INTERNATIONAL ECONOMIC REVIEW
Vol. 38, No. 1, February 1997

PRODUCT INNOVATION AND THE BUSINESS CYCLE*

BY BOYAN JOVANOVIC AND SAUL LACH[1]

University of Pennsylvania, U.S.A.
The Hebrew University, Israel, the Federal Reserve Board and the National Bureau of Economic Research, U.S.A.

Microeconomic data show two important facts about new products. First, some products are more important than others. And second, it takes them years to penetrate the market significantly. Our calibrated model with these features overpredicts the autocovariance of U.S. GNP at long lags, but underpredicts it at short lags. The latter is not surprising, since the model leaves out other obvious high-frequency shocks. The puzzle is why the U.S. GNP data do not show stronger autocorrelation at higher lags. A surprising finding is that while the speed of diffusion has huge level effects, it plays a minor role in shaping the business cycle.

1. INTRODUCTION

To implement new technology, a producer must usually buy some equipment that embodies it. This suggests that a technology shock will induce a new type of equipment, a new capital good. Indeed, one can distinguish technology shocks by the different capital goods that they may give rise to: The vintage of the capital good then reveals the date the technology shock occurs, and the subsequent sales of the capital good reveal how "big" the technology shock was.

1.A. *Two Questions*: We focus on two questions. The first is: How important is product innovation in shaping the business cycle? When calibrated to some micro U.S. data on new products, our model generates about one quarter of the variance of GNP around trend. This finding matches that of Greenwood, Hercowitz, and Krusell (1994), who use a very different model and methodology, and come up with a number of twenty percent.

The second question is: If product innovations do indeed cause aggregate fluctuations, then at what frequencies do they do so? We answer this question by comparing the autocovariances of U.S. per-capita GNP at various lags to the autocovariances predicted by the model. The fluctuations implied by the model turn

* Manuscript received December 1994; revised December 1995.

[1] We thank the C.V. Starr Center for Applied Economics at New York University, U.S.A. for technical and financial help, Michael Gort for providing the data, Frank Diebold, Doug Dwyer, Jordi Gali, Jeremy Greenwood, Zvi Griliches, Arnulf Grubler, Jinyong Hahn, Zvi Hercowitz, Pete Klenow, and Paul Romer for comments, and Ray Atje and Ken Rogoza for research assistance. This paper presents the authors' own views and not those of the Federal Reserve System.

out to be too sluggish; the long-lag autocovariances are *overpredicted*, mainly because of the longevity of products. On the other hand, the short lag autocovariances are *underpredicted*, probably because other shocks are left out. Surprisingly, the speed of diffusion turns out to have little bearing on the character of fluctuations, at least as summarized by the autocovariance function.

1.B. *Methodology*: We ask how much of the business cycle can one explain with a particular type of technology shock. We use the "vintage capital" idea (Johansen 1959, Salter 1960, and Solow 1959) that equipment is largely technology-specific. If this is correct, a technology shock can be measured, and its effects over time can be traced, by looking at what happens to the equipment that embodies the technology in question. We take the variation in quality of new products and their speed of diffusion as exogenous and then predict the autocovariance structure of GDP. We then compare it with what is observed using the metric of the autocovariance function. To do this, we need to use micro-evidence to justify our assumptions about the key parameters—how variable the quality of innovation is, how fast it spreads, and so on. We therefore look at the evolution of the twenty or so capital goods included in the Gort-Klepper (1982) sample. We then add these effects by assuming an aggregate production function of the kind that Romer (1987) and other growth theorists have proposed. In this production function, output depends on a weighted sum of capital goods. We then assume that the number of these capital goods increases over time.

1.C. *Other Approaches to Technology and the Business Cycle*: There are three other approaches to the same question. The first is the Real Business Cycle approach (Prescott 1986). It uses a model shocked by a first-order autoregressive technology process and, with just a handful of parameters, manages to produce realistic-looking cycles. A drawback of the approach, however, is that it does not measure the technology shock *directly*. Instead, it treats technology as a residual—an unobservable. It does not use micro data on technology in order to determine their size and the frequency of shocks.

In the second approach, Greenwood, Hercowitz and Krusell (1994) *measure* investment-specific technological change by looking at the price of capital goods relative to that of consumption goods, a series constructed by Gordon (1990). They report a correlation between the Hodrick-Prescott detrended relative price series and equipment-investment series of -0.46. When they simulate their model using a shock with the properties of the detrended relative price, they can explain only about 20 percent of the variation in output. Our own simulations use information from entirely different micro data, and yet we shall reach a similar quantitative conclusion. Of course, the relative price of capital goods can change for reasons other than technology shocks, so its unclear how much of the covariation between the price series and aggregate activity actually *is* caused by movements in technology.[2]

[2] A difficulty with the Gordon price series is that it falls so fast during the 70s and 80s that it implies tremendous technological *regress* in the consumption goods sector over the same period.

The third approach, like the second, also tries to measure technology directly. It focuses on innovation "counts," thereby pursuing Schumpeter's suggestion that business cycles are driven by a bunching of innovations. Kleinknecht (1987) studies aggregate patent statistics which do fluctuate substantially. Klenow (1994) looks at advertisements of new products, and finds them to be mildly procyclical. Geroski and Walters (1995) look at an innovation "count" measure for the U.K., and find that this measure is procyclical, as are patents. A problem with this approach is that the economic value of patents and "inventions" may itself change in such a way as to offset, at least in part, the effect that fluctuations in their *number* have on aggregates. Still, this approach is complementary to, and easily combined, with ours.

1.D. *Our Modeling Approach*: We do not try to explain how new products are invented, or why they differ in quality; presumably there is some randomness in the invention process—a process that is still largely a black box in theoretical work. But the diffusion of inventions has been much analyzed. The origins of diffusion lags have been primarily linked to vintage physical capital (originally Johansen 1960, Salter 1960, Solow 1959, extended to a cyclical context by Caballero and Hammour 1994, and by Bouccekine et al. 1995). But they also have been linked to vintage human capital (Chari and Hopenhayn, 1991), to lack of awareness (Griliches 1957, Jovanovic and MacDonald, 1994), to second-mover advantages (Jovanovic and Lach, 1989), and to the profitability of the inventions (Griliches 1957, Jovanovic and Rob 1990).

We do not build a structural model along any of these lines. Rather, we assume a steady arrival of innovations, in the form of intermediate products. Each generation of products is of different average quality. The randomness of the quality of each vintage produces aggregate fluctuations. We postulate a particular outcome for the diffusion curve for new inventions, which presumably are the most important component of technology shocks. 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. From this we infer how much the recurring invention of new products contributes to fluctuations of aggregate output.

1.E. *Summary of Results*: 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 8.1% per year. The diffusion of new products is, in other words, quite slow.

Given these estimates, we ask how much of the fluctuation in U.S. per-capita GNP stems from this type of shock. In our model, the exact nature of the persistence of the effects on output of the technology shock depends on the speed of diffusion and on the degree of persistence in how new products evolve. We find that since new products stay around for a long time, the shock explains (indeed it overpredicts) the long lag autocovariance component of the business cycle. The shock, however, does not generate large enough autocovariances at short lags. 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 relatively slow variation in aggregates. Indeed, when its parameters are estimated from micro data, the model substantially *overpredicts* the long-term autocovariance of de-trended GNP. The underprediction of short-lag autocovariances was expected—there are shocks that the model leaves out: policy shocks, legal shocks, shocks to import prices, shocks to management techniques, and so on. They presumably can account for the discrepancy between our model and the data at high frequencies. More puzzling, however, is the model's substantial overprediction of the longer-term autocovariances. This overprediction stems from two clear properties of the microeconomic data: new products evolve in a persistent manner, and the intergenerational variation in the quality of new products is substantial. Surprisingly, the rate of diffusion of new products plays a negligible role in what the model predicts for both amplitude and persistence of the business cycle. This challenges Schumpeter's view that a business cycle is a wave whose shape depends on the pattern of diffusion.

All this notwithstanding, we do find an effect of the speed of diffusion on the *level* of output, and the effect is *big*. A society in which technologies spread quickly will be far more developed than another in which technologies spread slowly.

1.F. *General Purpose Technologies*: We have not assumed any qualitative difference between the technologies—some are bigger than others, and there is a continuum of types. An alternative view is that there are two essentially different types of technological innovations. Most innovations are small, but occasionally we see what Bresnahan and Trajtenberg (1995) have called "general purpose technologies" (GPTs). *Their* diffusion speed matters because they interact directly with other technologies. Models that use this as a centerpiece are Jovanovic and Rob (1990), Andolfatto and MacDonald (1994), and Lippi and Reichlin (1994). The logic is that if there are big shocks that hit the economy once every few years, then diffusion lags must surely matter. The trouble with this view is that in fact, such big technologies take far too long to spread—much longer than the length of the typical business cycle. For example, the time from 10% to 90% diffusion of the railway was 54 years in the U.S.A., and 37 years in the U.S.S.R., and of the steam locomotive (a much smaller invention) 12 years in the U.S.A., and 13 years in the U.S.S.R (Grübler 1991, Table 1). Such long delays are found in a whole host of other inventions (Grübler 1991, p. 177), and in the Gort-Klepper sample that we study here.

1.G. *What Next?*: The contrasting views of the technology shock are mainly differences in emphasis. One can view cycles as the result of the occasional arrivals of GPTs and their subsequent spanning of application sectors. Or one can view growth as the result of the continuous arrival of *symmetric* innovations—some big, some small—and their subsequent improvements. Far from ruling out GPTs, we instead assume that they will recur with some statistical regularity. Similarly, one can view cycles as the result of a "bunching" of innovations; here more important products are the embodiment of a larger number of innovations, so that, for example, the computer embodies a large "bunch" of smaller innovations. Indeed, if every once in a while, some truly discrete changes do happen, such as the information technology revolution, our approach, and that of Greenwood et al. and Klenow ought to work. For instance, the Greenwood et al. approach to accounting for technical change should work if the investment deflator is correctly calculated to include all investment-specific technological changes, large and small. The Klenow "counting" approach will work if GPTs are associated with a wave of patents, or a wave of product introductions. And our approach should work if the sample of products is representative, and if a product's sales are proportional to its contribution to the aggregate production function.

Done correctly, these different approaches to accounting for the business cycle with technology shocks will reveal a stream of arrivals of new technologies, fluctuating in number and in quality. To be truly convincing, this literature will need to trace through the effects of various technological changes on the economy at large. Grübler and his coworkers are trying to do just this, with different tools than we have used here. We have added to this effort, but much more needs to be done to answer the essential question, which is: *which* technological developments caused *which* business cycles? A logical next step is to look at some sectorial and industry-of-use indicators (such as inventory, investment, and productivity), and try to organize the results in a way that aggregates to yield macroeconomic implications.[3]

1.H. *Plan of the Paper*: Section 2 presents the model. In order to get as quickly as possible to the crux of the matter, we describe the model's business cycle implications in Section 3. Section 4 then describes how the microeconomic data were used to generate the parameter values that underlie the exercise in the previous section. Section 5 presents a discussion of the results, and Section 6 concludes the paper. Technical details appear in the Appendix.

2. THE MODEL

2.A. *The Production Function*: Assume the following aggregate production function:

$$Y_t = L_t^{(1-\alpha)} \int_0^{A_t} q_{i,t}^{\alpha}\, di. \tag{1}$$

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. In this additive specification, no

[3] There is also a literature that tries to generate unpredictable fluctuations with no aggregate shocks at all—see Jovanovic (1987), and Scheinkman and Woodford (1994).

intermediate product is essential. An obvious alternative to (1) is a multiplicative specification, analyzed in Jovanovic and Rob (1990), but that is less plausible in that it says that each input is essential.

2.B. *The Diffusion of Intermediate Inputs*: Let $v(i)$ denote the vintage (i.e., date of birth) of product i. We shall make the following three assumptions about the diffusion of the q_i:

A.1. (No scale effects): Its diffusion by date t is proportional to L_t.
A.2. Its diffusion depends on its age, $t - v(i)$.
A.3. Its diffusion depends on shocks described by the stochastic process $\{\theta_{v(i),t}\}_{t=v(i)}^{\infty}$.
These shocks reflect inventions that "refine" the "basic" invention of product i.

These assumptions imply that

$$q_{i,t} = L_t F\left[t - v(i), \theta_{v(i),t}\right] \tag{2}$$

so that

$$\frac{Y_t}{L_t} \equiv y_t = \int_0^{A_t} \left[F\left(t - v(i), \theta_{v(i),t}\right)\right]^{\alpha} di \tag{3}$$

Now suppose that A_t grows exogenously at the rate λ, and normalize A_0 to equal unity. Then $A_t = e^{\lambda t}$. There is only one product per vintage; only product i arrives exactly at date $v(i)$. Changing the variable of integration from the product name i to its vintage v, where $i = e^{\lambda v}$, yields,

$$y_t = \lambda \int_{-\infty}^{t} e^{\lambda v}\left[F(t - v, \theta_{v,t})\right]^{\alpha} dv = \lambda e^{\lambda t} \int_0^{\infty} e^{-\lambda \tau}\left[F(\tau, \theta_{t-\tau,t})\right]^{\alpha} d\tau = e^{\lambda t} X_t \tag{4}$$

where

$$X_t \equiv \lambda \int_0^{\infty} e^{-\lambda \tau}\left[F(\tau, \theta_{t-\tau,t})\right]^{\alpha} d\tau.$$

Equation (4) is intuitive: There are $\lambda e^{\lambda(t-\tau)}$ products of age τ, and each contributes $F(\tau, \theta_{t-\tau})$. We then just add over all ages.

2.C. *Three Assumptions on the Shocks*: To keep the model manageable, we make the following assumptions about the product-innovation shocks:

A.4. (Independent shock sequences on different vintage products): If $v \neq v'$, $\{\theta_{v,t}\}_{t=v(i)}^{\infty}$ is independent of $\{\theta_{v',t}\}_{t=v'(i)}^{\infty}$.

A.5. (Autocorrelated shocks within a vintage): For given v, $\theta_{v,t}$ follows the law of motion $\log(\theta_{v,t+1}) = \beta \log(\theta_{v,t}) + u_{v,t,+1}$, with $u_{v,t}$ i.i.d. over v and t, with zero mean and variance σ_u^2. We assume that $|\beta| < 1$.

A.6. For each v, $\theta_{v,0}$ is drawn independently from the stationary distribution of the above process—the distribution implied by the representation $\log(\theta) = \Sigma_0^\infty \beta^j u_{v,-j}$.

3. THE NATURE OF THE BUSINESS CYCLE

We are now in a position to derive some implications for the business cycle

3.A. *Per Capita GNP is Trend-stationary*: This implication does not depend on the functional form of $F(.)$. We state it formally:

PROPOSITION 1. *If F is bounded,* $\log y_t$ *is stationary around the trend* λ.

PROOF. In view of A.5, X is a moving average of a stationary variable, and so it too must be stationary. □

The long-run growth rate of income is λ; diffusion lags therefore have only level effects in the long run.

Trend stationary processes can be highly persistent, and our model will provide an example of this. As such it will resemble a model where the trend is stochastic. A recent paper by Diebold and Senhadji (unpublished data, 1995) uses long spans of annual data (1875–1993) to argue that trend-stationarity is the appropriate formulation. But the debate about the properties of the time series of aggregates will no doubt continue.

3.B. *A Specific Form for F*: We seek a parsimonious, yet flexible form. Assume that

$$[F(\tau, \theta)]^\alpha = (1 - e^{-\rho\tau})\theta. \tag{5}$$

The parameter ρ measures the speed of diffusion—it is the same for each product.[4] The parameter θ now captures (through a slight abuse of notation) two things: the quality of the product (the original definition of θ), as well as its demand in production, as originally measured by α. This interpretation allows for the possibility that variations in sales may reflect production technology downstream rather than the intrinsic quality of the new product. Eq. (4) then implies

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

[4] We show that this parameter has surprisingly little effect on the business cycle. Allowing it to differ over products would make little difference to the model's implications.

3.C. *The Level Effects of the Speed of Diffusion*: Although ρ has no growth-rate effects, its level effects can be very large. Let $\mu \equiv E(\theta)$ be the unconditional mean of the $\{\theta\}$ process. The level effect of the diffusion rate ρ can be measured by the long-run mean of X_t:

$$E(X_t) = \frac{\rho\mu}{\rho + \lambda}. \tag{7}$$

So while the speed of diffusion does not affect the growth rate, it has a potentially huge level effect. This message should not be lost when we later conclude that the speed of diffusion has relatively little bearing on the nature of the business cycle.

3.D. *The Autocovariances, Autocorrelations, and Impulse Responses*: Appendix Section A shows that for each $k \geq 0$,

$$C_k \equiv \mathrm{Cov}[\log(X_t), \log(X_{t+k})] \approx \frac{\psi^2(\lambda + \rho)\lambda^2\beta^k e^{-\lambda k}}{2\lambda(2\lambda + \rho)}\left[1 + \frac{\lambda(1 - e^{-\rho k})}{\rho}\right] \tag{8}$$

where $\psi^2 \equiv \mathrm{Var}[\log(\theta)]$ is the unconditional variance of $\{\log(\theta)\}$. Since $\log X_t$ equals detrended $\log y_t$, this is the predicted autocovariance function for de-trended log per capita output. The amplitude of the business cycle is just the variance of $\log(X_t)$:

$$C_0 \equiv \frac{\psi^2(\lambda + \rho)\lambda}{2(2\lambda + \rho)}.$$

The autocorrelation coefficient of $\log(X_t)$ is $r_k \equiv C_k / C_0$:

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

The impulse response of invention shocks: If a high-quality basic invention is made at date t, this is represented by a high value of $\theta_{t,t}$, which will, by A.5. be correlated with $\theta_{t,t+k}$, since $\partial E(\theta_{t,t+k})/\partial\theta_{t,t} = \beta^k$. Now in equation (6), the weight on $\theta_{t,t+k}$ is $\lambda e^{-\lambda k}(1 - e^{-\rho k})$. So the response of y_{t+k} to a unit increase in $\theta_{t,t}$ is:

$$I_k = \lambda\beta^k e^{-\lambda k}(1 - e^{-\rho k}) \tag{10}$$

The function I_k is unimodal, and peaks at lag k^* given by

$$k^* = \frac{1}{\rho}\log\left(1 + \frac{\rho}{\lambda - \log\beta}\right) \tag{11}$$

where, of course, $\log(\beta) \leq 0$. We summarize the properties of C_k, r_k, I_k and k^* as follows:

PROPOSITION 2.
(*i*) C_k *is increasing in* β, *and it rotates clockwise when* λ, *or* ρ *rise.*
(*ii*) r_k *is increasing in* β, *and decreasing in* λ *and* ρ.
(*iii*) I_k *is increasing in* β *and* ρ, *and it rotates clockwise when* λ *rises,*
(*iv*) k^* *is increasing in* β, *and decreasing in* λ, *and* ρ.

PROOF. In Appendix Section B.

3.E. *Our Empirical Strategy*: We shall plot I_k and C_k (r_k is just a re-scaled version of C_k), evaluated at the estimates for λ, ρ, β, and ψ. Our strategy will be to estimate the parameters from microeconomic data, and then see how well the model explains the GNP data. In the tradition of real business-cycle analysis, we shall ask how well the predicted second moments match the second moments in the data. That is, we shall compare C_k to the autocovariance of de-trended per-capita real GNP.

We focus on the autocovariances in (8) and not the autocorrelations in (9) because the model omits noise that is in fact present in the y_t series, noise that might be autocorrelated at short lags only. Such noise, when added to the model, could affect r_k at *all* lags, raising it for low k and reducing it for large k. On the other hand, such noise affects the autocovariances only up to whatever length defines their high-frequency character. That is to say, if the omitted noise is $MA(q)$, then only the first q autocovariances are affected. Indeed this is why the model will underpredict the empirical autocorrelation for low k.

Our approach is in the calibration tradition of business cycle research, but we do depart from it in that we actually estimate some of the parameters from nontraditional microeconomic data. Our method is relatively straightforward, however, and there are more sophisticated alternatives. Our estimation does not take into account the aggregate model and data. In contrast, Diebold et al. (1995) present a more ambitious approach where, among other things, one can in principle estimate such parameters by seeking the estimates that minimize the divergence between the aggregate data and model. More specifically, their framework delivers the parameter configuration that minimizes divergence between the data spectrum and the model spectrum over frequencies of interest.

3.F. *The Parameter Estimates*: We move on to a summary of how well the model does in explaining the autocovariance function of per capita GNP. The parameter estimates are: $\rho = 0.081$, $\beta = 0.975$, $\psi^2 = 0.31$, and $\sigma_u^2 = 0.015$. How they were estimated is described in (the next) Section 4, where standard errors and sources are provided and discussed. We now present plots of I_k (in Figure 1) and C_k (in Figure 2) based on these estimates. The striking feature of Figure 1 is the length of time that it takes the impulse response to peak (3 years) and to die out.[5]

[5] It takes even longer (5 years) for the patent-citations impulse responses to peak (Jaffe and Trajtenberg, unpublished data, 1995].

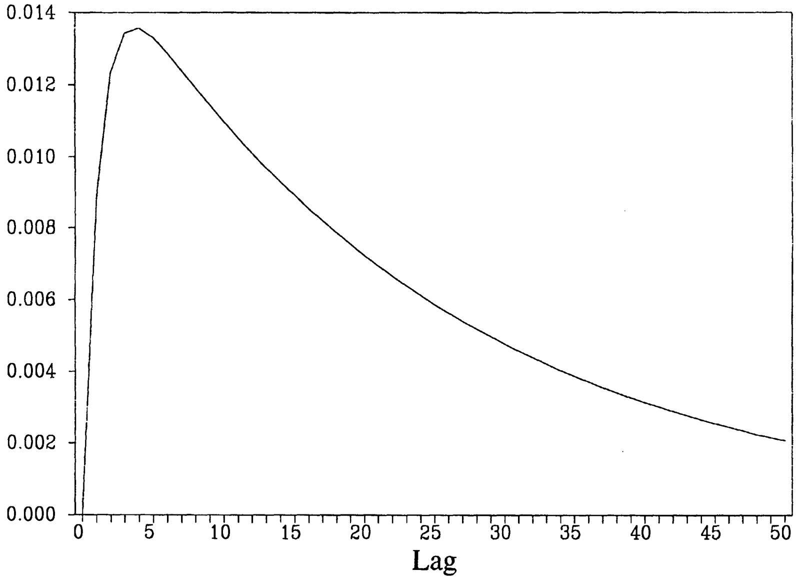

FIGURE 1

IMPULSE RESPONSE FUNCTION

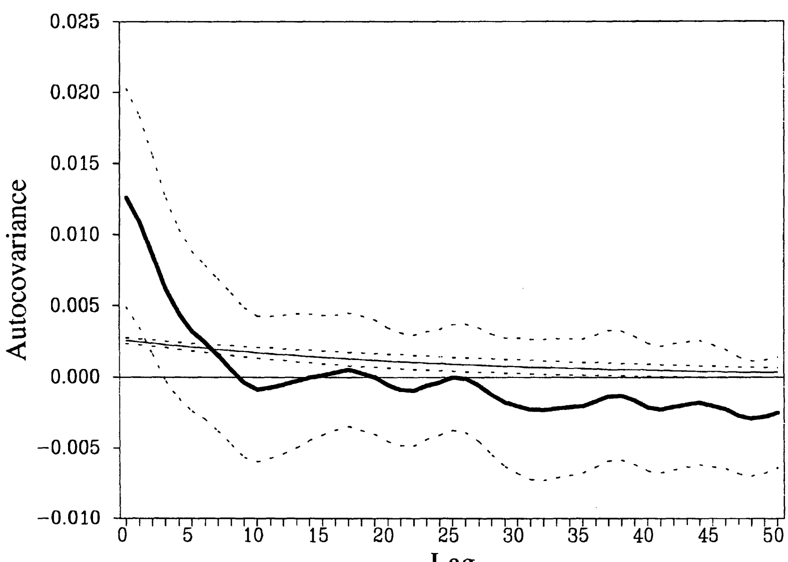

FIGURE 2

ACTUAL AND PREDICTED AUTOCOVARIANCES OF DE-TRENDED LOG y

Figure 2 depicts the actual autocovariances with the predicted ones, and the 1.96 standard-deviation bands around them (for discussion of the technicalities see Section 4).[6] For short lags, the model underpredicts the autocovariance, which makes sense because other shocks are omitted. For long lags, the model overpredicts the autocovariance, although not significantly. The reason is the estimate $\beta = 0.975$; there is extremely high persistence in $\log(\theta)$ within products, and in spite of the independence of shocks *across* products, this is enough to yield a lot of persistence.

What conclusions can we draw? We divide the issues into two sets: the first has to do with the amplitude of the business cycle, as measured by C_0, and the second has to do with the slope of the autocovariance.

3.G. *The Amplitude of the Business Cycle*: We make three points. First, shocks to product innovation explain about one quarter of the variance of GNP around trend. Given that there surely are shocks that the model leaves out, this is not bad.

Second, the amplitude of the predicted business cycles is largely *independent* of the rate of diffusion ρ:

$$\frac{\partial C_0}{\partial \rho} = \frac{\lambda}{(\lambda + \rho)(2\lambda + \rho)} C_0 = [0.024] C_0, \tag{12}$$

where the second equality obtains after substituting in the estimated values of λ and ρ. Hence the elasticity of C_0 with respect to ρ is roughly $[0.024]\rho$, which is about one fifth of one percent!

Third, we estimate that $C_0 = [0.083]\psi^2$, which means that the predicted variance of log GNP around trend is about *120 times smaller than* the variance of the log of the technology shock of each individual invention.

3.H. *The Slope of the Autocovariance Function*: We make two points. First, the model predicts a much flatter C_k than the data show. Since other shocks are probably present in the data, the underprediction of the short lag autocorrelations is to be expected. But the overprediction of C_k at longer lags is a problem for the model. The problem would be "solved" if we could somehow raise our estimate of λ

[6] The confidence interval around the actual autocovariance does not appear to increase with the lag length k, as one might expect, given that the number of observations used in the estimation of each autocovariance decreases with k. This is due to our use of a biased autocovariance estimator. This estimator divides the sum of $(N-k)$ cross-products by N, rather than $N-k$, where N is the number of observations. It is therefore biased but it usually has a lower mean squared error than the unbiased estimator (Priestley 1981, p. 323). It turns out that the exact variance of the biased estimator does not have to increase with k. Because the variance of the unbiased estimator is $(N/(N-k))^2$ times higher than the variance of the biased estimator, its variance (as expected) is more likely to be an increasing function of k.

or lower our estimate of β. Raising λ would make the absolute impact of a shock to θ fade faster relative to y, and lowering β would shorten the duration of the impulse effects of a shock (see Proposition 2.iii, and Figure 1).

Second, the flatness of C_k is *not* due to slow diffusion. Although Proposition 2.i says that a higher ρ makes C_k steeper, quantitatively the effect is small. Indeed, suppose that all products diffused *instantaneously*, so that $\rho \to \infty$, and so that

$$C_k = (\psi^2/2)\lambda e^{-\lambda k}\beta^k. \tag{13}$$

The intercept, C_0, would still be only 0.0026, which still underpredicts amplitude by a factor of four. And the model would still continue to substantially overpredict persistence at longer lags.

4. ESTIMATION

We present a more detailed discussion of our data sources and procedures, and make some remarks qualifying our parameter estimates.

4.A. *The Gort-Klepper Data*: The parameters ρ, β and ψ were estimated from the Gort-Klepper data. These data document the historical development of 46 products in terms of their sales, price, output, and numbers of producers over (part of) the life-cycle of each product. Gort and Klepper chose the products on the basis of the following three criteria:

1. To allow sufficient diversity by including consumer, industrial, and military products,
2. To include only products that were "basic" inventions,
3. To include products with adequate data on net entry.

We shall later discuss how these selection criteria may affect our results.

Table 1 lists the 21 products for which we have sales data. Only 11 or so of them seem to qualify as intermediate inputs. For example, there is no output in GNP that corresponds to the use of missiles, penicillin, DDT, blankets, and shavers. But because the number of products is small, we analyze them all. Column 1 of Table 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 product's life. The age range for which there are data appears in the second column, and the average volume of sales per product in the third. Sales, averaged over the 499 observations in the sample, were 580 million U.S. (1967) dollars.

4.B. *Procedure*: The production function in eq. (1) treats intermediate products as exchangeable inputs: One unit of product i and two units of product j can

TABLE 1
THE GORT-KLEPPER PRODUCTS

Product	Initial Year	Age Range	Sales[a]	Sample Size
1. Computers	1935	20–36	8567.6	17
2. Crystals	1936	25–36	45.6	12
3. DDT	1943	1–27	23.3	27
4. Electrocardiographs	1914	47–58	8.8	12
5. Electric Blankets	1911	35–61	61.5	27
6. Electric Shavers	1930	1–42	83	39
7. Fluorescent Lamps	1938	0–34	94.6	35
8. Freezers, Home & Farm	1929	18–43	335	26
9. Gyroscopes	1911	52–61	15.1	10
10. Lasers	1960	3–11	41.9	9
11. Missiles, Guided	1942	9–30	2150.8	22
12. Motors, Outboard	1908	42–64	138.6	23
13. Penicillin	1943	2–28	72.7	27
14. Pens, Ballpoint	1945	6–27	92.1	22
15. Records, Phonograph	1887	34–85	372.7	52
16. Streptomycin	1945	1–27	9.5	27
17. Styrene	1935	8–36	85.8	29
18. Tapes, Recording	1947	14–25	159.7	12
19. Television, Apparatus, Parts	1929	17–43	1355.1	27
20. Transistors	1948	6–24	266.5	19
21. Tubes, Cathode Ray	1922	26–50	157.5	25
Total				499

[a] Sales are deflated to 1967 dollars by the Wholesale Price Index, in millions.

produce as much final output as two units of product i and one unit of product j. As computers and ballpoint pens are not exchangeable in this sense, we bring them into common units by expressing everything in units of the 1967 "consumption good", so that for $q_{i,t}$ we shall use product i's sales at t, deflated to 1967 dollars by the Wholesale Price Index.

Gort and Klepper point out that on average there is a rapid decline in the rate at which 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), which, together with equation (2), implies that

$$(14) \qquad q^*_{i,t} \equiv \left[\frac{q_{i,t}}{L_t}\right]^{\alpha} = (1 - e^{-\rho\tau_{i,t}})\,\theta_{v(i),t}$$

where $q_{i,t}$ are sales in 1967 dollars, L is the population of the US, $\alpha = 1/3$, and $\tau_{i,t} = t - v(i)$ is the age of product i in year t. Then A.5 and (14) imply

$$(15) \quad \log q^*_{i,t} = \beta \log q^*_{i,1t-1} + \log(1 - e^{-\rho\tau_{i,t}}) - \beta \log(1 - e^{-\rho(\tau_{i,t}-1)}) + u_{v(i),t}$$

The parameters ρ, β and σ_u^2 can be consistently estimated from equation (15) by nonlinear least squares. Since the time series for each product is not too long and we assume that ρ, β and σ_u^2 are the same over products, we shall estimate them by

TABLE 2
NONLINEAR LEAST SQUARES ESTIMATES

Parameter[a, b]	Estimate
ρ	0.081 (.1683)
β	0.975 (0.106)
σ_u^2	0.015
ψ^2	0.311

[a] The number of pooled observations is 478. Asymptotic standard errors in parentheses.
[b] The regression includes a constant and 2 dummies for the pre- and post-war periods.

pooling the data. The estimates and their asymptotic standard errors are in Table 2.

4.C. *Other Estimates in the Literature*: The one point we have in common with much other empirical work is our estimate of diffusion speed. 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 $\rho = 0.081$ implies that our Δt is about 15 years. The most comprehensive study of diffusion (of 265 innovations) in the U.S. seems to be the one by Grübler (1991). He reports that the largest number of diffusion processes have Δt's of between 15 and 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 his sample does not contain information about the value of the inventions.

4.D. *How λ was Estimated*: The parameter λ is estimated from equation (4) using aggregate per-capita GNP data to measure y.[7] Taking the natural logarithm of (4) results in

$$\log y_t = a + \lambda t + \log X_t \qquad = 1, \ldots, T \tag{16}$$

where the constant term a is added to pick up the possibly nonzero mean of $\log X$. The assumed stationarity of the $\{\theta_{v,t}\}$ process implies that $\log X$ is now a zero mean stationary variable. The OLS estimator of λ in a regression of $\log y$ on a constant and on time is therefore consistent.

Further details on the estimation of λ: Since $\log X_t$ is serially autocorrelated the standard formula for the estimator of the variance of the OLS estimator is incorrect.[8] The correct covariance matrix is $(W'W)^{-1}W'VW(W'W)^{-1}$, where W is the $T \times 2$ matrix of regressors and V is the $T \times T$ autocovariance matrix of $\log X_t$. An estimate of V can be obtained since, under the assumption that the model is

[7] The data on annual GNP are from Balke and Gordon (1986) and cover the period 1869–1983. Population data are from the Historical Statistics Series and Statistical Abstracts of the U.S.A.

[8] The OLS estimator is also not efficient since the error is serially correlated and its autocovariance matrix is also a function of λ.

correct, we know the autocovariance function of log X_t. This is C_k in equation (8). We use the consistent OLS estimator of λ and the estimates of ρ, β and ψ from the Gort-Klepper data to estimate V and the correct variance of λ.[9]

Proposition 1 implies that the predicted long-run growth of y is λ. The average growth rate of per-capita income over the period is 0.01745, which is very close to the OLS estimate of 0.01668.[10] This result does not constitute a test of the model; it merely says that per-capita income is well approximated by an exponential trend.

Finally, the actual autocovariances in Figure 2 are based on de-trended log per-capita income which equals $\log y_t - \hat{a} - \hat{\lambda}t$; $\hat{a}$ and $\hat{\lambda}$ being the OLS estimates of a and λ in equation (16). Appendix Section C provides further detail on the computation of the confidence bands, and on the calculation of the estimates.

5. OVERVIEW OF THE RESULTS, THEIR LIMITATIONS, AND SOME EXTENSIONS

5.A. *Data Limitations and Their Bearing on the Estimates*: We make three points: First, the Gort-Klepper sample is selected to include only successful products. This suggests that our observations come from a truncated distribution that includes only the right tail, whereas the model deals with *all* innovations. As far as the predicted autocovariance goes, the critical parameter is ψ^2. Assuming that our estimate pertains to the left-truncated distribution of θ, whether ours is an overestimate or an underestimate of the population coefficient of variation depends on the form of the underlying distribution. If this distribution is exponential, our estimate of ψ^2 is an underestimate. The same is (somewhat surprisingly) true if the distribution is normal (Johnson and Kotz 1970, p. 81)

Second, 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 real sales, especially for computers (Gordon, 1990). This underestimate of quality change was the largest for those products for which sales were the largest, such as computers, television, and transistors. On these grounds, σ_u^2, and hence, ψ, is probably underestimated.

Third, ψ could be overestimated for the following reason: several inventions may be lumped together and labeled as just one invention. Perhaps computers should count as several inventions. Arbitrary classification errors therefore would lead to an overestimate of ψ. These are three reasons why it is hard to assign a standard error to our estimate of ψ reported in Table 2.

[9] We do not use the standard error of λ directly, only in the estimation of confidence intervals. Given the estimate of V we also computed a GLS estimator. Note, however, that this estimator is still not efficient since it does not estimate V and λ jointly. The OLS and GLS estimates of λ do not differ significantly.

[10] The average growth rate is the OLS estimate of the constant term, λ, in a first-differenced version of equation (16).

5.B. *Cyclical* λ: We concluded that the business cycle is shaped not by ρ, but by ψ and β, and secondarily by λ. This finding is at odds with Schumpeter's notion that the business cycle is a wave whose shape depends on the pattern of diffusion. How much does our conclusion depend on the assumption that λ is fixed? If innovations arrive in "bunches", then λ is random, and amplitude of the predicted GNP series would be higher. Klenow (1994) finds that advertised product introductions have an annual coefficient of variation of 0.19. What sort of adjustment would this call for? Suppose that this is a good estimate of the coefficient of variation of λ. Suppose, moreover, that the innovations to λ are uncorrelated with the innovations to quality of vintages, θ. Then with some algebra one can show that to a first approximation the formula for C_0 would be unchanged, except that in it the appropriate value of ψ^2 would be obtained by adding the number 0.19 to our own estimate of ψ^2, so that instead of 0.31, we would have $\psi^2 = 0.50$. Under this interpretation, product innovations explain around 42% of the amplitude of the business cycle. Similar adjustments could be make to incorporate the findings of Geroski and Walters (1995).

5.C. *Creative Destruction and External Effects*: We ignore the possibility that (a) some intermediate goods are substitutes, and a new product may displace an old one, (in this case a product's sales are an overestimate of the effect that the product's introduction has on GNP), or that (b) some intermediate goods are complements (the appearance of a product may raise the sales of another). These two effects work to bias our results in opposite ways. We hope that to a first-order approximation they can be ignored.

This relates to whether ideas "feed into each other"—and there is no doubt that they do. Caballero and Jaffe (1994) allow for both substitution of new technologies for old (obsolescence) and complementarity (current ideas breed future ideas). Economic historians have emphasized the occasional arrival of major inventions and GPTs, such as the steam engine, the diesel engine, the factory system, electricity, hybrid seeds, and semiconductors. These GPTs induce a correlation between other products that use them in production. A technology that uses a GPT as an input may itself be of any vintage, and hence in general the possible arrival of GPT's will induce a correlation in the sales of products of different vintages. Assumption A.4., however, allows such a correlation only among products of the same vintage.

6. CONCLUSION

This paper has sought to quantify the contribution of innovation to the business cycle. We measure technology directly and not as a residual. We use the Gort-Klepper data to measure fluctuations in the value of inventions. Although Gort-Klepper products are clearly "successes", and hence a biased sample, they are biased in terms of their importance. But it is not clear that this sample should provide a biased estimate of the two critical microeconomic parameters: ψ, which has the dimensions of a coefficient of variation, and β, which is an autocorrelation.

We conclude that product innovations can explain fluctuations at lower frequencies, but they underpredict fluctuations at higher frequencies. Since the diffusion of new products is rather persistent in character (β is high), the predicted autocovariance function is fairly flat—as in Figure 2. This finding appears to be robust. But we are less sure about the intercept of the predicted autocovariance because we are not too sure of 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 might cause us to revise the predicted amplitude, and perhaps persistence as well. Finally, it was a surprise to learn that the speed diffusion plays virtually no role in shaping the business cycle.

We definitely do find a big effect of the speed of diffusion on the *level* of output. A society in which technologies spread quickly will have a big level advantage over another in which technologies spread slowly.

We have used the identifying assumption that a technology shock is completely characterized by the new capital good, or equipment, that it induces. Sales of the capital good reflect only the direct effect of the shock. They do not capture the productivity rise of the user of the new capital good unless the inventor can extract all the rent. That is why the next step should be to look at sectoral or industry indicators systematically, rather than solely at output effects. In the case of an intermediate capital good such as the computer, one could look at what is happening in industries and sectors that use computer services. Our approach is to look only at sales of computers, but as we remarked earlier, computer sales may be underrepresented in aggregate impact relative to other technologies. Other sectoral indicators might be used—variables such as inventories, investment, and productivity. If we find big sectoral effects (in a major sector), then there must be effects on aggregate output. Such sectoral data will tell us more about the effects of particular technologies than aggregate data can, and the challenge will be to organize sectoral results in a coherent manner that can tell us something precise about the effect that technology shocks have on the movement of aggregates.

APPENDIX

A: *Calculating C_k in Equation* (8). Let $a_\tau \equiv e^{-\lambda\tau}(1-e^{-\rho\tau})$. Then

$$\begin{aligned}\operatorname{Cov}(X_{t+k}, X_t) &= \lambda^2 \operatorname{Cov}\left[\int_0^\infty a_\tau \theta_{t+k-\tau,t+k}\, d\tau, \int_0^\infty a_\tau \theta_{t-\tau,t}\, d\tau\right] \\ &= \lambda^2 \operatorname{Cov}\left[\int_0^\infty a_\tau \theta_{t+k-\tau,t+k}\, d\tau, \int_0^\infty a_\tau \beta^k \theta_{t-\tau,t-k}\, d\tau\right] \\ &= \lambda^2 \beta^k \operatorname{Cov}\left[\int_{-k}^\infty a_{s+k} \theta_{t-s,t-k}, \int_0^\infty a_\tau \theta_{t-\tau,t-k}\, d\tau\right]\end{aligned}$$

The second equality follows because $\theta_{t-\tau,t} = \beta^k \theta_{t-\tau,t-k} + \xi$, where, by A.4., ξ is independent of each other θ under the second integral after the first equality, and the third equality follows after changing variables in the first integral from τ to $s = \tau - k$, and factoring out β^k from the second integral. Now $\text{Cov}(\theta_{t-s,t-k}, \theta_{t-\tau,t-k}) = \sigma_\theta^2$ if $s = \tau$, and zero otherwise. Hence the first integral over the region $s \in [-k, 0]$ is uncorrelated with the second integral, and

$$\text{Cov}[X_t, X_{t+k}] = \lambda^2\beta^k\sigma_\theta^2 \int_0^\infty a_{s+k} a_s \, ds = \frac{\sigma_\theta^2 \lambda \rho^2 \beta^k e^{-\lambda k}}{2(\lambda + \rho)(2\lambda + \rho)} \left[1 + \frac{\lambda(1 - e^{-\rho k})}{\rho}\right]$$

Moreover, for any two variables u and v, $\text{Cov}[\log(u), \log(v)] \approx \text{Cov}(u, v)/E(u)E(v)$, which implies that $\psi^2 \approx \sigma_\theta^2/[E(\theta)]^2$. This approximation underlies the expression in equation (8).

A second way to obtain the above expression for the covariance is to first use A.4. to justify changing the limits of the integral defining X_{t+k} from $[0, \infty)$ to $[k, \infty)$, then change the variable in the first integral from τ to $s = \tau - k$.

B: *Proof of Proposition* 2. By "clockwise rotation" of, say C_k, in response to a parameter, say λ, we mean that $\partial C_k/\partial\lambda > 0$ for all k sufficiently small, and $\partial C_k/\partial\lambda < 0$ for all k sufficiently large. The derivations are straightforward, but signing some of them requires that the inequality $1 + x < e^x$ for $x \neq 0$ be invoked.

The only claim that is proved in a different way is claim (iv). The strategy here is to let $f(z) \equiv \log(z)$, so that $k^* = f(1 + \rho/B)/\rho$, where $B \equiv \rho/[\lambda - \log(\beta)]$. Then

$$\partial k^*/\partial\rho = [-f(1+B) + Bf'(1+B)]/\rho^2.$$

By Taylor's theorem with remainder,

$$f(1) = f(1+B) + f'(1+B)(-B) + [f''(x)/2]B^2,$$

for some $x \in [1, 1+B]$. Substituting for f into the previous equation, and observing that $f(1) = 0$ while $f'' < 0$ proves the claim.

C: *Computation of Confidence Intervals.* Confidence intervals of the *predicted* autocovariances are computed in the standard way using an asymptotic approximation. Let $\delta = (\lambda, \rho, \beta, \psi)$ and let "hats" denote estimated values. The variance of $C_k(\hat{\delta})$ is obtained from the variance of the linear Taylor expansion of $C_k(\hat{\delta})$ around δ. This results in an *estimated* variance given by $g(\hat{\delta})'\Omega g(\hat{\delta})$, where $g(\hat{\delta})$ is the gradient of $C_k(\delta)$ evaluated at $\hat{\delta}$, and Ω is the 4×4 estimated covariance matrix of $\hat{\delta}$.

Ω is specified as follows. Its diagonal elements are taken from the second column of Table 2 and from the assumption that $\hat{\psi}$ has zero variance which means that, on this account, the reported bands are tighter than they should be. The only nonzero off-diagonal element is the covariance between the estimated ρ and β that comes out from the estimation procedure; the other covariances are set to zero.

Under these assumptions, the estimated variance of $C_k(\hat{\delta})$ is,

$$A\,\hat{\text{var}}\big(C_k(\hat{\delta})\big) = \hat{\text{Var}}(\hat{\lambda})\big(\partial C_k(\hat{\delta})/\partial\lambda\big)^2 + \hat{\text{Var}}(\hat{\rho})\big(\partial C_k(\hat{\delta})/\partial\rho\big)^2$$
$$+ \hat{\text{Var}}(\hat{\beta})\big(\partial C_k(\hat{\delta})/\partial\beta\big)^2 + \text{Cov}(\hat{\rho}, \hat{\beta})\big(\partial C_k(\hat{\delta})/\partial\beta\big)\big(\partial C_k(\hat{\delta})/\partial\rho\big),$$

the partial derivatives being evaluated at the point estimates.

The confidence interval for the predicted autocovariance is based on the use of the normal distribution as an approximation to the true distribution of $C_k(\hat{\delta})$. The 95% confidence interval is therefore,

$$C_k(\hat{\delta}) \pm 1.96\left[A\,\mathrm{v\hat{a}r}\left(C_k(\hat{\delta})\right)\right]^{1/2}. \tag{C.1}$$

The "actual" autocovariances of de-trended $\log y_t$ are estimated in the standard way (e.g., equation (5.3.13) in Priestley, 1981). We also build confidence intervals around them, having the same form as those appearing in equation (C.1), but using the exact variance of these autocovariances under the assumption of normally distributed de-trended $\log y_t$ (see eq. (5.3.23) in Priestley, 1981).

Further Notes on the Estimation Procedure: The objective is to find values of ρ, β and γ that minimize

$$s(\rho, \beta, \gamma) = \sum_{i=1}^{21} \sum_{\tau} \Big[\log q^*_{i\tau_i} - \beta \log q^*_{i\tau_i - 1} - \log(1 - e^{-\rho\tau_i}) + \beta \log(1 - e^{-\rho\tau_i - 1}) - z_{it}\gamma\Big]^2$$

where z_{it} are product and period specific dummies. The estimation procedure is as follows: Choose values of β and ρ over the intervals $[-1,1]$ and $[0.01, 1.5]$, respectively, evaluate the regressors involving ρ and β and compute the sum of squared errors associated with the OLS estimator of γ, $S(\rho, \beta, \gamma(\rho, \beta))$. Select the values of ρ and β that minimize the latter concentrated sum of squared errors.

Asymptotic standard errors are estimated and computed in the standard way: $A\mathrm{cov} = \sigma_u^2(g'g)^{-1}$ where g is the gradient of the right-hand side of (15) with $z\gamma$ added to the end, $g = (g_\beta, g_\rho, z)$,

$$g_\beta = \log q^*_{i\tau_i - 1} - \log(1 - e^{-\rho\tau_i - 1}),$$

$$g_\rho = \frac{\tau e^{-\rho\tau}}{1 - e^{-\rho\tau}} - \frac{\beta(\tau - 1)e^{-\rho(\tau-1)}}{1 - e^{-\rho(\tau-1)}}$$

and where σ_u^2 is estimated by the minimized value of $S(\rho, \beta, \gamma(\rho, \beta))$, divided by the number of observations. The gradient is evaluated at the point estimates of ρ and β. The covariance of (β, ρ) is the 2×2 top-left matrix of $A\mathrm{cov}$.

REFERENCES

ANDOLFATTO, D. AND G. MACDONALD, "Endogenous Technological Change, Growth, and Aggregate Fluctuations," Unpublished paper, University of Rochester, February 1994.

BALKE, N. AND R. GORDON, "Appendix," in R. Gordon, *The American Business Cycle* (Chicago: University of Chicago Press, 1986).

BOUCEKKINE, R., M. GERMAIN AND O. LICANDRO, "Creative Destruction and Business Cycles," Working paper No. 95-16, Universidad Carlos III, May 1995.

BRESNAHAN, T. AND M. TRAJTENBERG, "General Purpose Technologies: Engines of Growth?" *Journal of Econometrics* 65 (1995): 83–108.

CABALLERO, R. AND M. HAMMOUR, "The Cleansing Effect of Recessions," *American Economic Review* 84 (1994): 1350–1364.

—— AND A. JAFFE, "How High Are the Giants' Shoulders: An Empirical Assessment of Knowledge Spillovers and Creative Destruction in a Model of Economic Growth," in O. Blanchard and S. Fischer, eds., *NBER Macroeconomics Annual*, (Cambridge, MA: MIT Press 1993, pp. 15–73).

CHARI, V. V. AND H. HOPENHAYN, "Vintage Human Capital, Growth, and the Diffusion of New Technology," *Journal of Political Economy* 99 (1991): 1142–65.

DIEBOLD, F., L. OHANIAN AND J. BERKOWITZ, "Dynamic Equilibrium Economies: A Framework for Comparing Models and Data", NBER Technical Working paper No. 174, February 1995.

—— AND A. SENHADJI, "Deterministic vs. Stochastic Trend in U.S. GNP, Yet Again," Unpublished paper, University of Pennsylvania, November 1995.

GEROSKI, P. AND C. WALTERS, "Innovative Activity over the Business Cycle," *Economic Journal* 105 (1995): 916–928.

GORDON, R., *The Measurement of Durable Goods Prices*, NBER, Chicago: University of Chicago Press, 1990.

GORT, M. AND S. KLEPPER, "Time Paths in the Diffusion of Product Innovations," *Economic Journal*, 92 (1982): 630–653.

GREENWOOD, J., Z. HERCOWITZ AND P. KRUSELL, "Macroeconomic Implications of Investment-Specific Technological Change," Working Paper No. 6-94, The Sackler Institute of Economic Studies, Tel Aviv University, March 1994.

GRILICHES, Z. "Hybrid Corn: An Exploration in the Economics of Technological Change," *Econometrica* 25 (1957): 501–22.

GRÜBLER, A., "Diffusion: Long Term Patterns and Discontinuities," *Technological Forecasting and Social Change* 39 (1991): 159–80.

JAFFE, A. AND M. TRAJTENBERG, "Flows of Knowledge from Universities and Federal Labs," Unpublished paper, Brandeis University, 1995.

JOHANSEN, L. "Substitution versus Fixed Production Coefficients in the Theory of Economic Growth," *Econometrica* 27 (1959): 157–176.

JOHNSON, N. AND S. KOTZ *Distributions in Statistics* (New York: Wiley 1970).

JOVANOVIC, B. "Micro Shocks and Aggregate Risk" *Quarterly Journal of Economics* (1987): 395–409.

—— AND S. LACH, "Entry, Exit, and Diffusion with Learning by Doing," *American Economic Review* 79, no. 4 (1989): 690–99.

—— AND R. ROB, "Long Waves and Short Waves: Growth through Intensive and Extensive Search," *Econometrica* 58, no. 6 (1990): 1391–1409.

—— AND G. MACDONALD, "Competitive Diffusion," *Journal of Political Economy* 102, no. 1 (1994): 24–52.

KLEINKNECHT, A. *Innovation Patterns in Crisis and Prosperity: Schumpeter's Long Cycle Reconsidered* (London: MacMillan, 1987).

KLENOW, P., "New Product Introductions," Unpublished paper, Graduate School of Business, University of Chicago, January 1994.

LIPPI, M. AND L. REICHLIN, "Diffusion of Technical Change and the Identification of the Trend Component in Real GNP," *Review of Economic Studies*, 61 (1994): 19–30.

PRESCOTT, E., "Theory Ahead of Measurement in Business Cycle Research," *Carnegie Rochester Conference on Public Policy* 25 (1986): 11–44.

PRIESTLEY, M. *Spectral Analysis and Time Series* (New York: Academic Press 1981).

ROMER, P. "Crazy Explanations for the Productivity Slowdown," *NBER Macroeconomics Annual* (1987).

SALTER, W. E. G., *Productivity and Technical Change* (New York: Cambridge University Press, 1960).

SCHEINKMAN, J. AND M. WOODFORD "Self-Organized Criticality and Economic Fluctuations," *American Economic Review* (*Papers and Proceedings*) 84 (1994): 417–421.

SOLOW, R. "Investment and Technical Progress," in K. Arrow, S. Karlin and P. Suppes, eds., *Mathematical Methods in the Social Sciences* (Stanford: Stanford University Press, 1959).