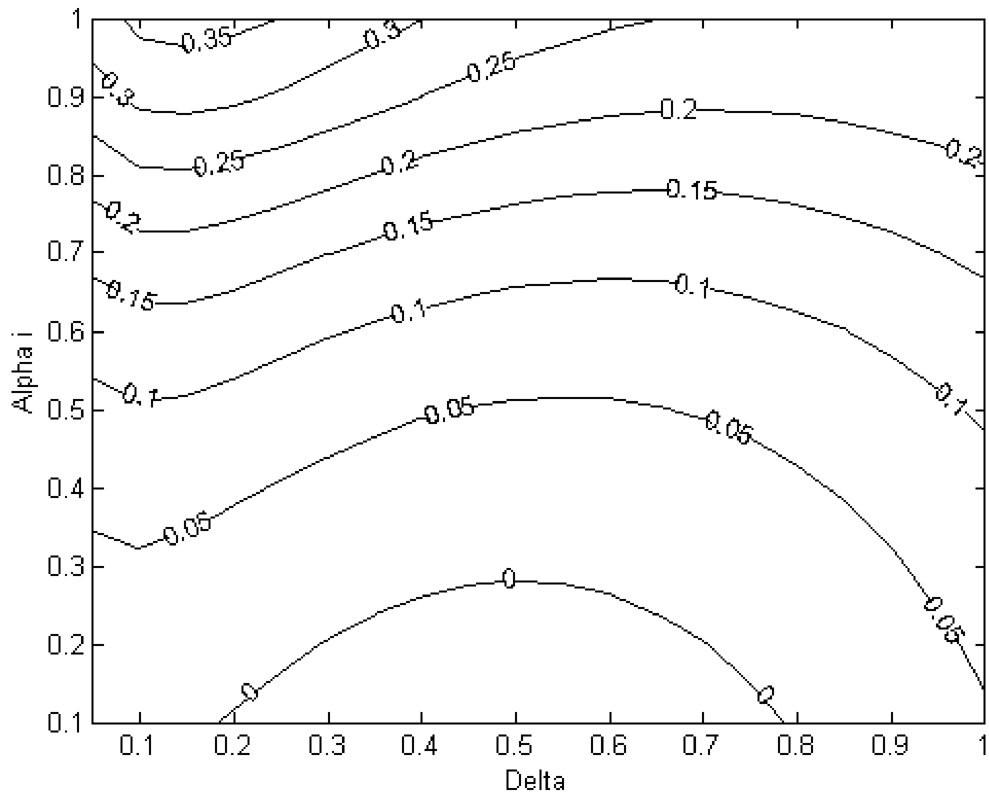

Figure 1: Contemporaneous correlation between the growth rates of the patent stock and the constructed stock of ideas, for different values of α_i and δ_T.

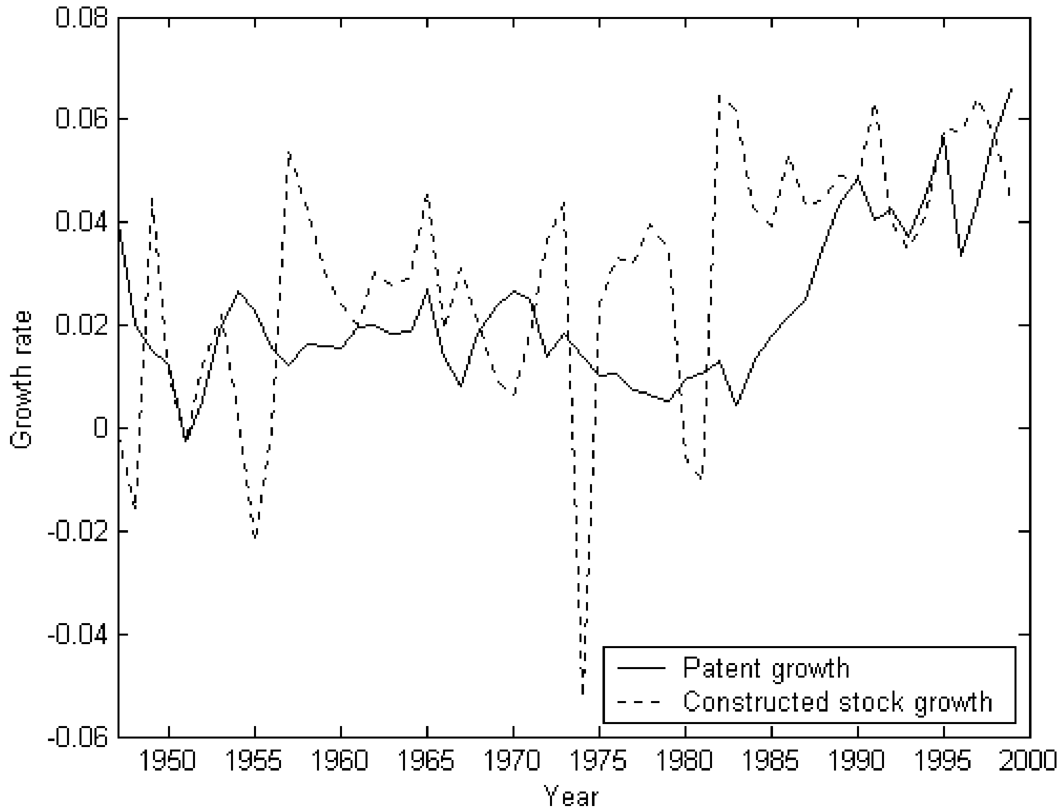

Figure 2: Patent stock and implicit knowledge stock derived from price data; growth rates. $\alpha_i = 1, \delta_T = 0.12$. Correlation: -35%.

the two series once more, this time Hodrick-Prescott filtered with smoothing parameters that are "low" to remove very high-frequency movements that might reflect noise due to changes in patent quality or measurement error – the Cummins and Violante (2002) price data is constructed under assumptions that arguably make it unsuitable for high frequency analysis, so the results using smoothed data may be more reliable. Notably, very little smoothing is necessary to raise the correlations between the two series even further. Observe that both measures of the knowledge stock display an increase in growth starting 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 likely not spurious but rather reflects an actual increase in the rate of innovation in the US economy.[16]

In the data, per-capita GDP growth has averaged 2.1% over the period. Suppressing neutral

[16] I also repeated the exercise of Figure (1) using the smoothed series. When the series are subjected to very little smoothing ($\lambda = 1$), results are similar and correlations are larger, up to almost 60% for $\alpha_i = 1$ and $\delta_T \approx 15\%$. Interestingly, when $\lambda = 10$, most of the high-frequency variation is smoothed out and the results continue to point to a high value of α_i but are no longer informative about δ_T. This is to be expected, as the absence of any short-term fluctuations in the filtered series implies an inability to distinguish how quickly or slowly they decay.

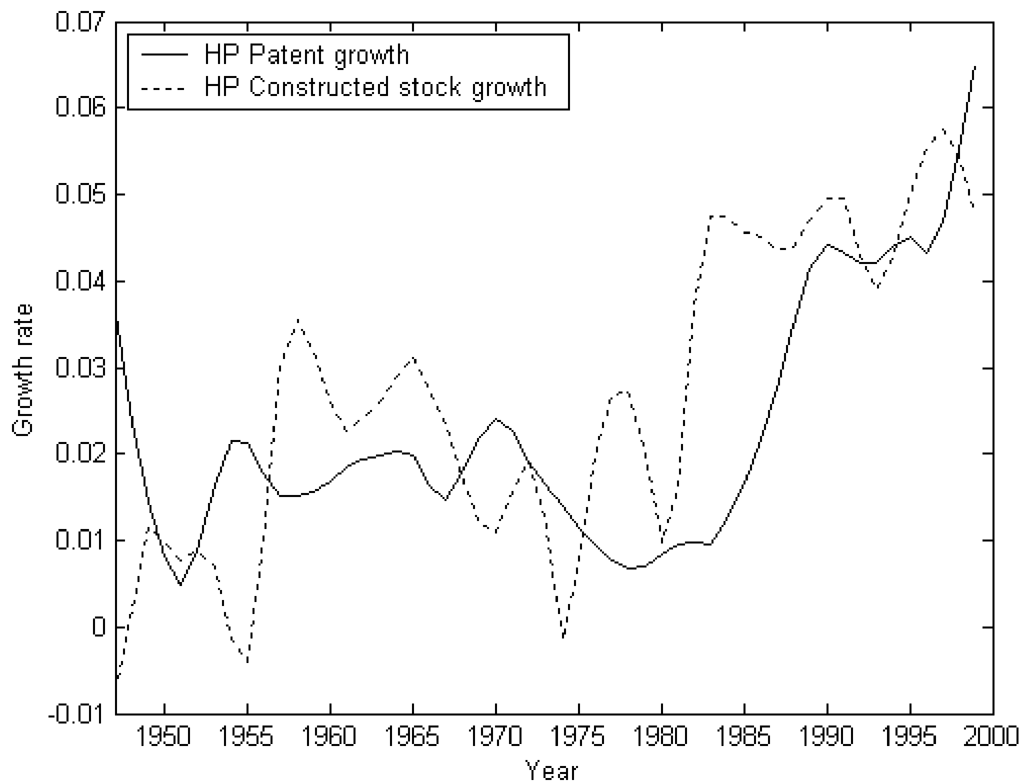

Figure 3: Patent stock and implicit knowledge stock, Hodrick-Prescott filtered growth rates with smoothing parameter $\lambda = 1$. $\alpha_i = 1, \delta_T = 0.12$. Correlation: -57%.

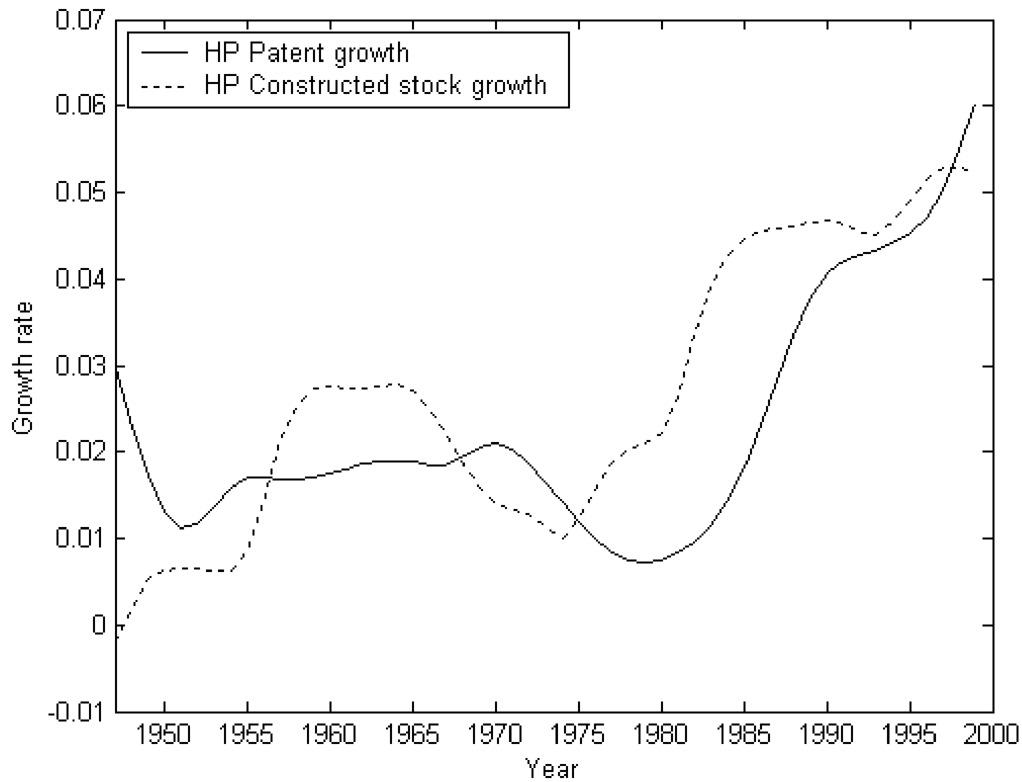

Figure 4: Patent stock and implicit knowledge stock, Hodrick-Prescott filtered growth rates with smoothing parameter $\lambda = 10$. $\alpha_i = 1, \delta_T = 0.12$. Correlation: -71%.

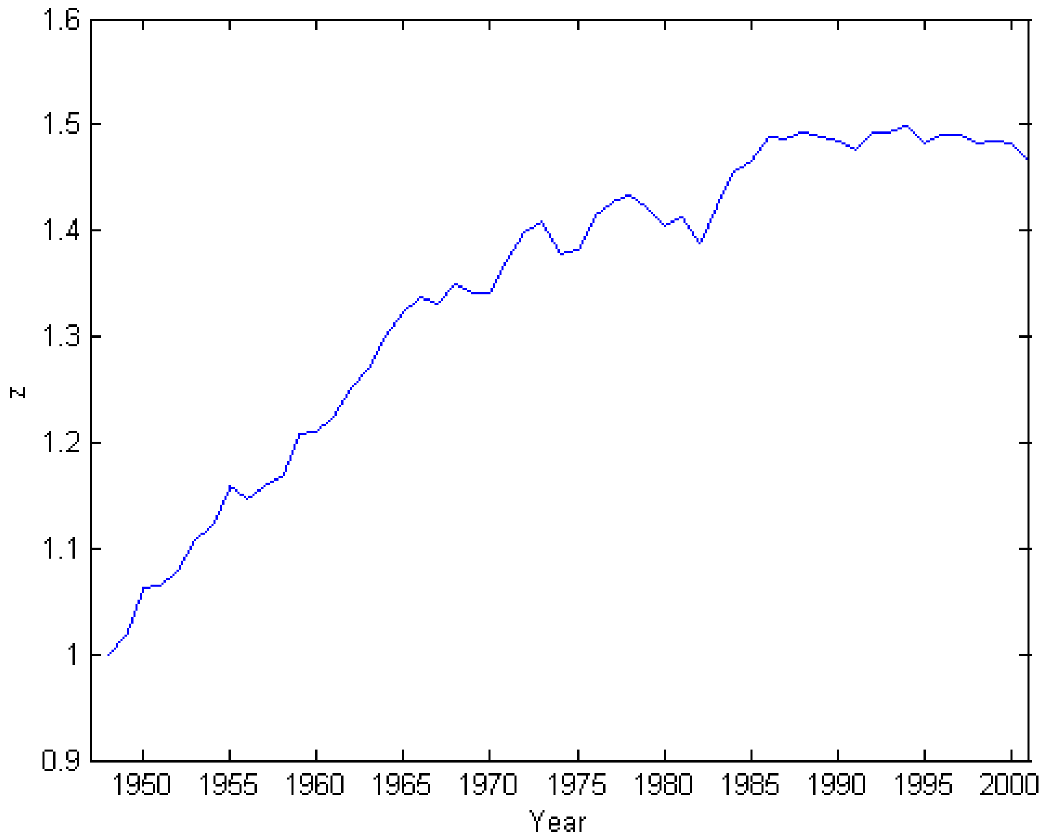

Figure 5: Neutral technical change z_t.

technical change γ_z, this number would have been 1.12%. On the other hand, neutral technical change alone would have led to 0.95% annual growth. Thus, investment-specific technical change accounts for about 55% of US economic growth over the post-war era. This is very close to the value of 58% 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. That this number is so robust to different approaches is remarkable.[17]

5.3 Total Factor Productivity

I now use the model to construct a series for neutral technical change. I do this by iterating on the capital accumulation equation (12) to impute a capital stock measure for each date, and then using equation (3) and data on output and labor to obtain $\{z_t\}_{t=1947}^{2000}$.

Figure (5) plots "true" total factor productivity z_t for the US, normalized by its initial value. It is worth noting that, based on this constructed series, $\gamma_z = 1.0072$, which is almost exactly

[17] Greenwood et al (1997) distinguishes between equipment and structures, so their growth accounting is sensitive to factor shares as well as the assumption that investment-specific technical change does not affect structures. I use a price index that includes structures, so this distinction is unnecessary and the factor share is more reliable.

the same as when a BGP is assumed.[18]

6 Economic Growth and the Production of Ideas

In this section, I discuss the implications of the model for the empirical structure of the ideas production function. Before doing so, however, I address some measurement issues that will provide the background for this task.

6.1 Patent "quality" and Time Trends

Griliches (1990) points out that a potential problem with the use of patent data is that patents may vary in terms of quality over time. The results of Hall, Jaffe and Trajtenberg (2001) suggest that patent quality may indeed vary 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, there is reason to believe that there may be discrepancies between the stock of economically useful knowledge and the patent stock *over time* as well as in cross-section.

To be precise, 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

$$q_t = m_t b_t, \tag{26}$$

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. Following the terminology of Griliches (1986), I refer to m_t as average "patent quality". On the basis of patent data, 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, 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. In this case – as I argue in Section (6.2.2) – estimates based on the patent stock will be biased.

To see this, 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:

$$\log b_t = \mu t + \phi \log T_t + \psi \log x_t + \varepsilon_t, \tag{27}$$

[18] As mentioned, the price-based series may be subject to measurement problems in the short run. Hence further comments on the short run behavior of z_t and T_t are left for an Appendix C.

Griliches (1990) and Abdih and Joutz (2003) 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%.[19]

Equation (27) is equivalent to the present setup, net of a re-labelling of variables. In the ideas production function (8), set $q_t = m_t b_t$. Taking logarithms, this becomes simply

$$\log b_t + \log m_t = \log s_t + \phi \log T_t + \psi \log x_t \tag{28}$$

Combining (27) and (28) yields $m_t = s_t e^{-(\mu t + \varepsilon_t)}$ so that

$$\mu = \log \gamma_s - \log \gamma_m. \tag{29}$$

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.

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? Empirically, the latter interpretation 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 towards secrecy as a means of protecting intellectual property among US manufacturing firms. Lanjow 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.

In addition, it is difficult to see any 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. A deterioration in the institutions of research might lead to a trend; however, this seems at odds with the evidence of Kortum and Lerner (1998), and it is not clear why this should result in a decrease in s_t. Griliches (1990) interprets the time trend as reflecting a tendency away from patenting due to the increasing opportunity cost of the patenting process – in this case, the trend does not represent a slump in innovativeness s_t, but an increase in unmeasured ideas m_t precisely as argued here.

Finally, the model can be used to derive a series for the discrepancy between patents and new ideas m_t. From equations (7) and (26), it follows that

$$m_t = \frac{T_{t+1} - T_t (1 - \delta_T)}{b_t}. \tag{30}$$

[19] 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 which is not very different from the values found by other authors.

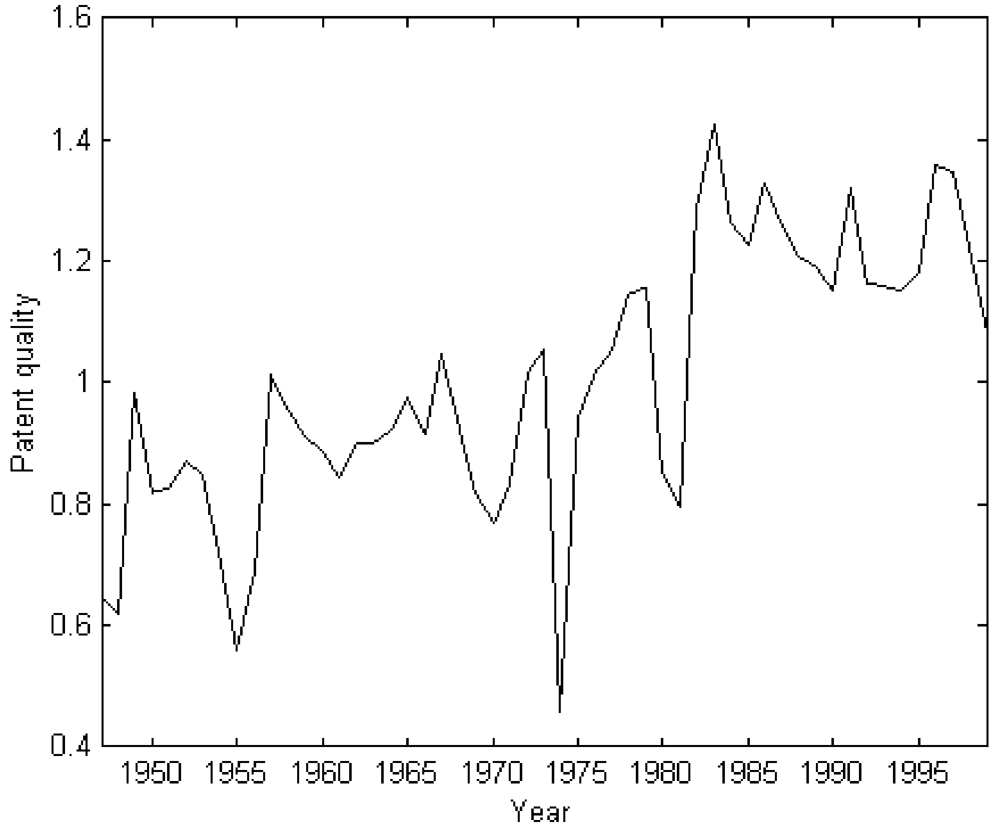

Figure 6: Patent quality m_t.

Iterating on (30) with the price-based series for T_t and setting b_t to equal patent applications yields the desired series, which is displayed in Figure (6). 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 time trends μ as the extent of measurement error underlying the use of patent data.

6.2 Intertemporal ideas spillovers

I now turn to the parameters of the ideas production function. First, I extend the calibration procedure to match these parameters. Second, I discuss the sensitivity of the empirical results of other authors to the measurement issues just raised.

6.2.1 Calibration

Equation (18) is the growth accounting relationship between parameters and aggregates, under the assumption that γ_T affects growth through ISTC. It can be rearranged as follows:

$$\phi = 1 - \frac{\alpha_k \psi \log \gamma_y + \alpha_k \log \gamma_s}{(1-\alpha_i \alpha_k) \log \gamma_y - \log \gamma_z - (1-\alpha_k) \log \gamma_n} \tag{31}$$

At this point, equation (31) contains three unknowns: ψ, γ_s and ϕ. The microeconomic estimates of ψ surveyed by Griliches (1990) range from 0.3 to 0.6, which turns out to be consistent with the macroeconomic estimates of other authors.[20] I focus upon values within this range.

Observe that there is a one-to-one mapping between ϕ and γ_s. Taken at face value, the measured trend μ suggests a range of values for γ_s between 0.97 and 0.98. However, as argued, these measurements are unable to distinguish between a time trend γ_s and patent quality or measurement error γ_m, and the latter interpretation seems more likely both for empirical and theoretical reasons. Hence I assume for now that $\gamma_s = 1$.

Depending on the value of ψ, the intertemporal spillover ϕ falls in the range $[0.23, 0.62]$. Like the empirical estimates, the calibration points to values of ϕ that are positive, suggesting that "fishing out" is not a feature of the data. However, the estimates are typically close to or even larger than 1.

6.2.2 Reconciling the estimates

Why are the empirical estimates of ϕ so much larger?

[20] 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:

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

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. Both Abdih and Joutz (2003) estimate an equation of the form $q_t = e^{\mu t} T_t^{\phi} k_t^{\eta} n_t^{\lambda}$, 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 $\psi \in [0.3, 0.64]$.

As mentioned, there is evidence that the patent stock underestimates growth in the stock of ideas ($\gamma_m > 1$). As a result, estimates of ϕ using patent data will be biased upwards. To see this, recall that the ideas' production function and the ideas' accumulation equation can be written:

$$\begin{aligned} \log b_t + \log m_t &= \log s_t + \phi \log T_t + \psi \log x_t, & (32)\\ T_{t+1} &= b_t m_t + (1-\delta_T) T_t. & (33) \end{aligned}$$

Observe that m_t appears in (32) as a time trend, and also in (33) as a "correction" to measuring the quantity of new ideas using patents.

On the other hand, the empirical implementation of these equations is:

$$\begin{aligned} \log b_t &= \mu t + \hat{\phi} \log B_t + \psi \log x_t + \varepsilon_t, & (34)\\ B_{t+1} &= b_t + (1-\delta_T) B_t. & (35) \end{aligned}$$

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

Writing the long-run forms of equations (32) and (34), and equating their right hand sides,

$$\begin{aligned} \phi \log \gamma_T &= \mu + \hat{\phi} \log \gamma_b - \log \gamma_s + \log \gamma_m, \\ \phi &= \hat{\phi} \frac{\log \gamma_b}{\log \gamma_T}, \end{aligned}$$

since $\mu = \log \gamma_s - \log \gamma_m$. Finally, $\gamma_T = \gamma_b \gamma_m$. Hence, the "true" value of ϕ underlying their estimates $\hat{\phi}_{measured}$ is:

$$\phi = \hat{\phi}_{measured} \left(\frac{\log \gamma_b}{\log \gamma_b + \log \gamma_m} \right). \quad (36)$$

If patent quality growth $\gamma_m > 1$, then $\hat{\phi}_{measured}$ is biased upwards.

Table (3) 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" using equation (36). Since γ_b is somewhat sensitive to the period of measurement, when discussing the estimates of any given paper I compute γ_b based on data for the corresponding years. Their point estimates turn out to be consistent with values of ϕ that are well below unity, and close to the range suggested by calibration.

Source	Period	γ_b	μ	$\phi_{measured}$	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

Table 3 – Patent and trend data. γ_b is the growth rate of patent applications over the relevant sample period. μ is drawn from the results of each paper.

Another approach towards reinterpreting the results is to use the ideas' production function to obtain bounds on ϕ that are consistent with the estimates, again without imposing the identification of ideas with investment-specific technical change.

On a balanced growth path, $\gamma_q = \gamma_T$, so the ideas production function can be re-written:

$$\phi = 1 - \frac{\psi \log \gamma_y + \log \gamma_s}{\log \gamma_T}. \tag{37}$$

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 (37) delivers a value for ϕ that assumes that the measured trend represents decreases in s_t. Denote this value $\phi_{\max}$.

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, which I denote $\phi_{\min}$.

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

Table 4 – Predicted and measured ϕ.

Results are reported in Table (4). Again, the column for $\phi_{\max}$ 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. On the other hand, if patent data systematically underestimate growth in the stock of ideas, the resulting values are all below one.[21]

Finally, the same procedure can be applied to equation (31), which does assume the structure imposed by linking ideas and ISTC. Results are displayed graphically in Figure (7). 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.

Aside from the values themselves, the following results should be emphasized. First, both estimates and calibration 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, the calibration suggests that there are likely to be *decreasing returns* in the ideas' production function: $\phi < 1$. The model is in fact consistent with values of ϕ that exceed unity: however, this hinges upon 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 from the use of patent data rather than a trend per se. 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. The paper shows how taking this into consideration in the context of an ISTC-based growth model points to values of ϕ that tally much better with theory.

[21] Growth in reported R&D spending has amounted to over 4% per year: if this is replaced for γ_y in equation (37) then the values of ϕ obtained are even lower.

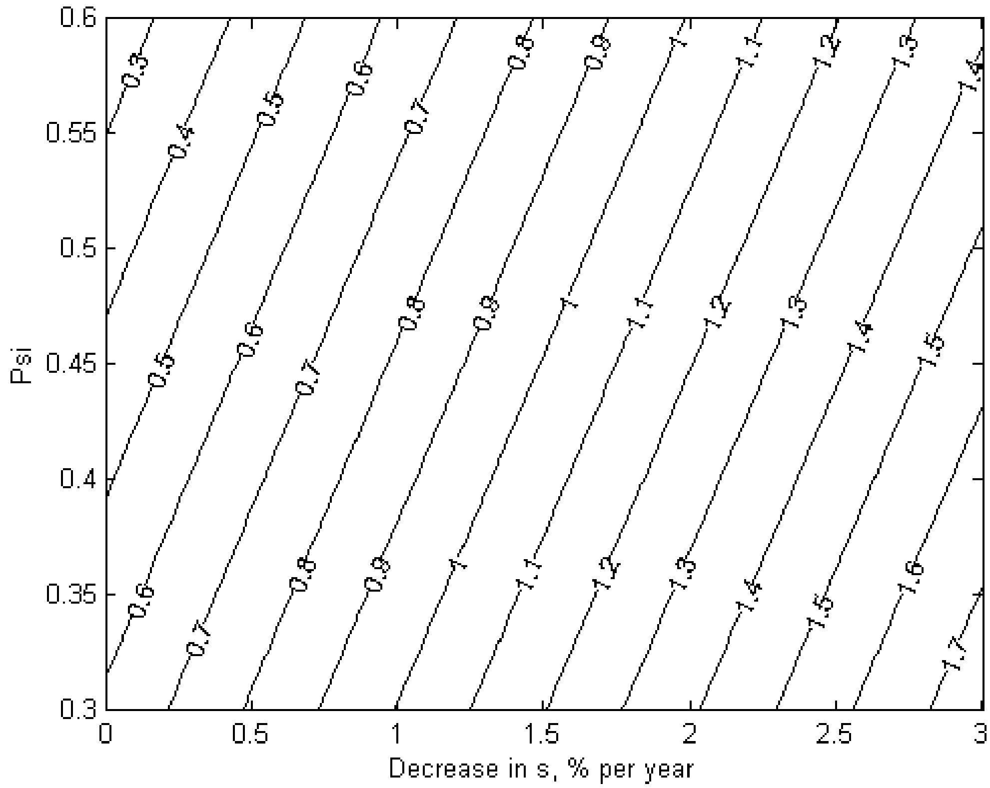

Figure 7: Values of ϕ consistent with the estimates, given assumptions upon ψ and γ_s.

6.3 Concluding Remarks

The paper proposes an aggregate model of economic growth that is consistent with the 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 ISTC that provides evidence of a link between ideas and growth through this channel. The calibration of the model economy then suggests that there are decreasing returns to old knowledge in the production of new knowledge.

A full simulation or estimation of the model is beyond the scope of this paper. However, given the theoretical link between p_t and T_t of the model, its business cycle properties are likely similar to those of models that link ISTC shocks to the business cycle using the price series alone. Fisher (2003) finds that they account for about 50% of the time-series variation in labor productivity.

An interesting extension of the model would be to investigate the policy implications of current results. The model economy is characterized by underinvestment in research, and the extent of this underinvestment is more severe to the extent that ideas persist over time. Given that the depreciation rate of ideas appears to be empirically low, there may be significant scope for research subsidies to increase output. This extension is left for future work.

7 Bibliography

Abdih, Yasser and Joutz, Frederick. "Relating the Knowledge Production Function to Total Factor Productivity." Mimeo: George Washington University (2003).

Caballero, Ricardo and Jaffe, Adam. 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. Cambridge, MIT Press (1993).

Cohen, Wesley; Nelson, Richard and Walsh, John. "Protecting their Intellectual Assets: Appropriability Conditions and why US Manufacturing Firms Patent (or Not)." NBER Working Paper w7552 (2000).

Cooley, Thomas F.; Prescott, Edward C. "Economic Growth and Business Cycles." *Frontiers of business cycle research*, Princeton University Press (1995) 1-38.

Cummins, Jason and Violante, Giovanni. "Investment-Specific Technical Change in the US (1947-2000): Measurement and Macroeconomic Consequences." *Review of Economic Dynamics* vol. 5(2), April 2002, 243-284.

Fisher, Jonas. "Technology Shocks Matter." Mimeo: Federal Reserve Bank of Chicago, (2003).

Greenwood,Jeremy, Hercowitz, Zvi and Krusell, Per. "Long-Run Implications of Investment-Specific Technological Change", *American Economic Review*, 87:3 (1997).

Greenwood, Jeremy; Jovanovic, Boyan. "Accounting for Growth." *New developments in productivity analysis*, NBER Studies in Income and Wealth, vol. 63. Chicago and London: University of Chicago Press (2001) 179-222.

Griliches, Zvi. "Productivity, R and D, and Basic Research at the Firm Level in the 1970's" *American Economic Review*, 76 (1996) 141-154.

Griliches, Zvi. "Patent Statistics as Economic Indicators: A Survey" *Journal of Economic Literature*, Vol. 28, No. 4. (1990) 1661-1707.

Hall, Bronwyn H., Jaffe, Adam B. and Trajtenberg, Manuel. "Market Value and Patent Citations: A First Look" NBER Working Paper w7741 (2000).

Harhoff, Dietmar; Scherer, Frederic M. and Vopel, Katrin. "Citations, family size, opposition and the value of patent rights." *Research Policy*, 32(8)1343-1363 (2003).

Howitt, Peter. "Steady Endogenous Growth with Population and R&D Inputs Growing." *Journal of Political Economy* 107(4) (1999) 715-30.

Huffman, Gregory W. "Propagation Through Endogenous Investment-Specific Technological Change." Vanderbilt University Economics Department Working Paper 02-W23 (2002).

Johansen, Soren. "Estimation and Hypothesis Testing of Cointegration Vectors in Gaussian Vector Autoregressive Models." *Econometrica*, Vol. 59, No. 6. (1991) 1551-1580.

Jones, Charles. "R&D-Based Models of Economic Growth." *Journal of Political Economy* 103 (1995) 759-84.

Kortum, Samuel and Lerner, Josh. "Stronger Protection or technological revolution: what is behind the recent surge in patenting?" *Carnegie-Rochester Conference Series on Public Policy* 48 (1998) 247-304.

Krusell, Per. "Investment-Specific R&D and the Decline in the Relative Price of Capital", *Journal of Economic Growth*, 3 (1998) 131-141.

Kydland, Finn E. and Prescott, Edward C. (1982) "Time to Build and Aggregate Fluctuations." *Econometrica*, 50(6) 1345-1370.

Lanjow, Jean and Schankerman, Mark. "The Quality of Ideas: Measuring Innovation with Multiple Indicators." NBER Working Paper w7345 (1999).

Leiva Bertrán, Fernando. "Pricing Patents Through Citations." Mimeo: University of Iowa (2004).

Maddison, Angus. "Growth and Slowdown in Advanced Capitalist Economies: Techniques of Quantitative Assessment." *Journal of Economic Literature* 25 (1987) 649-698.

Nadiri, M. Ishaq; Prucha, Ingmar R. "Estimation of the Depreciation Rate of Physical and R&D Capital in the U.S. Total Manufacturing Sector." *Economic Inquiry* 34 (1996) 43-56.

Pakes, Ariel; Schankerman, Mark. "The Rate of Obsolescence of Patents, Research Gestation Lags, and the Private Rate of Return to Research Resources."; *R&D, Patents, and Productivity*, NBER Conference Report. Chicago and London: University of Chicago Press (1984) 73-88,

Porter, Michael E. and Stern, Scott. "Measuring the "Ideas" Production Function: Evidence from International Patent Output" NBER Working Paper w7891 (2000).

Romer, Paul. "Endogenous Technological Change." *Journal of Political Economy* 98 (1990) S71-S101.

Scherer, Frederic M. and Harhoff, Dietmar. "Technology policy for a world of skew-distributed outcomes." *Research Policy*, 29(4) 559-566, (2000).

Stevens, Greg and Burley James, "3,000 Raw Ideas = 1 Commercial Success". *Research Technology Management*, 40, (1997), 16-27.

Wilson, Daniel J. "Is Embodied Technology the Result of Upstream R&D? Industry-Level Evidence," *Review of Economic Dynamics* 5 (2002) 285-317.

A Autocorrelation

A consequence of equation (24) which links traditionally-measured TFP ($\hat{z}_t$) to the two forms of technical change (z_t and T_t) is that, even if in reality shocks to z_t are not autocorrelated, it is still possible for $\hat{z}_t$ to be autocorrelated – so long as T_t is. Indeed, the autocorrelation of growth in the derived z_t series is merely 2%. This suggests that the persistence of US economic growth is *not* due to the persistence of neutral technical change.

Identifying x_t with data on R&D spending, one can construct a series for s_t by iterating on the ideas production function for given values of ϕ and ψ. Interestingly, this yields series for s_t that also display very little autocorrelation: apparently, *neither* type of technology shock z_t nor s_t is persistent. How then can the model be consistent with the highly persistent nature of measured TFP innovations?

The answer is that the model contains a powerful propagation mechanism: the ideas production function itself.[22]

Consider the following setup, in which for simplicity I suppress the endogenous response of x_t to shocks.[23] Suppose that the economy is indeed on a BGP, except that, at some date 0, there is a one-off shock to patent productivity so that $s_0 = e^{\varepsilon_0}$, where ε_0 is small and may be positive or negative. Thereafter, iterate on the ideas production and accumulation equations (7) and (26), while keeping all other variables on the BGP. In particular this means that, for $t > 0$, $s_t = 1$: the shock to s_t has *no persistence* by construction.

Let $\hat{\rho}$ be the measured autocorrelation of detrended traditionally-measured TFP $\hat{z}$:

Claim 1 *When s_t shocks are iid, $\hat{\rho} = \phi + \frac{(1-\phi)(1-\delta_T)}{\gamma_T}$.*

Argument Let $s_0 = e^{\varepsilon_0}$, so that $\varepsilon_0 > 0$ be the shock to quality at date $t = 0$. Thereafter, $s_t = 1$. Let $\kappa_t = s_0 \gamma_x^{t\psi} x_0^{\psi}$, and for simplicity I keep x constant in what follows. Also, define $\kappa_0 = s_0 x_0^{\psi}$ and $\gamma_\kappa^t = \gamma_x^{\psi t}$, so that $\kappa_t = \kappa_0 \gamma_\kappa^t$. This change of variables simplifies the algebra.

The evolution of new ideas is described by:

$$\begin{aligned} q_0 &= \kappa_0 T_0^{\phi} e^{\varepsilon_0}, \\ q_t &= \kappa_t T_t^{\phi} \text{ for } t > 0. \end{aligned}$$

Since $\alpha_i = 1$, $\hat{z}_t(\varepsilon_0) = z_t T_t^{\alpha_k}$. Define the proportion relative to trend $\xi_t \equiv \frac{\hat{z}_t(\varepsilon_0)}{\hat{z}_t(0)}$. Statistically, the autocorrelation $\hat{\rho}$ is estimated from the equation

$$\log \xi_{t+1} = \hat{\rho} \log \xi_t + \epsilon_t.$$

Analytically, the autocorrelation will in principle differ by the date. Define

$$\begin{aligned} \hat{\rho}_t &\equiv \frac{\log \xi_{t+1}}{\log \xi_t} = \frac{\log \hat{z}_{t+1}(\varepsilon_0) - \log \hat{z}_{t+1}(1)}{\hat{z}_t(\varepsilon_0) - \log \hat{z}_t(1)} \\ &= \frac{\log T_{t+1}(\varepsilon_0) - \log T_{t+1}(1)}{\log T_t(\varepsilon_0) - \log T_t(1)}, \end{aligned} \tag{38}$$

[22] Huffman (2002) also makes the point that the ideas-production function may operate as an endogenous propagation mechanism. However, his model lacks a notion of s_t, assuming instead that z_t is the source of shocks.

[23] This will bias my results *against* finding significant propagation. The endogenous response of x_t is likely to quantitatively small, in any case, as ψ is itself relatively small.

where the last step is because α_k and z_t cancel out of all the equations.

In the special case in which $\delta_T = 1$, the sequence $\{\xi_t\}_{t=0}^{\infty}$ has a simple expression that depends neither on t nor ε_0.

$$\begin{aligned}
\xi_0 &= 1, \xi_t = \left(e^{\varepsilon_0}\right)^{\phi^{t-1}\alpha_k}, t > 0 \\
&\Rightarrow \hat{\rho}_t = \frac{\log \xi_{t+1}}{\log \xi_t} = \frac{\phi^t \alpha_k \varepsilon_0}{\phi^{t-1}\alpha_k \varepsilon_0} = \phi.
\end{aligned}$$

For $\delta_T < 1$, there is a caveat. The analytical autocorrelation $\hat{\rho}_t$ will not be constant over time, as $\frac{\log \xi_{t+1}}{\log \xi_t}$ will not be either. Simulations show that, first of all, ρ_t varies very little over time, tending monotonically towards a limit and with the initial deviation depending on the size of the shock (for the calibrated values, the deviation is on the order of 0.1% for a 15% shock). The limit does not depend on the size of the shock, however. Hence, I compute the value of $\hat{\rho}_t$ as $\varepsilon_0 \to 0$. It will turn out that this expression does not depend on t.

Observe that, in the long run, $\gamma_T^{1-\phi} = \gamma_\kappa$. On the BGP, the following relationship holds between κ_0 and I_0 :

$$\begin{aligned}
\kappa_0 &= \frac{\gamma_T T_t - (1-\delta_T) T_t}{\gamma_x^{t\psi} . I_t^{\phi}} = \frac{\gamma_T T_0 \gamma_T^t - (1-\delta_T) T_0 \gamma_T^t}{\gamma_x^{t\psi} . T_0^{\phi} \gamma_T^{t\phi}} \\
\kappa_0 &= \frac{\gamma_T T_0 - (1-\delta_T) T_0}{T_0^{\phi}}.
\end{aligned}$$

The limit of the numerator and denominator of (38) as $\varepsilon_0 \to 0$ is zero, so we must apply L'Hôpital's rule. Define

$$\hat{\rho} \equiv \lim_{\varepsilon_0 \to 0} \hat{\rho}_t = \lim_{\varepsilon_0 \to 0} \left(\frac{\frac{dT_{t+1}(\varepsilon_0)}{d\varepsilon_0}}{T_{t+1}(\varepsilon_0)} \right) \div \left(\frac{\frac{dT_t(\varepsilon_0)}{d\varepsilon_0}}{T_t(\varepsilon_0)} \right),$$

which conjectures that $\hat{\rho}$ does not depend on t. It is easily shown that

$$\begin{aligned}
\frac{dT_{t+1}(\varepsilon_0)}{d\varepsilon_0} &= \frac{dT_t(\varepsilon_0)}{d\varepsilon_0} \left[\kappa_t \phi T(\varepsilon_0)_t^{\phi-1} + (1-\delta_T) \right] \\
&\Rightarrow \hat{\rho} = \lim_{\varepsilon_0 \to 0} \hat{\rho}_t \frac{T_t(\varepsilon_0) \left[\kappa_t \phi I_t^{\phi-1} + (1-\delta_T) \right]}{T_{t+1}(\varepsilon_0)}.
\end{aligned}$$

Solving,

$$\begin{aligned}
\hat{\rho} &= \lim_{\varepsilon_0 \to 0} \frac{\kappa_t \phi T_t(\varepsilon_0)^{\phi} + (1-\delta) T_t(\varepsilon_0)}{(1-\delta_T) T_t(\varepsilon_0) + \kappa_t T_t(\varepsilon_0)^{\phi}} = \lim_{\varepsilon_0 \to 0} \frac{\kappa_t \phi + (1-\delta_T) T_t(\varepsilon_0)^{1-\phi}}{(1-\delta_T) T_t(\varepsilon_0)^{1-\phi} + \kappa_t} \\
&= \lim_{\varepsilon_0 \to 0} \frac{\kappa_t - \kappa_t(1-\phi) + (1-\delta_T) T_t(\varepsilon_0)^{1-\phi}}{(1-\delta_T) T_t(\varepsilon_0)^{1-\phi} + \kappa_t} = 1 - \lim_{\varepsilon_0 \to 0} \frac{\kappa_t(1-\phi)}{(1-\delta_T) T_t(\varepsilon_0)^{1-\phi} + \kappa_t}.
\end{aligned}$$

Finally, as $\varepsilon_0 \to 0$, $T_t \to T_0\gamma_T^t$, which yields a value of $\hat{\rho}$ that is indeed independent of time.

$$\begin{aligned}\hat{\rho} &= 1 - \frac{(1-\phi)}{(1-\delta_T)\frac{\left(T_0\gamma_T^t\right)^{1-\phi}}{\kappa\gamma_\kappa^t}+1} = 1 - \frac{(1-\phi)}{(1-\delta_T)\frac{(T_0)^{1-\phi}}{\kappa}+1} \\ &= 1 - \frac{(1-\phi)}{(1-\delta_T)\frac{T_0^{1-\phi}T_0^{\phi}}{\gamma_T T_0-(1-\delta_T)T_0}+1} = \phi + \frac{(1-\phi)(1-\delta_T)}{\gamma_T}.\end{aligned}$$

When $\varepsilon_0 \neq 0$, $\hat{\rho}_t$ will not be constant. However, $\hat{\rho}_t \lesseqgtr \hat{\rho} \iff \varepsilon_0 \lesseqgtr 0$, so estimates of $\hat{\rho}$ will tend to measure the correct value, since innovations to detrended quality average zero by definition. Again, simulations show that these deviations are negligible.

First, notice that $\hat{\rho}$ is bounded below by ϕ so that, even the economy forgets new ideas immediately ($\delta_T = 1$), there will be endogenous persistence. Observe also that $\hat{\rho}$ is decreasing in δ_T: the undepreciated past stock of ideas provides another channel for the influence of the innovation ε_0 to perdure. It is also decreasing in γ_T, as larger values of γ_T imply that new knowledge more rapidly crowds out the effect of past innovations. On the other hand, it is increasing in ϕ.

Even using the lowest value of ϕ consistent with the calibration ($\phi = 0.23$) yields $\hat{\rho} = 0.89$, and an intermediate value of $\phi = 0.62$ is sufficient to yield $\hat{\rho} = 0.97$, in line with the values of Cooley and Prescott (1995) that are typical of the real business cycle literature. The entire autocorrelation of measured TFP can be accounted for by endogenous propagation through the ideas' production function – even when shocks are *iid*. What matters is not so much the value of ϕ as the fact that the empirically relevant range for δ_T is sufficiently low that it provides significant feedback after any innovation to s_t.

Another interesting observation regards variance. An investment-specific shock of fully 15% is necessary to generate an increase in detrended $\hat{z}_t$ of 0.7%, which is the standard deviation reported in Cooley and Prescott (1995). This is for two reasons. First, it is the stock – not the flow – of new ideas that matters, so that even a large investment-specific shock will lead to only a small change in $\hat{z}_t$. Second, the effect of knowledge upon aggregates is intermediated by the capital share α_k, which dampens the aggregate effect of investment-specific shocks.

There is no reason to believe that the outcome of research and development should be easily predictable – indeed, by the quiddity of discovery itself, the yield of research activity is likely to be highly variable,[24] which is consistent with the dynamics of p_t and s_t Nonetheless, the endogenous propagation mechanism implies that the extreme variability inherent to scientific discovery is consistent with very smooth aggregate dynamics.

This point is of quantitative relevance. Depending on the values of ϕ and ψ assumed, the standard deviation detrended s_t ranges from 4 up to 14%, consonant with the finding that a 15% shock leads to a shock to measured TFP of 0.7% – also one standard deviation. Thus, shocks to s_t of empirically plausible size and no persistence can account for the entirety of the magnitude and persistence of measured TFP shocks. By contrast, shocks to z_t are small and display no persistence at all.

[24] See Schrerer and Harhoff (2000). Stevens and Burley (1997) argue that the ratio of potential research projects to actual innovation-products is on the order of 3000:1.

B Notes on the Data

There are two alternative measures of the output of research activity: patent grants, and patent applications. I use patent *applications* for the following reasons. The "application-to-grant" lag is about 2 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 which, 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 and/or otherwise economically useful ideas. Kortum and Lerner (1998) argue that there has been an acceleration of innovative activity in the US 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 non-scientists may have economically useful ideas. It too displays very little high-frequency variation.

An further data issue is that y_t does not equal the NIPA measure of output unless $\alpha_i = 1$. Using (11), the resource constraint (4) becomes:

$$y_t \geq c_t + \alpha_i \int p_{jt} u_{jt} + x_t. \tag{39}$$

However, given α_i this is easily corrected.

Equation (25) computes two measures of the knowledge stock, once using patents and one using p_t. To initialize the two stocks, I assume that the economy was on the balanced growth path in 1947: small deviations in this assumption do not affect results. Patent data is available as far back as the late 18th century: constructing patent data based on the entire series does not change results.

C z_t and T_t in the short run

Although the current paper differs from Greenwood et al (1997) in amalgamating structures and equipment, the z_t series resembles theirs. In particular, it preserves the result that, since the mid-1970s, "true" TFP growth has decreased considerably – although it does not display the TFP *decline* that they detect. Before 1970, z_t grows by 1.7% per annum, dropping to under 0.4% per annum after 1970. The reason why the slowdown is not as precipitous here is because

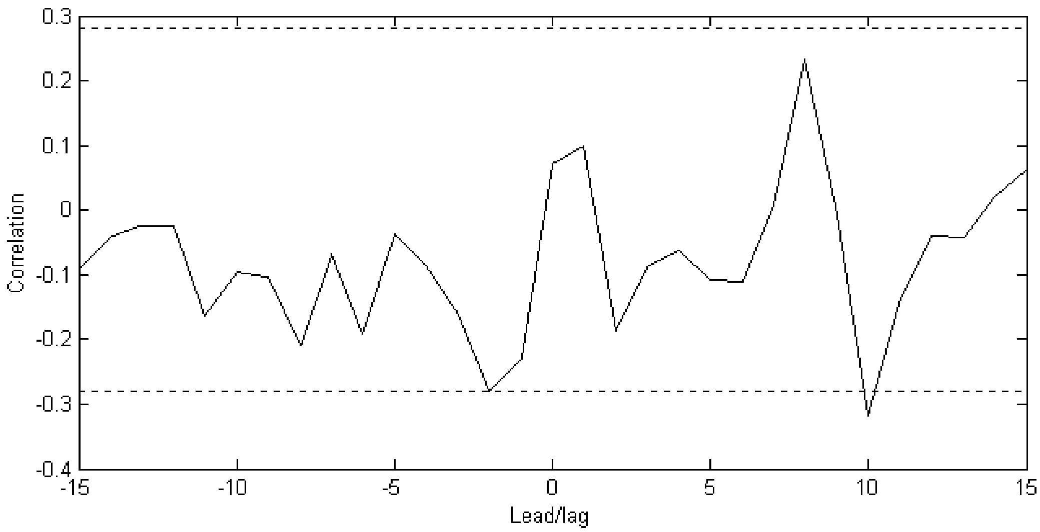

Figure 8: Cross-correlogram for growth in the price-based knowledge stock (t) and neutral technical change $(t+k)$.

considering structures lowers the degree to which the accumulation of quality-adjusted capital accelerates after 1970.

In general, one might argue that perhaps not *all* ideas are investment-specific, so that the use of ideas may be reflected in the behavior of z_t. However, the cross-correlogram of growth in z_t and T_t^{price} in Figure (8) shows no significant correlation at any lead or lag[25] – in contrast to Figure (9), which displays the correlogram for T_t^{price} and T_t^{patent}. This suggests that z_t primarily reflects aspects of productivity – such as institutional change, or the quality of fiscal or monetary policy – that are very different from the use of ideas.

Figure (9) displays the cross-correlogram of the growth rate of the two knowledge stocks. As can be seen, there is some co-movement at low leads and lags. This is to be expected if knowledge takes 1-2 years to diffuse through the economy, or if knowledge is an input into its own production. However, innovation apparently leads to implementation fairly quickly, as found by Caballero and Jaffe (1993).

[25] The same is true if one takes the patent stock rather than the price-based series for T_t.

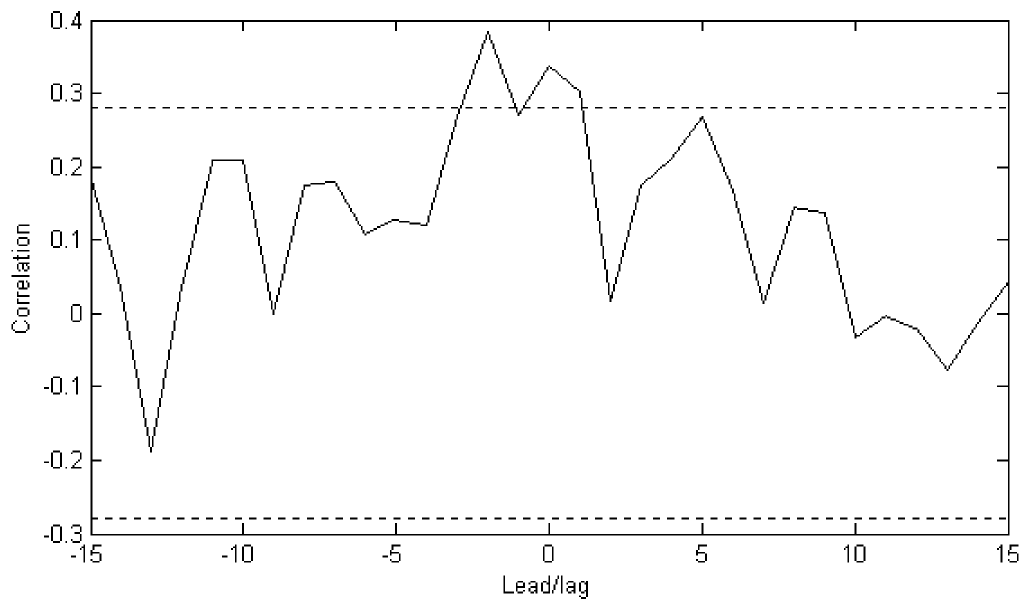

Figure 9: Cross-correllogram for growth in the patent-based knowledge stock (t) and the price-based stock $(t+k)$.