Center for Research on Economic and Social Theory
and
Department of Economics
Working Paper Series

Constant Returns and Small Markups in U.S. Manufacturing

Susanto Basu
John G. Fernald

April, 1993
Number 93-19

DEPARTMENT OF ECONOMICS
University of Michigan
Ann Arbor, Michigan 48109-1220

CONSTANT RETURNS AND SMALL MARKUPS IN U.S. MANUFACTURING

Susanto Basu
University of Michigan

and

John G. Fernald
Harvard University

April 23, 1993

Please address correspondence to:

Susanto Basu
Department of Economics
University of Michigan
611 Tappan Street
Ann Arbor MI 48109-1220.

Abstract

We reexamine Robert Hall's recent finding of increasing returns in U.S. manufacturing. With gross output data at roughly the two-digit SIC level, we estimate that returns to scale are close to constant. We show why, with imperfect competition, value-added data lead to biased estimates of returns to scale, and hence why Hall's results are easily explained. We show how to control for value-added bias by combining Hall's data with data on intermediate input use: using Hall's data with this correction, returns to scale again appear to be constant. We also estimate that the average markup of price over marginal cost is about 15 percent.

This paper takes Robert Hall seriously.

In a series of influential and controversial papers, Hall (1986, 1988, 1990) interprets the observed procyclicality of the Solow productivity residual as an indication of markup pricing, increasing returns to scale, or both. In his 1986 and 1988 papers, he proposes a model where firms produce with constant returns to scale but price output above marginal cost. However, this line of work implies that the procyclicality of the Solow residual should disappear if one uses cost shares instead of revenue shares to weight input growth, and also predicts that firms make large pure profits. Since neither appears to be true in the U.S. manufacturing data that Hall (1990) uses, he reinterprets his work as evidence of increasing returns to scale. Under this interpretation, increasing returns and markup pricing largely offset each other, leaving little pure profit.

In this paper, we reexamine Hall's finding of increasing returns in two-digit U.S. manufacturing, reaching conclusions very different from his. Hall uses industry data on real value added and primary inputs of capital and labor. We show why value-added data are inappropriate for estimating returns to scale, and hence why Hall's procedure can lead to a spurious finding of increasing returns. We, by contrast, apply his technique to data on industry gross output, and find that U.S. manufacturing is characterized, not by increasing returns, but by approximately constant returns. When we combine Hall's data with data on intermediate input use, we can correct for value-added bias, and recover the true degree of returns to scale. Even Hall's data, suitably augmented, imply that there are constant returns. In fact, our estimates with both gross-output and value-added data imply that the model proposed by Hall (1988) is roughly correct: we find that constant returns to scale and markup pricing can explain much of the short-run behavior of the Solow residual. But our estimates of the markup are much smaller than those found by Hall and others; thus, we do not face the problem of reconciling enormous markups with low profit rates.

We take seriously Hall's hypothesis that output is produced by sectors that may have non-constant returns and may price their output above its marginal cost of production. But when considering issues of increasing returns and markup pricing, one must be careful to specify the appropriate concept of output.

We start from a model where Hall's hypothesis applies to the production of gross output. We view gross output as the natural primitive concept of production. Firms produce gross output not real value added, which is an economic index number without physical interpretation. Similarly, firms sell gross output (with or without a markup); they do not sell real value added. Starting from Hall's model of production and pricing, now applied to gross output, we then derive analytically the relationship between the true degree of returns to scale of the gross output production function and returns to scale as estimated from value-added data.

Intuitively, the reasons these differ are exactly the reasons proposed by Hall to explain the lack of invariance of the Solow residual--increasing returns and markups. Real value added is constructed from gross output by subtracting the contribution of intermediate inputs from gross production. To do this, one must know or infer the marginal product of intermediate inputs. Real value added is constructed assuming that this marginal product is observable from factor payments to intermediate goods. But this assumption is violated under precisely the conditions that Hall claims are important and for which he is testing, namely markup pricing and increasing returns. As long as Hall confines his work to testing for the failure of the joint hypothesis of perfect competition and constant returns, he is justified in using value-added data: these data are correctly constructed under the null hypothesis. But it is inappropriate and misleading to use value-added data to estimate the degree of increasing returns under the maintained hypothesis that the data are generated by an imperfectly competitive sector. By construction, these data cannot (except in special cases) give an unbiased estimate of the degree of sectoral returns to scale. Hall's finding of increasing returns is an artifact of his use of data that his own analysis holds to be unsuitable.

Our critique differs considerably from other objections to Hall's work. Various authors have attacked Hall's approach because they claim it neglects other important sources of procyclical productivity. We briefly review three that are frequently mentioned, in order to show how our approach differs.

First, an early criticism of Hall's work made by Abbott, Griliches, and Hausman (1988) is that his results reflect cyclical measurement error. This is the standard interpretation of procyclical productivity and Okun's law: labor hoarding and procyclical worker effort combine to make labor input more procyclical than hours worked. Since official statistics do not reflect the increase in work effort, productivity appears procyclical. The issue of cyclical measurement error has recently been comprehensively examined by Gordon (1992). He reviews the various sources of measurement error that might lead to a spurious finding of procyclical productivity and concludes that a small amount of each is sufficient to explain Hall's result.

Second, Hall's work has also been attacked for neglecting the distinction between short- and long-run returns to scale. As explained by Berndt and Fuss (1986) and others, divergence between the two is caused by the presence of quasi-fixed factors. When factors of production, such as capital, are sunk in the short run, measured productivity growth as conventionally calculated is biased. The bias comes from the fact that the shadow value of the sunk factors, which are the right prices for calculating the shares of these factors in production, in general diverge from their long-run user costs.

Third, a recent line of work reinterprets Hall's findings of increasing returns to scale. Caballero and Lyons (1992) and Bartelsman, Caballero, and Lyons (1991) argue that Hall's finding of increasing returns at the industry level results from an omitted variable: productive externalities from aggregate output to the productivity of individual sectors. Their estimates imply that individual sectors have constant or diminishing returns to scale, but the economy as a whole enjoys increasing returns. In recent work, Hall (1991) interprets these findings as evidence for "thick market" or "agglomeration" effects. We have shown in other work (Basu and Fernald, 1993) that the statistical model estimated by Caballero and Lyons is subject to the same data difficulties that we identify here.

These various attacks or reinterpretations of Hall's work basically attribute his results to specification error. They do not dispute Hall's findings but rather his interpretation, arguing that his results have economic explanations other than increasing returns. By contrast, we work under the hypothesis that Hall's basic procedure is exactly correct. We show, however, that his method for measuring the degree of returns to scale is flawed by his use of incorrect data. As we demonstrate, in the presence of imperfect competition, Hall's procedure is in general biased, and may plausibly lead us to find increasing returns even when there are none. Controlling for the biases, Hall's findings disappear.

The paper is structured in four sections. First, we briefly review Hall's basic method, and we show why Hall's use of value-added data biases his estimates of returns to scale. Second, we discuss the data we use and why they are appropriate for estimating returns to scale. Third, we reestimate Hall's equations and show that Hall's findings are in fact biased in exactly the way that theory predicts. Fourth, we conclude with a summary and thoughts on directions for future research.

I. Method for Estimating Returns to Scale and Markups

In this section, we begin by reviewing the theory behind Hall's (1990) method of estimating returns to scale. This method essentially involves regressing output growth on a cost-weighted sum of the input growth rates.

We then discuss the relationship between estimates of returns to scale using gross-output data and using value-added data. We show that with imperfect competition, intermediate input use directly shifts a value-added function. This occurs because measures of real value added are calculated by subtracting the productive contribution of intermediate inputs from gross output, assuming that these intermediate inputs are paid the value of their marginal products. With imperfect competition, intermediate inputs are paid less than the value of their marginal product. As a result, some of the productive contribution of intermediate goods is incorrectly attributed to real value added.

After we discuss the biases that result from attempting to estimate returns to scale using data on value added and primary inputs only, we then show how one can use data on intermediate input use to control for these biases. An important point is that one requires the same data to estimate returns to scale using value-added output as one requires using gross output.

We begin with the following production function for an industry:

$$Y = F(K, L, M, T). \tag{1}$$

Y is gross output, not value added. K and L are primary inputs of capital and labor, while M is intermediate inputs of energy and materials. T is an index of the state of technology. We assume that the production function is homogeneous of degree γ in capital, labor, and intermediate goods, and that there are no fixed factors of production.

Differentiating the production function (1), we can express the growth rate of output as:

$$dy = \left(\frac{F_K K}{Y}\right)dk + \left(\frac{F_L L}{Y}\right)dl + \left(\frac{F_M M}{Y}\right)dm + dt. \tag{2}$$

Lower-case letters represent logs of their upper-case counterparts, so all of the quantity variables in (2) are log differences, or growth rates. For convenience, we normalize to unity the elasticity of output with respect to technology T. The sum of the output elasticities equals the degree of returns to scale γ, so that

$$\gamma = \left(\frac{F_K K}{Y}\right) + \left(\frac{F_L L}{Y}\right) + \left(\frac{F_M M}{Y}\right). \tag{3}$$

We allow firms to have some degree of monopoly power in the goods markets, though we assume that they are price takers in factor markets. We assume also that firms act as if they face a sequence of one-period static problems; this abstracts, for example, from any considerations of dynamic monopoly or investment behavior. Under these assumptions, the first-order conditions for profit maximization imply that the elasticity of output with respect to any factor J equals a markup μ multiplied by the share of that input in total revenue:

$$\left[\frac{F_J J}{Y}\right] = \mu\left[\frac{P_J J}{PY}\right] \qquad J = K,\ L,\ M. \tag{4}$$

P is the price of gross output, F_J is the marginal product of input J, and P_J is the price of the Jth input as perceived by the firm. Note that the price of capital, P_K, must be defined as the rental cost of capital. It does not include possible profits, which generally are also payments to capital. With perfect competition, where $\mu = 1$, equation (4) just states that the elasticity of output with respect to any input equals the input's share in revenue. Under imperfect competition, the elasticity of output exceeds the revenue share.

Equations (3) and (4) together imply that

$$\gamma = \mu\left[\frac{P_K K + P_L L + P_M M}{PY}\right]. \tag{5}$$

We can use equations (4) and (5) to eliminate the markup, μ, from the expression for the total differential of output. We find

$$\begin{aligned} dy &= \gamma\cdot\left[c_L\,dl + c_K\,dk + (1 - c_L - c_K)\,dm\right] + dt \\ &\equiv \gamma\cdot dx + dt. \end{aligned} \tag{6}$$

The shares of capital and labor inputs in total costs are c_L and c_K. dx is a cost-weighted sum of the growth rates of the various inputs. Intuitively, equation (6) says that the growth rate of output equals the growth of inputs multiplied by the degree of returns to scale, plus the contribution of productivity shocks. If there are constant returns and perfect competition--so the cost shares are the revenue shares--then equation (6) is just the standard equation defining the Solow residual.

As noted above, we wish to allow for the possibility that firms make positive profits. Hence, we follow Hall (1990) by computing the required cost of capital. We create a required rate of return series for capital and compute the rental price of capital as the product of this series with the value of the capital stock. We weight factors by their cost shares rather than revenue shares, since as Hall (1988) shows, with markup pricing the measured Solow residual is procyclical, even with constant returns and no change in technology.

Note that in the production function (1) the measure of output is gross output, which explicitly includes the contribution of intermediate goods to production. From the perspective of a firm or an industry, all inputs are symmetric: the firm faces the possibility of substituting among <u>all</u> factors of production--capital, labor, and intermediate inputs of energy and materials--to produce gross output. A number of researchers (including Hall) instead use NIPA data on real value added in order to estimate markups and returns to scale. Measures of value added attempt in some way to subtract from gross output the productive contribution of intermediate goods. This generates a measure of "net output" which, it is hoped, does not depend at all on intermediate inputs. That is, this approach assumes that we can write the value-added function as

$$V = G(K, L, T). \tag{7}$$

We now show, however, that in the presence of markups, intermediate input use directly shifts this value-added function. As a result, estimates of returns to scale from value-added data are in general biased.

The growth rate of the Divisia index of value added is defined as[1]

$$\frac{dV}{V} \equiv dv = \frac{1}{1-s_M}dy - \frac{s_M}{1-s_M}dm, \tag{8}$$

where s_M is the share of materials in revenue. Substituting from equation (2) for dy, we find:

$$dv = \frac{1}{(1-s_M)}\left[\left(\frac{F_K K}{Y}\right)dk + \left(\frac{F_L L}{Y}\right)dl\right] + \left[\left(\frac{F_M M}{Y}\right) - s_M\right]\left[\frac{1}{1-s_M}\right]dm + \frac{dt}{1-s_M}. \tag{9}$$

Substituting the first-order condition for intermediate input use, and combining capital and labor inputs into a cost-weighted index dx^V, we find:

$$dv = \frac{1}{(1-s_M)}\left[\frac{F_K K+F_L L}{Y}\right]\left[\left(\frac{F_K K}{F_K K+F_L L}\right)dk + \left(\frac{F_L L}{F_K K+F_L L}\right)dl\right] + (\mu-1)\left(\frac{s_M}{1-s_M}\right)dm + \frac{dt}{1-s_M}, \tag{10}$$

or,

$$dv = \gamma\cdot\left[\frac{1-\frac{\mu}{\gamma}s_M}{1-s_M}\right]dx^V + (\mu-1)\left(\frac{s_M}{1-s_M}\right)dm + \frac{dt}{1-s_M}. \tag{11}$$

Equation (11) provides the basis for our discussion of what it is that value added measures. We see that the growth rate of the index of value added depends on returns to scale, markups, and intermediate

[1] See Sato (1977) or Arrow (1974).

input growth, in addition to technological progress and the growth of primary inputs. If there are constant returns to scale and no markups, then the growth rate of value added equals the growth rate of primary inputs plus technological progress. That is,

$$dv = dx^V + \frac{dt}{(1-s_M)}. \tag{12}$$

One implication of equation (12) is that if s_M is constant, as Hall (1990) notes, under competition and constant returns the productivity residual ($dv - dx^V$) calculated from a Divisia index of value added is uncorrelated with any variable that neither causes productivity shifts, nor is caused by productivity shifts. In addition, dx^V can be calculated with either cost or revenue shares, since there are no profits.

In the presence of markups, however, equation (11) shows that the growth of materials and energy directly shifts the value-added production function. Intuitively, value added is calculated by subtracting from gross output the productive contribution of intermediate goods, assuming that the elasticity of output with respect to intermediate inputs equals its revenue share. With markups, this elasticity exceeds its revenue share. Hence, some of the contribution of materials and energy is attributed to value added. As a result, the value-added productivity residual is correlated with any variable that is correlated with intermediate-goods use, regardless of whether it is correlated with technology. Thus, Hall's argument that under constant returns the cost-based Solow residual should be invariant is true in gross-output data, but, apart from special cases, is not true with value-added data. In particular, if markups exceed 1, a rise in energy prices, which causes energy use to fall, shifts the value-added function downwards.

Suppose we attempt to estimate returns to scale using value-added data. There are two sources of bias, one upward and one downward.

First, if the markup exceeds 1 and we do not control for intermediate-input use, then the estimate of γ suffers from omitted-variable bias. Because intermediate input use is positively correlated with primary input use, this omitted variable bias is positive.

Second, consider estimates of returns to scale, γ^*, defined as

$$\gamma^* = \gamma \cdot \left[\frac{1 - \frac{\mu}{\gamma} s_M}{1 - s_M} \right]. \tag{13}$$

If markups μ exceed returns to scale γ--that is, if there are positive economic profits--then γ^* is less than γ. Hence, controlling for changes in intermediate-input use, estimates of returns to scale with value-added data are biased downwards.

With imperfect competition, it is thus difficult to interpret results using value-added data. In particular, Hall's parameter estimates of returns to scale are largely meaningless. These estimates can be biased either up or down.[2]

If we have data on the price and quantity of intermediate input use, however, we can recover the underlying gross-output parameters. Rearranging equation (11), we find:

$$dv + \left(\frac{s_M}{1-s_M}\right) dm = \gamma \left(\frac{dx^v}{1-s_M}\right) + \mu \left(\frac{s_M}{1-s_M}\right)(dm - dx^v) + \frac{dt}{1-s_M}. \tag{14}$$

We estimate this equation in Section III.

So far, our discussion of why value-added results are difficult to interpret has assumed that value added is calculated as a Divisia index. This assumes that the weights on output and intermediate inputs are adjusted continuously. The NIPA data instead uses the method of double-deflation to estimate real value added: gross output and intermediate inputs are deflated separately. If we normalize the base-year prices of gross output and intermediate inputs to one, the double-deflated estimate of real value added, V^{DD}, is

$$V^{DD} = Y - M. \tag{15}$$

Let n equal M/Y, the share of output going to materials in base-year prices. If we differentiate

[2] In one special case, where there are constant returns and dm equals dv, these biases exactly cancel, as Hall (1990) shows. This special case is obviously not relevant for estimating the degree of returns to scale.

equation (15), we find that the growth rate of double-deflated value added is:

$$dv^{DD} = \frac{1}{1-n} dy - \frac{n}{1-n} dm. \tag{16}$$

Note that the double-deflated index (16) differs from the Divisia index (8) by using weights computed with constant prices, rather than current prices.

With some algebraic manipulation, we can show that the growth rate of double-deflated value added equals the growth rate of the Divisia index, plus an additional term:

$$dv^{DD} = dv + \frac{n}{(1-n)(1-s_M)} \left(1 - \frac{P_M}{P}\right)(dy - dm). \tag{17}$$

The second term on the right-hand-side of equation (17) is the double-deflation bias. As Bruno (1978) and Bruno and Sachs (1985) point out, this bias term disappears in two special cases. First, it disappears if intermediate inputs grow at the same rate as gross output. For this to be true in general, the production technology must be Leontief between intermediate inputs and value added, so that intermediate inputs are used in fixed proportion to output. Second, the bias term disappears if the price of intermediate inputs relative to the price of output is always equal to one. This is the case if relative prices are constant over time; it is also the case if these base-year prices are allowed to change each period, so that the double-deflated value-added index becomes essentially the Divisia, shifting-weight index.

The relative price of intermediate inputs (especially energy) is certainly not constant over time, and NIPA adjusts the base-year prices approximately every ten years. Hence, even with constant returns and competition, double-deflated value added provides an unbiased measure of productivity growth only if intermediate inputs are used in fixed proportions to output. This is not only a strong assumption to impose, it is rejected by the data.[3]

To summarize, in this section we show the relationship between estimates of returns to scale and markups using data on gross output and value added. With imperfect competition, the data required to

[3] See, for example, Jorgenson, Gollop and Fraumeni (1987).

estimate returns to scale and markups with value-added data are the same as the required with gross output data. In particular, one cannot measure returns to scale and markups using data on primary inputs of capital and labor alone. One requires data on intermediate inputs as well.

II. The Data

We combine two sources of data on inputs and output for 21 manufacturing industries, at roughly the two-digit SIC level.[4] First, we use Robert Hall's data on real value added, capital stock, and hours worked by industry.[5] This is a double-deflated measure of real value added from NIPA. The two input measures--the BEA net capital stock and NIPA labor hours--are not adjusted for changes in quality over time.

Second, we use unpublished data provided by Dale Jorgenson and Barbara Fraumeni on industry-level inputs and outputs. In the following paragraphs we highlight a few key features of the data; in the appendix we discuss the underlying sources and methods at greater length. For a complete description, see Jorgenson, Gollop, and Fraumeni (1987) or Jorgenson (1990).

These sectoral accounts seek to provide accounts that are, to the extent possible, consistent with the economic theory of production. Output is measured as gross output, and inputs are separated into capital, labor, energy, and materials. For our purposes, an essential aspect of the data is their inclusion of intermediate inputs. From equation (14), we can combine this information with Hall's data to assess the importance of the value-added biases that arise from omitting intermediate inputs. In addition, we can use the Jorgenson data to estimate returns to scale directly, as in equation (6).

Jorgenson's input data are available both with and without an adjustment for input quality. Without the quality adjustment, the Jorgenson measures of labor and capital input are essentially the same

---

[4] All industries are 2-digit SIC industries with the exception of transport equipment (SIC 37), which is subdivided into "motor vehicles" (SIC 371) and "other transport equipment" (SIC 372-379).

[5] Hall (1990) describes the data sources. We obtained these data from Ricardo Caballero, who uses them in Caballero and Lyons (1992).

as Hall's measures--hours worked and the capital stock.[6] From the perspective of a firm, however, the relevant measure of its input of, say, labor is not merely labor hours. The firm also cares about the relative productivity of different workers. In creating a series for labor input, Jorgenson, Gollop, and Fraumeni assume that wages are proportional to marginal products. This allows them, in essence, to calculate quality-adjusted labor input by weighting the hours worked by different types of workers (distinguished by various demographic and occupational characteristics) by relative wage rates. Note that the validity of this procedure is unaffected by markup pricing, because markups affect the level of wages, not relative wages. Hence, labor input can increase either because the number of hours worked increases, or because the "quality" of those hours increases. Similarly, Jorgenson, Gollop and Fraumeni adjust inputs of capital and intermediate goods for changes in quality.

In order to allow a direct comparison between results using Hall's and using Jorgenson's data, we use Jorgenson's data to calculate a Divisia index of real value added, as in equation (8). Divisia indices are correctly given in continuous time. Letting $\bar{s}_m$ be the average value in periods T and T-1 of the share of intermediate inputs in revenue, we approximate equation (8) in discrete time as[7]

$$Ln\left[\frac{V_i(T)}{V_i(T-1)}\right] = \frac{1}{1-\bar{s}_m} Ln\left[\frac{Y_i(T)}{Y_i(T-1)}\right] - \frac{\bar{s}_m}{1-\bar{s}_m} Ln\left[\frac{M_i(T)}{M_i(T-1)}\right]. \tag{18}$$

To estimate the payments to capital, we follow Hall and Jorgenson (1967), Hall (1986, 1990), and Caballero and Lyons (1992), and compute a series for the user cost of capital r. The required payment for any type of capital is then $rP^K K$, where $P^K K$ is the current-dollar value of the stock of this type of

---

[6] The capital stock data differs in two ways. First, Jorgenson and Fraumeni include land and inventories. Second, in creating stock estimates from the underlying investment data, Jorgenson and Fraumeni assume that the underlying assets depreciate geometrically. The BEA net capital series that Hall uses assumes straight-line depreciation over the estimated service life of the asset, where the service life follows a modified Winfrey distribution.

[7] This is the most common method used for approximating Divisia indices in discrete time. Diewert (1976) provides a justification for this by showing that it is an exact index for a translog production function, which in turn provides a flexible second-order approximation to any production function.

capital.[8] In each sector, we use data on the current value of the 51 types of capital, plus land and inventories, distinguished by the BEA in constructing the national product accounts. Hence, for each of these 53 assets, we compute the user cost of capital as

$$r_s = (\rho + \delta_s)\frac{(1 - ITC - \tau d)}{(1 - \tau)}, \quad s = 1 \ to \ 53. \tag{19}$$

ρ is the required rate of return on capital, and δ_s is the depreciation rate for this asset.[9] ITC is the investment tax credit, τ is the corporate tax rate, and d is the present value of depreciation allowances. We assume that the required return ρ equals the dividend yield on the S&P 500. We obtained data on ITC, τ, and d from Alan Auerbach.

Because input use is likely to be correlated with technology shocks, we seek demand-side instruments for input use. We use the instruments recommended by Caballero and Lyons (1992): the growth rate of the price of oil deflated by the price of manufacturing durables; the growth rate of the price of oil deflated by the price of manufacturing nondurables; the growth rate of real government defence spending; and the political party of the President. We use the current value of each instrument, as well as one lag of defence spending and the party of the President.

III. Results.

Our results confirm both the importance of using gross-output data and the details of the biases from using value added that we discuss above. To summarize our results, using the theoretically-preferred quality-adjusted, gross-output data we find that there is no evidence for increasing returns. Hall's findings are an artifact of his use of value-added, non-quality-adjusted data and, in an odd way, his use of instruments.

---

8 Note that the Berndt and Fuss (1986) critique holds that instead of using p^K, the price of investment goods, one should use tax-adjusted q, the shadow value of installed capital.

9 The 51 types of capital, along with their estimated depreciation rates, are listed in Jorgenson (1990), Table 3.6. We assume land does not depreciate, while inventories depreciate at a 25 percent rate.

The most important results are reported in Table 1. The first line of Table 1 shows estimates of returns to scale from Hall's basic regression run with his data. For his output data, Hall uses the NIPA series on real value added. The second line reports estimates of returns to scale using our constructed Divisia index of value added. We use these data both with and without quality adjustment of the input series. The third line does the same, but now with the Jorgenson gross-output data. In all cases, the estimation technique is SUR, with the degree of returns to scale constrained to be equal in the 21 industries. All regressions include unreported constants for each industry. The sample period is 1959 to 1984.

Using the instruments described in Section II (essentially the instruments Hall recommends) and using Hall's data set, we find strong evidence of increasing returns to scale. The estimate of γ is 1.17, with a standard error of 0.03, so we can reject the hypothesis that $\gamma = 1$ at any reasonable level of significance. Interestingly, however, the uninstrumented regression estimates γ to be 0.89 with a standard error of 0.01--equally strong evidence of decreasing returns to scale. These results contrast with those obtained using our preferred data set: we find the OLS estimate of returns to scale using gross-output data substantially exceeds the instrumental variables estimate of 1.03.

This pattern might at first seem puzzling. It is easily explained, however, by the value-added biases that we identify in Section I. As we note in the discussion following equation (13), there are two biases from using value-added data: the direct estimate of γ is biased down, but the indirect effect of omitting materials growth tends to bias it up. A priori we cannot say which effect will dominate. In addition, the outcome need not be the same in both the instrumented and the uninstrumented cases. Indeed, we might expect that the second effect will be stronger when we instrument with oil prices, because oil prices are an "instrument" that is very strongly correlated with the error term (including the omitted variable of materials growth).

In fact, as the following example demonstrates, the instrumented results are likely to show higher returns to scale than the uninstrumented ones under fairly general conditions. Consider the case of an industry which is small in its factor markets, so that the change in the relative price of factors is exogenous noise (possibly with a non-zero mean). In periods when there are technology shocks, the industry can produce more output without needing a proportionate increase in input. The correlation between technology shocks and input use is in general positive (although as we discuss below, this does depend on the demand elasticity). Nevertheless, the correlation need not be high.

In the limiting case where the price elasticity of demand approaches 1, there is no correlation between technology shocks and input use.[10] In this case, assuming that technical progress is Hicks-neutral and the production function is homothetic, periods in which the change in output results solely from supply shocks also show no correlation between changes in capital and labor input and changes in materials inputs. So in the admittedly special case of unit elasticity of demand, when output movements result from supply shocks there is no upward omitted-variable bias on estimates of returns to scale in the uninstrumented, value-added regression. There is only the second, downward bias to the estimated returns to scale.

On the other hand, the instruments are designed to isolate periods when output movements are driven by demand shocks. At these times, since there are no technology shocks, increases in output can come only from increases in capital, labor, and materials. Since by homotheticity the use of all three inputs generally moves together, there will be a strong correlation between increases in capital and labor input and the omitted variable, the increase in materials input. In this example, the correlation between the first-stage predicted growth in primary inputs and materials growth (that is, the correlation between the projection of primary input use dx^V on our instruments and the growth in materials) is stronger than the correlation between actual primary input use and the composite error term, the sum of materials growth and technology shocks. Hence, if the world behaves in this fashion, we would estimate smaller returns to scale in the uninstrumented regression, where output changes result from both supply and demand

10 A one percent improvement in technology causes the price of output to fall by one percent at any given level of output. When the elasticity of demand equals 1, a one percent reduction in price stimulates a one percent increase in the quantity demanded. But since the demand shock allows the industry to produce one percent more output with the same inputs, the increase in demand is met without hiring additional factors.

shocks, than in the instrumented regression, where we consider only output movements that are due to demand shocks. Another way to put this is that the positive omitted-variable bias is larger in the instrumented than the uninstrumented regressions, and this effect can swamp the (also positive) bias caused by the endogeneity of factor inputs.

We have discussed this possibility in some depth because it is exactly the pattern that one finds. In Hall's basic value-added regressions, γ is much smaller than 1 when we do not instrument, and it becomes correspondingly larger than 1 when we instrument. The pattern (though not the exact magnitudes) are the same for the other regressions using value-added data. This result is puzzling if one thinks of the instruments as being used to solve the "transmission problem" of simultaneity between productivity shocks and factor demands. One expects higher productivity to lead to increased demand for factors. Now it is theoretically possible for the OLS bias to be negative: an industry with an inelastic demand curve reacts to a positive, neutral, productivity shock by decreasing the use of all its inputs. But in the benchmark case of imperfect competition that Hall considers--static monopolistic or monopolistically competitive behavior--the price elasticity of demand at equilibrium must be greater than 1. (That is why in our example above we refer to unit elasticity as a limiting case.) In that case, a positive technology shock does result in greater input use.

Even if, for some reason, the equilibrium elasticity of demand were less than 1, we should expect to see the instrumented estimates also exceed the uninstrumented ones when we estimate the same equation with the gross-output data. Here, by contrast, we see the opposite and expected pattern for both the quality-adjusted and non-quality-adjusted regressions: the OLS estimates are uniformly larger than the 3SLS estimates. So we interpret the unusual behavior of the instrumented estimates in the value-added regressions as another indication of the flaws of value-added data.[11]

11 We conjectured that the odd behavior of the instrumented results comes from the use of oil prices as an instrument. However, we find the same pattern of higher value-added estimates when we use only the defense spending and the party in power instruments. The higher correlation between the instrumented right hand side variable and the composite error term in value-added data does not seem to be driven primarily by our choice of instruments but rather by the very fact of instrumenting.

Our discussion of the biases from double-deflation is also relevant for interpreting the results of Table 1. We see that the estimates using Hall's data are different from those obtained with the non-quality-adjusted inputs and value-added data constructed as a Divisia index. For example, the uninstrumented estimate of γ using Hall's data is 0.89. Our non-quality-adjusted data gives an estimate of γ of 1.04. The significant difference between the two sets of output series is that our value-added data are correctly computed as a Divisia index, while the NIPA estimates of real value added are double-deflated, with base prices that are changed about once a decade. As we point out, double-deflated value added has all the biases of a Divisia index of value added, plus an additional one. The extra bias comes from the fact that the relative price of materials may change over time, thus changing the materials/output ratio. Since one of the materials is energy and there have been large changes in energy prices over the sample period, the double-deflation bias is potentially large. The fact that there are significant differences between the results obtained using Hall's NIPA data and the correctly constructed value-added data indicate that this bias is quite important.[12]

In order to confirm that the failure of value-added data is due to the causes we identify in equation (11), we added the growth of materials use as an additional right hand side variable in the value-added regressions. If value added is correctly measured then materials use should not enter significantly in this regression: the whole idea of value-added data is to create a concept of real output that is independent of materials input, and depends only on capital and labor used in the sector. The results are presented in Table 2. We decisively reject the hypothesis that value added is correctly measured. In all our regressions, materials use enters significantly and with a large positive coefficient (large relative to the estimated returns to scale). At the same time, we see that our prediction about the degree of returns to

[12] There are small differences between the right hand side variables used by Hall and those that we use in the non-quality-adjusted regressions. For example, in his calculation of the required rate of return to capital Hall assumes an uniform capital depreciation rate of 12.7% for all sectors. We use sector-specific depreciation rates in constructing each sector's required rate of return. In order to eliminate this source of difference between the two sets of estimates, we also regressed our value-added variable on Hall's cost-weighted inputs. We obtained results that are very similar to the ones we report, thus confirming that the difference between the two sets of estimates results from the differences in the output series.

scale is confirmed. Once we include materials use in the regression, the estimate of γ that we obtain from the value-added regressions is always biased down relative to its true value as obtained from our preferred gross output regressions. That is, as we predicted in our discussion following equation (13), γ^* is less than γ.

We can use equation (14) to estimate γ and the markup, μ, by combining data on value added and primary inputs with data on materials input. The estimates of equation (14) are found in Table 3. The first set of regressions combine Hall's data with Jorgenson's data on the intermediate input share and growth rate. The second set uses our constructed Divisia index along with this intermediate input data. The results are striking. In all cases the estimated returns to scale are economically indistinguishable from constant returns. In all but one case the estimates are also not statistically different from 1. This is exactly the result we obtain with gross-output data. The only puzzle remaining is why the estimate of returns to scale using non-quality-adjusted data is not significantly larger than its counterpart obtained from the quality-adjusted data. We expect this pattern a priori, and we do find it when we estimate γ alone in Table 1.

Equally interesting are our estimates of the markup. In only one case is the estimate as large as 30 percent; in all other cases it is about 20 percent or lower. In our preferred regression, using quality-adjusted data with instrumenting, we find that on average prices are only 15 percent higher than marginal cost. These figures are considerably smaller than others that have been estimated under more restrictive assumptions.

Here we find that our theoretical predictions are borne out in detail: the biases in the construction of value-added data stemming from the presence of small markups are sufficient to yield estimates of returns to scale that are greater or less than 1, even when the true degree of returns to scale is approximately constant. These results show that the different results we obtain using value-added as opposed to gross-output data stem not just from a general failure of the concept of value added, but from the specific problems that we identify in the Section I.

The final question we consider is whether we should use quality-adjusted data. We find that the difference between using quality-adjusted and non-quality-adjusted data is quantitatively important. As one predicts, returns to scale seem much larger with Hall's data and the non-quality-adjusted data than with the quality-adjusted data, either value-added or gross-output. From our discussion of the data above, it is clear that we should adjust inputs for quality; otherwise, an increase in output resulting from an increase in the general quality of inputs will be erroneously interpreted as reflecting a higher degree of returns to scale. Even with the theoretically correct gross-output data, returns to scale are significantly greater than 1 using non-quality-adjusted data. When we do not instrument, the estimate of returns to scale is 1.10. Instrumented it falls, but only to 1.07. In both cases, the estimates are significantly greater than 1.

It is only when we use the theoretically-preferred quality-adjusted gross-output data that we eliminate the evidence for increasing returns. As we noted above, the "transmission problem" operates in its expected direction with the gross output data: even with the quality-adjusted data, the uninstrumented estimate of γ is 1.07 with a standard error of 0.01, which is significantly greater than 1. But when we instrument to solve the endogeneity problem we obtain our preferred estimate of γ of 1.02--not significantly different from constant returns to scale.

Thus we see that both the flaws of value-added data and the use of quality-adjusted data are empirically important. About two-thirds of the difference between the estimated returns to scale between Hall's data and our preferred estimate comes from our use of gross output data; the remaining difference is due to the quality adjustment.

Here we see evidence that, even for a given estimation procedure, the choice of the right data is critical. We see also that there are clear theoretical grounds for choosing quality-adjusted gross-output data. Using these data, there is no evidence to support Hall's claim of increasing returns to scale.

IV. Conclusion

We draw two main conclusions, one technical and one substantive.

The technical lesson is that value-added data are inappropriate for use under the maintained hypothesis of non-constant returns to scale and markup pricing. Value added is not an atheoretic construct; it is an economic concept whose measurement assumes perfect competition and constant returns. Since the construction of value added is based on these two assumptions, it is not suitable for measuring the extent of departures from these benchmarks. This point is a subtle one, because value-added data (if calculated as a Divisia index) are appropriate for testing the joint hypothesis of perfect competition and constant returns: in that case, these assumptions are part of the null hypothesis.

We have identified three quantitatively important failures of value-added data. The first is a general point; the second and third are specific to the commonly used NIPA data.

Our first point is that estimates of returns to scale using value-added data can be biased up or down. We provide examples and empirical results that indicate both outcomes are possible, and under certain conditions may even be likely. Our empirical results show that we have correctly identified the sources of the failures of value-added data. In fact, we can even correct for these failures. But this is a Pyrrhic victory. The information needed to fix the problems of value added is exactly the information needed to estimate using gross-output data in the first place, so one might as well use the theoretically preferable gross output from the start. Regardless of what one thinks of Hall's (1990) estimation procedure as a method for detecting increasing returns, it is sobering to realize that a choice between different concepts of output can produce such large differences in results, even with identical estimation techniques. We hope our results stimulate other researchers to think about the appropriateness of value added for their empirical work, and to substitute gross output in situations where value added is unsuitable.

Second, we show that double deflation can lead to substantial biases. It has long been known that value added calculated as a Divisia index has better index number properties than double-deflated value added. Here we show that this distinction is quantitatively important; we obtain quite different results using these two different measures of value added.

Third, we also show that whether or not the input series are adjusted for changes in quality can produce significant changes in the results. Quality-adjusted data are always preferable for the reasons we have discussed; using these data we find substantially lower estimates of the degree of returns to scale.

These technical points should not obscure the importance of our substantive results. Correcting for the three problems we identify above, we find that U.S. manufacturing industries are characterized by essentially constant returns to scale and a mild degree of imperfect competition, leading to markups of price over marginal cost of about 15 percent. These results paint a very different picture of the macroeconomy than that proposed by Hall (1988, 1990) and his followers, and they are important for rethinking the micro underpinnings of many macroeconomic models.

Our finding that returns to scale are approximately constant is appealing because it is consistent with the cross-sectional evidence. Baily (1990) notes the conclusion of Griliches and Ringstad (1971) that estimated firm-level returns to scale are typically in the range of 1.03-1.05. As these authors note, on a priori grounds it seems unlikely that the huge size disparities between firms that is a feature of many industries could exist if returns to scale were not roughly constant.

If production technology exhibits constant returns, however, then one is also forced to the conclusion that markups must be small. Domowitz, Hubbard, and Petersen (1988) estimate the markup using gross-output data at the four-digit level, under the maintained hypothesis of constant returns to scale. They find a considerably higher estimate of the markup: on the order of price exceeding marginal cost by 60 percent. We are puzzled by the size of their markup. It is true that our preferred estimate is obtained using quality-adjusted data. It is easy to show, however, that our much smaller estimate of the markup is also more reasonable.

A finding of markups in excess of 60 percent implies that observed labor and materials elasticities are understated by more than a third. But under constant returns, the output elasticities of the three inputs must sum to 1. Given that the sum of the observed labor and materials shares is about 0.9, this implies that the true output elasticity of capital is roughly -0.4! Even our preferred results--returns to scale of 1.02 and a markup of 1.15--imply that the true capital elasticity is about zero. This is surely too low as well, but it is a step in the right direction.

Our finding of constant returns to scale is important for papers that stress increasing returns as the source of cyclical unemployment (e.g. Weitzman, 1982) or long-run growth (e.g. Romer, 1986). They also bear on the plausibility of the coordination failure literature (Diamond, 1982; Cooper and John, 1988). These papers cite productive externalities as the source of increasing returns. Our finding that there are no increasing returns is problematic for this literature. We pursue this point more deeply in a companion paper (Basu and Fernald, 1993), and conclude that productive externalities to output in two-digit manufacturing are at most small, and quite possibly nonexistent.

That we find only small markups is also important for models of imperfect competition in product markets. In recent years, it has become an accepted "stylized fact" that U.S. industries are characterized by high markups but low rates of pure profit. The two have been reconciled by assuming that there are large fixed (overhead) costs of operation that consume much of the operating profit. But by this metric, if markups are small, fixed costs must also be small. This conclusion can have large economic significance. For example, Rotemberg and Woodford (1991) assume a high value of the average markup (on the order of the 60 percent markup over marginal cost found by Domowitz, Hubbard, and Peterson, 1988) to generate their claim of countercyclical markups. (Their empirical procedure cannot estimate the average markup but must take it as known a priori.) Our finding is that markups are closer to 10 percent, which Rotemberg and Woodford term their "extremely conservative" case. But as Ramey (1991) points out, given Rotemberg and Woodford's (1991) other findings, an average markup of 10 percent actually implies a procyclical markup. This is one example of the many issues that depend on an accurate measure of the size of markups, and hence of fixed costs.

More research is still needed to produce a tight fit between theory, data, and estimation. One major step towards that goal would be to incorporate fixed costs into the analysis in a rigorous fashion. The current empirical work focusing on increasing returns (including our own) assumes that increasing returns must take the form of declining marginal cost. The benchmark case is that of a Cobb-Douglas production function where the sum of the exponents of capital and labor exceeds one. But increasing returns are correctly defined as declining average cost. Fixed costs are one reasons why we might see increasing returns without decreasing marginal cost: with fixed costs and constant marginal cost, average costs fall throughout the range of production. Much of our intuition about the form of production technology says that fixed costs should be a major source of increasing returns. For example, fixed costs may take the form of overhead capital and labor. Given the number of non-production workers in U.S. manufacturing, overhead labor may be quantitatively significant.

Fixed costs present major problems for consistent estimation of returns to scale. For example, with fixed costs the degree of returns to scale is no longer constant. Also, if the ratio of fixed to variable inputs is not the same for different inputs, the production function is not homothetic. Nevertheless, it is clear that fixed costs must be incorporated into the analysis. It is simply inconsistent for our theory and empirical work to assume that markup profits are dissipated by fixed costs when our estimates of returns to scale and markups are calculated under the hypothesis that fixed costs do not exist. We are focusing on these and similar issues in current work.

Appendix: The Jorgenson Data Set

The data provided by Jorgenson and Fraumeni are part of Jorgenson's long-term research effort, with a number of coauthors, aimed at creating a complete set of national accounts for both inputs and outputs at the level of individual industrial sectors as well as the economy as a whole. Their purpose is to allow Jorgenson to allocate U.S. economic growth to its sources at the level of individual industries. Because the goal is to account for growth in production, these data seek to provide measures of inputs and output that are, to the extent possible, consistent with the economic theory of production. A complete description of the data is contained in Jorgenson, Gollop, and Fraumeni (1987), henceforth referred to as JGF.

Output is measured as gross output, and inputs are separated into capital, labor, energy, and materials. Each of these inputs is, in turn, a composite of many more subcategories of these inputs. For example, capital input includes the contribution of hammers, computers, office buildings, and many others forms of physical capital. For both national accounting and productivity measurement, an important issue is how these many subcategories should be combined to form a small, manageable set of categories, in this case inputs of capital, labor, energy, and materials.

Gross output by industry, in current and constant prices, comes from the Office of Economic Growth of the Bureau of Labor Statistics. The main thing to note is that output is not measured as value added. From the perspective of a firm or industry, the proper measure of output is gross output, so that intermediate inputs are treated symmetrically with primary inputs. Value added by industry, as published in the National Income and Product Accounts, is useful for national accounting, since it sums to national expenditure. It does not, in general, have an interpretation as a measure of output.

Conceptually, JGF divide labor input into hours worked, and average labor quality. NIPA provides hours worked by industry. From the point of view of a producer, however, the proper measure of labor input is not merely labor hours: firms also care about the relative productivity of different workers. Fraumeni and Jorgenson use data from household surveys to disaggregate total hours into hours worked by different types of workers, categorized by demographic variables such as sex, age, and education. JGF then assume that workers are paid proportionately to the value of their marginal products. This assumption allows them to calculate labor input as essentially a weighted sum of the hours worked by different types of workers, weighting by relative wage rates. For the economy as a whole, labor-input growth from 1947 to 1985 averaged 1.81 percent; the growth of labor hours averaged 1.18 percent, while the growth of labor quality averaged 0.63 percent.[13]

JGF measure capital input analogously to labor input, attempting to weight the input of different types of capital by relative efficiencies. The BEA provides industry data for 27 categories of producer durables, including trucks and autos, and 23 categories of nonresidential structures. In addition to these, Jorgenson and Fraumeni include data on the stock of inventories and land by industry. A simple capital stock measure, with no adjustments for quality, can be calculated as an unweighted sum of the stocks of all types of capital. As an input measure, this is analogous to labor hours. Capital input from the point of view of a producer, however, should weight the different types of capital by relative productivities. Just as we need wage rates to calculate labor input, we require rental rates to calculate capital input. These rental rates are not directly observed. JGF assume that there are either constant returns and competition or monopolistic competition, however, so that total payments to capital are observed as property compensation, a residual after all other factors have been compensated. JGF use this to back out the implied rental rates for each type of capital, based on knowledge of the stock of each type of capital and its depreciation rate, as well as tax parameters such as the corporate income tax and investment tax credits.

Note that the largest determinant of relative rental rates is differences in depreciation rates. Computers, for example, depreciate at an estimated rate of 27 percent per year; office buildings depreciate at an estimated rate of 2.5 percent per year. Hence, a one-dollar investment in computers must provide a higher flow of services than a one-dollar investment in office buildings.

The data on intermediate inputs of energy and materials is the most problematic. Conceptually,

13 Calculated from Jorgenson (1990), Table 3.1.

the task is straightforward: for each year, we observe payments to intermediate factors as the difference between nominal gross output and nominal value-added; we want to divide this nominal intermediate input into indices of price and quantity. Note that we also observe the output price for each of the types of intermediate input, from the gross output price data. Thus, to construct an appropriate deflator for intermediate goods, we simply need to know how to weight the price of each component of the intermediate goods index.

The difficulty arises from the low quality of the underlying data: the BEA compiles comprehensive input-output tables only about every five years (1947, 1958, 1963, 1967, 1972, 1977, and 1982). Intermediate inputs into any sector include inputs from all sectors. As with capital and labor, these disaggregated inputs should be weighted by marginal productivities in order to calculate a composite intermediate input. This requires consistent annual input-output tables in current and constant prices. In brief, for the benchmark years the data is adjusted to make the definitions consistent over time; they are then aggregated to the 35-industry level. These benchmarks are converted to shares of industry output, and then these shares are interpolated from benchmark to benchmark. This gives an estimated input-output table for each year.[14] This then allows us to create an appropriate price deflator for nominal payments to intermediate factors in each year.

Is it a problem that we use data constructed from these interpolated input-output tables for regressions at annual frequencies? Probably not. First, the weights on the different components of the intermediate-goods index do not change much, even over a five to ten year period. Second, even when the weights on some component do change, they tend to change only gradually. Hence, a linear interpolation probably provides a reasonable approximation to the true weights.

14 Wilcoxen (1988) discusses in detail the steps required for all of these calculations.

BIBLIOGRAPHY

Abbott, T., Z. Griliches, and J. Hausman, 1988, Short run movements in productivity: Market power versus capacity utilization. Mimeo, Harvard University.

Arrow, K., 1974, The measurement of real value added, in: P. David and M. Reder, eds., Nations and households in economic growth: Essays in honor of Moses Abramovitz (Academic Press, New York).

Baily, M.N., 1990, Comment, in: Peter Diamond, ed., Growth/Productivity/Employment (MIT Press, Cambridge).

Bartelsman, E., R. Caballero, and R. Lyons, 1991, Short and long-run externalities, NBER Working Paper No. 3810.

Basu, S. and J. Fernald, 1993, Productive externalities in U.S. manufacturing: Do they exist and are they important? Mimeo, University of Michigan.

Berndt, E. and M. Fuss, 1986, Productivity Masurement Using Capital Asset Valuation to Adjust for Variations in Utilization, Journal of Econometrics, 33, 7-30.

Bruno, M., 1978, Duality, intermediate inputs, and value-added, in: M. Fuss and D. McFadden, eds., Production economics: A dual approach to theory and practice, Volume 2 (North-Holland, Amsterdam).

Bruno, M. and J. Sachs, 1985, The economics of worldwide stagflation (Harvard University Press, Cambridge).

Caballero, R. and R. Lyons, 1992, External effects in U.S. procyclical productivity, Journal of Monetary Economics 29, 209-226.

Cooper, R. and A. John, 1988, Coordinating coordination failures in Keynesian models, Quarterly Journal of Economics 84, 347-385.

Diamond, P., 1982, Aggregate demand management in search equilibrium, Journal of Political Economy 90, 881-94.

Domowitz, Ian, R. Glenn Hubbard, and Bruce C. Petersen (1988). "Market Structure and Cyclical Fluctuations in U.S. Manufacturing." Review of Economics and Statistics 55-66.

Gordon, R.J., 1992, Is procyclical productivity a figment of measurement error? Mimeo.

Griliches, Z. and V. Ringstad, 1971, Estimating production functions and technical change from micro data. (Aschehoug, Oslo).

Hall, R., 1986, Market structure and macroeconomic fluctuations, Brookings Papers in Economic Activity 2, 285-322.

Hall, R., 1988, The relation between price and marginal cost in U.S. industry, Journal of Political Economy 96, 921-947.

Hall, R., 1990, Invariance properties of Solow's productivity residual, in: Peter Diamond, ed., Growth/Productivity/Employment (MIT Press, Cambridge).

Hall, R., 1991, Labor demand, labor supply, and employment volatility, in: O. Blanchard and S. Fischer, eds, NBER Macroeconomics Annual 1991 (MIT Press, Cambridge).

Hall, R. and D. Jorgenson, 1967, Tax policy and investment behavior, American Economic Review 57, 391-414.

Jorgenson, D., 1990, Productivity and economic growth, in: E. Berndt and J. Triplett, eds., Fifty years of economic measurement (University of Chicago Press, Chicago).

Jorgenson, D., F. Gollop, and B. Fraumeni, 1987, Productivity and U.S. economic growth (Harvard University Press, Cambridge).

Ramey, V., 1991, Comment, in: O. Blanchard and S. Fischer, eds. NBER Macroeconomics Annual 1991 (MIT Press, Cambridge).

Romer, P., 1986, Increasing returns and long-run growth, Journal of Political Economy 94, 1002-37.

Rotemberg, J. and M. Woodford, 1991, Markups and the business cycle, in: Blanchard, O. and S. Fischer, eds. NBER Macroeconomics Annual 1991 (MIT Press, Cambridge).

Sato, K., 1976, The meaning and measurement of the real value added index, Review of Economics and Statistics 58, 434-442.

Weitzman, Martin, 1982, Increasing returns and the foundations of unemployment theory, Economic Journal 92, 787-804.

Table 1

$$dy = \gamma \cdot dx + dt$$

Data Set	Uninstrumented		Instrumented	
	Non-Qual-Adjusted	Quality-Adjusted	Non-Qual-Adjusted	Quality-Adjusted
Hall/NIPA VA	0.89 (0.013)	--	1.17 (0.031)	--
Divisia VA	1.04 (0.025)	0.66 (0.024)	1.12 (0.041)	0.85 (0.041)
Gross Output	1.10 (0.009)	1.07 (0.007)	1.07 (0.014)	1.02 (0.011)

Standard errors in parentheses. Sample is 1959 through 1984.

Table 2

$$dv_i = \gamma^* \cdot dx_i + \beta \cdot dm_i + dt_i$$

Data Set	Param	Uninstrumented		Instrumented	
		Non-Qual-Adjusted	Quality-Adjusted	Non-Qual-Adjusted	Quality-Adjusted
Hall/NIPA VA	γ^*	0.64 (0.018)	--	0.48 (0.051)	--
	β	0.09 (0.012)	--	0.56 (0.031)	--
Divisia VA	γ^*	0.64 (0.033)	0.47 (0.028)	0.86 (0.064)	0.63 (0.062)
	β	0.20 (0.022)	0.39 (0.018)	0.19 (0.041)	0.18 (0.042)

Standard errors in parentheses. Sample is 1959 through 1984.

Table 3

$$dv + \left(\frac{s_m}{1-s_m}\right)dm = \gamma\left(\frac{dx^v}{1-s_m}\right) + \mu\left(\frac{s_m}{1-s_m}\right)(dm - dx^v) + \frac{dt}{1-s_m}.$$

Data Set	Param	Uninstrumented		Instrumented	
		Non-Qual-Adjusted	Quality-Adjusted	Non-Qual-Adjusted	Quality-Adjusted
Hall/NIPA	γ	0.99 (0.009)	--	1.03 (0.012)	--
	μ	1.19 (0.011)	--	1.30 (0.023)	--
Divisia VA	γ	0.99 (0.011)	1.0 (0.007)	0.98 (0.019)	0.96 (0.014)
	μ	1.21 (0.014)	1.23 (0.012)	1.21 (0.023)	1.15 (0.026)

Standard errors in parentheses. Sample is 1959 through 1984.