NBER WORKING PAPER SERIES

DIFFUSION LAGS AND AGGREGATE FLUCTUATIONS

Boyan Jovanovic
Saul Lach

Working Paper No. 4455

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 1993

We thank the C.V. Starr Center for Applied Economics at NYU for technical and financial help. Michael Gort and John Campbell for providing the data, Ray Atje and Ken Rogoza for capable research assistance, and the referees for useful comments. This paper is part of NBER's research programs in Productivity and Economic Fluctuations. Any opinions expressed are those of the authors and not those of the National Bureau of Economic Research.

NBER Working Paper #4455
September 1993

DIFFUSION LAGS AND AGGREGATE FLUCTUATIONS

ABSTRACT

This paper studies how random product innovations affect the time series properties of aggregates. It posits that recurring inventions of new intermediate goods differ in quality, and that their usage spreads gradually through the economy. It examines how fluctuations in per capita GNP are affected by these features of the innovation process. Micro data from the U.S. show, first, that the dispersion of products' qualities is quite large: Its coefficient of variation is 0.56. More importantly, they also show that the rate of diffusion of new products is relatively slow: Only 4.3% of the potential market size is realized in every year. Because diffusion is so slow, the model explains only low frequency movements in per capita GNP in the G-7 countries.

Boyan Jovanovic
Department of Economics
New York University
New York, NY 10003
and NBER

Saul Lach
Department of Economics
Hebrew University
Mount Scopus, Jerusalem 91905
Israel
and NBER

1. Introduction.

This paper identifies a particular type of technology shock and measures its contribution to generating fluctuations in aggregates. The shock to technology we focus on is the invention of new products, some of which are intermediate goods and can therefore be interpreted as shocks to the production function for final goods. Microeconomic data tell us how fast these products spread after they are invented, and how important they eventually become. From this we infer how much recurring invention of new products contributes to fluctuations of aggregate output.

Romer (1987) and other growth theorists have considered models in which the number of intermediate goods increases over time. In these models, goods are symmetric, and they penetrate the market instantaneosly. The same is true in models in which the number of goods fluctuates over the business cycle, models such as Devereux et al. (1993); in these models too, products are all the same, and there are no diffusion lags. In our paper we relax these two assumptions. We then use some U.S. data on newly-invented products to estimate the speed of their diffusion, and the degree of heterogeneity in the products' importance.

We find that neither symmetry of products nor their instantaneous diffusion are good approximations to reality. First, new products differ greatly in importance: The coefficient of variation of the distribution of quality over products is estimated at 0.56. And second, if there is something that one can call "eventual market penetration" of a product, the typical product approaches this value very slowly -- at the rate of 4.3% per year. The diffusion of new products is, in other words, quite slow.

Given these estimates, we then ask how much of the fluctuation in per-capita GNP of the advanced countries is attributable to this type of shock. In our model, the persistence of the output effects of the technology shock depends on the speed of diffusion. The faster the usage of a new product diffuses throughout the economy, the less persistent its effect on output will be. Since we find that new products are slow to penetrate the market, the shock explains only the relatively highly persistent component of the business cycle. It does not generate movements at high frequencies. A really great new product like the computer will eventually raise output by a lot once it is in widespread use. But by the time it has spread, newer products will have appeared on the scene; the cumulative effect of these shocks is therefore a combination of many independent influences, and these are subjected to too much averaging to have an aggregate impact at high frequencies. In sum, the invention of new products can explain only low frequency movements in aggregates.

High frequency movements in aggregates can still result from process innovations that affect many goods simultaneously. We have not ruled out the possibility that widely applied process innovations may explain some high frequency movements in aggregates. And then there are other shocks that the model leaves out: Policy shocks, legal shocks, shocks to import prices, shocks to management techniques, and so on. Many of these occur at higher frequencies, and they presumably can account for the discrepancy between our model and the data.

2. Diffusion Lags and Aggregate Output.

Assume the following aggregate production function:

$$Y_t = L_t^{1-\alpha}\int_0^{A_t} q_{it}^{\alpha}\,di \ . \tag{1}$$

Here L_t is the labor input, q_{it} is the quantity of the i^{th} intermediate input, and A_t is the number of intermediate inputs available at t. Romer (1987) uses such a production function. We depart from his formulation in that we identify the intermediate inputs q_i with inventions of different quality that are gradually adopted by producers.

If market size does not affect a product's rate of adoption, q_{it} will be proportional to the country's number of producers, or population, L_t. And if product i was introduced at date s_i, its output at t should also depend on its age, $t - s_i$. Thus we shall write

$$q_{it} = L_t h_i(t - s_i) \ .$$

New products differ in their importance; to capture such differences let

$$h_i(t - s_i)^{\alpha} \equiv \phi(t - s_i,\ \theta_{s_i},\ \epsilon_{it}) \ . \tag{2}$$

The shock θ_s affects all products of vintage s, while the shock ϵ_{it} is product-specific, independent over i and t, with CDF $G(\epsilon)$. The average output at t of products of vintage s is

$$\int \phi(t-s, \theta, \epsilon) dG(\epsilon) \equiv f(t-s, \theta) . \qquad (3)$$

Because a continuum of products of measure 1 arrives at each date, the ϵ's will wash out, and average output coincides with total output. Per capita income then is

$$\frac{Y_t}{L_t} \equiv Y_t = \int_0^{A_t} f(t - s_i, \theta_{s_i}) di .$$

Now suppose that A_t grows exogenously at the rate λ, and normalize A_0 to equal unity. Then $A_t = e^{\lambda t}$. Changing the variable of integration from the product name i to its vintage s, where $i = e^{\lambda s}$, yields.

$$y_t = \lambda \int_{-\infty}^{t} e^{\lambda s} f(t - s, \theta_s) ds = \lambda e^{\lambda t} \int_0^{\infty} e^{-\lambda \tau} f(\tau, \theta_{t-\tau}) d\tau = e^{\lambda t} X_t , \qquad (4)$$

where $X_t \equiv \lambda \int_0^{\infty} e^{-\lambda \tau} f(\tau, \theta_{t-\tau}) d\tau$ is a moving average of past θ's. Equation (4) is intuitive: There are $\lambda e^{\lambda(t-\tau)}$ products of age τ, and each contributes an amount $f(\tau, \theta_{t-\tau})$. We then just add over all ages.

Proposition 1: If the $\{\theta_t\}$ process is stationary, $\log y_t$ is stationary around the trend λ.

The long-run growth rate is λ; it does not depend on f. Diffusion lags therefore have only level effects in the long run.

To see how diffusion lags affect the cyclical properties of y_t, the X_t component, assume that

$$\phi(\tau, \theta, \epsilon) = \begin{cases} (1 - e^{-\rho\tau + \epsilon})\theta & \text{for } \tau > 0 \\ 0 & \text{for } \tau \leq 0 . \end{cases} \tag{5}$$

The parameter θ denotes the product's eventual ($\tau = \infty$) market size.[1] A large θ_s indicates that products of vintage s, and that vintage only, will eventually become more important. The parameter ρ measures the speed of diffusion -- it is the same for each product.

Substitute for ϕ into eq. (3), and assume that $\int e^{\epsilon} dG(\epsilon) = 1$. Then $f(\tau, \theta) = (1 - e^{-\rho\tau})\theta$. Eq. (4) then implies

---

[1] Therefore products last for ever. This parametrization seems apt given the level of aggregation: The Gort-Klepper products show little tendency of disappearing as they age. But at lower levels of aggregation, there are bound to be many products that eventually disappear.

$$X_t = \lambda \int_0^{\infty} e^{-\lambda\tau}(1 - e^{-\rho\tau})\theta_{t-\tau} d\tau \ . \qquad (6)$$

The process $\{\theta_t\}$ is assumed to be stationary and serially uncorrelated:

$$E(\theta_t) = \mu \quad \text{and} \quad Cov(\theta_t, \theta_s) = \begin{cases} \sigma^2 & \text{if} \quad t = s \\ 0 & \text{if} \quad t \neq s \end{cases} \qquad (7)$$

The level effect of the diffusion rate can be measured by the long-run mean of X_t, which is increasing in the speed of diffusion ρ:[2]

$$EX_t = \frac{\rho\mu}{\rho + \lambda} .$$

The appendix shows that for each $k \geq 0$,

$$Cov(\ln X_t, \ln X_{t-k}) \approx \psi^2 \frac{\lambda(\lambda + \rho)}{2(2\lambda + \rho)} e^{-\lambda k} \left[1 + \frac{\lambda}{\rho}(1 - e^{-\rho k}) \right], \qquad (8)$$

where $\psi \equiv \sigma/\mu$ is the coefficient of variation of the distribution of θ. The

[2] The expressions that we shall now present are valid (in roughly the same form) in discrete time and with discrete numbers of innovations. We use continuous time here because it leads to simpler formulas.

autocorrelation coefficient of $\ln X_t$ is

$$r_k = e^{-\lambda k}\left[1 + \frac{\lambda}{\rho}(1 - e^{-\rho k})\right]. \tag{9}$$

While r_k decreases with ρ for all k, it decreases with λ only for large k, and it increases with λ for small k. As $\rho \to \infty$, $r_k \to e^{-\lambda k}$, and the variance of ln y tends to $\lambda\sigma^2/2\mu^2$. Thus faster diffusion reduces persistence, but it raises the steady-state variance around trend.

Since $\ln X_t$ equals de-trended $\ln y_t$, we shall evaluate the model by comparing the predicted autocovariances in eq. (8) to those of de-trended output. To do so, we need estimates of ρ, ψ, and λ. The first two are estimated using micro data, the third using aggregate data.

3. Estimates of ρ and ψ: The Gort-Klepper Data.

The data assembled by Gort and Klepper (1982) document the historical development of 46 new products in terms of their sales, price, output, and numbers of producers over (part of) the life-cycle of each product. Table 1 lists the 21 products for which we have sales data. Most of them seem to qualify as intermediate inputs in the sense of eq. (1). Since the number of products is not that large, we analyze them all. Column 1 tells us when each product was introduced into the market. There are old products such as records, dating from 1887, as well as relatively new ones such as lasers, which became available in 1960. The last year for which data were collected was 1972 and, in general, sales and quantity of output figures were available for only a part of the life of the product. The age range for which there are data appears in the second column.

The last three columns of table 1 report our estimates of the product-specific θ's. The production function in (1) treats intermediate products as exchangeable inputs: One unit of product i and two units of product j can produce as much final output as two units of product i and one unit of product j. Now computers and ballpoint pens are surely not exchangeable in this sense, and something must be done to bring them into common units. We shall do this by expressing everything in units of the 1967 "consumption good", so that for q_{it} we shall use product i's sales at t, deflated to 1967 dollars by the Wholesale Price Index. Gort and Klepper discuss their data at length. One property that they point to is that on average there is a rapid decline in the rate at which

Product	Initial Year	Age Range	$\hat{\theta}$	S.D. of θ	$\hat{\theta}^0 = \max_t \left(\frac{q_{it}}{L_t}\right)^\alpha$
1. Computers	1935	20-36	3.50	1.28	4.81
2. Crystals, Piezo	1936	25-36	0.65	0.05	.69
3. DDT	1943	1-27	0.58	0.10	.63
4. Electrocardiographs	1914	47-58	0.36	0.05	.38
5. Electric Blankets	1911	35-61	0.76	0.07	.80
6. Electric Shavers	1930	1-42	0.87	0.21	.97
7. Fluorescent Lamps	1938	0-34	0.87	0.17	.98
8. Freezers, Home and Farm	1929	18-43	1.39	0.19	1.47
9. Gyroscopes	1911	52-61	0.48	0.08	.52
10. Lasers	1960	3-11	0.59	0.23	.80
11. Missiles, Guided	1942	9-30	2.37	0.62	2.77
12. Motors, Outboard	1908	42-64	1.00	0.10	1.06
13. Penicillin	1943	2-28	0.95	0.30	1.03
14. Pens, Ballpoint	1945	6-27	0.83	0.11	.89
15. Records, Phonograph	1887	34-85	1.58	0.44	1.77
16. Streptomycin	1945	1-27	0.67	0.39	.71
17. Styrene	1935	8-36	0.85	0.08	.88
18. Tapes, Recording	1947	14-25	0.99	0.09	1.03
19. Television, Apparatus, Parts	1929	17-43	2.08	0.50	2.37
20. Transistors	1948	6-24	1.17	0.31	1.35
21. Tubes, Cathode Ray	1922	26-50	1.09	0.24	1.21
Average					

Notes: For each product, the size of the sample is the number of years included in the age range, including the end years. So for computers, for example, the size of the sample is 17.

Table 1: The Gort-Klepper Products.

sales and quantity of output grow with the age of the product, and that their growth rates asymptote to zero. The functional form with this property is in eq. (5); it implies that

$$(q_{it}/L_t)^{\alpha} = \left[1 - \exp(-\rho(t-s) + \epsilon_{it})\right]\theta_s \qquad \text{for } t \geq s, \tag{10}$$

where the q_i are sales in 1967 dollars, L is the population of the US, and α = 1/3. Since the time series for each product is not too long, and since we assume that the parameter ρ is the same over products, we shall estimate ρ by pooling the data.

Estimating ρ and θ_i (for each i) requires a non-linear procedure. We proceed in two steps: first we obtain a consistent estimator of θ_i and then use it, and the value of α, to transform the nonlinear equation into a linear one. The second step is to estimate ρ by linear OLS on a transformed version of (10). Given an estimate of ρ, we can reestimate θ and iterate this two-step procedure till convergence. It turned out, however, that the initial and first-round estimates of θ were quite similar so that we did not iterate.

Our initial estimate of θ_i was the maximal value over t of $(q_{it}/L_t)^{\alpha}$. Let T_i be the largest age for which data are available for product i. Since $(q_{it}/L_t)^{\alpha}$ never exceeds θ_s and approaches it with probability one as t gets large, this is a consistent estimator of θ_i as T_i increases. Denote the initial estimate by θ_i^0. It is reported in the last column of table 1. The second step is to estimate the OLS regression $p_{it} = \gamma - \rho AGE_{it} + u_{it}$, where p_{it}

$= \log\,[1 - (q_{it}/L_t)^\alpha/\hat{\theta}_i^0]$, AGE = t - s, u_{it} equals ϵ_{it} plus an additional error term generated by estimating θ_i and γ captures the nonzero mean of ϵ.

The OLS estimate of ρ from the pooled sample, 499 observations, is 0.043 with a standard deviation of 0.013. The first-round estimator of θ_i is obtained by solving for $\hat{\theta}_{it} = \dfrac{(q_{it}/L_t)^\alpha}{[1 - \exp(\hat{p}_{it})]}$, where $\hat{p}_{it}$ is the predicted p_{it} from the pooled regression. We then average over all ages t to obtain an estimate of θ_i. This estimate and its standard deviation within each product are in columns 3 and 4 of the Table. Finally, in column 5 we present the initial estimate of θ_i, namely $\hat{\theta}_i^0$; the two sets of estimates do not differ much.

For μ we used the average of the individual $\hat{\theta}_i$'s, namely 1.16, although this does not differ much from the pooled estimate. The estimate of σ^2 is the between-products variance of θ_i. If we regard the within-variance component as the variance of a measurement error then this should be subtracted from the estimated total variance of θ.[3] Its coefficient of variation, ψ, therefore is 0.56. The frequency distribution of $\hat{\theta}$ is in Figure 1.[4]

---

[3] This adjustment is needed because the $\hat{\theta}_i$ are estimates so that their variance partly reflects sampling variability.

[4] Being based on BLS figures, the Gort-Klepper sales data do not fully control for quality change that is passed on to the consumer and therefore not fully reflected in sales. Adjusting for this could make a huge difference to our estimates of some of the θ's, especially for computers (Gordon, 1990). This underestimate of quality change was the largest for those products for which our $\hat{\theta}_i$ is the largest, such as computers, television, and transistors. On these grounds, ψ is probably underestimated.

But a second consideration suggests that ψ is overestimated. Several inventions may be lumped together and labelled as just one invention. Perhaps computers should count as several inventions. Arbitrary classification errors therefore would lead to an overestimate of ψ. These are two reasons why it is hard to assign a standard error to our estimate of ψ reported in table 3.

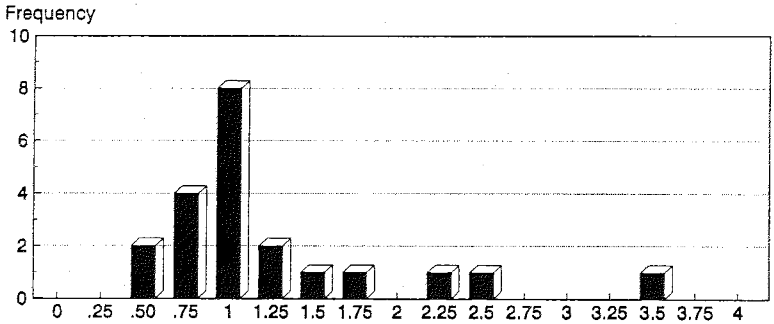

Figure 1: Frequency Distribution of $\hat{\theta}$ (in 1967 $'s per capita)

Figure 2 provides information that bears on our assumption that θ is stationary and, indeed, serially uncorrelated. The figure plots the size of $\hat{\theta}$ on its vintage, and it reveals neither trend nor autocorrelation, which is consistent with the assumptions in eq. (7).

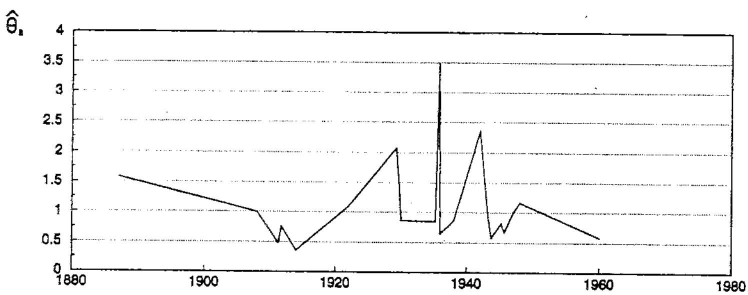

Figure 2: The estimated time series $\hat{\theta}_s$

4. Evaluating the Model: The G-7 Data.

In our model, only one shock drives per capita GNP. But actual data respond to other influences as well. Our aim is to find out how important our technology shock is relative to these other factors. Specifically, we want to measure its contribution to the time-series properties of per capita GNP in the G-7 countries. These are countries for which applying our US-based estimates of ρ and ψ seems justifiable to us.

The G-7 income data are from the International Monetary Funds's International Financial Statistics. These are post-war quarterly series for Gross National product (GNP) in four countries (Canada, Germany, Japan and the United States), and for Gross domestic product (GDP) in three countries (France, Italy and the United Kingdom). All series are seasonally adjusted. The sample periods are shown in Table A.1.

Proposition 1 states that the predicted long-run growth rate of y is λ, and so we have set λ equal to the average growth rate of per-capita incomes, which is 2.9 percent per year. The Gort-Klepper data have provided us with estimates of ρ and ψ summarized in Table 2.

Parameter	Estimate	Standard Error
λ	0.029	0.002
ρ	0.043	0.013
ψ	0.56	n.a.

Table 2: Parameter Estimates and their Standard Errors.

The model precisely predicts the autocovariance and autocorrelation of detrended $\ln y_t$, in equations (8) and (9). We shall focus on the autocovariance function because it easily reveals what our model accomplishes and where it falls short. It falls short at high frequencies because new products diffuse slowly and because the model omits all other temporary shocks.

The autocovariance function of detrended $\log y_t$ is based on the OLS residuals from a regression of $\ln y_t$ on time. Averaging over all 7 countries we get $\gamma(k) = \Sigma_i \gamma_i(k)/7$. Figure 3 plots $\gamma(k)$ on k ranging from 0 to 10 years. Using the estimates of λ, ρ and ψ we compute the predicted $\gamma(k)$, also plotted in figure 3.[5]

The explanatory power of the model rises with the length of the lag: it explains 7 percent at lag zero, 16 percent at lag 5, and most of it at lag 10. Evidently, it leaves most high frequency variation unexplained.

<u>The response of y to an impulse in θ</u>. In equation (6), the weight on $\theta_{t-\tau}$ is $\lambda e^{-\lambda\tau}(1 - e^{-\rho\tau})$. Dividing it by our estimate of σ yields the unimodal impulse response function shown in figure 3. The peak takes place at $t = \rho^{-1}\ln[(\lambda+\rho)/\lambda]$ after which it declines and converges to zero. The estimates imply that the impulse response function peaks at about twenty seven years, a surprisingly large number that underlies the high persistence of $\ln y_t$ in this trend-stationary model. This differs from the conventional model in which the response decays much faster.[6]

---

[5] The figure also reports a 95 percent confidence interval around the predicted $\gamma(k)$, as explained in the appendix.

[6] Technological shocks do not have permanent effects on y_t because the permanent <u>absolute</u> effect of θ_t on output shrinks relative to y_t as the latter grows without bound. Marco Lippi and Lucrezia Reichlin (1990) also discuss the

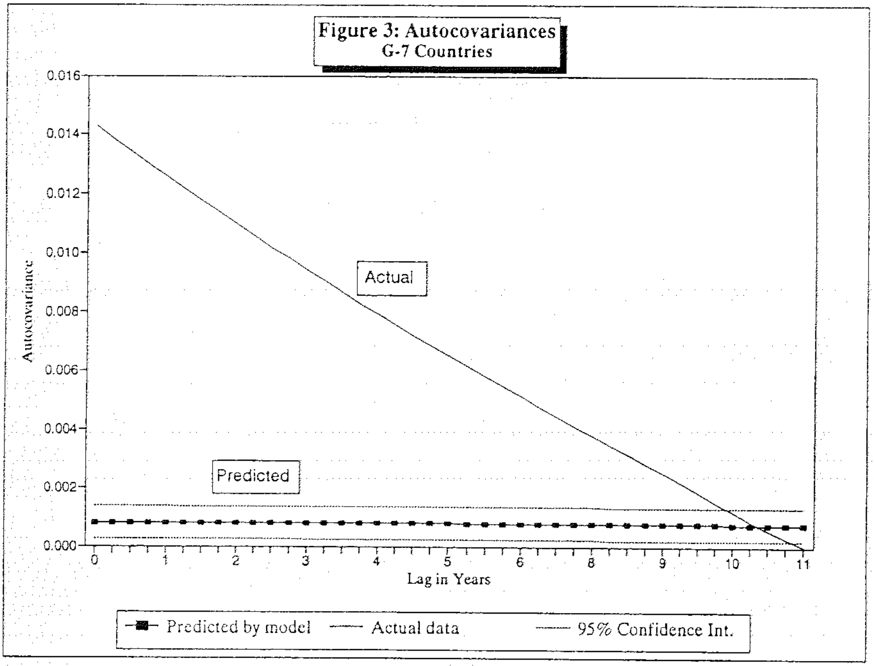
Figure 3: Autocovariances
G-7 Countries
0.016
0.014
0.012
0.010
0.008
0.006
0.004
0.002
0.000
Autocovariance
Actual
Predicted
0
1
2
3
4
5
6
8
9
10
11
Lag in Years
Predicted by model
Actual data
95% Confidence Int.

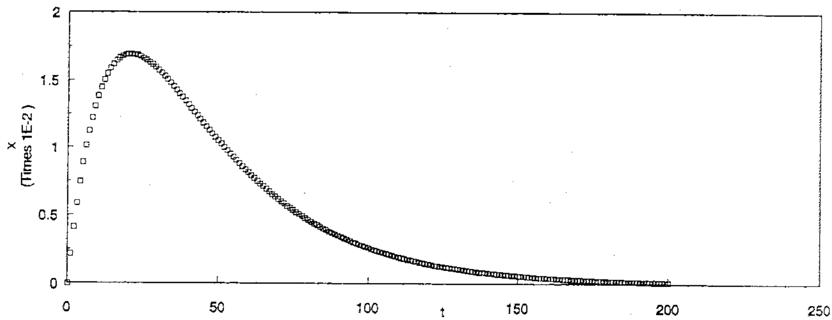

Figure 4: The response to a one standard-devation impulse in θ at date zero.

5. Comparison with Other Approaches.

The dominant approach to analyzing the effect of technology shocks on aggregates is the real business cycle model summarized by Prescott (1986). It treats the state of technology as an unobservable process, and aims to produce reasonable-looking business cycle behavior with a simple model shocked by a first-order autoregressive technology process. In spite of its success in this regard, a drawback of the approach is that it treats technology as a residual.

A second approach pursues Schumpeter's suggestion that business cycles may be driven by "bunching" of innovations. In our terminology this amounts to fluctuations in λ. Since direct measures of the frequency of inventions such as aggregate patent statistics and direct counts of innovations do fluctuate, Kleinknecht (1987) and others have argued that bunching of innovations does indeed occur, although they have had less success in tying the invention process to the business cycle. A problem with this approach is that the economic value

---

effect that diffusion lags might have on the time series properties of GNP.

of patents and "inventions" may itself change in such a way as to offset, at least in part, the effect that fluctuations in their <u>number</u> have on aggregates. While the Gort-Klepper data do not solve this problem entirely (see note 4), they do provide information on the economic value of the innovations.

Our approach incorporates elements of both of the above. As in the first approach, we quantify the implications of the moel, and compare them to the data in second moment form. Moreover, we derive an exact relation between the speed of diffusion of innovations and the persistence of output fluctuations. But as in the second approach, we measure technology directly and not as a residual.

6. <u>Conclusion</u>

We find that product innovations explain fluctuations at lower frequencies fairly well, but that it has little to do with fluctuations at higher frequencies. This finding is based on the Gort-Klepper data and on post-war aggregate data from the G-7 countries. One must therefore resist generalizing it to situations where that evidence does not apply.

We feel, however, that our conclusions are unlikely to be overturned by other data. First, output fluctuations are both less variable and more persistent in the developed world than elsewhere, so that bringing in more countries into the process of estimating the autocovariance function displayed in figure 3 would only show an even steeper empirical autocovariance. For the model to a better job at higher frequencies <u>but not overpredict at lower ones</u>, we would need a much higher estimate of the speed of diffusion of new products. But whatever scattered evidence there is on the diffusion of other products (or, for that matter process inventions as well) indicates a speed of diffusion much like that of the Gort-Klepper products.[7]

Since the diffusion of new products is slow, we can be confident that the predicted autocovariance function must be fairly flat -- as in figure 3. On the other hand, we are less sure about the intercept of the predicted autocovariance. This is because, as table 3 indicates, we are not too confident about our estimate of ψ: The data may understate quality change of some key products, and the BLS may have misclassified or wrongly lumped products together. Therefore, further attempts to measure quality differences among distinct new products would shed more light on the link between product innovation and the business cycle.

[7] The literature on diffusion often calculates the statistic "Δt", defined to be the time it takes a new product or process to grow between 10% and 90% diffusion. Our estimate of $\rho = 0.043$ implies that our Δt is 30 years. The most comprehensive study of diffusion (of 265 innoavtions) in the U.S. seems to be the one by Grubler (1991). He reports that the largest number of diffusion processes have Δt's on the order of 15-30 years, and the sample mean Δt is 41 years. Our estimate of the speed of diffusion is therefore reasonable. Unfortunately we are unable to check our estimate of ψ against his sample because it does not contain information about the value of the inventions. This is why the Gort-Klepper sample is so valuable.

References:

Devereux, Michael, Allen Head and Beverly Lapham, "Exit and Entry, Increasing Returns to Specialization, and Business Cycles," Queen's University, June 1993.

Gordon, Robert, The Measurement of Durable Goods Prices, NBER, Chicago: University Press, 1990.

Gort, Michael and Steven Klepper, " Time Paths in the Diffusion of Product Innovations," Economic Journal, 92 (September 1982): 630-653.

Grübler, Arnulf, "Diffusion: Long Term Patterns and Discontinuities," Technological Forecasting and Social Change 39 (1991): 159-80.

Kleinknecht, A. Innovation Patterns in Crisis and Prosperity: Schumpeter's Long Cycle Reconsidered London: Macmillan, 1987.

Lippi, Marco, and Lucrezia Reichlin, "Diffusion of Technical Change and the Identification of the Trend Component in Real GNP," unpublished paper, O.F.C.E., Paris, January 1990.

Prescott, Edward, "Theory Ahead of Measurement in Business Cycle Research," Carnegie Rochester Conference on Public Policy, Autumn 1986, 25, 11-44.

Romer, Paul, "Crazy Explanations for the Productivity Slowdown," NBER Macroeconomics Annual (1987).

Appendix A: Calculating $\mathrm{Cov}(X_t, X_{t+k})$.

Let $a_\tau \equiv e^{-\lambda\tau}(1-e^{-\rho\tau})$. Then

$$\mathrm{Cov}(X_{t+k}, X_t) = \lambda^2 \mathrm{Cov}\left(\int_0^\infty a_\tau \theta_{t+k+\tau} d\tau,\ \int_0^\infty a_\tau \theta_{t-\tau} d\tau \right).$$

But if we change variables from τ to $s \equiv \tau - k$, then $\tau = 0 \Rightarrow s = -k$, $\tau = \infty \Rightarrow s = \infty$, $t+k-\tau = t-s$, and $\tau = k+s$. Therefore, since $ds = d\tau$,

$$\mathrm{Cov}(X_{t+k}, X_t) = \lambda^2 \mathrm{Cov}\left(\int_{-k}^\infty a_{s+k}\theta_{t-s} ds,\ \int_0^\infty a_s \theta_{t-s} ds \right)$$

$$= \lambda^2 \mathrm{Cov}\left(\int_0^\infty a_{s+k}\theta_{t-s} ds,\ \int_0^\infty a_s \theta_{t-s} ds \right)$$

$$= \lambda^2\sigma^2 \int_0^\infty a_{s+k} a_s ds \qquad = \lambda^2\sigma^2 \int_0^\infty e^{-\lambda k} e^{-2\lambda s}(1-e^{-\rho(k+s)})(1-e^{-\rho s}) ds$$

$$= \lambda^2\sigma^2 e^{-\lambda k} \int_0^\infty \left[e^{-2\lambda s}(1-e^{-\rho s}) - e^{-\rho k} e^{-2\lambda s} e^{-\rho s}(1-e^{-\rho s}) \right] ds$$

$$= \lambda^2\sigma^2 e^{-\lambda k} \left[\int_0^\infty a^{-2\lambda s} ds - \int_0^\infty e^{-(2\lambda+\rho)s} ds - e^{-\rho k} \int_0^\infty e^{-(2\lambda+\rho)s} ds + e^{-\rho k} \int_0^\infty e^{-2(\lambda+\rho)s} \right] ds$$

$$= \lambda^2\sigma^2 e^{-\lambda k} \left[\frac{1}{2\lambda} - \frac{1}{2\lambda+\rho} - e^{-\rho k} \left(\frac{1}{2\lambda+\rho} - \frac{1}{2(\lambda+\rho)} \right) \right]$$

$$= \lambda^2\sigma^2 e^{-\lambda k} \left(\frac{\rho}{2\lambda(2\lambda+\rho)} - \frac{e^{-\rho k}\rho}{(2\lambda+\rho)2(\lambda+\rho)} \right) = \frac{\rho\lambda^2\sigma^2 e^{-\lambda k}}{2(2\lambda+\rho)} \left(\frac{1}{\lambda} - \frac{e^{-\rho k}}{\lambda+\rho} \right).$$

$$= \frac{\lambda\sigma^2\rho^2 e^{-\lambda k}}{2(\lambda+\rho)(2\lambda+\rho)} \left[1 + \frac{\lambda}{\rho}(1-e^{-\rho k}) \right].$$

Moreover, for any two variables u and v, $\mathrm{Cov}(\ln u, \ln v) \approx \mathrm{Cov}(u, v)/E(u)E(v)$. This approximation underlies the expression in eq. (8).

Computation of Confidence Intervals. Write $\mathrm{Cov}(\ln X_t, \ln X_{t-k}) \equiv \gamma(k; \delta)$ where δ is the vector of parameters (λ, ρ, ψ). Let hats denote estimated values. The variance of $\gamma(k; \hat{\delta})$ is gotten from the variance of the linear Taylor expansion of $\gamma(k; \hat{\delta})$ around δ. This results in an estimated variance given by $g(\hat{\delta})'\Omega g(\hat{\delta})$, where $g(\hat{\delta})$ is the column vector of partial derivatives of $\gamma(k; \hat{\delta})$ evaluated at $\hat{\delta}$ and Ω is the 3 x 3 covariance matrix of $\hat{\delta}$. The diagonal elements of Ω were taken from the third column of Table 2, assuming that the variance of ψ is zero, which means that the reported bands are tighter than they would be. [Footnote 4 mentions two reasons -- unmeasured quality change and classification error -- why some of the θ_i may be mismeasured, and why it, therefore, is hard to assign a standard error to our estimate of ψ.] The off-diagonal elements of Ω are assumed to be zero. The confidence interval for the true autocovariance is based on using the normal distribution as an approximation to the true distribution of $\gamma(k, \delta)$. The 95 percent confidence interval is then $\gamma(k; \hat{\delta}) \pm 1.96\ [g(\hat{\delta})'\Omega g(\hat{\delta})]^{1/2}$.

Country	Sample	Number of Observations
Canada (GNP)	1957:1-1991:2	138
France (GDP)	1970:1-1991:2	86
Germany (GNP)	1960:1-1991:2	126
Italy (GDP)	1960:1-1991:1	125
Japan (GNP)	1957:1-1991:1	137
United Kingdom (GDP)	1957:1-1991:1	137
United States (GNP)	1957:1-1991:2	138

Table A.1: Sample Periods and Number of Observations.