Nordic Journal of Political Economy

Volume 25 | 1999 | Pages 15-32

On the Macroeconomic Effects of Major Technological Change

Philippe Aghion Peter Howitt

Philippe Aghion* and Peter Howitt**

On the Macroeconomic Effects of Major Technological Change

Why is it that the adoption of new (more advanced) technological paradigms often entails *cyclical* growth patterns including long recession periods?

Among various attempts to account for Schumpeterian waves,[1] one that appears particularly promising and fruitful is the approach based on the notion of "General Purpose Technologies" (GPTs), that is technologies whose introduction affects the entire economic system. More precisely, whilst each new GPT raises output and productivity in the long-run, it can also cause cyclical fluctuations while the economy adjusts to it. Examples of GPTs include the steam engine, the electric dynamo, the laser and the computer (see David 1990)).

An interesting model of cyclical growth based on GPTs is Helpman and Trajtenberg (1995). The basic idea of this model is that GPTs do not come ready to use off the shelf. Instead, each GPT requires an entirely new set of intermediate goods before it can be implemented. The discovery and development of these intermediate goods is a costly activity, and the economy must wait until some critical mass of intermediate components has been accumulated before it is profitable for firms to switch from the previous GPT. During the period between the discovery of a new GPT and its ultimate implementation, national income will fall as resources are taken out of production and put into R&D activities aimed at the discovery of new intermediate input components.

There are two aspects of this theory which

* EBRD, London and Nuffield College, Oxford

** Ohio State University and CIAR

1. Precusory contributions include Jovanovic and Rob (1990) and Cheng and Dinopoulos (1992) which tried to generate Schumpeterian waves based on the dichotomy between fundamental and secondary innovations, with each fundamental innovation being followed by a sequence of more and more incremental innovations. Of particular interest as a macroeconomic model is the Cheng-Dinopoulous (1992) paper in which Schumpeterian waves obtain as a unique [non steady-state] equilibrium solution, along which the current flow of monopoly profits follows a cyclical evolution. "Because the economy's wealth is equal to the discounted present value of aggregate monopoly profit, fluctuations in profits generate procyclical fluctuations in wealth, the interest factor, consumption (...) and aggregate R&D investments". (Cheng-Dinopoulos.)

may call its empirical relevance into question. The first is the likely *size of the slump* that it might cause. All of the decline in output is attributable to the transfer of labor out of manufacturing and into R&D. But since the total amount of R&D labor on average is only about two and a half percent of the labor force, it is hard to see how this can account for change in aggregate production of more than a fraction of a percent. (The size of the slump would be even smaller if we assumed, as Helpman and Trajtenberg do, that some national income is imputed to the R&D workers even before their research pays off in a positive stream of profits in the intermediate sector).[2]

The second questionable aspect of this theory has to do with the *timing* of slowdowns: the Helpman and Trajtenberg model implies an *immediate* slump as soon as the GPT arrives. This in turns follows from the assumption that: (i) agents need to see the new GPT before investing in research in order to discover the complementary components, and: (ii) these research activities are sufficiently profitable that they always divert some labor resources away from manufacturing. In fact, as Paul David (1990) argues it may take several decades before major technological innovations can have a significant impact on macroeconomic activity (Paul David talks about a pre-paradigm phase of 25 years in the case of the electric dynamo). Then it is hard to believe that labor could be diverted on a large scale into an activity which will pay off only in the very distant future. The fact that so much secondary knowledge has to be accumulated before anyone will know what to do with the new GPT means instead that most firms will choose to ignore it, so treat it as just an academic theoretical discovery with no foreseeable practical significance, and that no significant speed-up R&D activity, and hence no significant slump, will take place for a long time.

The first of these problems is relatively easy to deal with (at least conceptually), as one can think of a number of reasons why the adjustment to a massive and fundamental technological change would cause adjustment and coordination problems resulting in a slump. For example, as Atkeson and Kehoe (1993) analyse, the arrival of a new GPT might induce firms to engage in *risky experimentation* on a large scale with startup firms, not all of which will succeed. The capital sunk into these startup firms will not yield a competitive return right away except by chance; meanwhile national income will drop as a result of that capital not being used in less risky ways using the old GPT. Also, an increase in the pace of innovation aimed at exploiting the new GPT may well result in an increased rate of job turnover, and hence in an increased rate of unemployment. Greenwood and Yorukoglu (1996) present an analysis in which the costs of learning to use equipment embodying the new GPT can account for a prolonged productivity showdown. Howitt (1996) shows how the arrival of a new GPT can cause output growth to slow down because it accelerates the rate of obsolescence of existing physical and human capital.

The second problem is more challenging

2. Helpman and Trajtenberg find that a measured slump occurs when the GPT arrives even if the full cost of R&D is imputed as national income. The reason is that the discovery induces workers to leave a sector where their marginal product is higher than the wage (because the intermediate sector is imperfectly competitive and pays according to the marginal revenue product of labour rather than the marginal value product), and to enter a sector – research – where their (imputed) marginal product is just equal to the same wage.

to deal with. The question is, if the exploitation of a new GPT is spread out over a period of many decades why should it not result in simply a slow enhancement in aggregate productivity, as one industry after another learns to use the new technology?

Again, several answers come to mind and we actually think of the following three explanations as being complementary. First are the *measurability* problems: as already stressed by David and others, it may take a while before the new products and services embodying the new GPT can be fully accounted for by the conventional statistics. (This, however, does not explain the possibility of delayed *slumps*). Second, the existence of *strategic complementarities* in the adoption of new GPTs by the various sectors of the economy may generate temporary lock-in effects, of a kind similar to the implementation cycles in Shleifer (1986). It may then take real labor costs or other "exogenous" economic parameters to reach a minimum threshold before a critical number of sectors decide to jump on the bandwagon of the new GPT. A third explanation, which will be the main focus of our analysis below, lies in the phenomenon of *social learning*. That is, the way most firms learn to use a new technology is not to discover everything on their own but to learn from the experience of other firms in a similar situation: that is, for a firm to learn from other firms for whom the problems that must be solved before the technology can successfully be implemented bear enough resemblance to the problems that must be solved in this firm, that it is worth-while trying to use the procedures of those successful firms as a "template" on which to prepare for adoption in this firm. Thus at first the fact that no one knows how to exploit a new GPT means that almost nothing happens in the aggregate. Only minor improvements in knowledge take place for a long time, because successful implementation in any sector requires firms to make independent discoveries with little guidance from the successful experience of others. But if this activity continues for long enough, a point will eventually be reached when almost everyone can see enough other firms using the new technology to make it worth their while experimenting with it. Thus even though the spread of a new GPT takes place over a long period of time, most of the costly experimentation through which the spread takes place may be concentrated over a relatively short subperiod, during which there will be is a cascade or snowball effect resulting sometimes in a (delayed) aggregate slump.

The paper is organized as follows. We present a simplified version of the Helpman and Trajtenberg model of GPT which fits nicely into the basic Schumpeterian framework developed in Aghion and Howitt (1992), and which also permits us to endogeneize the long-run growth rate which Helpman and Trajtenberg take as given. We then extend this to introduce social learning considerations with a view to addressing the objections concerning the timing of economic slow-downs. Finally we illustrate how the objection concerning the size of slow-down might be addressed by introducing three alternative features into the basic social learning model namely skill-difference, job search and obsolescence.

A simplified presentation of the Helpman-Trajtenberg model of GPTs

A brief reminder of the basic Schumpeterian growth model

Let us first recall the main features of the basic Schumpeterian growth model as

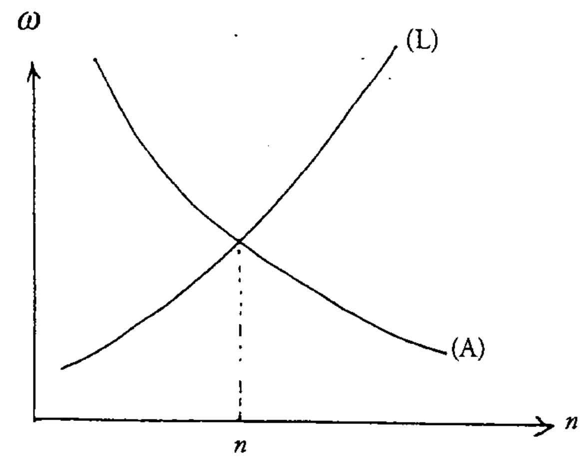

Figure 1

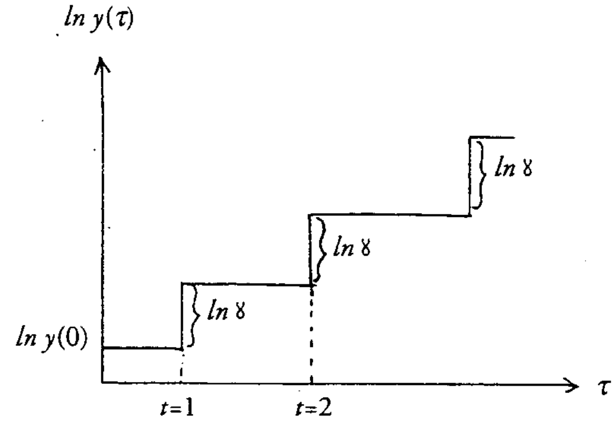

Figure 2

developed in Aghion-Howitt (1992). Final output is produced according to the flow production function:

$$y = A \cdot x^{\alpha}$$

where x is the flow of intermediate input and A is a productivity parameter measuring the quality of intermediate input x. (Intermediate input itself is produced with labor according to a one-to-one linear technology, so that x corresponds also to the flow of manufacturing labor). In this economy where population is constant (equal to L, which is also the total flow of labor supply under the assumption that each individual is endowed with one unit flow of labor per unit of time), growth will entirely result from vertical innovations, that is from quality improvements in A. That is, each innovation will augment current productivity by the multiplicative factor $\gamma > 1 : A_{t+1} = \gamma A_t$. Innovations in turn are the [random] outcome of research activities, and are assumed to arrive discretely with Poisson rate λn, where n is the current flow of research activities.

In steady-state the allocation of labor between research and manufacturing remains constant over time, and is determined by the arbitrage equation:

$$\omega = \lambda \gamma v \qquad \text{(A)}$$

where the LHS of (A) is the productivity-adjusted wage rate $\omega = \frac{\omega}{A}$ which a worker earns by working in the manufacturing sector; and $\lambda \cdot \gamma \cdot v$ is the expected reward from investing one unit flow of labor in research. The productivity-adjusted value v of an innovation is determined by the Bellman equation:

$$rv = \pi(\omega) - \lambda n v$$

where $\pi(\omega)$ denotes the productivity-adjusted flow of monopoly profits accruing to a successful innovator and where the term $(-\lambda n v)$ corresponds to the capital loss involved in being replaced by new subsequent innovators.

The above arbitrage equation, which can be reexpressed as:

$$\omega = \lambda \gamma \frac{\pi(\omega)}{r + \lambda n} \qquad \text{(A)}$$

together with the labor market clearing equation:

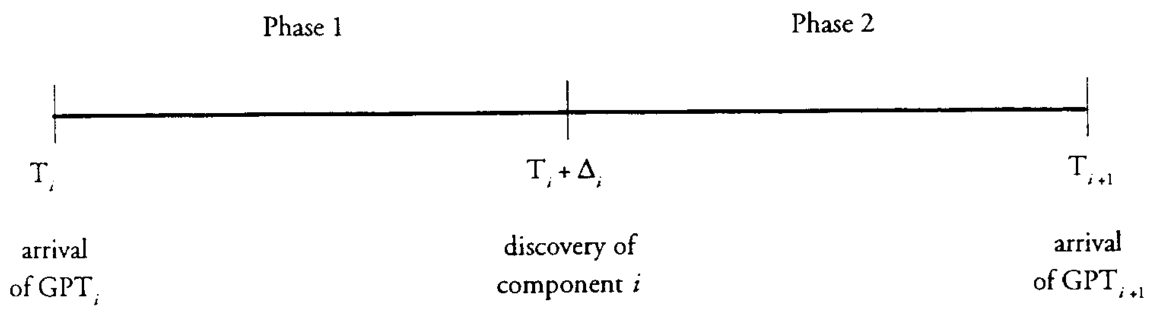

Figure 3

$$x(\omega)+n=L \qquad \text{(L)}$$

where $x(\omega)$ is the manufacturing demand for labor,[3] jointly determine the steady-state amount of research n as a function of the parameters $\lambda, \gamma, L, r, \alpha$. Figure 1 above depicts the two curves (A) and (L), and shows the straightforward comparative statics results.

In steady-state the flow of consumption good (or final output) produced between the t^{th} and the $(t+1)^{th}$ innovation is:

$$y_t = A_t(L-n)^{\alpha}$$

which implies that in real time (whith we denote by τ), the log of final output will increase by $\ln \gamma$ each time a new innovation occurs. Thus in this *one-sector* economy where each innovation corresponds by definition to a *major* technological change (i.e. to the arrival of a new GPT), growth will be *uneven* (see Figure 2) with the time path of $\log y(\tau)$ being a random step function.

The average growth rate will be equal to the size of each step, that is $\ln \gamma$, times the average number of innovations per unit of time, that is λn : i.e, $g=\lambda n \ln \gamma$.

Although it is uneven, the time path of aggregate output as depicted above does not involve any slump. Accounting for the existence of slumps requires an adequate extension of the basic Schumpeterian model, for example of the kind developed by Helpman and Trajtenberg to which we now turn.

The Helpman-Trajtenberg model revisited

As before, there are L workers who can engage either in production of existing intermediate goods or in research aimed at discovering new intermediate goods. But each intermediate good is now linked to a particular GPT. We follow Helpman and Trajtenber in supposing that before any of the intermediate goods associated with GPT can be used profitably in the final goods sector, some minimal number of them must be available. But we lose nothing essential by supposing that this minimal number is one. Once the good has been invented, its discoverer profits from a patent on its exclusive use in production, exactly as in the basic model above.

3. $$x(\omega) = \arg\max_x (\underbrace{\alpha \cdot x^{\alpha-1}}_{p(x)} \cdot x - \omega \cdot x)$$

and

$$x(\omega) = \max (\alpha \cdot x^{\alpha-1} \cdot x - \omega \cdot x).$$

Thus the difference between this model and our basic model is that now the discovery of a new generation of intermediate goods comes in *two* stages. First a new GPT must come, and then the intermediate good must be invented that implements that GPT. Neither can come before the other. You need to see the GPT before knowing what sort of good will implement it, and people need to see the previous GPT in action before anyone can think of a new one. For simplicity we assume that no one directs R&D towards the discovery of a GPT. Instead, the discovery arrives as a serendipitous byproduct of the collective experience of using the previous one.

Thus the economy will pass through a sequence of cycles, each having two phases, as indicated in Figure 3 below. GPT_i arrives at time T_i. At that time the economy enters phase 1 of the i^{th} cycle. During phase 1, the amount n of labor is devoted to research. Phase 2 begins at time $T_i+\Delta_i$ when this research discovers an intermediate good to implement GPT_i. During phase 2 all labor is allocated to manufacturing, until GPT_{i+1} arrives, at which time the next cycle begins. Over the cycle output is equal to $A_{i-1}F(L-n)$ during phase 1 and $A_iF(L)$ during phase 2. Thus the drawing of labor out of manufacturing and into research causes output to fall each time a GPT is discovered, by an amount equal to $A_{i-1}[F(L)-F(L-n)]$.

A steady state equilibrium is one in which people choose to do the same amount of research each time the economy is in phase 1, that is where n is constant from one GPT to the next. As before, we can solve for the equilibrium value of n using a research arbitrage equation and a labor market equilibrium curve. Let ω_j be the wage, and v_j the discounted expected net value of profits in the intermediate goods sector in phase j, each divided by the productivity parameter A of the GPT currently in use. In a steady state these productivity-adjusted variables will all be independent of which GPT is currently in use.

Since research is conducted in phase 1 but pays off when the economy enters into phase 2 with a productivity parameter raised by the factor γ, the usual arbitrage condition must hold in order for there to be a positive level of research in the economy:

$$\omega_1=\lambda\gamma v_2 \tag{1}$$

Suppose that once we are in phase 2, the new GPT is delivered by a Poisson process with a constant arrival rate equal to μ. Then the value of v_2 is determined by the Bellman equation:

$$rv_2=\pi(\omega_2)+\mu(v_1-v_2) \tag{2}$$

By analogous reasoning, we have:

$$rv_1=\pi(\omega_1)-\lambda n v_1 \tag{3}$$

Combining (1) – (3) yields the research arbitrage equation:

$$\omega_1=\lambda\gamma\,[\,\pi(\omega_2)+\frac{\mu\pi(\omega_1)}{r+\lambda\gamma}]\,/[r+\mu] \tag{4}$$

Since no one does research in phase 2, we know that the value of ω_2 is determined independently of research, by the market clearing condition: $L=x(\omega_2)$. Thus we can take this value as given and regard equation (4) as determining ω_1 as a function of n. The value of n is determined, as usual, by this equation together with the labor-market equation:

$$L-n=x(\omega_1) \tag{5}$$

As in the basic model, the level of research n is an increasing function of the productivity of research λ, the size of improvement created by each GPT γ, and the population L; and a decreasing function of the rate of interest r. The arrival rate μ of GPTs, can be shown to have a negative affect on research;[4] intuitively, an increase in μ discourages research by reducing the expected duration of the first of the two phases over which the successful researcher can capitalise the rents from an innovation. The size of the slump ln $(F(L))$ –ln $(F(L-n))$ is an increasing function of n, and hence will tend to be positively correlated with the average growth rate.

The average growth rate will be the frequency of innovations times the size $\ln \gamma$, for exactly the same reason as in the basic model. The frequency, however, is determined a little differently than before because the economy must pass through *two* phases. An innovation is implemented each time a full cycle is completed. The frequency with which this happens is the inverse of the expected length of a complete cycle. This in turn is just the expected length of phase 1 plus the expected length of phase 2:

$$1/\lambda n+1/\mu=\frac{\mu+\lambda n}{\mu\cdot\lambda n}$$

Thus we have the growth equation:

$$Eg=\ln\gamma\cdot\frac{\mu\cdot\lambda n}{\mu+\lambda n} \tag{6}$$

Thus the expected growth rate will be positively affected by anything that raises research, with the possible exception of a fall in μ. In the limit, when μ falls to zero, growth must also fall to zero as the economy will spend an infinitely long time in phase 2, without growing. Thus for small enough values of μ, Eg and n will be affected in opposite directions by a change in μ.

One further property of this cycle worth mentioning is that, as Helpman and Trajtenberg point out, the wage rate will rise when the economy goes into a slump. That is, since there is no research in phase 2 the normalized wage must be low enough to provide employment for all L workers in the manufacturing sector, whereas with the arrival of the new GPT, the wage must rise to induce manufacturers to release workers into research.

As already discussed in the introduction, the Helpman and Trajtenberg model may not quite fit the empirical and/or anecdotal evidence on the macroeconomic impact of major technological changes, to the extent that it predicts immediate slumps of very small magnitude. Whilst the introduction of social learning considerations in the next section will contribute to explaining the observed *delays* in the macroeconomic response to new GPTs, other considerations such as skill differentials, job search and obsolescence introduced in the following

4. To show this, it suffices to show that an increase in μ shifts the research arbitrage curve to the left. By applying the implicit function theorem to (4) we see that the sign of this shift is:

$$sgn\frac{dn}{d\mu}=-sgn\,(\omega_1-\frac{\lambda\gamma\pi(\omega_1)}{r+\lambda n}).$$

Since no research is done in phase 2, labour market equilibrium requires $\omega_2<\omega_1$, and hence $\pi(\omega_2)>\pi(\omega_1)$.
Applying this to (4) yields:

$$\omega_1>\frac{\lambda\gamma\pi(\omega_1)}{r+\lambda n}\cdot\frac{r+\lambda n+\mu}{r+\mu}>\frac{\lambda\gamma\pi(\omega_1)}{r+\lambda n}.$$

section can help account for the macroeconomic *significance* of GPT – driven fluctuations.

A model of major technological change through social learning

Basic set-up

We consider the following dynamic model of the spread of technology, which is similar to the sorts of models used by epidemiologists when studying the spread of disease, which also takes place through a process of social interaction between those who have and those who have not yet been exposed to the new phenomenon. The setting of model is like the model we have just described, with a continuum of sectors, uniformly distributed on the unit interval, except that now each sector must invent its own intermediate good in order to exploit the GPT. We study here the nature of the cycle caused by the arrival of a single GPT, under the assumption that the arrival rate μ is so small that there is insignificant probability that the next GPT will arrive before almost all sectors have adopted the one that has just arrived. In order to simplify the analysis even further we suppose that the amount of research in each sector is given by a fixed endowment of specialized research labor. Thus all the dynamics will result from the effects of social learning on the payoff rate to experimentation. This is the phenomenon that we believe to be at the heart of the timing of the delayed cyclical response to GPTs. Endogenizing the allocation of labor between research and manufacturing would just accentuate the effects we find, as it would draw more labor into research, hence augmenting the intensity of experimentation, just when the informational cascade we focus on is already having the same effect.

Aggregate output at any point in time is produced by labor according to the constant returns technology:

$$Y=\{\int_0^1 A(i)^\alpha\, x(i)^\alpha di\,\}^{1/\alpha} \qquad (7)$$

where $A(i)=1$ in sectors where the old GPT is still used, and $A(i)=\gamma>1$ in sectors that have successfully innovated, while $x(i)$ is manufacturing labor used to produce the intermediate good in sector i.

We assume now that in each sector an innovation requires *three* breakthroughs rather than the two breakthroughs of the previous model. First, the economy wide GPT must be discovered. Second, a firm in that sector must acquire a "template", on which to base experimentation. Third, the firm must use this template to discover how to implement the GPT in its particular sector. (This third stage is equivalent to the component finding stage in the Helpman and Trajtenberg model, whilst the second stage is new). Thus all sectors are in one of three states. In state 0 are those sectors who have not yet acquired a template. In state 1 are those who have a template but have not yet discovered how to implement it. In state 2 are those sectors who have succeeded in making the transition to the new GPT. We let the fraction of sectors in each state be represented by n_0, n_1, n_2, and suppose that initially $n_0=1, n_1=n_2=0$.

A sector will move from state 0 to state 1 if a firm in that sector either makes an independent discovery of a template or if it discovers one by "imitation" that is by observing at least k "similarly located" firms that have made a successful transition to the new GPT (firms in state 2). The Poisson arrival rate of independent discoveries to such a sector is $\lambda_0<<1$. The Poisson arrival rate of opportunities to observe m similarly located firms is assumed to equal unity. The probability that such an observation will pay off (in other words the probability that at least k among the m similar firms will have successfully experimented the new GPT) is given by the cumulative Binomial:

$$\varphi(m,k,n_2)=\sum_{j=k}^{m}\binom{m}{k} n_2^j \cdot (1-n_2)^{m-j}$$

since n_2 is the probability that a randomly selected firm will be in state 2. Thus the flow of sectors from state 0 to state 1 will be n_0 times the flow probability of each sector making the transition: $\lambda_0+\varphi(m,k,n_2)$.

For a sector to move from state 1 to state 2, the firm with the template must employ at least N units of labor per periods (the equivalent of n in the Helpman and Trajtenberg model). We can think of this labor as being used in formal R&D, informal R&D, or in an experimental starting firm. In any case it is not producing current output. Instead it is allowing the sector access to a Poisson process that will deliver a workable implementation of the new GPT with an arrival rate of λ_1. Thus the flow of sectors from states 1 to 2 will be the number of sectors in state 1, n_1, times the success rate per sector per unit of time λ_1.

We can summarize the discussion to this point by observing that the evolution over time of the two variables n_1, and n_2 is given by the autonomous system of ordinary differential equation:

$$\begin{aligned} \dot{n}_1 &= [\lambda_0+\varphi(m,k,n_2)]\,(1-n_1-n_2)-\lambda_1 n_1 \\ \dot{n}_1 &= \lambda_1 \cdot n_1 \end{aligned} \qquad \text{(S)}$$

with initial condition: $n_1(0)=0, n_2(0)=0$. (The time path of n_0 is then given automatically by the identity $n_0 \equiv 1-n_1-n_2$.)

Figure 4 depicts the solution to the above system (S). Not surprisingly, the timepath of n_2 follows a logistic curve, accelerating at first and slowing down as n_2 approaches 1, with the maximal growth rate occurring somewhere in the middle. Likewise the path of n_1 must peak somewhere in the middle of the transition, since it starts and ends at zero. If the arrival rate λ_0 of independent discoveries is very small then both n_1 and n_2 will remain near zero for a long time. Figure 4 shows the behaviour of n_1 and n_2 in the case where $\lambda_0=.005$, $\lambda_1=.3$, $m=10$ and $k=3$. The number of sectors engaging in experimentation peaks sharply in year 20 due to social learning.

The solution to the system (S) can be used with the aggregate production function (7) and the market clearing condition for labor to determine the time path of aggregate output. Using the symmetry of the production technology (7), which implies that all the sectors using the same GPT (either old or new) will demand the same amount of manufacturing labor, we can reexpress the flow of aggregate output as:

$$Y=\left\{\int_0^{1-n_2} x_0(i)^\alpha\, di+\gamma^\alpha \cdot \int_{1-n^2}^{1} x_N(i)^\alpha di\right\}^{1/\alpha} \qquad (8)$$

where x_0 (resp. x_N) denotes the flow of labor demand by a sector using the old (resp the new) GPT.

In this Cobb-Douglas world the local monopolists in sectors in state 0 and 1 who use the old technology will demand labor according to the demand function[5]

5. This follows from profit-maximization: for any sector i, $x(i)= \arg_x \max \{p_i(x) \cdot x-wx\}$, where:

$$\begin{aligned} p_i(x)= \frac{\partial Y}{\partial x} &= \alpha x^{\alpha-1} \cdot Y^{1-\alpha} && \text{if sector } i \text{ uses } \textit{old} \text{ technology} \\ &= \alpha \cdot \gamma^\alpha x^{\alpha-1} \cdot Y^{1-\alpha} && \text{if sector is uses the } \textit{new} \text{ technology.} \end{aligned}$$

The corresponding first-order conditions, respectively for old and new sectors, yield the above equations (9) and (10)..

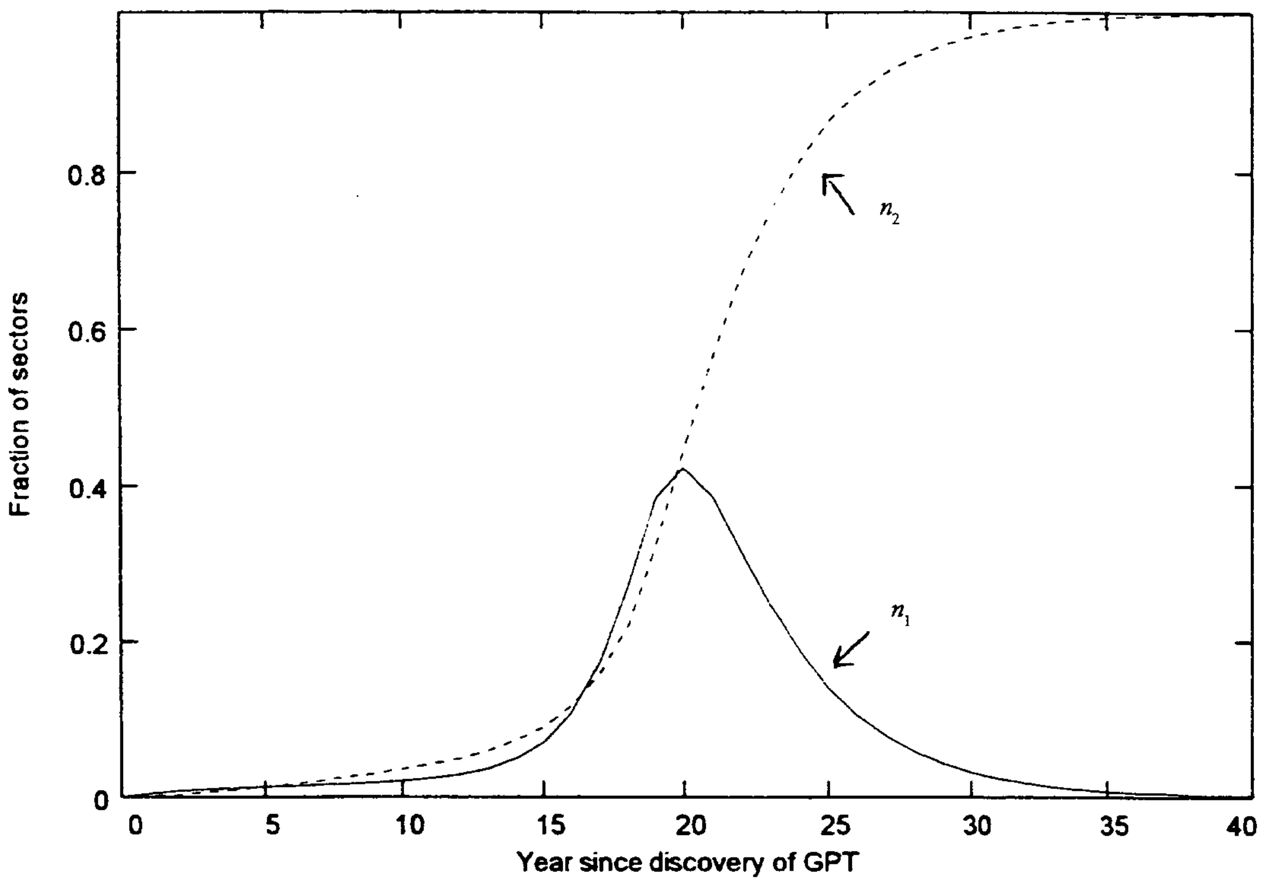

Figure 4

$$x_0=(w/\alpha)^{\frac{1}{\alpha-1}} \cdot Y \qquad (9)$$

while those in sectors in state 2 will demand labor according to:

$$x_N=(w/\alpha \cdot \lambda^{\alpha})^{\frac{1}{\alpha-1}} \cdot Y \qquad (10)$$

where w is the real wage rate.

Now using the market clearing condition:

$$\underbrace{(1-n_2)x_0+n_2 \cdot x_N}_{\text{manufacturing labor demand}} + \underbrace{n_1 N}_{\text{experimenting labor}} = L \qquad (L)$$

one can solve for the real wage w as a function of Y, n_1 and n_2. Substituting this solution into the above expressions for x_0 and x_N and then substituting the resulting values of x_0 and x_N into (8) yields the following reduced form expression for output:

$$Y=(L-n_1 N) \cdot (1-n_2+n_2\gamma^{\frac{\alpha}{\alpha-1}})^{\frac{\alpha-1}{\alpha}} \qquad (11)$$

Figure 5 shows the time path of output, which results from the above dynamics in n_1 and n_2, in the benchmark case where N=6, L=10 and α =.5. As expected, output is not much affected by the new GPT for the first decade and half, but then it enters a severe recession precisely when the number of sectors engaging in experimentation increases sharply as a result of social learning: output reaches a trough in year 19, after a 10.5% drop in output. From there output begins to grow, ultimately attaining a value of γ (=1.5) times its original value.

The delay in the slump caused by the GPT could not have occurred without the impact of social learning.[6] That is, suppose

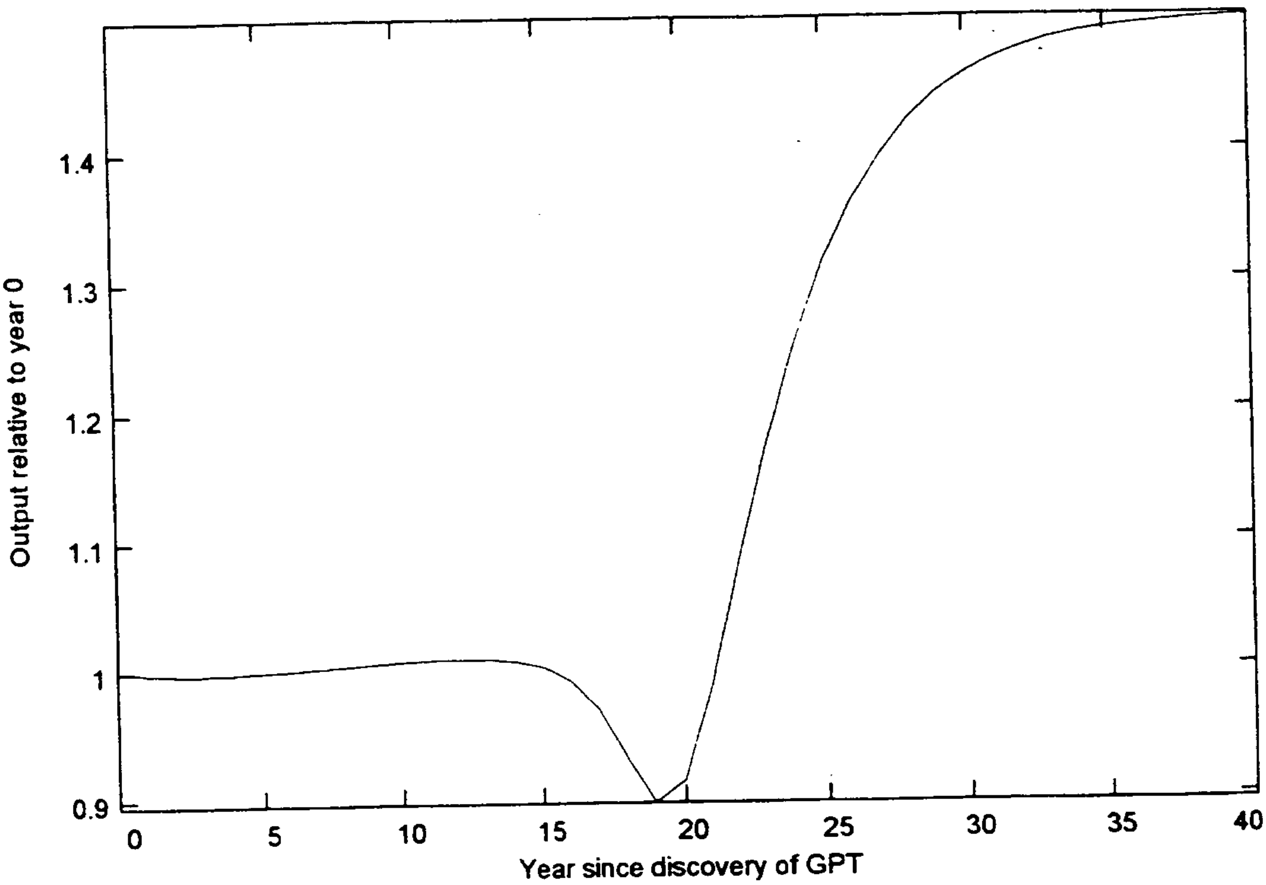

Figure 5

that the function $\varphi(m,k,n_2)$ that embodies the effects of social learning were replaced by a constant value, $\varphi_0=\int_0^1\varphi(m,k,n_2)dn_2$, whose average value was the same but which was not affected by the process of observing other sectors that have succeeded in implementing the new GPT. Then n_2 would still follow a mild logistic curve but the intensity of experimentation n_1 would rise immediately following the arrival of GPT and would fall monotically from then on. Output could go through a slump but the maximal rate of decline would occur immediately at year 0. This benchmark case of no social learning is illustrated in Figure 6.

Intuitively, the reason why the slump cannot be delayed in this case is as follows. In order for output to be falling there must be a positive flow of sectors into state 1, which is drawing workers out of manufacturing. But without social learning this flow must be diminishing whenever a slump is underway, since the rise in the level of n_1 (and in n_2) will be reducing the rate of growth of n_1. (See the above equation $\dot{n}_1$ in (S)). That is, the rise in n_1+n_2 reduces the number of sectors n_0 from which new experimentation can arrive, while the rise in n_1 increases the flow of successful innovators out of the state of experimentation. Thus either the slump starts right

6. Greenwodd and Yorukoglu (1996) assume a private learning process that also produces diffusion according to a mild logistic curve.

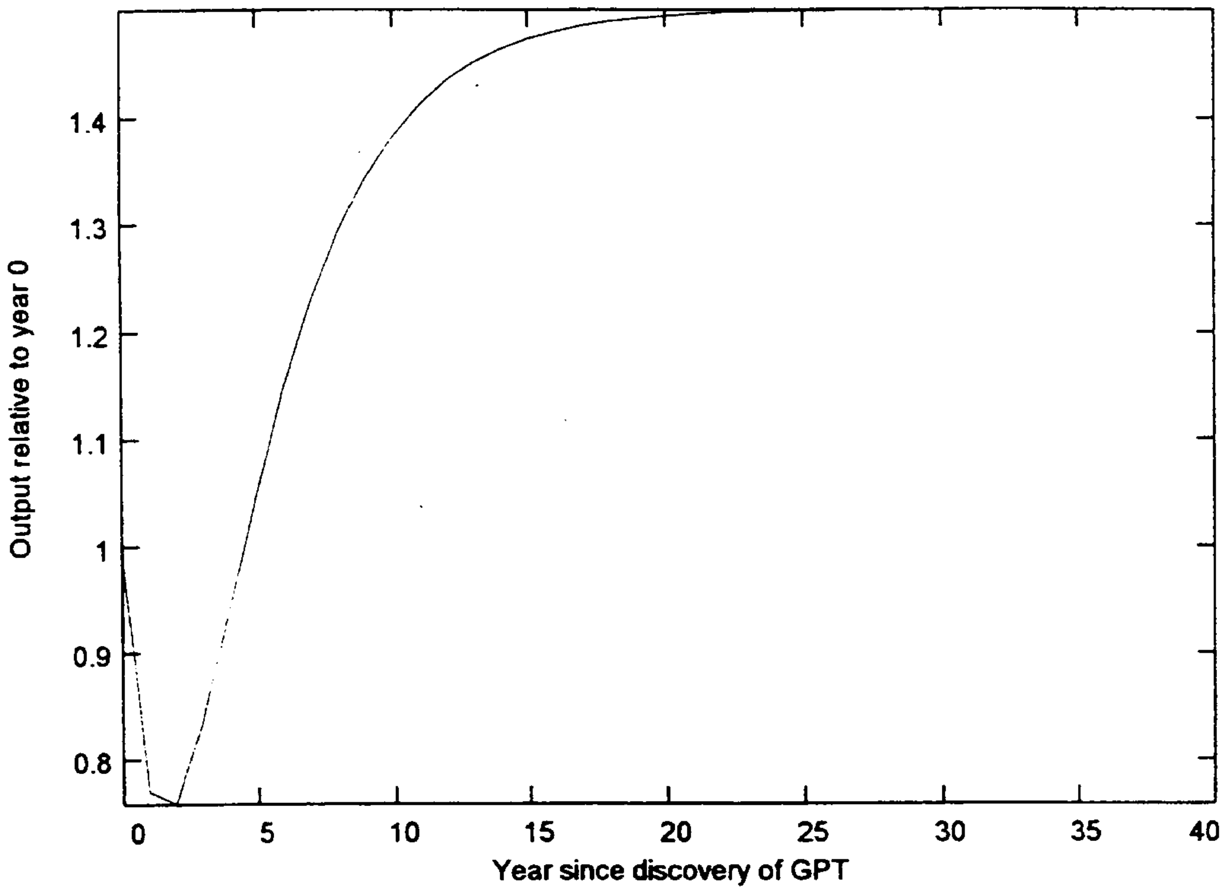

Figure 6

away, in which case its intensity will diminish steadily, or it never starts at all.[7] What social learning does is to reverse the effect of n_2 on the rate of growth of n_1; that is, as n_2 rises, the resulting increase in the likelihood of imitation counterbalances the fall in the number of possible imitators, thus causing the cascade at the heart of our analysis.

7. To see this more formally, suppose that the output function (11) can be approximated by its first-order Taylor expansion around $n_1 = n_2 = 0$:

$$Y \approx L - N \cdot n_1 + \xi \cdot n_2; \quad \xi \equiv \frac{1-\alpha}{\alpha} (\gamma^{1-\alpha} - 1)\, L > 0.$$

Because $\dot{n}_2$ is always positive, $\dot{n}_1$ must also be positive whenever Y is not rising. Thus:
If $Y \leq 0$, then:

$$\dot{Y} \approx -N \cdot \dot{n}_1 + \xi\, \dot{n}_2$$

$$\dot{Y} \approx -N [(\lambda_0 + \varphi_0)(1 - n_1 - n_2) - \lambda_1 n_1] + \xi \lambda_1 n_1$$

$$\ddot{Y} \approx [N(\lambda_0 + \varphi_0 + \lambda_1) + \xi \lambda_1]\, \dot{n}_1 + N(\lambda_0 + \varphi_0)\, \dot{n}_2$$

$$\ddot{Y} > 0 \text{ (because } \dot{n}_1 > 0,\ \dot{n}_2 > 0).$$

Hence a delayed slump, with $\dot{Y}$ turning negative or becoming more negative at some date $t > 0$, is impossible.

Table 1

Parameter	Value	Slump	Size	Peak date	Trough date
Benchmark		yes	11 %	12	19
α	0.2	yes	12 %	12	19
(0.5)	0.8	yes	8 %	15	19
γ	1.1	yes	22 %	0	20
(1.5)	3.0	no			
k	1	yes	23 %	0	5
(3)	5	yes	4 %	37	42
m	3	no			
(10)	30	yes	22 %	4	10
N	2	no			
(6)	8	yes	20 %	11	19
λ_0	.001	yes	11 %	32	40
(.005)	.025	yes	10 %	5	10
λ_2	0.1	yes	22 %	13	29
(0.3)	1.0	no			

Some comparative dynamics
Table 1 below shows how the time path of aggregate output responds to variations in the basic parameters of the model, namely:

- α, which measures the degree of substitutability across intermediate inputs.
- γ, which measures the size of productivity improvements brought about by the new GPT.
- N, the number of workers taken out of manufacturing by each experimenting firm.
- m, the number of sectors potentially "similar" to a given sector.
- k, the required number of observations of successful experimentations in order to acquire a template "by imitation".
- λ_0, the arrival rate of independent ideas for new templates.
- λ_1, the arrival rate of success for experimenting firms.

In all cases the simulation produces either a marked slump, as in Figure 5 above, or a monotonic increase in output. When there is no slump there is an initial period of relatively slow growth followed by a sharp acceleration coming just after the peak in experimentation. When there is a slump it almost always comes after a period of mild growth, which itself is often preceded by a very mild (less than half a percentage point) recession. The size of slump reported in Table 1 is the percentage shortfall from the peak attained at the end of the period of mild growth (or from year 0 if no such period exits) to the trough. From Table 1 we can see:

(a) the magnitude of slumps increases as α decreases, that is when intermediate inputs become less substitutable. This is fairly intuitive: as α decreases, the downsizing of old manufacturing sectors which results from labor being diverted away into experimentation, is less and less substituted for by the new – more productive – intermediate good sectors.

(b) the magnitude of slumps decreases as γ increases, and for sufficiently large γ the slump even disappears. Again, this result is intuitive: the bigger productivity of new sectors compensates for the reduction in output in old sectors caused by experimentation (and by the resulting wage increase), thereby reducing the scope for aggregate slumps.

(c) If m is too small, output grows steadily: indeed the lower m, the lower the scope for social learning and for the resulting snow-ball effects on aggregate output.

(d) An increase in k leads to bigger delays but smaller slumps: as k increases it will take longer for "imitation" and social learning to become operational and by the time it becomes so, a higher number of sectors will have already moved into using the new and more productive GPT, hence the smaller the size of aggregate slumps.

(e) An increase in N leads to larger slumps. This is straightforward: the bigger N, the more labor will be diverted away from manufacturing into experimentation by firms in state 1, and therefore the bigger the size of slumps when social learning causes the fraction of experimenting sectors to sharply increase.

(f) An increase in the arrival rate of independent ideas λ_0 speeds up the macroeconomic response to the new GPT. This is not surprising, for the larger λ_0 the faster the conditions will be created for social learning to operate.

(g) An increase in the success rate of experimentation λ_1 reduces the size of slumps: this is again easy to understand, for the larger λ_1 the faster the emergence of sectors using the new GPT, which compensate for the downsizing of manufacturing activities induced by experimentation activities.

Accounting for the size of slowdowns

Skill differentials

The last five years or so have witnessed an upsurge of empirical papers on skill differentials and wage inequality, and their relationship with technological change (in particular, see Juhn et al. (1993). It turns out that a straightforward extensions of our GPT model can immediately account for the observed positive correlation between the acceleration of technological progress resulting from the introduction of new GPT, and the increasing skill differential.[8] The

8. Our explanation of both the differential and the slowdown is similar in spirit to that of Greenwood and Yorukoglu (1996) who also emphasize the role of skilled labor in implementing new technologies.

same extension can also magnify the slump.

More formally, suppose that the labor force L is now divided into skilled and unskilled workers, and that the implementation of the new GPT requires *skilled* labor whereas old sectors can indifferently use skilled or unskilled workers to manufacture their intermediate inputs. Also, let us assume that the fraction of skilled workers is increasing over time, e.g. as a result of schooling and/or training investments which we do not model here:

$$L_s(t)=L\,(1-(1-\tau)e^{-\lambda_2 t}),\quad \tau<1,$$

where τ is the initial fraction of skilled workers and λ_2 is a positive number measuring the speed of skill acquisition.

The transition process from the old to the new GPT can then be divided into two subperiods. First, in the early phase of transition (i.e. when t is low) the number of sectors using the new GPT is too small to absorb the whole skilled labor force, which in turn implies that a positive fraction of skilled workers will have to be employed by the old sectors at the same wage as their unskilled peers. Thus, during the early phase of transition the labor market will remain "unsegmented', with aggregated output and real wage being determined exactly as before.[9]

Second, in the later state of transition, where the fraction of new sectors has grown sufficiently large that it can absorb the whole skilled labor force, the labor market will become segmented, with skilled workers being exclusively employed (at a higher wage) by new sectors whilst the unskilled workers remain in old sectors. Let w_u and w_s denote the real wages respectively paid to unskilled and skilled workers. The demand for manufacturing labor by the old and new sectors are still given by:

$$x_0=\left(\frac{W_u}{\alpha}\right)^{\frac{1}{\alpha-1}}\cdot Y$$

and

$$x_N=\left(\frac{W_s}{\alpha\gamma^{\alpha}}\right)^{\frac{1}{\alpha-1}}\cdot Y$$

except that we now have: $w_s > w_u$, where the two real wages are determined by two separate labor market clearing conditions, respectively:

$$L_2=n_1N+n_2\cdot x_N \rightarrow w_s \qquad (12)$$

and

$$L_1=L-L_2=(1-n_2)\cdot x_0 \rightarrow w_u \qquad (13)$$

These equations yield:

$$w_s=\gamma^{\alpha}\alpha\left(\frac{n_2 Y}{L_2-n_2N}\right)^{\alpha-1}$$

and

$$w_u=\alpha\left(\frac{(1-n_2)Y}{L-L_2}\right)^{\alpha-1}$$

which, after substitution for w_s and w_u in the above expressions for x_0 and x_N and after

9. i.e. by equations (L) and (11), which yield:

$$Y=(L-n_1N)(1-n_2+n_2\gamma^{\frac{1-\alpha}{\alpha}})^{\frac{1-\alpha}{\alpha}}$$

and

$$w=\alpha[1-n_2+n_2\gamma^{\frac{1-\alpha}{\alpha}}]^{\frac{1-\alpha}{\alpha}}.$$

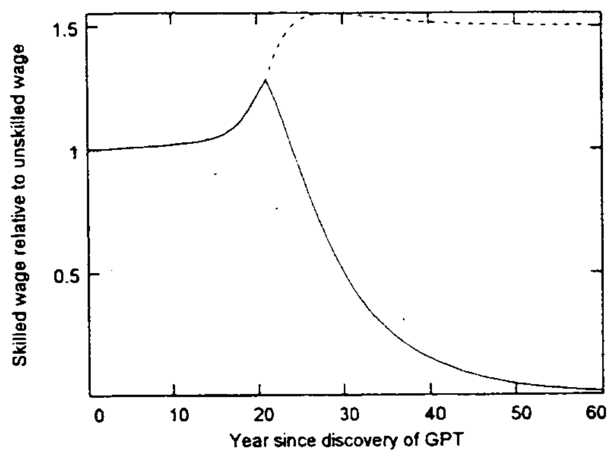

Figure 7a

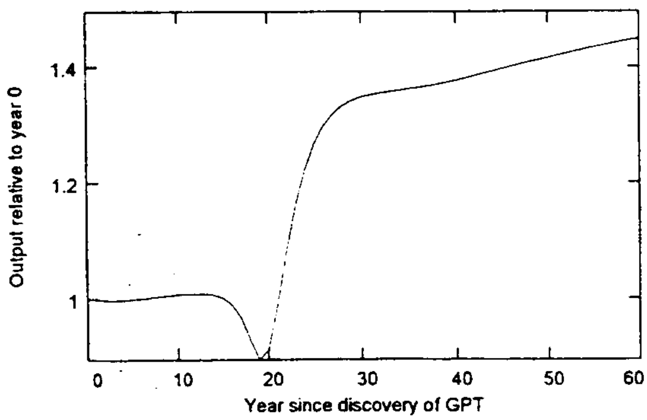

Figure 7b

substituting the resulting values of x_0 and x_N in equation (11), yields the following expression for aggregate output during the segmented phase of transition:

$$Y=[(1-n_2)^{1-\alpha}(L-L_2)^{\alpha}+n_2^{1-\alpha}\gamma^{\alpha}(L_2-n_1N)^{\alpha}]^{\frac{1}{\alpha}}$$

(The cut-off date between the unsegmented and segmented phases of transition to the new GPT is simply determined by:

$$w_s(t_0)=w_u(t_0).)$$

Figure 7a depicts the time-path of real wages and Figure 7b the time-path of aggregate output in the benchmark case of the previous section with λ_2=.05 and τ =.0.25. Two interesting conclusions emerge from this simulation.

(a) The skill premium (w_s/w_u) starts increasing sharply in the year n=21 when social learning is accelerating the flow of new sectors in the economy, and then the premium keeps on increasing although more slowly during the remaining part of the transition process.[10] Since everyone ends up earning the same (skilled) wage, standard measures of wage inequality first rise and then fall.

(b) Compared to the benchmark case *without* skill differentials and labor market segmentation, the *magnitude of the slump* is the same (11%) but the recovery is slower: the reason for this is simply that high productivity sectors are "constrained" by the short supply of skilled labor; in simulations with other parameter values we see that the slump is exacerbated by the skill shortage if the market becomes segmented near the peak of experimentation.

Job search

Let us now extend the basic set-up in another direction namely by introducing other costly job search which, together with the

10. The acceleration in the premium, with w_s increasing and w_u decreasing sharply at the beginning of the segmented phase, has to do with the high demand for skilled experimentation labor during this time-interval where social learning peaks. The skilled real wage w_s starts tapering off thereafter where most sectors are already in phase 2 and the supply of skilled labor keeps on increasing over time.

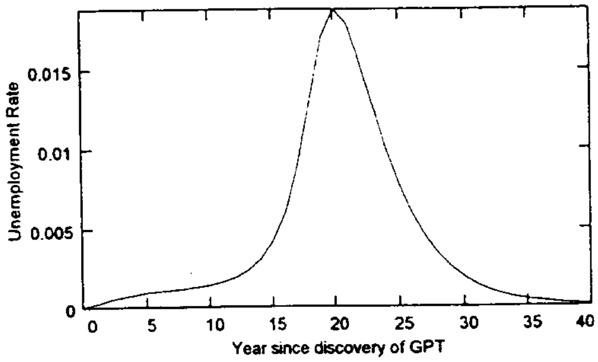

Figure 8a

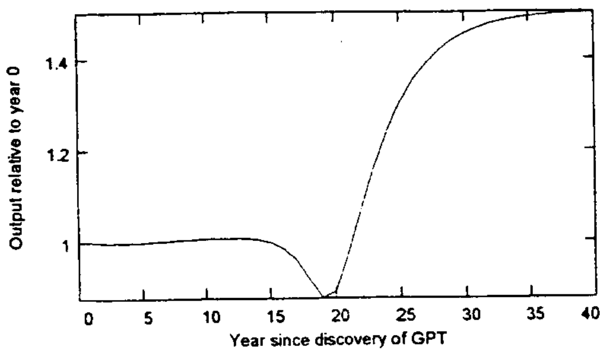

Figure 8b

destruction of jobs by new sectors generates unemployment on the transition path. Unemployment in turn diverts a higher fraction of the labor force out of manufacturing activities, thereby *increasing* the size of slumps relative to the benchmark case simulated above. Indeed slumps can now occur even if the labor N needed to perform experiments is negligible.

More formally, suppose that, the fraction β of workers in each sector that adapts the GPT (and moves in to n_2) go into temporary unemployment, because they are unable to adapt to the new GPT in the sector where they were formally unemployed. Suppose also that the fraction λ_3 of the unemployed per period succeed in finding a new job. Then the evolution of U, the number unemployed, is governed by:

$$\dot{U} = \underbrace{\beta x_0(w)\lambda_1 n_1}_{\text{job destruction}} - \underbrace{\lambda_3 \cdot U}_{\text{job creation}}.$$

Output and the real wage are determined exactly as in the basic model except with the "effective labor force" $L-U$ instead of L.[11] Putting this real wage into the demand function (9) and substituting for Y yields the equilibrium quantity:

$$x_0(w) = \frac{L-U-n_1 N}{1-n_2+n_2^{\frac{\alpha}{1-\alpha}}}$$

Figure 8 depicts the time paths of unemployment and aggregate output with the benchmark parameter set from Section 3 together with β=0.5 and λ_3=2. The unemployment rate reaches a sharp peak in year 20, just after experimentation reaches its peak, with the predictable effect of increasing the size of the slump (from 11% to 13%).

Obsolescence

Our analysis in the previous section has already discussed various mechanisms that may potentially account for the significant size of macroeconomic fluctuations caused by the arrival of a new GPT: in particular the

11. For simplicity, we identify flows into unemployment with flows out of the labor force. This allows us to bypass the technical complications involved in modelling explicitly the bargaining game between new sectors and workers. Taking the latter more traditional modelling route would signicantly complicate the algebra without adding much in terms of economic insights.

existence of labor market segmentations; or the occurrence of errors in the experimentation process which, together with labor market frictions, will generate unemployment fluctuations on the transition path to the new GPT. There is however another and maybe more straightforward explanation for the slow-downs or slumps induced by major technological changes, one that should immediately occur to anyone remotely familiar with Schumpeter's ideas: namely, the (capital) obsolescence caused by the new wave of (secondary) innovations initiated by a new GPT.

To capture this idea, we first reinterpret the basic model by supposing that the factor used in both production and research is not labor but capital, either physical or human. Each time an innovation arrives implementing the GPT in a sector it destroys a fraction δ of the capital that had previously been employed in that sector, because all capital must be tailor made to use a specific technology in a specific sector, and some of the capital is lost when it is converted to use in another sector or with another technology.[12] For simplicity, suppose that people are target savers, that is, they save a constant fraction s per period of the gap between the desired capital stock L and the actual stock K. Then the rate of net accumulation of capital is:

$$\dot{K} = \underbrace{s \cdot (L-K)}_{\text{gross saving}} - \underbrace{\delta x_0(w) \cdot \lambda_1 n_1}_{\text{obsolescence}}.$$

Output and the real wage (that is, the real rate of return to capital) are determined as in the basic model, but with L replaced by K. The initial stationary state with $n_1=0$ has $K=L$.

It is easy to see that this modification of the basic model is formally equivalent to that of the previous section, with the gap $L-K$ replacing the number of unemployed U, the saving rate s replacing the job-finding rate λ_3, and the obsolescence fraction δ replacing the job-destruction fraction β. Thus, for the same reasons as in the previous section, the capital shortfall will peak sharply around the same time as the peak in experimentation, and the slump will be larger than if there were no obsolescence.

12. This is the assumption made in Howitt (1996).

References

Aghion, P. and P. Howitt: "A Model of Growth through Creative Destruction", *Econometrica*, 1992.

Atkeson, A. and P. Kehoe: "Industry Evolution and Transition: The Role of Information Capital", unpublished, Univ of Pennsylvania, 1993.

Cheng, L. and E. Dinopoulos, "A Schumpeterian Model of Economic Growth and Fluctuations", Mimeo, University of Florida, 1992.

David, P, "The Dynamo and the Computer: An Historical Perspective on the Productivity Paradox", *American Economic Review*, 1990, Vol 80, 2, pp. 355-361.

Greenwood, J. and M. Yorukoglu: "1974", unpublished, University of Rochester, 1996.

Helpman, E. and M. Trajtenberg, "A Time to Sow and a Time to Reap: Growth Based on General Purpose Technologies", CIAR Working Paper, 1994, No 32.

Howitt, P: "Measurement, Obsolescence, and the Adjustment to a New General Purpose Technology", in preparation for this volume 1996.

Jovanovic, B and R. Rob, "Long Waves and Short Wave: Growth Through Intensive and Extensive Search", *Econometrica*, 1990, 58, 1391-1409.

Juhn C., Murphy K., and Pierce B.,"Wage inequality and the Rise in Returns to Skill", Journal of Political Economy, 1993, 101, No3.

Shleifer, A., "Implementation Cycles", *Journal of Political Economy*, 1986