Some New Economy Lessons for Macroeconomists

Karl Whelan*

Division of Research and Statistics, Federal Reserve Board

1 Introduction

This may seem like a strange time to write about the implications of the “New Economy” for macroeconomists. At the time of writing (May 2001) it has become clear that many of the more extreme notions associated with the term are turning out to be incorrect. In the space of a year, the NASDAQ stock index – perhaps the most potent symbol of the New Economy – has tumbled from the giddy heights of 5000, and is currently hanging around 2000. Many high-profile “dot-com” firms, some of which were recently touted as the future giants of the economy, are now thought to be based on flawed business models. And the commonly-heard argument that the technological advances associated with investments in information technologies signalled the end of business cycles also looks pretty far fetched right now, given that the U.S. economy has slowed noticeably over the past year, leading to widespread concerns about the possibility of recession.

Most macroeconomists are probably not too surprised by the recent turn of events. Stock valuations for technology companies in early 2000 had far surpassed what most consider reasonable yardsticks, and the business cycle is a phenomenon with a long history that always seemed unlikely to be eradicated by the Internet.¹ So, recent events may seem to confirm skepticism over whether there really is anything new about the New Economy.

* Mail Stop 80, 20th and C Streets NW, Washington DC 20551. Email : kwhelan@frb.gov. The views expressed are my own and do not necessarily reflect the views of the Board of Governors or the staff of the Federal Reserve System.

¹ See Campbell and Shiller (2001) for a discussion of valuations that captures the typical viewpoint. As always, though, one can find opposing arguments. See, for instance, Glassman and Hassett (1999).

While skepticism about the idea that the traditional laws of economics have somehow just been overturned is, of course, always justified, in this paper I want to emphasize that there is still much to be learned from analyzing the U.S. experience with the New Economy, by which I mean analyzing the role that information technologies played in the U.S. expansion of the 1990s. While the lessons may not of the “revolutionary” variety stressed by many New Economy advocates, they could still prove fundamental in changing the way macroeconomists think about a number of important issues.

The paper begins with a brief review of the evidence on U.S. investment in high-tech equipment and labor productivity in the 1990s. The rest of the paper focuses on three areas.

First, I discuss how efforts to correctly capture the role of information technologies have raised a number of important measurement issues, and how these issues led to a change in the construction of aggregate real series in the U.S. national accounts, such as real GDP. These measurement issues are worth highlighting because they have significant implications for the interpretation of U.S. macroeconomic data, and, as of yet, they are not well understood by many mainstream macroeconomists. Second, I argue that the behavior of the U.S. economy in the 1990s provided an important confirmation for traditional neoclassical theories of business investment and productivity. Third, I suggest that the recent experience has important implications for what type of theoretical and empirical models of economic growth are likely to prove helpful in the future.

2 The U.S. Economy in the 1990s

Figure 1 shows that, behind all the hype, there has been some substance to the idea that the U.S. economy has experienced a technology-related pickup in productivity growth. The top panel shows the growth rate of real business investment in high-tech equipment, as measured by the National Income and Product Accounts (NIPA) series on business investment in “information-processing equipment”. This series mainly consists of outlays on computer hardware, computer software, and communications equipment. It shows that real high-tech business investment accelerated throughout the economic expansion of the 1990s.

The middle panel shows the rate of change of prices for high-tech equipment. These prices fell at an accelerating rate for most of the 1990s, the result of continuing technological improvements in the production of these goods. The year 2000, however, proved to be an exception to this pattern, as price declines fell to a rate of only 2 percent, more in line with the late 1980s than with the expansion of the 1990s. It is too early to know if this change signals a reversal of the pattern of rapid technological improvements in the

Figure 1 : U.S. High-Tech Investment and Productivity

high-tech industry. However, anecdotal reports of chip shortages and long lead times suggested that, for much of 2000, exceptionally strong demand for high-tech equipment may have been restraining the traditional pattern of supply-driven price declines.

The bottom panel displays productivity growth for the U.S. private business sector. It shows a steady improvement in productivity growth throughout the second half of the 1990s, an unusual event given the ad-

vanced stage of U.S. business cycle at that point. As I will discuss further below, a number of studies have shown that a relatively large fraction of the improvement in U.S. productivity growth over this period can be assigned to the efficiency gains that resulted from investments in high-tech equipment.

3 Some Measurement Issues

We have seen how the U.S. price indexes for high-tech equipment declined substantially in the 1990s. Most of this decline was due to falling prices for computing equipment; the price indexes for software and communications equipment have not fallen nearly as fast as the index for computer hardware. In this section, I first discuss the approach to measurement of computer prices in the U.S. national accounts, which since 1985 has been based on the hedonic price index method.

I then describe how the introduction of hedonic indexes for computing equipment has had a profound effect on the statistical properties of measures of U.S. real GDP, and how this led to a change in the measurement of all aggregate real variables in the U.S. national accounts. Finally, there is a discussion of how calculations that do not account for the methodology used to construct real aggregate series can substantially overstate the role that information technologies play in the U.S. economy.

3.1 Hedonic Price Indexes for Computers

A good starting point when thinking about the measurement of computer prices is to acknowledge the innate complexity of computing equipment when compared with the apples, oranges, and widgets of textbook economic theory. A quick glance through Dell's online catalog reveals that one can buy "a computer" for \$1000, or for \$2000, or for \$3000. And presumably, those who choose to spend \$3000 do so for a good reason: Their \$3000 computer is quicker, has better speakers, a screen with more pixels, and so on. The point to emphasize here is that "a computer" is not a well-defined economic commodity, but rather something harder to measure – call it *computing power* – is what firms and businesses actually care about.

While the idea that we can measure the price and quantity of computing power may seem somewhat abstract, the methodology behind such measures pre-dated the computer revolution. Zvi Griliches's famous 1961 paper "Hedonic Price Indexes for Automobiles: An Econometric Analysis of Quality Change" showed that, by using dummy variables to estimate the market value of measurable features of motor vehicles, one could develop a price index that acknowledged the value to consumers of all the various factors which differentiate a Mercedes from a Hyundai. The path-breaking

studies of Gregory Chow (1967), and later Roseanne Cole et al (1986), showed that this hedonic methodology could be feasibly applied to measure the price of computing power. These studies revealed remarkable rates of price decline. This was, of course, a reflection of the rapid pace of technological innovation in the computing industry, a phenomenon well-summarized by Intel founder, George Moore's famous "law" that the speed-per-dollar of microprocessors could double every 18 months.

Influenced by these academic studies, the U.S. Commerce Department's Bureau of Economic Analysis (BEA), which produces the NIPAs, decided in 1985 – after a somewhat acrimonious public debate – to adopt the hedonic computer price method. It is important to note, though, that this decision in no way affected the methodology for the measurement of *nominal* spending on computing equipment. Rather, it affected the *quantity* series used in the construction of real GDP. The U.S. series on "real computer spending" now refers, not to "the quantity of computers" but rather to the quantity of computing power that the nominal dollar outlays were able to acquire.

I will now turn to some of the measurement issues raised by the use of these hedonic or quality-adjusted computer deflators.

3.2 Measurement of Real GDP

When asked why they focus on "real GDP" as opposed to the nominal series, most macroeconomists will reply that real GDP is more useful because it "controls for the effects of inflation". While it is well known that there are many alternative approaches to constructing price indexes, and thus many different possible measures of real GDP, it is often assumed that each of these methods are roughly equivalent. And when the prices of all products in the economy change at roughly the same rate, then this assumption is a reasonable one. However, once some categories have prices that change at very different rates from the rest, as has occurred in the U.S. since the introduction of hedonic prices for computers, then it turns out that different approaches to constructing real GDP can result in series with very different statistical properties. The U.S. experience with this has been instructive, and is worth some discussion.

The simplest measure of real GDP is the fixed-weight or Laspeyres measure. Until 1996, U.S. real GDP was constructed according to this method. The fixed-weight approach starts with a set of prices from a specific base year, and uses these prices to weight the quantities of each category. The resulting series has the interpretation of "the value of period t 's output had all prices remained at their year- b level". If we also express each of the component quantity or "real" series in terms of what their cost would have been in year b , then we can express total real GDP as the arithmetic sum of the component real series.

While the fixed-weight methodology has the advantage of simplicity and ease of interpretation, it also has a number of undesirable features. Most importantly, the growth rate of a fixed-weight measure of real GDP depends on the choice of base year. Some simple calculations show that, for the U.S. data, this base-year dependence is an important phenomenon. Take 1998 as an example: The growth rate of fixed-weight real GDP for the U.S. in this year was 4.5 percent if we use 1995 as the base year; using 1990 prices it was 6.5 percent; using 1980 prices it was 18.8 percent; and using 1970 prices, it was a stunning 37.4 percent!

The reason we get higher growth rates when using earlier base years is the well-known problem of “substitution bias” associated with fixed-weight indexes. The categories with declining relative prices – most importantly in this case, computers – tend to have faster growth in quantities. The further back the base year, the larger is the weight on these fast-growing categories, and so the faster is the growth rate of aggregate real output. It is for this reason that the introduction of hedonic price indexes for computing equipment had important implications for the measurement of real GDP.

Because of the problem of base-year dependence with fixed-weight measures, the BEA abandoned this approach in 1996. Instead, it now employs a so-called chain index method to construct all real aggregates in the U.S. NIPAs, including real GDP. Instead of using a fixed set of price weights, chain indexes continually update the relative prices used to calculate the growth rate of the aggregate. Specifically, the growth rate of all real aggregates in the U.S. national accounts are now calculated using the so-called “ideal” chain index popularized by Irving Fisher (1922).

The Fisher chain index method calculates the gross growth rate (in other words the ratio of time t 's value to time $t - 1$'s) of the real aggregate at time t as a geometric average of the gross growth rates of two separate fixed-weight indexes, one a Paasche index (using period t prices as weights) and the other a Laspeyres index (using period $t - 1$ prices as weights.) Algebraically, the formula is

$$Q(t) = Q(t-1) \sqrt{\frac{\sum_{i=1}^n P_i(t)Q_i(t)}{\sum_{i=1}^n P_i(t)Q_i(t-1)} \times \frac{\sum_{i=1}^n P_i(t-1)Q_i(t)}{\sum_{i=1}^n P_i(t-1)Q_i(t-1)}} \quad (1)$$

where Q_i 's are the quantities of the individual categories, and the P_i 's are the prices.

The chain index approach has a number of important advantages over the fixed-weight method. Because the growth rate of the chain aggregate at time t depends only on the prices and quantities prevailing at times t and $t - 1$, there is no problem with substitution bias: The growth rate of this measure of real GDP does not depend on some arbitrary base year. In fact, the “base year” for chain-aggregates is simply the year chosen to equate the

real and nominal series, with the level of the series obtained by “chaining” forward and backward from there using the index.

The elimination of the base-year-dependence problem has been particularly important in recent years. The 1990s saw a combination of rapidly declining computer prices and large increases in nominal spending on computers. These developments would have made fixed-weight measures of GDP growth particularly subject to substitution bias. Prior to the adoption of the chain aggregation procedure, BEA’s practice had been to move the base year forward every five years; as our example comparing 1990-based and 1995-based fixed-weight measures showed, such a procedure would have resulted in predictable revisions to published real GDP growth of over two percentage points. By preventing the need for these large revisions, the move away from a fixed-weight approach avoided a problem that would have greatly complicated the interpretation of the recent macroeconomic performance of the U.S. economy.

Clearly, then, the chain aggregation approach greatly alleviates the interpretational problems associated with the fixed-weight measures of real output *growth*. Nevertheless, few improvements come without some cost, and the principal problem with chain aggregation is that it makes the interpretation of the *level* of real output more complex. BEA’s procedure has been to set real chain aggregates equal to their nominal counterparts in the same base year, b , used to define the published real series for the components series (the $Q_i(t)$ s). The published levels of real aggregates are then described as being in terms of “chained year- b dollars”. These series must be interpreted very carefully.

The level of chain-aggregated real GDP is the cumulation of period-by-period growth rates, where the growth rates are determined by continuously updated price weights. So, the “chained year- b dollar” terminology reflects only the year chosen to equate real and nominal output. Importantly, this measure of the level of real output *cannot* be interpreted as the cost of output had all prices remained fixed at their year- b levels. So, by definition, “chained year- b dollar” real GDP does not equal the simple sum of the real year- b dollar series of its individual components.

The non-additivity of chain aggregates may seem a little mysterious to those used to the fixed-weight approach, particularly since the theoretical models that frame most macroeconomists’ thinking usually feature aggregate “resource constraints” in which real output is simply defined as the sum of real consumption and real investment. However, the pattern of the non-additivity is actually quite simple and intuitive.

Note from equation (1) that the growth rate of a chain aggregate will be the same as that of a fixed-weight aggregate if relative prices do not change. But if relative prices are changing, then those products that decline in relative price—such as computers—will have a smaller impact on chained GDP growth after the base year, and a larger impact prior to the base year, than they would have in a fixed-weight calculation. Because

quantities of these products tend to grow fastest this means that, in general, chain aggregates grow slower than their fixed-weight counterparts after the base year, and faster prior to the base year. Both methods equate real and nominal output in the base year, so the difference between the levels of chain-weight and fixed-weight GDP follows an inverse-U shape, equalling zero in the base year and becoming more negative as we move away from the base year in both directions.

3.3 The Role of Information Technology

One issue that the introduction of the chain index method has complicated is the role that information technology plays in the determination of real GDP. Because many economists are unfamiliar with the chain aggregation methodology, it has become common to see calculations that are mistakenly based on the assumption of additivity. And because the non-additivity applies most to categories with large relative price changes, such as computing equipment, such mistakes are likely to be particularly misleading for these categories.²

For example, one might imagine that the effect of computer output on U.S. real GDP growth could be calculated by subtracting real computer output from real GDP, and then comparing the growth rate of this series with the published growth rate for real GDP. However, this calculation would only produce a valid estimate of non-computer real GDP if the published aggregate series had been constructed according to an additive, fixed-weight formula. And the result of this type of calculation can be to grossly overstate the effect of computer output on real GDP growth.

Currently, 1996 is the year used to equate real and nominal aggregate series. Disaggregated real series, such as computer output, are also expressed in terms of what their cost would have been in 1996. Aggregate real GDP growth is calculated by weighting real series according to their current and previous-quarter prices. Because the current prices for computers are significantly lower than their 1996 level, a simple subtraction of the 1996-based computer output series from the published series for real GDP will substantially overstate the direct effect of computer output in recent years. For example, for 2000, directly subtracting real investment and real consumption of computers from real GDP, and then calculating the growth rate of the resulting series, we get 3.8 percent, compared with 5.0 percent for total real GDP. However, the correctly-calculated series for non-computer real GDP grew 4.6 percent. Thus, the direct contribution of the computer sector to output growth in 2000 was 0.4 percent, as opposed to the 1.2 percent suggested by the incorrect calculation based on the assumption of additivity.

² Whelan (2000a) contains a more extensive discussion of the issues covered in this section.

Another series that can be misleading is the ratio of real computer output to real GDP. This series can usefully illustrate the fact that real purchases of computing power have grown much faster than real outlays on other goods and services, but it is unfortunately common to also see this ratio used to illustrate the “share in real GDP” of computer output, or the increasing effect of the high-tech sector on aggregate real output. The problem with the “share in real GDP” argument is that these are not shares at all, because the sum of the ratios across all categories will not equal one. And if the purpose of this type of calculation is to show the increasingly important role the numerator (in this case, computer output) plays in the determination of the denominator (chain-aggregated real GDP), then it can also give a misleading impression, because it fails to capture the changes over time in the weight that the numerator receives in the calculation of the denominator.

That real shares are an ill-defined concept with chain-weighted data may be a little frustrating for those used to performing such calculations. However, in many cases, there is an easy solution, which is to use nominal series. While inflation has an adverse effect on the use of nominal series for certain tasks, that doesn't mean they can't ever be used. In fact, if the question is about resource allocation, then nominal ratios can be interpreted as shares, and usually give an intuitive answer. For instance, suppose we want to know what proportion of output is being allocated towards capital investment. The ratio of nominal investment to nominal GDP gives a much cleaner answer than the corresponding real ratio: The nominal ratio tells us simply what fraction of each dollar spent is allocated to purchasing investment goods.

Nominal shares can also help correct some misleading impressions that real ratios may give about the changing role of information technology in the economy. While the ratio of real 1996-dollar business investment in high-tech equipment to real GDP goes from 0.003 in 1970 to 0.073 in 2000, the corresponding nominal ratio only changes from 0.016 to 0.053 over the same period. This shows that, in terms of actual dollars spent, the increase in the role of information technology has been more modest than one might think. Information technologies may have been extremely expensive in 1970, with large nominal expenditures buying small amounts of computing power; nevertheless, many firms were aware of the use of these technologies and were willing to allocate significant fractions of their capital spending budgets to them.

In highlighting the potential pitfalls when performing calculations with the U.S. data on the high-tech sector, my point is not to suggest that this sector has been unimportant in the U.S. economy's recent performance. In fact, the growth accounting studies that I discuss in the next section have all concluded just the opposite. Rather, the point is that it is easy to misunderstand the role of the high-tech sector if one uses the U.S. data in a fashion that is inconsistent with the way they have been constructed.

4 Investment and Productivity

The introduction of hedonic price indexes for computers in the U.S. national accounts, and the behavior of the economy in the 1990s, have together helped to underscore the usefulness of two empirical methods principally associated with Dale Jorgenson : the use of the user cost of capital as a tool for explaining investment behavior, and neoclassical growth accounting as a tool for understanding the determination of aggregate productivity.

4.1 The User Cost of Capital

As developed by Jorgenson (1963), the user cost of capital is a formula describing the required marginal productivity of a capital good as a function of its purchase price, p , the rate of interest, r , and the rate at which it depreciates, δ :

$$\frac{\partial F}{\partial K} = p \left(r + \delta - \frac{\Delta p}{p} \right)$$

Under stylized neoclassical conditions, this formula should, for example, summarize the effects that interest rates have on capital investment : As the rate of interest rises, the required marginal productivity of capital must also rise. Consequently, investment projects that do not meet the new higher hurdle rate do not get done, and the capital stock is reduced. More elaborated versions of the user cost formula have often been used to examine the effects of tax policy on investment.

Despite its common use in theoretical calculations, the user cost of capital has been widely believed to have one major problem : It didn't seem to help to explain investment. In a comprehensive survey of empirical studies of investment, Robert Chirinko (1993) concluded that "on balance, the response of investment to prices tends to be quite small and unimportant relative to quantity variables". Some researchers have questioned this conclusion, arguing that traditional econometric estimates of the effects of interest rates and tax policy on investment suffered from simultaneity bias : For example, positive shocks to investment may tend to produce higher interest rates, through a combination of the equilibrium response of the bond market and monetary policy reaction. In a well-known contribution, Jason Cummins, Kevin Hassett, and Glenn Hubbard (1994) argued that once the focus was restricted to "natural experiments" such as major tax reforms, then one could detect a large effect on investment of changes in the cost of capital. However, the absence of a systematic time series relationship between investment and the cost of capital has remained a constant question mark over the traditional, neoclassical theory of investment.

Recent research has shown that the behavior of investment in computers may provide an important time series example of the first-order importance of changes in the cost of capital. Unlike other types of capital, for

which small (and often temporary) changes in the rate of interest or the tax code provide the dominant source of variation, the user cost for computers is dominated by the price term, p , which has shown a steady and substantial decline over time. And these declines appear to have an important effect. Tevlin and Whelan (2001) have shown that there is a strong statistical relationship between real investment in computing equipment and the user cost of capital for computers.³ They show that two-equation regression models, which allow for a separate estimated effect of the cost of capital for computers, far outperformed standard aggregate models in explaining the boom in U.S. equipment investment in the 1990s.

The evidence of a strong relationship between computer investment and the cost of capital shows that—at least when the changes in the cost of capital are big and persistent—then we can see a significant reaction in firms' investment behavior. However, this still leaves unanswered the question of why traditional aggregate time series regressions have failed to detect a significant effect. One possibility is that, because the cost of capital for computers is dominated by the exogenous technological improvements associated with Moore's Law, it is free from the endogeneity problems that plague other measures. This suggests that Tevlin and Whelan's results may be a time series example of the "natural experiment" approach of Cummins, Hassett, and Hubbard.

Another possible explanation is that firms only respond significantly to variations in the cost of capital that they perceive as being permanent. This suggests that the key feature of computer prices that distinguishes them from other elements of the cost of capital is not their econometric exogeneity, but rather their persistence. Figuring out the relative merits of these two explanations would appear to be an important future research goal.

4.2 Growth Accounting

The user cost of capital has also been a useful tool in helping to understand the role that information technologies have played in determining aggregate productivity. To see why, note that to understand how investments in capital add to output, we need an estimate of the marginal productivity of these investments. Since Jorgenson and Griliches (1967), empirical growth accounting calculations have used Jorgenson's formula for the user cost of capital as a proxy for the marginal productivity of various types of capi-

³ One might be concerned that this relationship is merely a statistical mirage caused by measurement error for the hedonic computer price indexes. After all, a one percent mis-measured price decline will automatically produce a one percent increase in measured real investment, by virtue of the fact that real investment is measured as nominal investment divided by the price deflator. As this example suggests, the effect of such measurement error is to bias the estimated price elasticity of demand for computer capital towards minus one. However, Tevlin and Whelan found the estimated elasticity of demand to be greater than one in magnitude, so if there is measurement error, then it is likely that their results understated the true elasticity of demand.

tal. This neoclassical approach to growth accounting has provided a useful way to understand the effect of information technologies on aggregate U.S. productivity.

Perhaps surprisingly for those used to hearing about the New Economy, it is not that long since information technologies were seen as having had a disappointing productivity payoff. In the late 1980s, Robert Solow quipped that one could see computers everywhere except in the productivity statistics, thereby provoking a whole host of explanations for this so-called “Solow Paradox”. Some figured that firms were simply being too optimistic about the productivity benefits of IT; others such as Paul David (1990) used historical examples about diffusion of previous technologies as evidence that the benefits from high-tech investments would take time to show up, as firms and workers learned how to use the new technologies. To my mind, though, the most convincing answer came from my colleagues Stephen Oliner and Daniel Sichel (1994).

Oliner and Sichel used the neoclassical growth accounting methodology to show that one should not have expected investments in computers to have had a big payoff for aggregate productivity in the 1980s and early 1990s. While computers tend to depreciate rapidly, implying that a dollar spent on a computer needed to have a bigger productivity payoff today than a dollar spent on other types of capital, (recall the user cost formula from equation 2), Oliner and Sichel showed that the stock of computers was not very large when compared with the aggregate capital stock. So, despite qualitative impressions, in this *quantitative* sense, computers really were not “everywhere” at all in the early 1990s. Oliner and Sichel’s calculations implied that, at that time, investments in computers were probably boosting aggregate productivity growth by only about 0.2 percentage points per year.

The high-tech investment boom of the 1990s provided an important check on the neoclassical growth accounting methodology. As computer prices plunged, and investment accelerated, we might have expected that high-tech capital would become a more important part of total capital input, and that U.S. productivity growth would pick up. And this is exactly what happened. In updated research using the same methodology, Oliner and Sichel (2000) found that investments in high-tech equipment had played a crucial role in the improved productivity performance of the second half of the 1990s, adding about one percentage point per year to the growth rate of labor productivity over the period 1996-99. Jorgenson and Stiroh (2000) and Whelan (2000b) have presented alternative calculations that also suggested a key role for high-tech investments.

There are good reasons why the neoclassical growth accounting methodology has proved a successful empirical framework for understanding the effects of investments in information technology. Ultimately, the pattern of diffusion of the new technologies associated with the New Economy can be considered an excellent example of the logic underlying this approach. For example, the idea behind the Internet – networking computers toge-

ther, and using telephone lines to send text and images from one computer to another – had been around long before the 1990s. However, prior to the recent period, the cost of computing power and other communications technologies made the use of such technologies uneconomical for most: The marginal productivity of the Internet was simply not high enough for it to be worth the cost. As computing power and communication technologies became progressively cheaper in the 1990s, firms were able to invest in these new technologies and provide Internet services at prices that generated significant demand. And while the logic of the neoclassical model suggests that the high-tech investments of the late 1990s had a lower *marginal* productivity than previous high-tech outlays, there should be no confusion that these investments likely did add to the absolute level of productivity.

5 Implications for Theoretical Macroeconomic Models

An important theme in our discussion about information technologies has been that technological progress in the production of these goods has been faster than in the rest of the economy, and as a result, relative prices for these goods have fallen. While I have focused exclusively thus far on the high-tech sector, this pattern of faster productivity growth and falling relative prices has also applied to industries producing other durable goods, such as motor vehicles, consumer electronics, and non-high-tech capital equipment. In fact, according to U.S. NIPA data, while their share in nominal GDP has remained fairly stable, relative prices for durable goods have fallen steadily since the early 1960s. Thus, the growth rate of real output for the durable goods sector has consistently outpaced the rest of the economy.

This pattern of faster productivity growth in the production of durable goods is clearly incompatible with the standard one-sector, Solow-Ramsey model of economic growth. And it turns out that some of the standard “stylized facts” that are commonly cited as evidence for the one-sector growth model have not been holding up very well of late. For example, consider the model’s prediction that the economy should exhibit “balanced growth” such that the real series for consumption, investment, and output all tend to grow at the same rate in the long run.

While studies using data through the late 1980s, such as King, Plosser, Stock, and Watson (1991) found evidence in favor of the balanced growth hypothesis, in a recent paper, (Whelan, 2001), I document that once we extend the sample to incorporate the U.S. investment boom period of the 1990s, then one can strongly reject the hypothesis that real consumption and real investment have a common long-run statistical trend. The reason for the failure of the balanced growth prediction is the one-sector model’s inability to distinguish the behavior of the durable goods sector from that of

the rest of the economy. Most investment spending is on durable goods while most consumption spending is on nondurables and services. As a result, real investment has tended to grow faster than real consumption, a pattern that has been particularly evident since 1991.

An important implications of these facts is that the Solow-Ramsey model of economic growth, which remains the standard textbook model for long-run macroeconomic analysis (see, for instance, Barro and Sala-i-Martin, 1995) is actually a poor model of the U.S. economy. In Whelan (2001), I show how a simple two-sector approach, in which technological progress proceeds at different exogenous rates in the durable goods sector and the rest of the economy, can be used to model the long-run properties of aggregate U.S. data. Because the growth rate of a Fisher chain-aggregate is well approximated as a weighted average of the growth rates of its components, where the weights are the component's share in the corresponding nominal aggregate, I show that each of the major real aggregates (consumption, investment, capital stock, and output) can be expected to grow at different rates in the long run, depending on their nominal share for durable goods.

This type of model, with exogenous rates of technological progress in each sector, can be used to provide a better description of long-run U.S. data, and helps to explain the differential growth rates of consumption and investment seen in recent years. However, it also has its limitations. Ultimately, such models fail to tell us exactly *why* some industries have had exceptionally fast rates of productivity growth, and others have not.

While the 1980s and 1990s saw an enormous amount of work directed at explaining technological progress as an endogenous phenomenon resulting from the profit-maximizing actions of rational agents, almost none of this work focused on understanding the apparently large gap between technological progress in the durable goods sector and in the rest of the economy. (For example, there is no mention of this pattern in Barro and Sala-i-Martin's text.) There may be large gains to future empirical research in this direction. In particular, an exploration of whether this pattern can be explained by differential rates of R&D activity, or other "spillovers" stressed in the endogenous growth models of Romer (1990) and Jones and Williams (1998), would appear to be particularly worthwhile, and could yield important policy implications.

Bibliography

- Barro, Robert and Xavier Sala-i-Martin (1995), *Economic Growth*, New York : McGraw-Hill.
- Campbell, John and Robert Shiller (2001), "Valuation Ratios and the Long-Run Stock Market Outlook", NBER Working Paper, No. 8221.
- Chirinko, Robert (1993), "Business Fixed Investment Spending : A Critical Survey of Modelling Strategies, Empirical Results, and Policy Implications", *Journal of Economic Literature*, 1875-1911.
- Chow, Gregory (1967), "Technological Change and the Demand for Computers", *American Economic Review*, 57(5), 1117-30.
- Cole, Roseanne, Y. C. Chen, Joan Barquin-Stolleman, Ellen Dulberger, Nurthan Helvacian, and James Hodge (1986), "Quality-Adjusted Price Indexes for Computer Processors and Selected Peripheral Equipment", *Survey of Current Business*, 66, January, 41-50.
- Cummins, Jason, Kevin Hassett, and R. Glenn Hubbard (1994), "A Reconsideration of Investment Behavior Using Tax Reforms as Natural Experiments", *Brookings Papers on Economic Activity*, 2, 1-59.
- David, Paul (1990), "The Dynamo and the Computer : An Historical Perspective on the Modern Productivity Paradox", *American Economic Review*, 80, May, 355-361.
- Fisher, Irving (1922), *The Making of Index Numbers*, Boston : Houghton-Mifflin.
- Glassman, James K. and Kevin Hassett (1999), *Dow 36,000 : The New Strategy for Profiting from the Coming Rise in the Stock Market*, New York : Times Business.
- Griliches, Zvi (1961), "Hedonic Price Indexes for Automobiles : An Econometric Analysis of Quality Change", reprinted in Zvi Griliches (ed.) *Price Indexes and Quality Change*, 1971, Cambridge : Harvard University Press.
- Jones, Charles I. and John Williams (1998). "Measuring the Social Return to R&D", *Quarterly Journal of Economics*, 113, 1119-1135.
- Jorgenson, Dale (1963), "Capital Theory and Investment Behavior", *American Economic Review*, 53, 247-259.
- Jorgenson, Dale and Zvi Griliches (1967), "The Explanation of Productivity Change", *Review of Economic Studies*, 34, 249-280.
- Jorgenson, Dale and Kevin Stiroh (2000). "U.S. Economic Growth in the New Millennium", *Brookings Papers on Economic Activity*, 1, 125-211.
- King, Robert, Charles Plosser, James Stock, and Mark Watson (1991), "Stochastic Trends and Economic Fluctuations", *American Economic Review*, 81, No. 4, 819-840.

- Oliner, Stephen and Daniel Sichel (1994), "Computers and Output Growth Revisited : How Big is the Puzzle?", *Brookings Papers on Economic Activity*, 2, 273-317.
- Oliner, Stephen and Daniel Sichel (2000), "The Resurgence of Growth in the Late 1990s : Is Information Technology the Story?", *Journal of Economic Perspectives*, 14, Fall, 3-22.
- Romer, Paul (1990), "Endogenous Technological Change", *Journal of Political Economy*, 98(5), S71-102.
- Tevlin, Stacey and Karl Whelan (2001), "Explaining the Investment Boom of the 1990s", forthcoming, *Journal of Money, Credit, and Banking*.
- Whelan, Karl (2000a), "A Guide to the Use of Chain Aggregated NIPA Data", Federal Reserve Board, Finance and Economics Discussion Series Paper No. 2000-35.
- Whelan, Karl (2000b). "Computers, Obsolescence, and Productivity", forthcoming, *Review of Economics and Statistics*.
- Whelan, Karl (2001), "A Two-Sector Approach to Modeling U.S. NIPA Data", forthcoming in *Journal of Money, Credit, and Banking*.