A General Theory of Sustained Growth with Co-Existing GPTs*

Kenneth I. Carlaw[1] and Richard G. Lipsey[2]

1 University of British Columbia and University of Waikato, 2 Simon Fraser University, Email: Kenneth.carlaw@ubc.ca

Keywords: *General Purpose Technologies: sustained growth, logistic diffusion*

EXTENDED ABSTRACT:

In this paper we generalize the single-GPT growth model first developed for Lipsey and Carlaw (2006) and subsequently elaborated for Lipsey, Carlaw and Bekar (2005) to cover any number of GPTs. The model has three main sectors: a consumption goods sector that uses the results of applied research in its production function, an applied research sector that uses the GPTs invented by the pure research sector in its production function that is used to produce applied knowledge that is in turn used in the production of both consumption goods and pure knowledge, and a pure research sector that produces pure knowledge that occasionally results in the discovery of a new GPT. All three sectors also use the economy's scarce resources in their production processes. The full potential of any new GPT diffuses to the applied research sector according to its own logistic diffusion process. In this model, there can be as many pure research activities, as many GPTs, as many applied R&D activities, and as many consumption goods as desired. Simulation of the model matches the accepted stylized growth facts and demonstrates the model's potential for addressing important growth policy questions such as what is the appropriate mix between applied and pure R&D support for sustained long term economic growth.

I. INTRODUCTION

In this paper we incorporate the following empirically established characteristics of general purpose technologies (GPTs).[1] These characteristics have been detailed in Lipsey, Carlaw and Bekar (2005, hereafter LCB) but have been omitted from all previous formal GPT models.

1. The use of a new GPT spreads slowly through the economy and its full diffusion takes decades. During that process, its efficiency increases greatly.

2. GPTs occur in each of several "classes" of technology, such as materials, ICTs, power sources, transportation equipment, and organizational forms (e.g., the factory system).

3. At any one time, there are typically several GPTs in use, at least one "version" from each class and often there are several versions of any one class in simultaneous use.

4. Different versions of GPTs in any one class typically compete with each other while in contrast GPTs of different classes often complement each other.

5. Many of the sources of uncertainty in invention and innovation are modeled including: (i) how much potentially useful pure knowledge will be discovered by any given amount of research activity; (ii) the timing of the discovery of new technologies; (iii) how productive a newly innovated GPT will be over its lifetime; (iv) how well the new GPT will interact with GPTs of other classes; (v) how long a new GPT will continue to evolve in usefulness; (vi) when it will begin to be replaced by a new superior version of a GPT of the same class (vii) how long that displacement will take; and (viii) if the displacement will be complete.

Over the last several years there has been substantial activity in applying GPT theory to issues in economic history industrial organization and economic policy. These applications have been hampered by the nature of the first generation of GPT models found in Helpman (1998), which cover only the evolution of a single GPT that dominates the macro performance of the whole economy. This has led, for example, to attempts to infer the existence of a single GPT by examining the aggregate behavior of the economy as measured by its GDP.

In spite of the need for more empirically relevant models, there have been no further developments of GPT theory using standard techniques. One reason for this is that because these models all use dynamically stationary equilibrium concepts that become analytically intractable when more empirically relevant assumptions are introduced. We have circumvented this modeling problem by building models of GPT-driven growth that incorporate the complex structure of technology yet are analytically manageable because we forgo the use of a stationary foresighted equilibrium. Agents face a future that is uncertain and do the best they can with limited knowledge of current relations.

Knowledge is produced at two different levels: a pure research level that produces fundamental ideas and an applied research level that uses these ideas to produce applications. The pure research sector could be producing GPTs, or knowledge that is less universal but still provides fodder for the applied R&D sector. It could also be producing knowledge that allows GPTs developed abroad to be used to produce applications.[2]

II. THE SIMULTANEOUS GPT MODEL

In our model we initially incorporate the first four numbered characteristics listed in the introduction while the various sources of uncertainty outlined in point 5 are introduced as we specify our model formally in the next section. We introduce multiple activities in each of our three sectors. The pure research sector has X laboratory complexes, which we call 'labs' for short, each producing a distinct class of pure knowledge and making use of all types of applied knowledge. Each lab occasionally invents a new version of a GPT in its particular class. The applied R&D sector has Y 'facilities,' each producing a distinct type of useful knowledge, and employing one version of each class of GPT. The consumption sector has I industries, each producing a different consumption good, and making use of the stocks of applied knowledge. The productivity of each GPT, $\left(G_{n_x}\right)_t$, evolves according to a logistic diffusion process that is given in equation (8) below, until its full potential has been realized. In order to allow a new GPT to spread through the economy over time, the productivity of each GPT is allowed to differ in each of the applied R&D facilities. This is done by multiplying each GPT's productivity by a parameter v that is specific both to that GPT and the R&D facility in which it is operating.

Each R&D facility is initially seeded with one version of each class of GPT. When a new 'challenge' GPT is invented each applied R&D facility must decide whether to adopt it or to stay with the version of the 'incumbent' GPT of that class currently in use. Because of the uncertainties listed in (5) above we use a simple comparison of the challenger's current level of productivity with the current productivity of the incumbent. If a new challenging GPT has a lower current productivity than the incumbent in all applied R&D facilities, it is sent back to the pure research sector for further development and it is reconsidered for adoption in every subsequent period. If it has a higher productivity in one or more facilities, it is adopted there. Each subsequent period, the facilities that have not adopted this new GPT compare its evolving productivity with the evolving productivity of their incumbent and switch when the former exceeds that latter. It is possible that there may be several versions of GPT of any one class in use at any one time.

Some Terminology

The variable x indicates a specific *class* of GPT. The index n_x identifies a GPT in the sequence of *versions*, 1...n, of GPTs in class x that has been invented and adopted by at least one applied R&D facility. At any one time, the latest version adopted in that class is denoted by n_x and the previously adopted version by $(n\text{-}1)_x$. We refer to the productivity of the most recently invented GPT from lab x at time t as $(G_{n_x})_t$ and the previously invented GPT in that class as $(G_{(n-1)_x})_t$. The variable t_{n_x} refers to the invention date of the latest version of the class of GPT invented by lab x, while $t_{(n-1)_x}$ refers to the invention date of the previous version from lab x and $(t-1)_{n_x}$ refers to the period just prior to the invention of GPT of version n_x.

Relations Among GPTs

The productivity coefficient in each applied R&D facility is the geometric mean of the productivities of the GPTs that it uses, each pre-multiplied by an associated v value (see equation (2) below).

The vs can be set out in a YxX matrix where a row indicates the class of GPT and a column the research facility. We call this the 'operative v matrix'. When a new version of a particular class of GPT is invented, it brings with it its own matrix of potential new vs, which we call its 'potential v matrix'. When each R&D facility decides whether or not to adopt the challenger, it must form expectations over the potential vs required to make the calculation of potential output using the challenger. We denote these expected values by $\bar{v}$. Because the $\bar{v}$s vary across any row, different facilities will evaluate the productivity of the challenger differently and so some may adopt it while others do not. When a particular applied R&D facility adopts the challenger, the challenger's potential vs replace the existing ones in that facility's column vector in the operative v matrix. The resulting change in the vs associated with GPTs in classes other than the newly adopted GPT indicates whether the new GPT cooperates better or worse with the existing GPTs of other types than did the replaced incumbent.

The full identification of each v requires four indexes, $v_{y,z}^{n_x}$. The superscript n_x tells us that the v is modifying a GPT of class x, version n. The subscript tells us that the GPT is being used by research facility y and replaced the previous v when a new version of class z GPT was adopted. Thus, z indicates the class of GPT that was last adopted by that facility and it is the source of all the vs in that column, while n_x indicates the version of class-x GPT to which the v is being applied. Only the n and z change over time, the n indicating the version of the particular class-x GPT being used by the facility in question and the z indicating the class of the last GPT adopted by that facility (and hence the source of all the vs in that facility's column vector). The vs are determined randomly in this version of our model.

The Model

For simplicity, we assume a fixed supply of the composite resource, R, that is allocated by private price-taking agents in the consumption and applied R&D sectors and by a government that implicitly taxes the applied R&D and consumption sectors to fund pure research. We assume (1) agents are profit seeking and would maximize their profits if they could, (2) agents are price takers, (3) agents are operating under conditions of uncertainty for all of the reasons set out in characteristic 5 (4) because they cannot assign probabilities to alternative future consumption payoffs, agents seek to maximize their profits on the basis of current prices.

The constraint imposed by the composite resource is

(1) $$R_t = \sum_{i=1}^{I} r_t^i + \sum_{y=1}^{Y} r_t^y + \sum_{x=1}^{X} r_t^x$$

The output of applied knowledge from each applied R&D facility, y, depends on the amount of the composite resource it uses and its productivity coefficient, which is the geometric mean of each $\left(G_{n_x}\right)_t$ term multiplied by its corresponding v term, as shown in equation (2).

(2) $$a_t^y = \left[\prod_{x=1}^{X} (v_{y,z}^{n_x} (G_{n_x})_{t-1})^{\beta_x} \right]^{\frac{1}{X}} (r_t^y)^{\beta_{X+1}}$$

$\beta_x \in (0,1]\ \forall x \in X,\ \beta_{X+1} \in (0,1)$

In the consumption sector, we make the simplifying assumptions (1) that there are the same number of applied research facilities and consumption industries, $I = Y$, and (2) that the knowledge produced in each of the facilities, y, is useful only in the one corresponding consumption industry, $i = y$. The production function for each of the I industries in the consumption sector is then expressed as follows:

(3) $$c_t^i = (\mu A_{t-1}^y)^{\alpha_y} (r_t^i)^{\alpha_{Y+1}},$$

$\alpha_y \in (0,1]\ \forall y \in Y,\ \alpha_{Y+1} \in (0,1)$ and $i = y$

The stock of applied knowledge generated from each facility accumulates according to:

(4) $$A_t^y = a_t^y + (1-\varepsilon) A_{t-1}^y,$$

where $\varepsilon \in (0,1)$ is a depreciation parameter.

The Pure Knowledge Sector and the Endogenous Production of GPTs

We assume that the productivity coefficient in each lab is the geometric mean of the various amounts of the Y different kinds of applied knowledge that are useful in further pure research (one for each applied R&D facility and each raised to a power σ_y). The output of pure knowledge in lab x, g_t^x, is a function of this productivity coefficient and the amount of the composite resource devoted to that lab.

(5) $$g_t^x = \left[\prod_{y=1}^{Y} \left((1-\mu) A_{t-1}^y \right)^{\sigma_y} \right]^{\frac{1}{Y}} \left(\theta_t^x r_t^x \right)^{\sigma_{Y+1}},$$

$\sigma_y \in (0,1],\ \forall\ y \in Y$ and $\sigma_{Y+1} \in (0,1)$.

The term, θ_t^x, models the uncertainty surrounding the productivity of the composite resource devoted to pure research as indicated in uncertainty source 5 (i) listed above.

The stocks of *potentially useful* knowledge produced by each of the X labs accumulate according to:

(6) $$\Omega_t^x = g_t^x + (1-\delta)\Omega_{t-1}^x$$

where $\delta \in (0,1)$ is a depreciation parameter.

New GPTs are invented infrequently in each of the X labs and their invention date is determined when the drawing of the random variable $\lambda_t^x \geq \lambda^{*x}$ (uncertainty source 5 (ii)). For simplicity, we let the critical value of lambda for each of the X labs be the same: $\lambda^{*x} = \lambda^*\ \forall\ x \in X$. When at any time, t, $\lambda_t^x \geq \lambda^*$, indicating that a new version of class-x GPT is invented, the index t_{n_x} is reset to equal the current t, and n_x is augmented by one.

Because agents do not know how productive a new GPT will be over its lifetime (uncertainty source 5 (iii)), they must make their adoption decisions with incomplete information. Out of the many adoption criteria that we mentioned earlier, we use for this paper the rule to adopt the class-x challenger wherever its initial productivity is expected to exceed that of the class-*x* incumbent. The expected productivity from using the challenger is determined from its new $\left(G_{n_x}\right)_t$ term defined in equation (8) below, the productivities of the other classes of GPT used by that facility and the challenger's expected $\bar{v}$s , which alter those productivities. We assume that agents considering adopting the new GPT can correctly predict the potential v associated with the version of the challenger, the minimum knowledge that they need to make some kind of evidenced-based adoption decision, but that they predict the other $\bar{v}$s in their facility's column vector will be unchanged. This prediction may be falsified since the potential vs brought in by the challenger may differ from those in the current operative matrix (uncertainty source 5(iv)).

Since in each applied R&D facility the only v that agents expect to change is the one associated with the challenging x-class GPT, we can compare the productivities for any of the y facilities by simply comparing the $v_{y,z}^{(n-1)x}\left(G_{(n-1)x}\right)_{t_{n_x}}$ that would be produced if the incumbent were left in place with the $\bar{v}_{y,z}^{n_x}\left(G_{n_x}\right)_{t_{n_x}}$ that is expected to be produced if the challenger were adopted. This comparison is made in each of the Y applied R&D facilities at time $t = t_{n_x}$ so the test, stated generally for all applied R&D facilities, is:

$$(7)\quad \left[\bar{v}_{y,z}^{n_x}\left(G_{n_x}\right)_{t_{n_x}}\right] \geq \left[v_{y,z}^{(n-1)_x}\left(G_{(n-1)_x}\right)_{t_{n_x}}\right] \text{ for each}$$

$y \in [1,Y]$. If the test is passed, the new GPT is adopted in facility y.

If none of the y applied R&D facilities adopts the GPT, it is returned to its pure knowledge industry. The indexes t_{n_x} and n_x are incremented back to their previous values—it is as if the favorable drawing of $\lambda_t > \lambda^*$ had not occurred. Pure research then continues to improve the new GPT and it is reconsidered every period.

The diffusion process defined in equation (8) below starts for the newly arrived GPT.

$$(8)\quad \left(G_{n_x}\right)_t = \left(G_{(n-1)x}\right)_{(t-1)n_x} + \left(\frac{e^{\tau+\gamma(t-t_{n_x})}}{1+e^{\tau+\gamma(t-t_{n_x})}}\right)\left(\Omega_{t_{n_x}}^{x} - (G_{(n-1)x})_{(t-1)n_x}\right).$$

The equation shows actually useful general purpose knowledge, $\left(G_{n_x}\right)_t$, becoming available for use in applied research according to a logistic diffusion process in which t_{n_x} is the invention date of the version n_x, of the class-x GPT, $\Omega_{t_{n_x}}^{x}$ is the full *potential* productivity of the new version of GPT x, $\left(G_{(n-1)x}\right)_{t_{(n-1)x}}$ is the *potential* productivity of the version that it replaced, evaluated at the time at which that earlier version was last used, $t_{(n-1)_x}$ and γ and τ are calibration parameters that control the rate of diffusion.[3]

In addition to modeling uncertainty sources 5 (iii) and 5 (vi) which we discussed above, the interaction of the G and the v terms model uncertainty about when a new GPT will begin to be replaced by a challenger (source 5 (vi)), how long it will take to displace the incumbent GPT (source 5 (vii)), and whether a GPT will be completely displaced (source 5 (viii)). Thus, they model the many aspects of the general observation that the applied potential of a GPT cannot be precisely predicted when it is originally being developed.

In the subsequent periods, the test in equation (7) is modified to note the productivity changes that occur over time:

$$(7')\quad \left[\bar{v}_{y,z}^{n_x}\left(G_{n_x}\right)_t\right] \geq \left[v_{y,z}^{(n-1)x}\left(G_{(n-1)x}\right)_t\right]$$

for each $y \in [1, Y]$ that has not yet adopted GPT G_{n_x}.

In our model the economy's GDP is the current period's output of the consumption and applied R&D sectors plus resource costs of the pure knowledge sector.[4] We count the output of the applied R&D sector as additions to the capital stock because it is applied R&D that gets embodied in capital goods. Pure knowledge is an input into the applied research sector, being of no use in producing GDP until it is turned into applied knowledge. It is not, therefore, a part of the capital stock as usually measured. So for our purposes, the stock of applied knowledge is the capital stock and the flow of applied knowledge is investment.

Multi-GPT Resource Allocation

In the pure knowledge sector, the government pays for and allocates a fixed amount of the generic resource, R, to each of the pure knowledge producing labs. Producers in the applied R&D and consumption sectors maximize their profits each period taking prices as given. The prices for output from the I consumption industries are derived from the maximization of an aggregate utility function which we assume is additively separable across the I consumption goods.

$$(9)\quad U=\sum_{i=1}^{I}\left(c^{i}\right)^{\phi^{i}} \text{ and } \phi^{i}=\phi^{i'}=1, i\neq i'\,\forall i,i'\in I$$

Maximizing this utility function yields:

$$(10)\quad \frac{MU^{i=1}}{MU^{i\neq 1}}=\frac{P^{i=1}}{P^{i\neq 1}}=\frac{\phi^{i=1}\left(c^{i=1}\right)^{\phi^{i=1}-1}}{\phi^{i\neq 1}\left(c^{i\neq 1}\right)^{\phi^{i\neq 1}-1}}$$

Since $\phi^{i}=1\ \forall\ i\in I$ it follows that $P^{i=1}=P^{i\neq 1}$.

We assume a competitive equilibrium in the market for the composite resource so it earns the same wage, w, regardless of where it is allocated.

Each consumption industry's profits are:

$$(11)\quad \pi^{i}=P^{i}c^{i}-wr^{i}-P^{y}A^{y}$$

We suppress the time subscripts here because we assume each agent in the consumption industry re-evaluates its profit function every period. Profit maximization yields the following FOCs in each of the I consumption industries:

$$(12)\quad \begin{aligned} &P^{i}mp\,r^{i}-w=0\\ &P^{i}mpA^{y}-P^{y}=0\end{aligned}$$

where mp represents marginal product. From the first FOC, the assumption the $P^{i}=1$, and the definition of the production function for industry i we get the reduced form demand for the composite resource:

$$(13)\quad r^{i}*=\left[\frac{\alpha_{Y+1}}{w}\left(\mu A^{y}\right)^{\alpha_{y}}\right]^{\frac{1}{1-\alpha_{Y+1}}},$$

From the combination of both FOCs from the profit function for consumption industry i and the definition of the production function we get, $\frac{w}{P^{y}}=\frac{\alpha_{Y+1}}{\alpha_{y}}\frac{A^{y}}{r^{i}}$ which implies:

$$(14)\quad P^{y}*=\frac{w}{\alpha_{Y+1}A^{y}}\left[\frac{\alpha_{Y+1}}{\alpha_{y}}\left(\mu A^{y}\right)^{\alpha_{y}}\right]^{\frac{1}{1-\alpha_{Y+1}}}$$

Each applied R&D facility maximizes profits (15). The pure knowledge input in the form the currently adopted set of X GPTs is provided freely to the applied R&D facilities.

$$(15)\quad \pi^{y}=P^{y}a^{y}-wr^{y}$$

Maximization yields the following demand for the composite resource from each applied R&D facility:

(16)

$$r^{y*}=\left[\beta_{X+1}\left[\prod_{x=1}^{X}\left(v_{y,z}^{n_x}\left(G_{n_x}\right)_t\right)^{\beta_x}\right]^{\frac{1}{X}}\frac{P^{y*}}{w}\right]^{\frac{1}{1-\beta_{X+1}}}$$

With these resource demand equations we now have a complete description of the allocation of the composite resource across the three sectors.

Simulation of the Model

We restrict the model to three industries within the consumption sector, three facilities in the applied R&D sector and three labs in the pure knowledge sector ($I = Y = X = 3$). We choose $\alpha_{Y+1}=\beta_{X+1}=\sigma_{Y+1}$ and $\alpha_{y}=\beta_{x}=\sigma_{y}$ to impose symmetry across sectors and specific activities (i.e., industries, facilities and labs) within sectors.

We choose values of the parameters and initial conditions so that the model will replicate the accepted stylized facts of economic growth.[5] Some values are chosen to ensure consistency with observed data in the following ways: diminishing returns to our composite resource in all activities ($\alpha_{Y+1}=\beta_{X+1}=\sigma_{Y+1}=0.3$), an average annual growth rate between 1.5% and 2% ($\varepsilon=\delta=0.25$);[6] GPTs arriving on average every 35 years ($\lambda^{*}=0.66$). We choose $\gamma=0.07$ and $\tau=-7$ to so that 90% of a GPT's diffusion occurs over 130 years. We choose $\mu=0.95$ in order to set the income shares of the labor and capital (physical and human) at approximately 0.3 and 0.7. We set $\alpha_{y}=\beta_{x}=\sigma_{y}=1$ to ensure that knowledge had constant returns.

The Table 1 gives the parameter values and initial conditions used to simulate the results of the multi-GPT model as reported in the text and shown in the figures.

TABLE 1

NUMERICAL SIMULATION OF THE MULTI-GPT MODEL

Parameter values:

$$\alpha_y = 1 \quad \forall i \in [1,Y] \qquad \alpha_{Y+1} = 0.3 \qquad \beta_x = 1 \quad \forall x \in [1,X]$$

$$\beta_{X+1} = 0.3 \qquad \sigma_y = 1 \quad \forall y \in [1,Y] \qquad \sigma_{Y+1} = 0.3$$

$$\gamma = 0.07 \qquad \tau = -7 \qquad \phi_i = 1 \quad \forall i \in [1,I] \qquad \varepsilon = 0.025$$

$$\delta = 0.025 \qquad \mu = 0.95 \qquad I = 3 \qquad Y = 3 \qquad X = 3$$

$$\lambda^* = 0.66 \qquad R = \sum_{i=1}^{I} r_t^i + \sum_{y=1}^{Y} r_t^y + \sum_{x=1}^{X} r_t^x = 1000$$

Initial conditions:

$$\left(G_{n_x}\right)_0 = 1 \quad \forall x \in (1,X) \qquad A_0^y = 1 \quad \forall y \in [1,Y] \qquad n_x = 1 \qquad \theta_0^x = 1$$

$$t_{n_x} = 2 \qquad P_0^i = 1 \ \forall i \in [1,I] \qquad d_{y,x} = 1 \qquad w_0 = 1$$

The set of θ^x are random variables distributed uniformly with support [0.9, 1.1]. The λ's are derived from beta distributions, where each distribution is defined as $\text{beta}(x|\psi,\eta) = \frac{x^{(\psi-1)} x^{(\eta-1)}}{\text{Beta}(\psi,\eta)}$. Beta($\psi,\eta$) is the Beta function, and ψ and η are parameters which take on positive integer values. We choose $\psi = 5$ and $\eta = 10$.

Next we need to determine the *v*s in the for the challenger's potential matrix. They are calculated as the vs in the operative m matrix modified by the addition of a random variable π drawn from a uniform distribution with support [-0.05, 0.05]. For the example, the potential *v*s associated with a challenger of version $n+1$ from class x = 2 are derived from the incumbent's *v*s as follows:

$$(16) \quad v_{y,2}^{(n+1)2} = \begin{cases} v_{y,2}^{n2} + \pi & \text{if } (v_{y,2}^{n2} + \pi) \in [0.5, 1.5] \\ 0.5 & \text{if} \quad v_{y,2}^{n2} + \pi < 0.5 \quad y \in [1,\ 3] \\ 1.5 & \text{if} \quad v_{y,2}^{n2} + \pi > 1.5 \end{cases}$$

To demonstrate the some important properties of the model and its applicability for analysis of real world problems we do three things. First, we show that observable macroeconomic aggregates are not determined by a single GPT and therefore are not particularly useful for identifying the arrival and diffusion of a given GPT. Second, we demonstrate that the model is consistent with the accepted Stylized facts of growth. Third, we illustrate the model's ability to address certain important policy questions by posing and analysis the question what is the appropriate policy mix to support pure versus applied research.

One of our objectives in this paper is demonstrate that macro aggregates such as the growth rate of output cannot be used to identify the arrival and diffusion of a single GPT (or that a single GPT can determine the pattern of macro aggregates). Figure 2 shows the growth rates of output and of useful pure knowledge generated from the $X = 3$ classes of pure knowledge production.

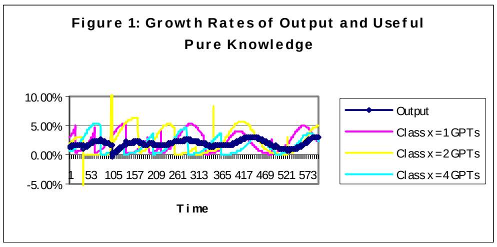

Figure 1: Growth Rates of Output and Useful Pure Knowledge

The figure indicates clearly that the arrival and evolution of any one GPT does not necessarily determine the pattern of growth of aggregate output.

At the beginning of this section we laid out our rational for calibrating the simulation of our model. Here we present some of the general results and we refer the reader to the following web URL http://www.sfu.ca/~rlipsey/ to see how our model meets the rest of the accepted stylized facts.

As shown in Figure 2, our model exhibits a positive long term trend growth of 1.9% but the growth rate fluctuates around this trend quite dramatically over time.

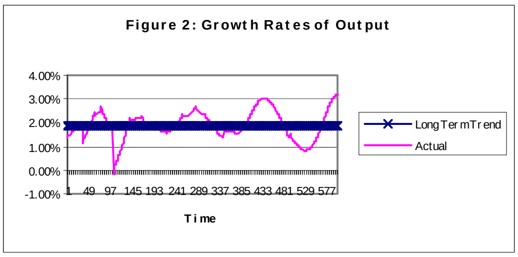

Figure 2: Growth Rates of Output

In Figure 2 there is no pattern of convergence between the long run trend and the growth rate. The variance is 0.0018. The spikes in the growth rate occur as a result of the slight jumps in the level productivity. The variance of our output growth series is not as large as the actual growth rates of GDP for Canada and the US calculated from Madison's data shown in Figure 3. However, spikes in the growth rates are present in both the real and the simulated data.

Figure 3: Anual Growth Rate of Real GDP

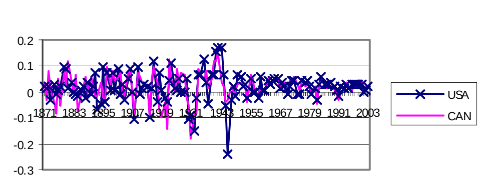

The variance of growth rate for Canada is 0.00264 which is much higher than that of our simulation. The key point is that growth rates are volatile for both the real and the simulated data.

Next, we use the model to address the policy question of what is the appropriate mix of support for pure versus applied R&D. We run four different simulations that model four different mixes of pure and applied R&D public funding support. We calibrate the model such that 3.4 % of total GDP is allocated by government to pure knowledge production. We choose 3.4% because Canada allocates 1.9% of total GDP to R&D and 2.5% (i.e., 4.4% in total) of GDP to post-secondary education and research, but these figures count both private and public contributions.[7] The simulations are run for an initial period of time with all of the government support going to pure knowledge production then we change the mix of pure and applied R&D that is funded publicly. Figure 4 shows the natural log of GDP from the point at which the policy change occurs forward.

In the first policy simulation all of the 3.4% of government allocated support goes to producing pure knowledge and the applied R&D sector privately produces applied knowledge. In simulations 2 – 4 some mix of pure and applied support is provided. In all cases the pure sector is wholly publicly funded and the applied sector produces using a mix of private and public funding. The second policy simulation evenly splits the government support between the pure knowledge and applied knowledge sectors. In the third policy simulation 87.5% of the government provided support to research goes to applied R&D and only 12.5% to pure knowledge production. In the fourth policy simulation 98.25% of the public support goes to applied R&C and a mere 1.25% goes to pure knowledge production.

Figure 4: Natural log of GDP

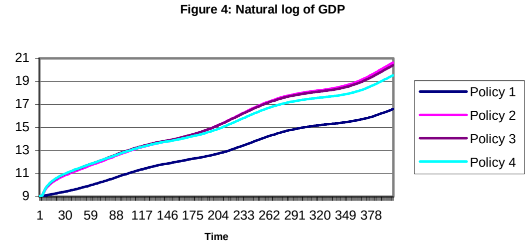

It is quite clear from Figure 4 that some mix of public support is preferable to government support going exclusively to the pure knowledge sector. However, the appropriate mix of applied and pure R&D that is publicly funded is not obvious from the last three simulations. Clearly there is an early gain to allocating a large amount of public support away from pure knowledge to applied R&D, provided that some amount of pure knowledge has been accumulated prior to the policy change. However, as time goes by the gain from allocating the bulk of public support to applied R&D is more than offset by the loss from not maintaining a sufficiently high proportion of support for pure knowledge production. This is shown by the fact that Policies 3 and 4 out perform Policy 2 initially but later 3 and 4 under perform relative to 2.

The policy analysis shown here is obviously crude and requires significant amounts of context specific data to calibrate to the model before explicit policy advice can be given in detail. However, what is clear in the qualitative results is that a "one size fits all" approach of solely funding pure or applied R&D is not optimal. And some mix of support to the two activities is preferred. It is also clear that within our model private activity on its own will under produce applied R&D relative to some mix of private and publicly supported applied R&D production.

BIBLIOGRAPHY

Carlaw, K. I. and R.G. Lipsey (2006) "GPT-Driven, Endogenous Growth", *Economic Journal* 116, 155-74.

Cohen and Levinthal (1990), "Absorptive capacity: A new perspective on learning and innovation", *Administrative Science Quarterly*, Volume 35, Issue 1 pg. 128-152.

Easterly and Levine (2001) "It's not factor accumulation: stylized facts and growth models", *World Bank Economic Review* 15(2).

Helpman, E. (1998) (ed.) *General Purpose Technologies and Economic Growth* (Cambridge Mass.: MIT Press)

Kaldor, N. (1961) "Capital Accumulation and Economic Growth" in Lutz, (ed.), *Theory of Capital.*

Lipsey, Richard G., Kenneth I. Carlaw and Cliff Bekar (2005) *Economic Transformations: General Purpose Technologies and Long-run Economic Growth* (Oxford University Press: Oxford, UK).

Statistics Canada (2003) "Education Indicators in Canada: Report of the Pan-Canadian Education Indicators program" *Catalogue 81-582-XIE* (Ottawa).

[*] The authors gratefully acknowledge funding support for this research from the Royal Society of New Zealand's Marsden Fund, grant number 04 UOC 020.

[1] "A GPT is a single generic technology, recognizable as such over its whole lifetime, that initially has much scope for improvement and eventually comes to be widely used, to have multiple uses, and to have many spillover effects." Lipsey Carlaw and Bekar (2005: 98)

[2] Cohen and Leventhal (1990) call this knowledge producing activity the building of "absorptive capacity".

[3] Agents in our model do not know the values of the logistic diffusion parameters and these can be assumed fixed or determined randomly in the model each time a new GPT arrives. As such, they model uncertainty source 5 (v).

[4] Because the pure knowledge sector only periodically produces a useful GPT, we adopt the accounting convention of valuing the output of that sector at its input costs in each period.

[5] These facts are the original ones from Kaldor (1961) and the more recently proposed facts from Easterly and Levine (2001).

[6] This is based on Madison's historical data set (see http://www.ggdc.net/maddison/). The average annual growth rate of GDP per person from 1870 to 2003 for the USA is 1.86% and for Canada is 1.96%.

[7] See "Education Indicators in Canada: Report of the Pan-Canadian Education Indicators program, Statistics Canada (2003) catalogue 81-582-XIE (Ottawa).