

Quantifying Embodied Technological Change¹

Plutarchos Sakellaris

University of Maryland, Athens University of Economics and Business, and IMOP

and

Daniel J. Wilson

Federal Reserve Bank of San Francisco

Abstract:

We estimate the rate of embodied technological change directly from plant-level manufacturing data on current output and input choices along with histories on their vintages of equipment investment. Our estimates range between 8 and 17 percent for the typical U.S. manufacturing plant during the years 1972-1996. Any number in this range is substantially larger than is conventionally accepted with some important implications. First, the role of investment-specific technological change as an engine of growth is even larger than previously estimated. Second, existing producer durable price indices do not adequately account for quality change. As a result, measured capital stock growth is biased. Third, if accurate, the Hulten and Wykoff (1981) economic depreciation rates may primarily reflect obsolescence.

Keywords: productivity growth, embodied technological change, equipment investment, plant, producer durable price index.

JEL: O3, D24, L60

1 Correspondence: Daniel.Wilson@sf.frb.org. This is a revised version of a paper previously circulated as “The Production-Side Approach to Estimating Embodied Technological Change.” We would like to thank Carol Corrado, Boyan Jovanovic, two anonymous referees and seminar participants at Athens University of Economics and Business, University of Ioannina, Koc University, University of Maryland, New York University, European Central Bank, Federal Reserve Banks of Boston and San Francisco, Federal Reserve Board, Center for Economic Studies (CES), U.S. Census Bureau for very useful comments. Plutarchos Sakellaris would like to thank the staff of the Industrial Output Section of the Federal Reserve Board for their hospitality while he was working on this paper. Daniel Wilson would like to thank Clopper Almon and the staff of Inforum for their generous support, financial and otherwise, of this research. This paper was in part revised when the first author visited the European Central Bank’s Directorate General Research as part of their Research Visitor Programme. The data used in this paper were collected under the provisions of Title 13 US Code and are only available at the CES. The research in this paper was conducted by the authors as Research Associates at the CES. The views expressed here do not represent those of the Federal Reserve System, its staff, or the U.S. Census Bureau.

1. Introduction

Johansen (1959) and Solow (1960) argued that more recent vintages of capital may embody technological advances that make them “better”, or equivalently “of higher quality”, than older vintages. In that case, investment is essential in order to reap benefits from the part of technological progress that is embodied in capital. The distinction between embodied and disembodied technological change is important in analyzing many issues related to long-run standards of living and growth.

Jovanovic and Rob (1997) show that distortions to investment incentives have a substantially larger effect when technological growth is embodied in equipment. In particular, the taxation of equipment plays a big role in explaining development gaps across countries. Eaton and Kortum (2001) attribute about 25 percent of cross-country productivity differences to variation in the relative price of equipment, about half of which they ascribe to barriers to trade in equipment that inhibit the diffusion of embodied technological change across borders. The decomposition of technological change to embodied and disembodied may be important for understanding medium-run growth experience, as well. From 1995 to 1999 U.S. labor productivity and total factor productivity (TFP) grew at rates about double those of the preceding 20 years. This dramatic change has led to some important questions. What are the reasons for this productivity surge? Is it going to be long lasting?

A possible explanation for this surge in productivity growth assigns a big role to equipment investment. The U.S. economy experienced an equipment investment boom in the middle and late 1990's. Investment in equipment and software averaged 9.6 percent of GDP

between 1995 and 1999, 3.2 percentage points higher than the average from 1987-94. Much of this new equipment was for Information Technology but other types of equipment were also purchased en masse. If new equipment embodies technology that is more productive on the margin than that of older equipment and official price indices do not fully reflect such quality changes then an equipment investment boom should be accompanied by an increase in measured TFP growth. TFP will grow faster than average for as long as the rates of investment are heightened.

There is another possibility for sustainably high TFP growth rates within this framework. If the rate of growth of technology that is embodied in new vintages of investment has increased then at any rate of equipment investment measured TFP growth will be higher.² Some have argued that the advent of Information Technology (the computer revolution) and its incorporation into production has slowly pushed the average rate of embodied technological change higher (see Greenwood and Yorukoglu, 1997, among others), especially after 1973.

A serious shortcoming in the above discussion has been the lack of reliable estimates of the rate of technological change that is embodied in equipment capital. Most of the literature on embodiment, including the papers of Hulten (1992) and Greenwood, Hercowitz, and Krusell (1997), has relied on an estimate of about 3 percent for the years 1954 to 1990. This estimate was derived by comparing the rates of growth of either a consumption deflator or the official price index for equipment investment (which turn out to be quite similar) with another one constructed by Gordon (1990), which is intended, *inter alia*, to capture change in the quality of

² Again it is important for this argument to go through that official price indices do not incorporate quality growth adequately.

new vintages of equipment. We call this method of deriving embodied technological change the price approach. However, as Gordon (1996) points out, the difference in the growth rates of his investment price index and either of the other two can also be affected by factors other than quality change. Furthermore, not all kinds of quality change are captured by the Gordon (1990) index.³

In this paper, we propose an estimation framework that arrives at estimates of the rate of growth of embodied technological change directly from observed production, input, and investment decisions at the plant level. We use data contained in the Longitudinal Research Database maintained by the Center for Economic Studies at the Bureau of Census. If there are vintage effects (embodiment) then plants with relatively newer equipment should be more productive, after controlling for materials, labor input and utilization of capital and labor. These effects may be estimated in a production function framework where the capital stock of equipment is not constructed using the perpetual inventory method. Instead, we include in the estimating equation the complete history of investment in equipment by vintage, deflated by an index that does not correct for quality change. The enormous cross-sectional variation in plant-level investment histories that is afforded by the LRD allows us to estimate vintage effects in a manner not possible with more aggregate data.

A simple example may help illustrate our basic idea. Consider two plants, A and B, both born in 1980 and observed continuously until 1989 (inclusive). Plant A expended 55% of its lifetime (up to 1989) investment in 1980 and 5% in each year thereafter. In contrast, plant B

³ See Section 3 for further discussion on the use of Gordon's index to capture embodied technological change.

invested 55% in 1989 and 5% in each prior year. The plants are otherwise identical. One can exploit the variation in these two plants' vintage distributions in investment (net of physical decay) to estimate embodied technological change by comparing their output in 1989. Doms & Dunne (1998) provide empirical evidence of this kind of extreme variation in investment. They find that over a 17-year period, 50% of the average plant's investment is concentrated in 3 years (not necessarily consecutive). Furthermore, lumpy investment activity is not synchronized across plants. Our basic idea is then to exploit the large cross-sectional variation in investment profiles as well as variation in productivity in order to estimate how much more productive is each new vintage of equipment investment.

The result that we obtain is that each vintage is about 12 percent more productive than the previous year's vintage (in the preferred specification).⁴ This number is far greater than the rate of 5 percent that Cummins and Violante (2002) estimate after extending Gordon's (1990) series. However, it is consistent with a number of other studies that, like ours, are not based on Gordon's series (e.g., Bahk and Gort (1993) and Hobijn, et al. (2002)). We evaluate the impact of embodied technological change on US manufacturing productivity growth between 1972 and 1996. We estimate that the effective capital stock of equipment grew about three times faster than commonly estimated and that the contribution of embodied technological change to total technological change was about two thirds.

These findings have several important implications. First, the role of investment-specific technological change as an engine of growth is even larger than previously estimated. Second, existing producer durable price indices substantially mismeasure quality change, yielding

⁴ This assumes that quality change in consumption goods is approximately zero.

downward biased measures of capital stock growth. Lastly, assuming Hulten and Wykoff's (1981) estimates of economic depreciation, which have since been adopted by the U.S. Bureau of Economic Analysis (BEA), are correct, our estimates suggest that obsolescence is the most important factor in the decline of the average capital asset's value over time.

2. Concepts

To better understand what we mean by embodied technological change and how we go about estimating it, consider the following general definition of capital stock in efficiency (quality) units (relative to their efficiency in the base year, t_0):

$$J_t = \sum_{s=1}^T I_{t-s} D_{t,t-s} (1 + \gamma)^{t-s-t_0} \quad (1)$$

where I_{t-s} is investment of vintage $t-s$ (i.e., s years prior to the current year), $D_{t,t-s}$ is the fraction of vintage $t-s$ investment still physically productive in year t , and $(1 + \gamma)^{t-s-t_0}$ represents capital-embodied technology, with γ denoting the rate of embodied technological change.

As shown in the embodied technological change models developed by Hornstein and Krusell (1996); Greenwood, et al. (1997); Hercowitz (1998); and others, if I_{t-s} is measured in foregone consumption units, then $(1 + \gamma)^{t-s-t_0}$ represents the relative productive efficiency of vintage $t-s$ investment, whose resource cost (in year $t-s$) was one unit of the consumption good. This vintage $t-s$ efficiency is measured relative to the efficiency of vintage t_0 investment (costing one consumption unit in t_0). In other words, $(1 + \gamma)^{t-s-t_0}$ represents embodied technology as long as I_{t-s} is deflated by a consumption price deflator that is not adjusted in any way to reflect

quality changes occurring in capital goods. In referring to this relative efficiency change as embodied technological change, we are implicitly assuming that there is no technological change in consumption goods. To the extent that there is technological change occurring in the consumption sector, $(1 + \gamma)^{t-s-t_0}$ would more properly be termed “investment-specific technological change,” defined as that technological change in investment goods above and beyond any technological change in consumption goods.

Equation (1) can be rewritten as

$$\begin{aligned} J_t &= (1 + \gamma)^{t-t_0} \sum_{s=1}^T I_{t-s} \left(\frac{D_{t,t-s}}{(1 + \gamma)^s} \right) \\ &= (1 + \gamma)^{t-t_0} \sum_{s=1}^T I_{t-s} \tilde{D}_{t,t-s} \end{aligned} \tag{2}$$

Here $\tilde{D}_{t,t-s}$ represents what is often termed *economic depreciation*. It captures the decline in the productive efficiency of capital due to the joint effects of embodied technological change (which is synonymous with obsolescence) and physical depreciation.

Hulten and Wykoff (1981) aim at estimating economic depreciation by estimating the rate at which used asset prices decline with the age of the asset. They find that the average piece of capital equipment experiences an age-related price decline of around 12%. Hulten and Wykoff were not able to separately identify physical depreciation and embodied technological change but rather their joint effect by exploiting the variation in used asset price data across age for a given cross-section year.

Similarly, in our approach, we are unable to separately identify these two effects econometrically. However, by exploiting cross-sectional variation in plant productivity (i.e., output after controlling for observable factor inputs) and plant investment history (the stream of

$I_t, I_{t-1}, I_{t-2}, \dots$), we can estimate their joint effect. In this way, our method is exactly the dual approach to Hulten and Wykoff's price approach.⁵

In order to isolate the effect of embodied technological change, we rely on physical depreciation estimates developed by the Federal Reserve Board (FRB). Because these estimates are non-geometric and vary by industry, we cannot simply estimate the joint effect and then subtract an exponential rate of physical depreciation. There is a complex interaction between investment histories and the non-geometric physical depreciation. As we show later, the estimate of γ would change substantially were we to assume a simple geometric rate for physical depreciation data fitted to the FRB pattern.

Our method to identify γ :

In Section 4, we propose a method of estimating the rate of growth of embodied technological change, γ , without the use of a quality-adjusted price index for equipment investment. This method requires cross-sectional data, say at the plant level, on gross output and current and past investment, in addition to data on labor and materials inputs. There are four important elements to this approach:

1) Deflate nominal investment by a price index that does not adjust for quality change in capital. We choose to use the Personal Consumption Expenditures (PCE) deflator, implicitly assuming that this price index does not adjust for quality growth embodied in capital. As noted

⁵ Hall (1967) demonstrates the fundamental inability to separately identify physical depreciation and embodied technological change with only a cross-sectional dimension and an age/vintage dimension. This fundamental identification problem implies that either physical depreciation or embodied technological change must be "observed" in some way. Measuring the quantity change in the technological characteristics in capital goods such as computers is an attempt to "observe" embodied technological change. The FRB and BLS "observe" physical

earlier, to the extent that quality growth occurs in consumption goods as well, our method will, strictly speaking, identify investment-specific technological change rather than total capital-embodied technological change.

2) Apply depreciation adjustments to capital that capture physical depreciation but not obsolescence. Thus, we construct $D_{t,t-s}I_{t-s}$. Here, we choose to use the methodology followed by the Federal Reserve Board (FRB) for capturing solely physical depreciation. We do not use the Hulten-Wyckoff/BEA rates since they measure economic depreciation and thus include the effect of embodied technological change.

3) Allow the marginal productivity of each vintage to be freely estimated or constrained to grow at the rate $(1+\gamma)$. This provides an estimate of embodied technological change. The parameter γ is identified from the cross-sectional variation in our data.

4) Use the time series variation in our data to identify disembodied technological change. In particular, we capture disembodied technological change with year dummies and industry-specific time trends.

It is important to emphasize here that the rates of growth of embodied and of disembodied technological change are identified conditional on use of accurate physical depreciation adjustments $D_{t,t-s}$ as well as the use of a deflator that does not incorporate quality growth. Later in the paper, we discuss in detail our choices and examine issues of robustness.

3. Related Literature

depreciation for particular capital goods based on engineering knowledge of the physical service lives of assets and the shape of physical depreciation pattern over age.

One could classify the set of related papers into three camps. Most of the recent contributions use Gordon's (1990) quality-adjusted price indices for PDE in order to identify embodied technological change and then answer important questions related to economic growth or fluctuations in the U.S. The second camp contains older contributions that estimated embodied technological change using data on production and capital stock age using an approach due to Nelson (1964). The third approach was pioneered by Hobijn (2000) and might be called the structural model-based approach. We review these three camps briefly here and point out the main differences of our approach.

Price-Based Estimates of Embodied Technological Change

Gordon (1990) is a major study aimed at correcting mismeasurement in equipment price indices due to quality change. He provides quality-adjusted price indices for 22 types of equipment and their components. Hulten (1992) was the first to use these series in order to identify embodied technological change. He constructed a single aggregate index from Gordon's indices as well as one for the corresponding price indices published by BLS. Taking the ratio of the two, he calculated the average annual growth rate of embodied technological change to be 3.44 percent for U.S. manufacturing during 1949-1983. As a result, he attributed about 20 percent of the residual growth of quality-adjusted manufacturing gross output to embodied technological change.

Various papers followed Hulten (1992) in using Gordon's data but differed in the methodology employed. Greenwood et al. (1997) argued that the baseline index for comparison should be the implicit price deflator for non-durable consumption goods. This had very little

effect on their estimate of embodied technological change. Three other differences in methodology, however, were important. Greenwood et al. (1997) advocated that output not be adjusted for quality change, that value added data be used in place of gross output and that a general equilibrium approach be used that accounts for input growth due to embodied technological change. They found that embodied technological change contributed about 58 percent of all output growth in the U.S. between 1954 and 1990.⁶

Hornstein and Krusell (1996) and Greenwood and Yorukoglu (1997), among others, note that the Gordon index points to a large increase in the rate of embodied technological change after 1973.⁷ They argue that increased technological change and the adjustment processes necessitated by it are largely responsible for the post-1973 measured slowdown in productivity growth. Greenwood et al. (2000) use Gordon's data to investigate the importance of embodied technological change for postwar U.S. aggregate fluctuations. They document a negative comovement between the relative price of equipment and equipment investment. Their analysis suggests that about 30 percent of output fluctuations are due to shocks in this relative price.⁸

The most recent addition to this literature is a paper by Cummins and Violante (2002). Using Gordon's quality-adjusted price indexes for 1947-1983 they estimate the quality bias implicit in the official (NIPA) indexes for the same period. They then extrapolate the quality bias in the official indexes for the period 1984 to 2000. This way they estimate that embodied technological change grew at an annual rate of about 3 percent up to 1975 and accelerated to

⁶ Disembodied change provided the rest. The authors actually use the terminology investment-specific versus neutral technological change.

⁷ This increase is about 25 percent. See, e.g., Hornstein and Krusell (1996) p. 231.

about 5 percent thereafter.

The practice of using Gordon's (1990) quality-adjusted price indices for PDE in order to identify embodied technological change is not uncontroversial. Gordon (1996) himself even takes issue with it. He claims that differences between his indices and the official ones are not entirely due to quality change. He offers as an example the traditional substitution bias introduced by a change in the relative prices of goods. Furthermore, he continues, quality adjustment may arise for reasons unrelated to embodied technological change. For example, an energy price increase may lead consumers to shift from inexpensive and energy-inefficient air conditioners to expensive and energy-efficient ones. The latter are also more costly to produce at any given level of technology in the sector that produces them and do not necessarily represent an increase in technology.

Finally not all kinds of quality change are captured by the Gordon index. Gordon was only able to create new price indices for durable goods for which sufficient data on model characteristics and prices existed. The Sears catalog was the primary source of this data. For a large number of goods there simply was no data.⁹ Given these problems, it seems fruitful to examine alternative methods of estimating embodied technological change.

⁸ There are several related papers. See Greenwood et al. (2000) pp. 110-2 for a partial review.

⁹ Gort and Wall (1998) argue that estimates of embodied technological change based on Gordon (1990) may be biased substantially towards zero (pp. 1658-9). They also point out another problem with the group of studies applying this methodology. While they adjust investment flows for quality change, they implicitly assume that economic depreciation rates, derived from Hulten and Wykoff (1981) and incorporated in the NIPA, measure only physical decay. This is unlikely to be true as these measures also incorporate obsolescence. It is clear that research effort has to be directed towards estimating rates of physical depreciation for assets. The only such example we are aware of is the work of Oliner (1993) for mainframes and peripherals.

Production-Based Estimates of Embodied Technological Change

Nelson (1964) developed a variant of Solow's (1960) embodiment model that illustrated the relationship between the rate of embodied technological change and the average age of capital. He showed that, to an approximation, the log of efficiency-adjusted capital ($\bar{k}$) is proportional to a time trend *plus* the difference of γa_t and the log of unadjusted capital (k), where a_t is the average age of the unadjusted capital stock. Thus, a standard production function estimation (in logs) provides an estimate of embodied technological change by dividing the coefficient on average age by the coefficient on capital stock. Bahk and Gort (1993) study a sample of young manufacturing plants and find that a 1-year drop in average age is associated with between a 2.5 and a 3.5 percent rise in the plant's gross output (See their Table 1 and p. 571). Assuming a one-sixth share weight for capital in the production function of gross output, these results correspond to a 15-21 percent annual rate of growth of embodied technological change. This is three to four times higher than the baseline price-based estimates discussed above.

However, Bahk and Gort (1993) make the dubious assumption that maintenance outlays offset the effect of physical decay on the capital stock. Their capital stock construct is the sum of gross investments of all vintages without any adjustment for physical depreciation.¹⁰ Their estimates are best regarded as describing the joint impact of physical depreciation and embodiment. Our view is that physical decay occurs at a rate higher than zero.

Bahk and Gort's results should be contrasted with a series of earlier papers that used

¹⁰ While they admit that "this assumption is only at best an approximation of reality" (p.566) they do not provide concrete supporting evidence for it.

production-based techniques and found no evidence of embodied technological change (Wickens, 1970, You, 1976). It seems that the reason for this difference is the use in these papers of aggregate time series data. McHugh and Lane (1987) demonstrated that using cross-sectional information together with an adjustment for differential utilization of vintages provides ample evidence for embodied technological change.

Structural Model-Based Estimates of Embodied Technological Change

Hobijn (2000) measured embodied technological change by fitting a simple stochastic vintage capital model to aggregate US data. After calibrating its parameters he filtered out the implied path of embodied technological change and estimated the rate of growth of embodied technological change at about 4% for the period 1975-1999. However, this estimate is biased downward because, as the author admits, the physical depreciation rate used in this calibration (4.6%) is too high. A more reasonable rate of physical depreciation would probably put the estimate of embodied technological change from Hobijn's model at about 4.5% or even higher.

A related approach is taken in Hobijn et al. (2002). They modify the neoclassical model of investment with convex adjustment costs to include embodied technological change. They then estimate the Euler equations for investment using data for 4-digit US manufacturing industries. Their estimates of the rate of embodied technological change for equipment range between 9.5% and 32.6%. The stark difference in estimates between the two papers may be due to the higher level of aggregation of the data in the first paper. As mentioned earlier, papers using a production-side approach have already concluded that aggregation biases the rate of embodied technological change downward.

4. Our Methodology

We now formalize our approach. Consider a Cobb-Douglas function for the production of plant output:

$$y_{it} = \beta \cdot l_{it} + \theta \cdot m_{it} + \alpha \cdot \log(U_{it}^J J_{it}) + \eta \cdot \log(U_{it}^S S_{it}) + \zeta_{it} \quad (3)$$

where y is (log) gross output, l is (log) labor hours, S is the structures capital stock, m is (log) materials (including energy expenditures), J is the equipment capital stock, i indexes plants, and t indexes year.¹¹ The term, ζ_{it} , captures stochastic shocks to disembodied technology. It may contain both an idiosyncratic and a non-idiosyncratic component. In the regressions reported below, we estimate the non-idiosyncratic component of ζ_{it} with a set of year dummies, industry dummies, and industry-specific linear time trends. The idiosyncratic component is absorbed into the error term. We are thus assuming that it is uncorrelated with the factor inputs, though we also explore relaxing this assumption below.

The variables U^J and U^S in equation (3) denote that the utilization rates of equipment and structures capital, respectively. To measure utilization, we follow Petropoulos (1999) who shows that under certain reasonable conditions, the intensity of a plant's energy usage can be used as a proxy for capital utilization. Specifically, we assume:

$$U_{it}^J = (E_{it} / J_{it})^{\frac{1}{\tau_J}} \quad \text{and} \quad U_{it}^S = (E_{it} / S_{it})^{\frac{1}{\tau_S}} \quad (4)$$

where E denotes energy expenditures (fuel plus electricity). For purposes of identification, we

¹¹ All variables are in constant dollars unless stated otherwise.

assume that $\tau_J = \tau_S = \tau$. Substituting these expressions into the production function (3) and rearranging, we obtain

$$y_{it} = \beta \cdot l_{it} + \theta \cdot m_{it} + [\alpha(\tau - 1)/\tau] \cdot \log(J_{it}) + [\eta(\tau - 1)/\tau] \cdot \log(S_{it}) + [(\alpha + \eta)/\tau] \cdot e_{it} + \zeta_{it} \quad (5)$$

where e denotes $\log(E)$.

In this paper, we focus on the estimation of the rate of embodied technological change in *equipment* capital. Essentially, we assume that to the extent that there is embodied technological change in structures it is adequately captured in the official price indexes. Therefore, plants' structures capital stock is calculated according to the traditional perpetual inventory definition using historical structures investment from the LRD, price deflators from NBER-CES, and physical depreciation data from the FRB. Equipment capital, on the other hand, is a parameterized stream of past real equipment investment (net of physical decay, $D_{i,t,t-s}$):

$$J_{it} = \sum_{s=1}^T I_{i,t-s} D_{i,t,t-s} \phi_{t-s} \quad (6)$$

where T is the age of the plant, $I_{i,t-s}$ is real equipment investment of vintage $t-s$ capital goods, $D_{i,t,t-s}$ is the fraction of one dollar's worth of vintage $t-s$ investment that is still used in production in year t , and ϕ_{t-s} is the level of embodied technology in equipment capital of vintage $t-s$ relative to some numeraire year's technology.¹²

¹² We assume that investment in vintage $t-s$ capital goods is synonymous with new investment in year $t-s$. This is incorrect to the extent that there is investment in used capital of earlier vintages. Unfortunately, data on the vintage of a plant's used investment is not available in the LRD. This should have little effect on our results since used investment is typically a negligible part of total (new plus used) investment.

It is important to note that we deflate nominal equipment investment by the price index for Personal Consumption Expenditures.¹³ In so doing we are assuming that this price index does not incorporate any technological change that is embodied in new equipment. Equation (6) incorporates a one-year time-to-build assumption; that is, new investment is not put into operation until the following year. It takes time to build because of actual assembly requirements or because time is needed to train workers on how to use the new equipment.¹⁴

Substituting equation (6) into (5) yields the estimating model. Given the large number of parameters and the non-linearity of the model, we found that obtaining estimates (precise or otherwise) of ϕ_{t-s} was not possible in all but the most parsimonious of specifications. Thus, for the regressions reported in this paper, we simplify by assuming a constant geometric rate of embodied technological change. This changes the specification of J_{it} to:

$$J_{it} = \sum_{s=1}^T I_{i,t-s} D_{i,t,t-s} (1 + \gamma)^{t-s-t_0} \quad (7)$$

where γ is the rate of embodied technological change and t_0 is the numeraire year in which the level of embodied technology is normalized to 1. We choose 1996, the last year of our data, as the numeraire year.¹⁵

13 Some authors in the price-based literature suggest other consumption deflators such as PCE for nondurables and nonhousing services. We found that using such subaggregate PCE deflators have virtually no effect on our empirical results, which is not surprising given that these subaggregates are very highly correlated with the aggregate PCE deflator.

14 We estimated some specifications allowing no time to build but assuming that investment is distributed evenly throughout the year so that, on average, six months' worth of total investment in the newest vintage is used at any one time. The estimates of γ are uniformly higher under this specification.

15 Equations (5) and (6) provide a way of aggregating vintages that embody different technology levels. Fisher (1965) shows that a necessary and sufficient condition for the existence of a capital aggregate are that the marginal rate of substitution between any pair of inputs within the aggregate be independent of the inputs outside the aggregate. Under constant returns to scale, as Solow (1960) showed, this condition requires that 1) production with

Discussion of variable capital utilization

Unmeasured variation in the intensity with which plants utilize capital may lead to biases in production function estimation. We avoid this pitfall by proxying for capital utilization with energy use.¹⁶ The parameter τ in the above model is the elasticity of the rate of energy use with respect to capital utilization. It allows for energy use being proportional to capital services ($U^J J$ or $U^S S$, as in Burnside, Eichenbaum, and Rebelo, 1995) when $\tau = 1$. For values of $\tau > 1$ the marginal cost of capital services increases faster with utilization than with capital stock and at the extreme ($\tau = \infty$) there is no variation in utilization as it is too expensive. Similarly, when $\tau < 1$ the marginal cost of capital services increases slower with utilization than with capital stock and at the extreme ($\tau = 0$) any variation in utilization is achieved by an infinitesimal variation in energy use.¹⁷

Discussion of physical depreciation assumptions

Our estimates of the rate of growth of embodied technological change rely importantly on accurate measures of physical depreciation. This is the part of economic depreciation that is due

each vintage be additively separable, 2) total factor productivity be the same across all vintage production functions at a given year, and 3) investment in a “better” vintage of equipment is equivalent to a larger amount” of investment measured in constant quality.

16 This approach was suggested by Jorgenson and Griliches (1967) and followed by Burnside et al. (1995) for industry-level estimation and Petropoulos (1999) for plant-level estimation, among others. See Fernald and Basu (1999) for pitfalls arising from unmeasured factor utilization.

17 Petropoulos (1999) argues that often plants increase capital utilization by “dusting off” older, less efficient machines. Then, increases in utilization would require increasing rates of energy use per unit of capital utilized ($\tau > 1$).

to wear and tear resulting from using an asset in production. We do not use the exponential/geometric depreciation rates produced by BEA, which are largely based on the estimates of Hulten and Wykoff (1981), as these reflect both physical depreciation and embodied technological change. Instead, we employ the methodology used by BLS and FRB in constructing capital stocks adjusted for the effects of physical depreciation (see Mohr and Gilbert (1996)).

There are two important differences. First, the FRB-BLS methodology results in an age-efficiency schedule that is substantially different from geometric, especially in the early part of an asset's life. Second, the implied rate of physical depreciation is substantially lower than the BEA's rates of economic depreciation. These features can be seen in Figure 2, which graphs the average depreciation schedule by age in our primary sample for each source of depreciation data. Specifically, for the FRB-BLS depreciation data, we plot the average by age (s) over all years (t) and plants (i), of the $D_{i,t,t-s}$ in equation (7). The FRB-BLS depreciation data are available at the 3-digit SIC industry level. Thus, though $D_{i,t,t-s}$ has a plant subscript, empirically this variable is the same for all plants within a 3-digit industry.

Similarly, we compute and plot the average of the $D_{i,t,t-s}$ data that is implied by the BEA depreciation rates according to the identity $D_{i,t,t-s} = (1 - \delta_s^{BEA})^{t-s}$. Lastly, the depreciation schedule implied by a 12% geometric rate is shown as a frame of reference. Clearly, the average BEA (economic) depreciation pattern is quite similar to a 12% exponential decline. The FRB-BLS (physical) depreciation pattern, on the other hand, has a reverse-S shape as well as having a much slower decline in general. In the empirical section, we explore the robustness of our estimates of γ to different assumptions about physical depreciation.

5. Data

Sample

The plant-level data we use come from the confidential Longitudinal Research Database (LRD), which resides at the U.S. Census Bureau's Center for Economic Studies.¹⁸ The LRD consists of annual data on U.S. manufacturing establishments collected in the quinquennial Census of Manufactures (CM) and the Annual Survey of Manufacturing (ASM). At the time of this writing it covers 1963, 1967, and 1972-96.

We construct several data samples with which to estimate equation (4).¹⁹ Our primary sample is an unbalanced panel consisting of all plant-year observations from plants born in or after 1972 that survived for at least four consecutive years (including the birth year), with the last possible observation year being 1996. Thus, a single plant may have multiple observations in our sample, provided that we continuously observe the plant's equipment investment for every year from birth until the current observation year. Observations from plants born prior to 1972 are excluded as we cannot observe their entire investment history. These missing investment variables would likely have caused substantial biases in estimating the relative efficiency of different vintages of investment. An alternative sample was constructed in order to test the robustness of our results to the possibility of sample selection bias introduced by the exclusion of plants born prior to 1972. The results are robust to this change of sample (results from alternative samples are available from authors upon request).

¹⁸ Access to the LRD can be obtained for approved research purposes from the Center for Economic Studies.

¹⁹ We will report results for our primary sample only. Results using other samples are available from the authors upon request.

We keep only observations from plants that have survived at least four consecutive years in an effort to avoid two possible problems. First, the factor elasticities, particularly for capital, in the first three years of a plant's operations have been shown to be significantly different than those in later years (see Bahk and Gort (1993)). In fact, the elasticity of output with respect to capital is often found to be statistically insignificant in the first few years. Bahk and Gort attribute this phenomenon to learning by doing. Second, data in the LRD is imputed for a substantial portion of any given year's newly-born plants (with the exception of data on industry, employment, and payroll). Thus, data from the first year of plants is fraught with measurement error. Since we utilize the entire investment history of a plant, we are not immune to the measurement error introduced by imputation of the first year's equipment investment. However, the share of investment that is imputed becomes smaller and smaller as the plant ages and thus, it is hoped, its impact should be minimal by the fourth year. Using an alternative sample that retained observations of two- and three-year old plants, we found that our results do not appear to be biased by their exclusion (alternative sample results available from authors upon request).

Variable Construction

The definitions of gross output, labor, structures capital stock, materials, and energy are similar to those used elsewhere in the plant-level literature (see Dunne 2000). Real gross output was defined as the value of shipments plus inventory change deflated by the 4-digit SIC shipments deflator in the 4-digit SIC level NBER-CES Productivity Database (hereafter, NBER-CES)(see Bartelsman and Gray 1996). The concept of "production-worker equivalent hours" was used for the labor variable:

$$L = ph + (nw/pw)ph,$$

where ph is production worker hours, nw is total non-production worker salaries and wages, and pw is total production worker salaries and wages. Assuming workers are paid their marginal product, the second term should capture “production worker equivalent hours” contributed by non-production workers. Real energy expenditures is measured as the sum of the costs of fuel and electricity deflated by the NBER-CES energy deflator. Real materials is the sum of the costs of fuel, electricity, and parts, deflated by the NBER-CES materials deflator.²⁰ The measurement of real structures capital stock and real equipment investment are discussed above in Section 4, yet let us reiterate here that we deflate structures investment by the NBER-CES structures deflator (implicitly assuming it captures any technological change that may have occurred in structures) while we deflate equipment investment by the PCE deflator (and leave equipment-embodied technological change to be picked up in the estimated vintage weights).

Physical Depreciation Data

To obtain physical depreciation data, we employ the methodology used by BLS and FRB in constructing capital stocks adjusted for the effects of physical depreciation (for details see Mohr & Gilbert 1996). For each vintage of investment we repeat the following procedure. First, 3-digit industry-level capital expenditures are split among 35 asset categories. This is accomplished with an iterative matrix balancing (RAS) technique that employs the industry-level investment data as column controls and utilizes (aggregate-economy) NIPA data on asset-level

²⁰ A data appendix describing the construction of each of the variables and containing a description of the properties of our samples can be provided by the authors upon request.

capital expenditures as row controls for the 35 asset categories. Data derived from BEA's Capital Flows Tables (CFT) provide initial asset-by-industry investment shares for the iterative procedure.

Second, these asset expenditures by industry are transformed from a current-dollar to a constant-dollar basis to obtain estimates of real investment by asset type that can be compared across time. All assets are deflated using the PCE deflator (consistent with our use of the PCE deflator to deflate plant-level total investment for our regressions). Third, we use mean service lives that are specific to each asset type and adjusted to account for expected retirements from each asset-type investment bundle around its mean life. These mean service lives were supplied by BEA except for autos. A discard density function captures stochastic retirements around these mean lives, whereas a hyperbolic- (or beta-) decay function captures the effect of physical depreciation due to wear and tear. We adjust real industry investment by asset type for the joint effects of the decay and discard processes rather than just the latter. In effect, we create a series for each asset of the amount of each vintage that will still be productive at age a . Finally, we aggregate all assets of the same vintage to derive age-efficiency schedules specific to vintage and investing-industry.

Is our Sample Representative of Manufacturing?

Our primary sample consists of a total of 96,846 plant-year observations covering 24,404 plants. We performed a detailed comparison of this sample to the other samples we use and to the aggregate manufacturing sector (the results of which can be provided by the authors upon request). In terms of the dynamic behavior of its plants, this sample is quite representative of the

manufacturing sector as a whole. The average growth rates of investment and employment in our sample are quite similar to those of aggregate manufacturing over the 1975-96 period. This is true even though the sample, due to the nature of the LRD, consists of plants that are larger, on average, than the typical manufacturing plant (in terms of gross shipments, employment, or investment) as aggregate activity is primarily shaped by large plants. The sample distribution of shipments across 2-digit SIC industries is also quite representative of manufacturing though there is a tendency to under-represent mature-plant industries such as Petroleum and over-represent young-plant industries such as Communications.

6. Results

In this section we present the results of our estimation of the model consisting of equation (7) substituted into equation (5). We employ non-linear least squares (NLLS) with a heteroskedasticity- and autocorrelation-consistent (HAC) variance-covariance matrix. The HAC VC matrix is necessary because autocorrelation of the errors is likely due to the fact that the same plant can have multiple observations in our sample. In order to capture disembodied technological change (ζ_{it} in equation (5)), all regressions include year and industry dummies as well as industry-specific time trends. In addition, we include a dummy variable, which we call *Multi*, indicating whether the plant is part of a firm that operates multiple plants. Table 1 contains our main results. The results of sensitivity analysis with different samples are available separately upon request.

Initially, we do not allow for variable capital utilization by setting $\tau = \infty$. As may be seen in column (A), the rate of growth of embodied technological change, γ , is estimated to be a little

under 8 percent. This is higher than the rate calculated in the price-based approaches using Gordon's (1990) data. The coefficient on *Multi* indicates that plants in a multi-establishment firm have on average 8.7 percent higher TFP than single-establishment firms, other things equal. This is consistent with past results on plant-level production estimation (see, e.g., Baily, Hulten, and Campbell (1992)). The estimated factor elasticities are quite close to observed factor shares with the exception of the capital coefficients, which are lower.²¹

The structures elasticity (η) is particularly low. We estimate an elasticity that is about 1/10 that of equipment's elasticity; factor shares implied by NIPA data suggest it should be closer to 1/3. It is well known that unobserved utilization can lead to a downward bias on the output elasticities of quasi-fixed factors. Structures is likely the most fixed of factors. As to how the unobserved utilization in structures may bias gamma, one must answer: does the unobserved utilization have a higher correlation with newer investment or older investment? If the answer is newer investment, the bias on gamma is positive; if it's older investment, the bias is negative. It is unclear whether newer equipment tends to require more unobserved structures services than older equipment. However, next we attempt to at least partially control for unobserved capital utilization and we can then see how the estimates of η and γ change.

Controlling for Unobserved Capital Utilization

The data seem to support the necessity of allowing for variable capital utilization, at least

²¹ Allowing factor elasticities to vary by industry yields a slightly lower estimate of γ of 0.055 (0.028). However, many of the estimated elasticities are nonsensical, particularly those for structures. It appears allowing industry-varying elasticities may be asking too much from the data and therefore we maintain the assumption of common elasticities for the regressions to follow.

as proxied by energy use. As may be seen in Column (B), the estimate of τ is 2.08 (with a standard error of 0.09). This value implies that the marginal cost of capital services (UJ or US) increases faster with utilization than with capital stock. In particular, the Leontief assumption employed by Jorgenson and Griliches (1967) and Burnside et al. (1995), $\tau = 1$, is rejected by our data. The estimated elasticities of capital and labor change in the direction one would expect if plants vary their capital utilization significantly in response to shocks. Not taking such variation into account creates an omitted variable likely to be positively correlated with employment (biasing its coefficient upward) and negatively correlated with the stock of capital (biasing its coefficient downward).²²

Controlling for variation in utilization increases the estimate of γ to 17 percent. In addition, it reduces the coefficient of labor by 2 percentage points while it increases the two capital coefficients substantially. The structures elasticity doubles while the equipment elasticity increases by more than a third. These increases bring the ratio of the elasticity of labor to the sum of the elasticities of capital to about 2:1, as expected, and brings returns to scale to 0.98 which is not statistically significantly different from one. Though the structures elasticity is still below our *a priori* expected value given aggregate data on its factor share, the increase in the structures elasticity supports the hypothesis that this parameter is biased downward in the base specification (column A) due to unobserved variation in utilization. And the fact that γ increases when one (partially at least) controls for variation in utilization suggests that the any bias on γ due to unobserved capital utilization is likely to be negative.

22 For plant-level evidence on these patterns of correlation see Sakellaris (forthcoming).

Learning Effects

Several papers since Arrow (1962) have suggested that the installation of new equipment embodying improved technology may involve significant subsequent investments in training workers as well as experimentation in the organization of production. These would tend to reduce productivity initially while raising it eventually and Arrow (1962) termed this process “learning by doing.”²³ Ignoring such learning effects might bias our estimate of γ , though the direction of the bias is not clear.

To explore the extent of such problems we repeat our estimation including an indicator of whether the plant undertook an equipment-investment spike in the current year together with seven lags of this indicator.²⁴ Our idea is that spikes are associated with most instances of an introduction of a new and “better” vintage of technology and the inclusion of the spike indicators should control for learning effects on productivity. As may be seen in column C of Table 1, the estimate of γ is 11.6 percent and the rest of the coefficients do not change much.²⁵

Average Vintage Effect

In Section 3 we mentioned that an alternative production-based approach which obtains γ via the estimated coefficient on average age or vintage is due to Nelson (1964). In the most

23 See Greenwood and Jovanovic (1998) for an extensive analysis of the macroeconomic implications of learning-by-doing as well as references to some recent work.

24 We identify a spike with observations of equipment investment to capital ratios greater than 0.20. See Power (1998), Doms and Dunne (1998), Cooper et al. (1999) and Sakellaris (forthcoming) for justification of using this definition at the plant level.

25 The coefficients on the spike indicators show an approximately 6% drop in productivity in the first year after a spike with slow recovery thereafter. A similar effect was reported by Huggett and Ospina (2001) for Colombian plant-level data.

parsimonious case where one does not allow for unobserved utilization or learning effects, this approach leads to the following specification (ignoring the constant, the error term, and time and industry dummies):

$$y_{it} = \rho \cdot t + \beta \cdot l_{it} + \theta \cdot m_{it} + \eta \cdot s_{it} + \alpha \cdot j'_{it} + \alpha \cdot \log(1 + \gamma \bar{V}_{it}) \quad (8)$$

where j' is log equipment capital unadjusted for embodied technological change (i.e. j' is just j from equation (7) with γ set equal to zero) and $\bar{V}$ is the average vintage of equipment for plant i in year t .

Generally, $\gamma \bar{V}$ is assumed to be close to zero allowing the last term in the above equation to be approximated by $\alpha \gamma \bar{V}_{it}$. In our sample, however, the mean (over all plant-year observations) average vintage is approximately 12 (corresponding to 1984). The product of 12 and even a very small γ , say 0.01, will still be far from zero suggesting that serious specification errors are likely when we estimate (8). In fact, as shown in column (D) of Table 1, the Nelson method yields the rather implausible estimate of 27 percent for γ which is significantly different from the 12 percent found using our method.

Industry heterogeneity

One concern we have is that the dynamic behavior of output, inputs, and disembodied technological change may differ substantially across industries. A prominent example is the computer sector which has enjoyed rapid growth in its productivity relative to other manufacturing industries and a correspondingly rapid relative decline in its price index. In order to control for this we included a set of industry-specific time trends in each of our regressions.

All industry trends were estimated between ± 1 percent with the important exceptions of computers (8.3%) and communication equipment (3.3%).²⁶ Given that the computer industry has such a high rate of disembodied technological change, it is natural to wonder whether the computer industry is the primary driver of our high estimates of γ . Omitting the computer industry decreases our estimate from 0.116 to 0.070 (with a std. error of 0.053). Thus, it appears the computer industry is *a* but probably not *the* primary driver of economy-wide embodied technology.

Time variation in the parameters

One may also be concerned with whether γ has varied greatly over time. The price-based literature claims quite forcefully a sharp acceleration in γ around 1973. Since our data spans 1972 to 1996 this is not a hypothesis that we can test here. Within the price-based literature one paper, Cummins and Violante (2002), goes even further and claims an acceleration of embodied technological change in the 1990s associated with the diffusion of information technologies. To test this, we estimate our full specification allowing γ to vary by year and test the hypothesis that γ is constant. The F-statistic for this test is 0.17 compared to a 95% critical value of 1.00. Thus, one cannot reject the hypothesis of a constant γ over cross-sectional years. We hasten to add, though, that the reason for this rejection is very high standard errors around very variable parameters rather than precisely estimated parameters near a common value. Thus, the test is not very powerful in addressing a possible acceleration in the rate of embodied technological change.

²⁶ These estimates refer to the regression in Column C though the corresponding estimates from the other regressions are quite similar. We were able to reject the hypothesis that industries do have a common time trend via a Wald test at a 99% level of confidence. The F-statistic is 84.5, far greater than the 99% critical value of 1.8.

Another concern is whether our pooled estimate of γ is affected by the assumption that the elasticity of energy usage with respect to capital utilization, τ , is constant. It is possible that τ was higher in the 1970's when energy costs were high than it was in the 1980's and 1990's. Since the contribution of equipment capital to output, $\alpha(\tau-1)/\tau$, is increasing in τ , if the estimated τ is constrained to be constant whereas, in reality, it was higher early on, then the contribution of early vintage equipment to output would be understated. To compensate, least squares estimation will tend to overweight early vintage investment relative to later investment in order to minimize the sum of squared residuals. This implies an downward bias in γ when constraining τ to be constant. However, allowing τ to vary by year and estimating our full specification over the pooled sample yields a gamma of 0.127 (0.054) compared to the 0.116 (0.060) in Table 1, Column C. Furthermore, τ_t does exhibit a slightly rising trend over time, which is statistically significant according to a Wald test.²⁷

Alternative Depreciation Assumptions

The estimate of γ depends on the assumption for physical depreciation. In the above regressions, we used the FRB physical depreciation estimates. We now entertain the possibility that physical depreciation rates are near-geometric and amount to various fractions of the BEA's economic depreciation rates. Figure 2 shows the overall averages for both the FRB and the BEA estimates. Recall that economic depreciation rates include both physical depreciation and embodied technological change. We construct the parameterized capital stock assuming that

²⁷ The F-statistic for the test of the null hypothesis that τ is constant for all years is $8.48 > F_{J,T-K}(.95)=1.00$.

physical depreciation is a constant fraction of the BEA-provided rate for each observation's industry and year. We then estimate how γ varies.

Table 3 and Figure 3 contain the results. As expected, the lower is the assumed depreciation rate the higher is the estimate of γ , which ranges from 0 to 25 percent (with standard errors ranging from 4 to 6 percent). Clearly as we allow less of the decrease in productivity of early vintages of investment over time to be explained by physical depreciation, γ is left to explain more of this decrease. Total economic depreciation, the sum of physical depreciation and embodied technological change, ranges from 13 to 25 percent. Also shown in Figure 3 is the average BEA economic depreciation rate of 12% and the implied rates of embodied technological change as the physical depreciation rate is increased.

To conclude, even with geometric physical depreciation the estimated rate of embodied technological change is high except in the case where the BEA rates are assumed to reflect only physical depreciation and no embodied technological change. It is straightforward to show theoretically that in the latter case one excludes the possibility of embodied technological change by assumption.²⁸

Simultaneity Bias

It is well known that OLS estimation of production function relationships is subject to potential simultaneity biases.²⁹ Since the “independent” variables are production inputs that are

28 Deflating by the equipment-investment deflator instead of the PCE deflator and repeating the estimation does not change our conclusions.

29 Marschak and Andrews (1944) were the first to recognize this problem. Griliches and Mairesse (1995) provide a thorough discussion of the issue together with attempts to ameliorate the problem using plant-level data.

chosen optimally by the producers, the usual exogeneity assumptions that are required for the consistency of OLS may fail. A plant manager's input choices are determined by plant quality (or, equivalently, managerial efficiency, or disembodied technology) along with factor prices and product demand. Since part of this quality is unobserved by the econometrician, it is subsumed into the disturbance term of the production function. The result is that variable inputs may be correlated with the disturbance term.

To address the possibility of simultaneity bias, we attempted a non-linear instrumental variables (NLIV) estimation. An appropriate set of instrumental variables should include all exogenous variables in the model together with other exogenous and relevant (i.e. correlated with the Jacobian vector of first derivatives of the model with respect to the parameter vector) instruments. For identification it is necessary to have at least one instrument per parameter to be estimated. Unfortunately, it is very difficult to find good instruments (exogenous and relevant) at the plant level for our purposes. This limits severely the success of our NLIV estimations as we will see below. In order to address these problems partially, we try to minimize the number of parameters to be estimated by following a suggestion of Griliches and Ringstad (1971). We replace factor elasticities of some or all of the inputs (e.g. of materials and labor) with measures of their share in cost (using 4-digit industry-level data from the NBER-CES Productivity Database and 2-digit equipment and structures rental rates from the BLS). This leaves in some cases γ , and τ as the only parameters to be estimated.

For instruments, we use the set of 4-digit industry-level downstream demand indicators originally constructed by Bartelsman, Caballero, and Lyons (1994) and modified by Baily, Bartelsman, and Haltiwanger (2001). An industry's demand indicator is a weighted average of

the economic activity of manufacturing and service industries (downstream industries) that purchase the industry's output. The weights are the share of each downstream industry's purchases in the upstream industry's total output and the measure of economic activity is the sum of the cost-share weighted growth rates of each factor input (capital, labor, and materials). In order to filter out any endogenous effect that an upstream industry's productivity may have, through its output price, on downstream industries' activity, the indicator excludes the activity of downstream industries whose purchases from the upstream industry are greater than 5% of their total intermediate input purchases. The instrument set also includes the second through fourth powers of this downstream demand indicator, and investment lagged 3 through 24 years (24 is the highest possible non-zero investment lag that a plant in our sample can possibly have).³⁰

The results of these regressions are shown in Table 4. Column A shows the regression results from our preferred specification where labor's and material's elasticities are not estimated but rather are measured by their industry cost shares. Column B shows the results from the same regression except that all factor elasticities are measured by their industry cost shares and no utilization adjustment is made.³¹ Column C contains results from an average vintage specification also without a utilization adjustment and with all factor elasticities measured by cost shares.

The results of our IV estimations are inconclusive. Largely due to the inclusion of lagged investment in our instrument set, our first-stage R^2 's (not shown) are fairly high, particularly for

30 It may be argued that past investment, even lagged more than three years, is endogenous. Unfortunately, when we omit lagged investment from our set of instruments, our first-stage R^2 drops substantially and the standard errors rise making the estimates essentially meaningless.

31 The iterative estimation procedure was not able to converge when τ was included in the regression.

the first derivative of γ . Yet, despite the apparent relevance of the instruments, we are not able to estimate γ with any reasonable degree of precision. Of the three regressions shown, only the estimate of γ in column B is significant. Based on unreported NLLS results we conclude that the negativity of γ is most likely due to the imposition of the industry cost-shares rather than to the use of IV. In summary, despite our efforts to instrument for endogeneity, we are not able to rule out the possibility of simultaneity bias in our estimate of γ .

7. Productivity growth decomposition

Our findings have the important implication that the equipment capital stock is mismeasured considerably in studies of productivity growth or production function estimation. In particular, we estimate that the stock of equipment grows at much faster rates than is reflected in conventional methods. As may be seen in Table 5, the growth rate of equipment capital stocks depends crucially on the assumptions made on physical depreciation, the investment-price deflator, and the rate of growth of embodied technological change.³²

For both our primary plant-level sample and for the aggregate manufacturing data (from BEA), we calculate the average annual growth rate of equipment capital stock under six different sets of assumptions.³³ The growth rate of the stock implied by the results from our main

32 For the construction of these capital stocks, a numeraire year (t_0) of 1972 was used. The choice of numeraire year can affect whether the growth rate of capital is monotonically increasing or monotonically decreasing in γ . This can be seen by taking the derivative of (J_t/J_{t-1}) with respect to γ where J_t is defined according to equation (6). Choosing a numeraire year at the beginning of our sample ensures that the growth rate will increase with γ . What is important for comparing capital stock growth rates using different γ 's is not necessarily the sign of this derivative, but just that the derivative is monotonic in γ .

33 Specifically, we regress the log of equipment capital stock on t and an intercept. The coefficient on year gives the average annual growth rate.

regression (Column C of Table 1), with $\gamma=0.116$, is 14.7 percent in our sample and 14.0 percent in the aggregate data. This is substantially higher than the growth rate of the conventional capital stock ($\gamma=0$, nominal investment deflated by PPI investment deflators and adjusted for depreciation using the BEA/Hulten-Wyckoff rates), which is 4.2 percent in our sample and 3.0 in the aggregate. It is also far above the growth rates of the capital stocks endorsed by either Hulten (1992) or Greenwood, et al. (1997). The key element that generates the higher growth rate in our capital stock construct is the value of γ . When we set γ to zero but use the FRB physical depreciation data and the PCE deflator to construct the capital stock its growth rate is only 4.4 percent in our sample and 2.1 in the aggregate.

What is the relative importance of embodied and disembodied technological change for productivity growth in our sample of U.S. manufacturing plants? According to equation (3), the residual growth rate in output, after accounting for quality-adjusted input growth, is due to technological change that is not embodied in equipment investment. That growth rate in our sample is 0.57 percent annually.³⁴ Correspondingly, the rate of output growth that is due to embodied technological change is equal to the equipment elasticity times the differential growth rate of quality-adjusted and non-quality adjusted equipment capital stocks. This measure averages 1.17 percent annually in our sample.³⁵ The implication is that equipment-embodied technological change accounted for about two-thirds of total technological change between 1972 and 1996 in our sample of U.S. manufacturing plants.

³⁴ This number is the coefficient on time in a regression of the log of TFP, constructed using the equipment capital stock with a $\gamma=.116$, on time and an intercept.

³⁵ The equipment elasticity is 0.114 from Column C of Table 1. The approximate annual growth rate in the equipment capital stock due to embodied technological change is 10.3 percent. This number comes from subtracting

8. Conclusion

Assessing the rate of embodied technological change is of high importance. It is a necessary ingredient for productivity analysis and for constructing accurate measures of capital accumulation. Furthermore, it tells us how much of the decline in an asset's value as it ages can be attributed to obsolescence.

In this paper, we develop a production-side approach that provides alternative estimates of embodied technological change without relying on accurate measurement of price indices for producer durable equipment. We find that the rate of embodied technological change is between 8 and 17 percent with our preferred estimate being 12 percent. According to our preferred estimate of 12%, embodied technological change may account for as much as two-thirds of the total growth in TFP of manufacturing plants between 1972 and 1996, suggesting an important role for investment in spurring productivity growth above and beyond its traditional role of capital deepening.

Table 6 displays our estimates along with a few other estimates of embodied technological change found in the literature. Our estimates are consistent with those of other studies that use a production-side approach. However, they are much higher than those based on Gordon's (1990) equipment price deflators such as the 5% estimate obtained by Cummins and Violante (2002). Though we are not able to rule out the possibility of an upward simultaneity bias in our estimates, they are consistent with arguments made by Gordon (1996), Gort and Wall

the 14.7 percent estimate in Row A of Table 6 from the 4.4 percent estimate in Row C.

(1998), and others who argue that the price-based estimates are likely to understate true embodied technological change. For one thing, Gordon's quality-adjusted price indices cover only a subset of all types of equipment, hence an aggregate price index based on these indices may not capture all equipment quality change. Furthermore, Gordon (1996) points out that movements in the price ratio of investment to consumption can be caused not only by embodied technological change but also by substitution effects. An alternative approach based on estimation of an optimizing structural model (Hobijn, 2000 and Hobijn et al., 2002) does not seem to tilt the balance either way. Clearly, further research is important in refining our estimates. The potential gain is immense as accurate descriptions of productivity growth rest on our understanding of the extent of capital quality growth.

REFERENCES

- Arrow, Kenneth. "The Economic Implications of Learning by Doing." *Review of Economic Studies*, 29(3), June 1962, pp. 155-173.
- Bahk, Byong-Hyong, and Gort, Michael. "Decomposing Learning by Doing in New Plants." *Journal of Political Economy*, 101(4), August 1993, pp. 561-583.
- Baily, Martin; Bartelsman, Eric; and Haltiwanger, John. "Labor Productivity: Structural Change and Cyclical Dynamics." *Review of Economics and Statistics*, August 2001, 83(3): 420-33.
- Bartelsman, Eric, Caballero, Ricardo, and Lyons, Richard. "Customer- and Supplier-Driven Externalities." *American Economic Review*, September 1994, 84(4), pp. 1075-84.
- Bartelsman, Eric J., and Gray, Wayne. "The NBER Manufacturing Productivity Database." NBER Technical Working Paper 205, October 1996.
- Basu, Susanto and Fernald, John. "Why is Productivity Procyclical? Why Do We Care?" Board of Governors of the Federal Reserve System, IFDP no. 638, June 1999.
- Burnside, Craig; Eichenbaum, Martin; and Rebelo, Sergio. "Capital Utilization and Returns to Scale." In *NBER Macroeconomics Annual 1995*. Editors Julio J. Rotemberg and Ben S. Bernanke. Cambridge, MA: MIT Press, 1995.
- Cooper, Russell; Haltiwanger, John; and Power, Laura. "Machine Replacement and the Business Cycle: Lumps and Bumps." *American Economic Review*, December 1999, 89(4), pp. 921-46.
- Cummins, Jason; and Violante, Giovanni L. "Investment-Specific Technical Change in the United States (1947-2000): Measurement and Macroeconomic Consequences." *Review of Economic Dynamics*, 5(2), April 2002, pp. 243-84.
- Davis, Steven; Haltiwanger, John; and Schuh, Scott. *Job Creation and Destruction*. Cambridge, MA: MIT Press, 1996.
- Doms, Mark, and Timothy Dunne. "Capital Adjustment Patterns in Manufacturing Plants." Center for Economic Studies Discussion Paper 94-11, August 1994.
- Eaton, Jonathan and Samuel Kortum. "Trade in Capital Goods," NBER working paper # 8070, January 2001.
- Fisher, Franklin M. "Embodied Technical Change and the Existence of an Aggregate Capital Stock." *Review of Economic Studies*, 32, October 1965, pp. 263-88.
- Gordon, Robert J. *The Measurement of Durable Goods Prices*. Chicago: University of Chicago Press, 1990.
- Gordon, Robert J. "Can Technology Improvements Cause Productivity Slowdowns? Comment." In *NBER Macroeconomics Annual 1996*. Editors Julio J. Rotemberg and Ben S. Bernanke. Cambridge, MA: MIT Press, 1996.
- Gort, Michael, and Richard Wall. "Obsolescence, Input Augmentation, and Growth Accounting." *European Economic Review*, 42, 1998, pp. 1653-65.
- Greenwood, Jeremy; Hercowitz, Zvi; and Krusell, Per. "Long-Run Implications of Investment-Specific Technological Change." *American Economic Review*, June 1997, 87(3), pp. 342-362.
- Greenwood, Jeremy; Hercowitz, Zvi; and Krusell, Per. "The Role of Investment-specific

- Technological Change in the Business Cycle.” *European Economic Review*, 44(2), April 2000, pp. 91-115.
- Greenwood, Jeremy, and Jovanovic, Boyan. “Accounting for Growth.” Mimeo, June 1998.
- Greenwood, Jeremy, and Mehmet Yorukoglu. “1974.” *Carnegie-Rochester Conference Series on Public Policy*, 46(0), 1997, pp. 49-95.
- Griliches, Zvi, and Mairesse, Jacques. “Production Functions: The Search for Identification.” NBER Working Paper No. 5067, March 1995.
- Griliches, Zvi, and Ringstad, V. *Economies of Scale and the Form of the Production Function*. Amsterdam: North-Holland, 1971.
- Hercowitz, Zvi. “The ‘embodiment’ hypothesis.” *Journal of Monetary Economics*, 41(1), February 1998, pp. 217-24.
- Hobijn, Bart. “Identifying Sources of Growth,” mimeo, Federal Reserve Bank of New York, September 2000.
- Hobijn, Bart, Ana Maria Oviedo and Ashwin Vasan. “Embodied Technological Change in U.S. Manufacturing,” mimeo, Federal Reserve Bank of New York, December 2002.
- Hornstein, Andreas, and Krusell, Per. “Can technology improvements cause productivity slowdowns?” In *NBER Macroeconomics Annual 1996*. Editors Julio J. Rotemberg and Ben S. Bernanke. Cambridge, MA: MIT Press, 1996.
- Huggett, Mark and Sandra Ospina. “Does productivity growth fall after the adoption of new technology?” *Journal of Monetary Economics*, August 2001, 48, pp. 173-195.
- Hulten, Charles. “Growth Accounting when Technical Change is Embodied in Capital.” *American Economic Review*, September 1992, 82(4), pp. 964-980.
- Hulten, Charles and Frank Wykoff. “The Measurement of Economic Depreciation.” In *Depreciation, Inflation, and the Taxation of Income From Capital*. Editor Charles Hulten, pp. 81-125. Washington, D.C.: Urban Institute, 1981.
- Jorgenson, Dale W., and Griliches, Zvi. “The Explanation of Productivity Change.” *Review of Economic Studies*, 1967, 34, pp. 249-83.
- Johansen, Leif. “Substitution versus Fixed Proportion Coefficients in the Theory of Economic Growth,” *Econometrica*, 1944, 12(3/4), pp. 143-205.
- Jovanovic, Boyan and Rafael Rob. “Solow vs. Solow,” 1999, mimeo, NYU.
- Marschak, Jacob, and William Andrews. “Random Simultaneous Equations and the Theory of Production.” *Econometrica*, 1944, 12(3/4), pp. 143-205.
- McHugh, Richard and Julia Lane, “The Age of Capital, the Age of Utilized Capital, and Tests of the Embodiment Hypothesis,” *The Review of Economics and Statistics*, May 1987, v. 69, no. 2, pp. 362-67.
- Mohr, Michael, and Gilbert, Charles. “Capital Stock Estimates for Manufacturing Industries: Methods and Data.” Board of Governors of the Federal Reserve System, March 1996.
- Nelson, Richard. “Aggregate Production Functions and Medium-Range Growth Projections.” *American Economic Review*, September 1964, 54(5), pp. 575-606.
- Olley, Steven and Pakes, Ariel. “The Dynamics of Productivity in the Telecommunications Equipment Industry.” *Econometrica*, 55(4), 1996, pp.1263-98.
- Petropoulos, Wendy. “Industry Productivity Dynamics and Unmeasured Capacity Utilization.” Mimeo, University of Michigan, November 1999.

- Power, Laura. "The Missing Link: Technology, Investment, and Productivity." *Review of Economics and Statistics*, May 1998, 80(2), pp. 300-313.
- Sakellaris, Plutarchos. "Patterns of Plant Adjustment," *Journal of Monetary Economics*, forthcoming.
- Shea, John. "Do Supply Curves Slope Up?" *Quarterly Journal of Economics*, 1993, 108(1), pp. 1-32.
- Solow, Robert. "Investment and Technical Progress." In Kenneth Arrow, Samuel Karlin, and Patrick Suppes, eds., *Mathematical Methods in the Social Sciences*. Stanford, CA: Stanford University Press, 1960, pp. 89-104.
- Wickens, Michael R., "Estimation of the Vintage Cobb-Douglas Production Function for the United States 1900-1960," *The Review of Economics and Statistics*, May 1970, v. 52, no.2, pp. 187-93.
- You, Jon Keun, "Embodied and Disembodied Technical Progress in the United States, 1929-1968," *The Review of Economics and Statistics*, February 1976, v. 58, no.1, pp. 123-27.

Table 1

Parameter	A	B	C	D
γ	0.077 (0.029)	0.169 (0.049)	0.116 (0.060)	0.271 (0.059)
Intercept	2.298 (0.049)	2.421 (0.054)	2.369 (0.080)	-1.183 (1.077)
Multi	0.087 (0.008)	0.079 (0.008)	0.080 (0.008)	0.080 (0.005)
β	0.344 (0.005)	0.322 (0.005)	0.319 (0.005)	0.318 (0.002)
Θ	0.545 (0.005)	0.531 (0.005)	0.529 (0.005)	0.529 (0.002)
α	0.078 (0.003)	0.108 (0.005)	0.114 (0.005)	0.114 (0.002)
η	0.008 (0.002)	0.020 (0.004)	0.018 (0.004)	0.019 (0.002)
τ		2.076 (0.091)	2.219 (0.105)	2.232 (0.050)
Adjusted R^2	0.926	0.927	0.927	0.927

A: Base regression (does not allow for capital utilization or learning-by-doing effects).

B: Adjusts for capital utilization using energy expenditures (see equation (5)).

C: Full specification (adjusts for capital utilization and includes variables capturing learning-by-doing effects associated with investment spikes).

D: Average vintage regression (see equation (8)).

Note: Regressions A-D also included year and industry dummy variables as well as industry-specific time trends. The coefficients on these are not shown in order to conserve space. They can be obtained from the authors upon request.

Table 2

Parameter	A	B	C	D
γ	0.019 (0.025)	0.153 (0.029)	0.005 (0.042)	0.245 (0.065)
Intercept	2.218 (0.047)	2.417 (0.024)	2.229 (0.034)	2.534 (0.081)
Multi	0.087 (0.008)	0.086 (0.008)	0.080 (0.008)	0.080 (0.008)
β	0.344 (0.005)	0.344 (0.005)	0.319 (0.005)	0.319 (0.005)
Θ	0.545 (0.005)	0.545 (0.005)	0.529 (0.005)	0.529 (0.005)
α	0.079 (0.003)	0.079 (0.003)	0.115 (0.005)	0.116 (0.005)
η	0.007 (0.002)	0.007 (0.002)	0.017 (0.004)	0.016 (0.004)
τ			2.219 (0.106)	2.249 (0.109)
Adjusted R ²	0.926	0.926	0.927	0.927

A: Base regression (does not allow for capital utilization or learning-by-doing effects); physical depreciation measured by BEA economic depreciation rates.

B: Base regression; no physical depreciation allowed for.

C: Full specification (adjusts for capital utilization and includes variables capturing learning-by-doing effects associated with investment spikes); physical depreciation measured by BEA economic depreciation rates.

D: Full specification; no physical depreciation allowed for.

Note: Regressions A-D include year and industry dummy variables which are not shown in order to conserve space. They can be obtained from the authors upon request.

Table 3

Fraction of BEA rates	Estimated γ	Implied Economic Depreciation Rate ¹
0	0.245 (0.065)	0.245
0.25	0.205 (0.064)	0.236
0.5	0.159 (0.061)	0.220
0.75	0.087 (0.053)	0.178
1	0.005 (0.042)	0.126

Note: All of the estimates in column 2 were obtained using the same, preferred specification: adjustment to control for unobserved utilization and spike dummies included to control for learning-by-doing effects (i.e., the same specification as was used in Table 1, Column C).

¹(Column 1)*(Average BEA depreciation rate = 0.121) + (Column 2).

Table 4

	A	B	C
γ	0.332 (0.319)	-0.134 (0.048)	0.144 (0.098)
Intercept	-0.030 (0.980)	1.519 (0.029)	2.494 (1.006)
α	0.166 (0.033)		
η	0.112 (0.030)		
Multi	-0.269 (0.047)	0.058 (0.011)	0.049 (0.010)
Scale		0.939 (0.007)	0.947 (0.006)
P-value for Test of over-identifying restrictions	1.000	1.000	1.000
Adjusted R ²	0.905	0.905	0.906

A: Regression with labor and materials elasticities measured by industry cost-shares.

B: Regression with all factor elasticities measured by industry cost-shares.

C: Average vintage regression with all factor elasticities measured by industry-level cost-shares.

Note: Regressions A-C also include year and industry dummy variables.

Table 5
Annual Growth Rate from 1972-96

<u>Type of Equipment Capital Stock</u>	<u>Our Sample</u>	<u>Aggregate Manufacturing</u>
A. This paper -- $\gamma = 11.6\%$	14.7%	14.0%
B. This paper -- $\gamma = 7.7\%$	11.2%	10.0%
C. This paper -- $\gamma = 0$	4.4%	2.1%
D. Conventional	4.2%	3.0%
E. Hulten (1992)	7.0%	6.0%
F. Greenwood, et al. (1997)	7.0%	5.7%

A -- The equipment capital stock implied by the γ found in Table 1, Column C: $\gamma=0.116$ and the FRB physical depreciation schedules and the PCE deflator are used.

B -- $\gamma = 0.077$, FRB physical depreciation, and PCE deflator.

C -- $\gamma = 0$, FRB physical depreciation, and PCE deflator.

D -- $\gamma = 0$, BEA depreciation, and FRB/PPI investment deflators.

E -- $\gamma = 0.03$, BEA depreciation, and FRB/PPI investment deflators.

F -- $\gamma = 0.03$, BEA depreciation, and deflator for PCE of nondurables and nonhousing services.

Table 6

Alternative Estimates of Embodied Technological Change in the Literature

<u>Paper</u>	<u>Sample Period</u>	<u>Estimated Rate of Embodied Technological Change</u>
This paper	1972-1996	8 - 17%, preferred estimate = 12%
Hobijn (2000)	1947 - 1998	4%
Hobijn, Oviedo, and Vasan (2002)	1960 - 1996	Industry estimates range from 9.5% to 32.6%
Bahk and Gort (1993)	1973-1986	15-21%
Cummins and Violante (2002), based on Gordon (1990)	1948 - 2000	4%
	1975 - 2000	5%
Hornstein and Krusell (2000), based on Gordon (1990)	1973 - 1990	4%

FIGURE 1
PCE deflator vs. average FRB equipment deflator used in our sample from 1972-96




FIGURE 2
Average Depreciation Pattern




Figure 3

