Dark Matter: Measuring Unobserved Productivity Growth due to Computers through its Impact on Observables

Alwyn Young, London School of Economics, June 2025

"You can see the computer age everywhere
but in the productivity statistics." (Solow 1987)

Abstract

I find systematic deviations in the relationship between measured industry total factor productivity growth and price & downstream demand growth associated with the use of computer and electronics intermediate inputs in production. The effects are robustly negative, indicating an overstatement of quality adjusted output and productivity growth in using industries of about .003 per .01 expenditure share on computer & electronic intermediates. These estimates are confirmed by regressions of measured productivity growth on own and upstream computer input use. After adjustment for mismeasurement, there is no association between computer intermediates use and total factor productivity growth.

*I am grateful to participants in the LSE macro work-in-progress seminar for helpful comments.

I. Introduction

Unobserved objects can be quantified and measured by observed systematic discrepancies in the causal relationships between observables. Linearly and trivially, if $y = x\beta$, y_o & x_o are observed and β is known, then unobserved x is given by $x_{uo} = y_o/\beta - x_o$. Econometrically, in a world in which x is not the sole determinant of y and β is unknown, if an observable indicator z_o exists such that $x_{uo} = \gamma z_o + \eta$, then both β and γ can be consistently estimated by running the non-linear regression $y_o = (x_o + \gamma z_o)\beta + \varepsilon$. The key identifying assumptions are that the impact of x_{uo} on y_o (β) is the same as that of x_o and, implicitly, that z_o affects y_o only through its influence (γ) on x_{uo}, i.e. an exclusion restriction. The first assumption can be relaxed to allow x_{uo} to have an impact $f(\beta)$, where the function f() is known. This naturally arises in cases where the linear relationship is on an unobservable $y_{uo} = (x_o + x_{uo})\beta + \varepsilon$, and there is an additional known discrepancy between y_o and y_{uo} driven by x_{uo}, as is the case below.

This paper applies this idea to a topic that has troubled economists and statisticians for some decades, the question of whether we are properly measuring the quality-adjusted gains from the use of computer technology. The lefthand side observables are measured quality-adjusted prices and downstream input demand growth (y_o), the righthand side observables are measured total factor productivity growth (x_o) and computer factor income shares (z_o) and the unobservable of interest is mismeasurement of total factor productivity growth associated with the use of computer related inputs ($x_{uo} = \gamma z_o$). The key identifying assumptions are that: (1) the movements of true quality adjusted prices and quantities demanded (y_{uo}) are the same function of true total factor productivity growth (x_o+x_{uo}), regardless of its origin (a common β); and (2) the use of computer inputs has no impact on quality adjusted price and quantity demanded other than through its impact on true productivity growth (the exclusion restriction). These are motivated using downstream industry demand and the relationship between price and costs. The tested null hypothesis is that of no systematic mismeasurement associated with computer related inputs, $\gamma = 0$.

The linear relation, $x_{uo} = \gamma z_o$, is motivated by Solow's comment quoted above, which suggests that we fail to measure what we do with computers. This can be operationalized as mismeasurement of average rates of computer-factor augmenting technical change in computer using industries, so that the degree of mismeasurement is proportional to the expenditure share on computer inputs. Mismeasurement of the output of one sector translates into mismeasurement of total factor productivity growth in the opposite direction in downstream industries that use its output. Thus, the mismeasurement hypothesis actually implies mismeasurement in users and in users of users, with effects going in opposite directions, something that does not seem to have been considered.

I find empirical evidence of mismeasurement in what we do with computer and electronic intermediates. The point estimates are all decidedly negative, implying that we systematically overstate the factor augmenting gains associated with the use of these inputs. These results are arrived at using likelihood and bootstrap techniques shown in Monte Carlos to be extremely conservative, with rejection probabilities of true nulls well below nominal value, while retaining power to reject false nulls of zero effects. They are robust to multiple changes in the specification of ancillary variables, the systematic delete-one removal of individual industries from the sample, the use of varied likelihood techniques to account for heavy tailed data, the allowance for industry level heterogeneity in the association between productivity growth and price and quantity changes, adjustments for business cycle mismeasurement in total factor productivity, the use of simple long run industry means, and, even, a wide grid search over all possible combinations of the elasticity of supply and demand.

Put simply, the relationship between price and demand growth and total factor productivity growth differs systematically and robustly from that implied by elasticities of demand and supply in a manner correlated with both within and between industry variation in the quantity of computer intermediate inputs. Similarly statistically significant and robust deviations are not found for any other input. I interpret this deviation in terms of a model of overstatement of factor augmenting technical change, i.e. what we do with computer inputs. Point estimates indicate an overstatement of industry output and total factor productivity growth of about .003 per annum per percentage share of computer and electronics inputs in total factor payments.

The paper whose observations are closest in spirit to this one is Acemoglu et al (2014). While documenting the correlation of higher industry labour productivity growth with some measures of computer technology use, they note that it is peculiarly negatively associated with real output growth, i.e. inconsistent with expected changes in demand following a reduction in price brought on by total factor productivity growth. The present paper expands this emphasis on using confirmatory observables into a methodology that quantifies mismeasurement using estimated relationships between left and righthand side observables, finding, similarly, that there are in fact no supply and demand elasticities that can eliminate the discrepancy between price, quantity and total factor productivity growth associated with computer and electronics intermediates use. Unless, of course, one allows for the possibility that the gains from such use are overestimated.

The association of higher labour productivity growth with some measures of computer use in pre- and early millennial data has been documented by Stiroh (2002) and Acemoglu et al (2014). In that spirit, this paper also runs a simple linear regression of total factor productivity growth on the expenditure share of domestically produced computer and electronics intermediate inputs, documenting a similar positive relation in post-millennial data. However, it also finds a negative relation between industry total factor productivity growth and the use of computer and electronic intermediates in upstream industries. This positive own effect and negative supplier effect is consistent with overstatement of the output gains from intermediate input use, which would overstate productivity growth in users while understating it in users of users. The point estimates from these regressions are virtually identical to those found using the "dark matter" methodology that concentrates on the relationship between observable price, quantity and total factor productivity growth.

Once adjustment for estimated mismeasurement is made, there is basically no relation between industry productivity growth and computer intermediate input use. This echoes the pushback to Solow's comment given by Gordon (2000) and Oliner & Sichel (2000): while the growth of computer capital has contributed greatly, in a standard growth accounting framework, to growth outside that sector and there has also been rapid total factor productivity growth in computing industries, it is less obvious that there should be an additional productivity contribution from the use of those inputs. The computer productivity revolution most obviously lies in Moore's Law and the extraordinary fall in the price of computational power, and less obviously in what we accomplish with that power.[1]

The new millennium has witnessed a US productivity slowdown, with the Bureau of Economic Analysis' data described below showing private sector total factor productivity growth of .0084 per annum in 1987-2004 falling to .0032 per annum in 2004-2021. Opinion pieces in the popular press express the view that this is due to a failure to properly measure the gains from computer technology (e.g. Aeppel 2015, Alloway 2015), but in considered academic analyses Byrne, Fernald & Reinsdorf (2016) and Syverson (2017) persuasively argue that such unmeasured gains cannot explain the productivity slowdown. This paper argues that insofar as this mismeasurement is attached to the sale of particular products, it should show up as discrepancies in the relation between price, quantity and total factor productivity growth and, hence, is actually measurable. It finds evidence of mismeasurement, but unexpectedly in the opposite direction, suggesting that post-millennial growth is even slower than believed. At the aggregate level, this translates into about ⅛ of measured US output growth between 2000 and 2021, and ½ of total factor productivity growth during that period as well. Decades ago, Jones (1995, 2002) argued that the incremental costs of innovation and productivity growth rise with the level of technology, implying, absent ever increasing market size, declining growth. That view appears increasingly prescient.

The Boskin commission (1996) famously concluded that the Consumer Price Index (CPI) was biased upwards by 1.1 percent per year, of which about half could be attributed to a failure

[1]In this regard, it is sobering to reflect on the fact that the Apollo 11 command and lunar module guidance computers each had only 2KB of RAM, which would be insufficient for almost any phone app today.

to measure quality improvements. This, along with other well known studies finding unmeasured gains to quality improvement and variety (e.g. Gordon 1990, Bils and Klenow 2001), might lead readers to conclude that growth is unambiguously underestimated. There are a number of studies, however, pointing in the opposite direction. For example, Gordon (2009) and Gordon and VanGoethem (2007) find downward biases in the CPI of 3 percent per annum for women's apparel and 1 percent per annum in rental shelter over many decades, while Aizcorbe and Ripperger-Suhler (2024) estimate a negative chain drift in hedonic price indices in 2011-2020 of 6 and 8 percent per annum for desktop and notebook computers, respectively. On the theoretical level, Feenstra (1995) finds that in a discrete choice framework with pricing above marginal cost log-linear hedonic regressions, such as are used in the analysis of computer prices, would overstate price declines, Hobijn (2002) shows that if price per unit quality rises with quality both hedonic and model matching price indices will overstate price declines, Harper (2007) notes that durable goods obsolescence leads to an overstatement of quality change, and Aizcorbe and Copeland (2007) argue that with intermittent purchases price indices will tend to understate true movements in the cost of living index as consumers do not gain from price declines above their reservation value. While this paper is not about the methodology of price indices, to aid in the interpretation of its results I summarize some of these insights in a short explanation of how obsolescence and lifecycle differences in the reservation values of buyers can lead both hedonic and model matching price indices to overstate price declines and output growth in durable goods industries upgrading quality through the use of computer and electronics intermediates.

The paper proceeds as follows: Section II presents a model of systematic mismeasurement of factor augmenting technical change and discusses how data on price and downstream intermediate input demand can be used to identify the rate of mismeasurement. Special emphasis is given to explaining the sources of identification, the steps taken to avoid endogeneity bias, and the methods used to demonstrate it is not determining the results. Section III introduces the BEA industry level total factor productivity and input output data. The data are heavy-tailed and consequently I use both the standard multivariate normal and better fitting multivariate-t in estimation. Monte Carlos show bootstrapped standard errors achieve conservative control of size in this environment, but only the multivariate-t retains substantive power to reject erroneous nulls of zero effects. Consequently, while point estimates are similar across likelihoods, those using the multivariate-t are more consistently statistically significant.

Section IV presents the main results using the dark matter methodology, emphasizing the robustness of point estimates to the regression and likelihood specification and sample changes. The finding of overstated productivity growth in computer using industries is shown to remain when elasticities of demand and supply are taken as known and given any value the reader may like, emphasizing the robustness of the baseline results to any possible bias in the estimation of these elasticities. As confirmatory evidence, Section V regresses total factor productivity growth on computer factor shares. While such regressions might be interpreted as demonstrating the positive effects of computer use, in the light of the dark matter results the finding of a positive association with own computer intermediates use and negative association with upstream use suggests a different interpretation, i.e. that of mismeasurement. Section VI summarizes the implications for growth by industry and in the aggregate economy and concludes with a short discussion of potential sources of downward bias in hedonic and model matching durable goods price indices.

II. Estimating Mismeasurement Using Observables

(a) A Model of Systematic Mismeasurement in an Input-Output Framework

We assume throughout that national income accountants accurately measure nominal values, but have difficulty disentangling these into price and quantity components. Let true and measured gross output and price in perfectly competitive industry *i* in period *t* be given by

$$(1a) \quad Q_{it}^T = F_i(\, f_{1it}^T X_{1it}^T \,, \dots, \, f_{Jit}^T X_{Jit}^T \,) \quad P_{it}^T = C_i(\, W_{1t}^T / f_{1it}^T, \dots, \, W_{Jt}^T / f_{Jit}^T)$$

$$(1b) \quad Q_{it}^M = F_i(f_{1it}^M X_{1it}^T \,, \dots, \, f_{Jit}^M X_{Jit}^T) \quad P_{it}^M = C_i(\, W_{1t}^T / f_{1it}^M, \dots, \, W_{Jt}^T / f_{Jit}^M)$$

where *T* and *M* denote true and measured values, F_i production functions which are constant returns to scale in J inputs $X_1...X_J$, C_i cost functions which are constant returns to scale in J input prices $W_1... W_J$, and f_{jit} factor augmenting productivity parameters. (1a) are standard production and cost functions. (1b) is a formalization of *what* national income accountants implicitly measure, not *how* they actually measure output quantities and prices as, with rare exceptions,[2] these are not measured by examining input quantities and prices, let alone productivity parameters. For this reason, the factor quantity and price arguments in the functions on the righthand side of (1b) are true values, even though these might not be measured accurately. The failure to properly appreciate and quantify the degree to which technical change is allowing industries to use inputs in novel and more productive ways (f_{jit}) appears as implicit unobserved (UO) discrepancies between true and measured factor augmenting productivity, $f_{jit}^{UO} = f_{jit}^{T}/f_{jit}^{M}$.

Differentiating (1) and (2) with respect to time and using the equivalence between elasticities and expenditure shares implied by perfect competition, the difference between measured and true output quantity and price growth is seen to be

$$(2a) \quad \hat{Q}_{it}^{M} = \hat{Q}_{it}^{T} - \sum_{j=1}^{J} \theta_{jit} \hat{f}_{jit}^{UO} \qquad (2b) \quad \hat{P}_{it}^{M} = \hat{P}_{it}^{T} + \sum_{j=1}^{J} \theta_{jit} \hat{f}_{jit}^{UO},$$

where θ_{jit} is the expenditure share of input j and ^ denotes a proportional growth rate. Since nominal output is measured correctly, i.e. $P_{it}^{T} Q_{it}^{T} = P_{it}^{M} Q_{it}^{M}$, offsetting errors in prices and quantities arise equally whether national income accountants derive real quantity indices by deflating nominal values using constructed price deflators or price indices by dividing nominal values by constructed quantity measures. Furthermore, as the output of each sector is used as an input in others, the same errors in disentangling price and quantity are propagated through the input-output table. Thus, when industry n is used as an input in sector i, we have:

$$(3) \quad \hat{X}_{nit}^{M} - \hat{X}_{nit}^{T} = \hat{Q}_{nt}^{M} - \hat{Q}_{nt}^{T} \ \& \ \widehat{W}_{nt}^{M} - \widehat{W}_{nt}^{T} = \hat{P}_{nt}^{M} - \hat{P}_{nt}^{T}.$$

We order inputs so that the first N correspond to the industry indices i, with the remaining J - N consisting of primary factors.

While the mismeasurement of what we do with computer inputs translates naturally into

[2]Most notably government, which is not included as an industry in the analysis below.

the mismeasurement of the intermediate inputs of downstream sectors, for the purposes of our analysis here we assume that primary factor inputs are measured accurately. Recognizing that this assumption may be of concern to readers, in the on-line appendix I estimate mismeasurement of primary factor inputs using a similar methodology. Although point estimates suggest that the growth of computer and software capital may be overstated, the results are not statistically significant nor robust to changes in the sample. While one might think that the mismeasurement of domestic output found below would result in mismeasurement of the growth of capital input, this is not the case, as durable goods computer related capital investment is predominantly and increasingly based upon imports, as discussed more fully in the conclusion.

While true total factor productivity growth in industry i, i.e. the true growth of output minus the factor income share weighted growth of true factor inputs, is

$$(4) \quad \hat{A}_{it}^{T} = \hat{Q}_{it}^{T} - \textstyle\sum_{j=1}^{J} \theta_{jit} \hat{X}_{jit}^{T} = \sum_{j=1}^{J} \theta_{jit} \hat{f}_{jit}^{T},$$

measured total factor productivity growth, equal to measured output growth minus the weighted measured growth of inputs, is given by

$$(5) \quad \hat{A}_{it}^{M} = \hat{Q}_{it}^{M} - \sum_{j=1}^{J} \theta_{jit} \hat{X}_{jit}^{M} = \hat{A}_{it}^{T} + \hat{Q}_{it}^{M} - \hat{Q}_{it}^{T} - \sum_{j=1}^{J} \theta_{jit} (\hat{X}_{jit}^{M} - \hat{X}_{jit}^{T})$$

$$\Rightarrow \hat{A}_{it}^{M} = \hat{A}_{it}^{T} - \sum_{j=1}^{J} \theta_{jit} \hat{f}_{jit}^{UO} + \sum_{j=1}^{N} \theta_{jit} \sum_{k=1}^{J} \theta_{kjt} \hat{f}_{kjt}^{UO}.$$

Mismeasurement of what is done with inputs impacts own industry TFP estimates directly in one direction and the TFP estimates of downstream industries indirectly in the opposite direction.

(b) Estimation framework

We aim to identify the above by looking at the demand for intermediate inputs, where the moving demand and supply curves for the total use of intermediate input X_{it} follow:

$$(6a) \quad \hat{P}_{it}^{D} = \eta_{i}^{D} + \eta_{t}^{D} - \frac{1}{\sigma} \hat{X}_{it}^{D} + \varepsilon_{it}^{D} \quad (6b) \quad \hat{P}_{it}^{S} = \eta_{i}^{S} + \eta_{t}^{S} - \hat{A}_{it}^{T} + \frac{1}{\rho} \hat{X}_{it}^{S} + \varepsilon_{it}^{S},$$

where the η denote industry and time supply and demand fixed effects. Controlling for industry and time fixed effects, growth of input quantity demanded varies inversely with price growth with elasticity of substitution σ, while growth of input quantity supplied varies positively with price growth with elasticity of supply ρ. Under perfect competition total factor productivity growth lowers unit costs and the supply curve one-for-one, hence the "$-\hat{A}^T_{it}$" of (6b). Setting demand equal to supply yields the equilibrium solutions:

$$(7)\quad \hat{P}^T_{it} = \frac{\sigma(\eta^D_i + \eta^D_t + \varepsilon^D_{it})}{\sigma+\rho} - \frac{\rho(\hat{A}^T_{it} - \eta^S_i - \eta^S_t - \varepsilon^S_{it})}{\sigma+\rho}$$
$$\hat{X}^T_{it} = \frac{\sigma\rho(\hat{A}^T_{it} + \eta^D_i + \eta^D_t - \eta^S_i - \eta^S_t + \varepsilon^D_{it} - \varepsilon^S_{it})}{\sigma+\rho}.$$

As already stated, mismeasurement of total intermediate input use follows that of industry output, $\hat{X}^M_{it} - \hat{X}^T_{it} = \hat{Q}^M_{it} - \hat{Q}^T_{it}$.

If mismeasurement applies only to the use of input j, (2) & (5) can be simplified to

$$(8)\quad \hat{P}^M_{it} = \hat{P}^T_{it} + \theta_{jit}\hat{f}^{UO}_{jit},\quad \hat{X}^M_{it} = \hat{X}^T_{it} - \theta_{jit}\hat{f}^{UO}_{jit}$$
$$\&\quad \hat{A}^T_{it} = \hat{A}^M_{it} + \theta_{jit}\hat{f}^{UO}_{jit} - \sum_{n=1}^{N}\theta_{nit}\theta_{jnt}\hat{f}^{UO}_{jit},$$

so that (7) can be restated as the estimating equations on observables:

$$(9)\quad \hat{P}^M_{it} = \beta^P\left[\hat{A}^M_{it} + \gamma_j\left(\theta_{jit} - \Omega_{jit}\right)\right] + \gamma_j\theta_{jit} + \eta^P_i + \eta^P_t + \varepsilon^P_{it}$$
$$\hat{X}^M_{it} = \beta^X\left[\hat{A}^M_{it} + \gamma_j\left(\theta_{jit} - \Omega_{jit}\right)\right] - \gamma_j\theta_{jit} + \eta^X_i + \eta^X_t + \varepsilon^X_{it}$$
$$\text{where }\ \beta^P = \frac{-\rho}{\sigma+\rho},\quad \beta^X = \frac{\sigma\rho}{\sigma+\rho}\quad \&\ \ \Omega_{jit} = \sum_{n=1}^{N}\theta_{nit}\theta_{jnt},$$

and where γ_j is the economy-wide average rate of mismeasurement of factor augmenting productivity growth in the use of input j. (9) is a seemingly unrelated system of non-linear regressions with observed regressors $\hat{A}^M_{it}$, θ_{jit}, & Ω_{jit}, as well as industry and time fixed effects η_i and η_t. The parameter of interest is γ_j, the economy-wide average rate of mismeasurement of factor augmenting productivity growth in the use of input j, where we assume that

$$(10)\quad \hat{f}^{UO}_{jit} = \gamma_j + \zeta_{jit},\ \text{ with }\ E\left(\zeta_{jit}\right) = 0.$$

Variation ζ_{jit} of $\hat{f}^{UO}_{jit}$ from its average rate γ_j is implicitly included in the errors, which, as ζ_{jit} is multiplied by θ_{jit} and Ω_{jit}, makes them heteroskedatic. For consistent estimation the probability limit of the errors times the regressors must be zero, which in this case requires the additional assumption that the proportional rate of mismeasurement is independent of the factor shares. Similar assumptions are made in all regressions that do not deny the possibility of heterogeneous effects.

(c) Discussion of Identification

As noted in the Introduction, the identification of mismeasurement in (9) rests on two additional assumptions beyond those needed in any regression. First, that the impact of true total factor productivity growth on true price and quantity growth is the same regardless of its factor augmenting origin, i.e. a common β. Second, that expenditure shares have no impact on true price and quantity growth other than through their effects on true total factor productivity growth.

Total factor productivity growth impacts many observables within an industry, such as the relative use of factors, but these observables are likely to be heavily influenced by the form the total factor productivity growth takes. It is plausible, however, that the form total factor productivity growth takes within an industry is not relevant to the downward shift in the supply curve or the equilibrium demand in arms-length downstream industries. In keeping with this argument, the measure of X_{it} used below excludes own industry input demand. For computer expenditure shares not to influence price and quantity other than through total factor productivity growth, it is necessary that our conception of price, quantity and productivity include quality improvements. (9) allows for this, as it describes factor-augmenting-technical-change adjusted price and quantity as functions of true factor-augmenting-adjusted productivity. Here we follow standard national income accounting principles, re-expressing quality improvements as changes in quality adjusted prices per unit of quality adjusted quantity.

The identification of the slopes of supply and demand curves is a standard econometric problem, but of less importance here than the preceding exposition suggests. In (6) and (7) above total factor productivity growth shifts the supply curve, allowing the identification of the demand elasticity σ. The additional assumption that productivity growth shifts the supply curve one for one (under perfect competition) identifies the supply elasticity. To see this, note that if total factor productivity growth shifts the supply function by an unknown amount τ , then in moving from (6) to (9) we have

$$(11)\ \ \beta^P = \frac{-\rho\tau}{\sigma+\rho} \ \ \&\ \ \beta^X = \frac{\sigma\rho\tau}{\sigma+\rho}$$

and it is still the case that the ratio $-\beta^X/\beta^P$ identifies σ. The separate identification of the supply elasticity ρ from these two coefficients, however, requires that τ be known. However, as can be seen in (9), knowledge of neither ρ nor σ is needed to estimate the value of γ_j, which depends only on β^P and β^Q and not their decomposition. Estimates of these elasticities are given below as orthogonal matters of interest and a means of evaluating (under the assumption that $\tau = 1$) the estimates of β^P and β^Q.

There is also the conventional issue of the endogeneity of regressors, i.e. their correlation with the error terms in (9). With regards to the factor shares θ_{jit} and Ω_{jit}, we are regressing the *change* in prices and quantities between period t and t+1 on the *levels* of these regressors,[3] so any endogeneity that exists should rely mostly on the cumulative effects of correlated shocks. To this end, in the analysis below I include lagged values of the dependent variables as regressors to "whiten" the residuals. To avoid size distortions brought about by pre-testing, results with different lag structures are presented side-by-side and standard errors always clustered at the industry level to correct for any within industry correlation, and of course heteroskedasticity, left in any given specification. Key point estimates for computer intermediates do move somewhat towards zero with the addition of lagged dependent variables, and I treat these as more reliable, although differences across lag structures are not statistically significant.

Regarding total factor productivity growth, mismeasurement of this variable due to changes in capacity utilization brought about by demand and supply shocks do make it

[3]To remain consistent with the BEA's total factor productivity Tornqvist indices, I actually use the average value of factor shares in periods t and t+1. Results in the on-line appendix show that using period t factor shares as regressors, point estimates for computer intermediates shrink by about 10% towards zero, but are otherwise unchanged for other computer regressors, except in cases with very large standard errors using average factor income shares to begin with.

endogenous to price and quantity changes. As a baseline, I address this by including year fixed effects to account for business cycle movements of the regressors in general. These year fixed effects do play a role, as absent these the results are statistically less significant and point estimates closer to zero, as shown below. However, further corrections for productivity growth mismeasurement, in the form of industry level adjustments for business cycle fluctuations, have little effect on point estimates. Results using long run industry means, where mismeasurement due to capacity utilization should be much less of an issue, are also very similar to those found using year fixed effects. As endogeneity in total factor productivity growth will most directly bias the estimated βs, I also show that point estimates of the mismeasurement parameter γ_j are robust to very large exogenously imposed changes in the βs, i.e. taking them as known at different values. While the assumption of common βs is used to identify the degree of mismeasurement, this is not equivalent to saying that point estimates are very sensitive to the estimated values of those β. In practice they are not and hence any residual endogeneity of total factor productivity growth should be of little importance to the results.

III. Characteristics of the Data & their Implication for Estimation & Inference

I use the Bureau of Economic Analysis's industry level total factor productivity estimates covering 61 private sector industries from 1987 to 2021 and input-output tables covering the same from 1997 to 2021. One industry (social assistance) is not used as an intermediate input in any industry, while another (hospitals and nursing & residential care) is used as an intermediate input only in itself. Dropping these two and taking the intersection of the two data sets leaves 59 private sector industries for 1997 to 2021. The BEA productivity estimates provide factor income share and quantity data for 14 inputs, comprised of 9 classes of capital, college and non-college labour, and energy, service and materials intermediates. The input-output tables allow the more detailed calculation of the intermediate input shares of the domestic 59 private sector industries, giving a total of 73 input categories.[4] Our interest lies in those inputs most obviously associated with computer technology, namely (i) computer capital, (ii) software capital, (iii) computer and electronic intermediates and (iv) computer systems design & related intermediates. As the data are available, mismeasurement in the remaining 69 input categories is also examined, albeit parenthetically. The discussion of statistical significance below incorporates multiple testing bounds to control the family wise Type I error rate across the computer hypothesis tests.

Our dependent variables are the measured growth of total domestic intermediate input demand by industry, net of imports and own industry use, and intermediate input prices (taken as the growth of the sectoral domestic output deflator in the productivity accounts), while our regressors are measured total factor productivity growth by industry, factor and domestic intermediate input expenditure shares, and year and industry fixed effects. Regressing the dependent variables on industry & year dummies, I find that the residuals of annual industry price and non-own-use domestic intermediate input demand growth have kurtoses of 41 and 37, respectively. This indicates a large deviation from the normal distribution. The remainder of this section shows how alternative likelihoods combined with the bootstrap achieve conservative control of size while retaining power in this environment.

The multivariate t-distribution provides a computationally tractable way of modelling heavy-tailed multivariate data taking both positive and negative values on the reals. If the coefficient and covariance estimates for a normal model with common regressors for each dependent variable and multivariate normal errors are $\widehat{\boldsymbol{\beta}} = (\mathbf{X}'\mathbf{X})^{-1}(\mathbf{X}'\mathbf{Y})$ & $\widehat{\boldsymbol{\Sigma}} = \hat{\boldsymbol{\varepsilon}}'\hat{\boldsymbol{\varepsilon}}/\text{n}$, where $\mathbf{Y}$ is the n x m matrix of dependent variables, $\mathbf{X}$ the n x k matrix of common regressors, and $\hat{\boldsymbol{\varepsilon}}$ the n x m matrix of residuals, the corresponding measures for a multivariate t likelihood are weighted versions of the same, namely $\widehat{\boldsymbol{\beta}} = (\mathbf{X}'\mathbf{W}\mathbf{X})^{-1}(\mathbf{X}'\mathbf{W}\mathbf{Y})$ and $\widehat{\boldsymbol{\Sigma}} = \hat{\boldsymbol{\varepsilon}}'\mathbf{W}\hat{\boldsymbol{\varepsilon}}/\text{sum}(\mathbf{W})$,[5] where sum()

[4]In actuality, the BEA input-output tables are not a single matrix but rather separate supply and use tables for commodities. While the commodity-industry supply tables are largely diagonal, this is not entirely and completely the case. I calculate domestic input shares by dividing an industry's use of domestically produced commodities across industries in proportion to their shares of the supply thereof.

[5]The latter isn't actually the covariance (second central moment) matrix, which equals $\Sigma*\nu/(\nu-2)$, with ν equal to the degrees of freedom of the distribution.

denotes the sum of matrix entries and $\mathbf{W}$ is diagonal with elements $w_i = [\hat{\nu} + \hat{\boldsymbol{\varepsilon}}_i' \hat{\boldsymbol{\Sigma}}^{-1} \hat{\boldsymbol{\varepsilon}}_i]^{-1}$, with $\hat{\boldsymbol{\varepsilon}}_i$ the estimated residuals for observation i and $\hat{\nu}$ the estimated degrees of freedom. In estimating parameters, the multivariate-t systematically underweights outliers based upon their squared Mahalanobis distance $\hat{\boldsymbol{\varepsilon}}_i' \hat{\boldsymbol{\Sigma}}^{-1} \hat{\boldsymbol{\varepsilon}}_i$, with the underweighting increasing as the estimated degrees of freedom falls. This makes the point estimates less sensitive to error realizations in the heavy-tails of the distribution, as extreme deviations from point estimates are taken as being less unusual, and hence less relevant, than is the case with normal errors.

Table I below presents baseline maximum likelihood modelling of intermediate input price and quantity demanded (net of own use) using the multivariate normal and the multivariate-t. Each column includes industry and year fixed effects (η_i, η_t), differing only in the order T of lags of the dependent variables included in the vector auto regressions, namely

$$(12)\quad \hat{P}_{it}^{M} = \sum_{l=1}^{T} (L_{lP}^{P} \hat{P}_{it-l}^{M} + L_{lX}^{P} \hat{X}_{it-l}^{M}) + \eta_i^P + \eta_t^P + \varepsilon_{it}^P$$

$$\hat{X}_{it}^{M} = \sum_{l=1}^{T} (L_{lP}^{X} \hat{P}_{it-l}^{M} + L_{lX}^{X} \hat{X}_{it-l}^{M}) + \eta_i^X + \eta_t^X + \varepsilon_{it}^X.$$

I refer to the 2x2 matrix of l^{th} lag coefficients as $\mathbf{L}_l$. Standard error estimates, in parentheses, and covariance matrices are clustered and (/) cluster bootstrapped, both at the industry level. As seen, the multivariate t provides a much better fit to the data, with likelihoods that are 1300 to 1400 ln points higher for each lag structure. With the normal likelihood tests of the joint significance of the matrix of 1st order lag coefficients ($\mathbf{L}_1$) do not reject the null of the 2x2 0 matrix $\mathbf{0}_{2x2}$, but do reject the same null for 2nd and 3rd order lags ($\mathbf{L}_2$ and $\mathbf{L}_3$). In contrast, with the t-likelihood $\mathbf{L}_1 = \mathbf{0}_{2x2}$ and $\mathbf{L}_3 = \mathbf{0}_{2x2}$ are consistently rejected at the .05 level, but $\mathbf{L}_2 = \mathbf{0}_{2x2}$ is not when the bootstrap is used to estimate the covariance matrix. Specifications with $\mathbf{L}_4$, not reported in the table, do not reject the null of $\mathbf{0}_{2x2}$ with the t likelihood and either covariance estimate and with the normal when bootstrapped.[6] As noted earlier, if there is any endogeneity of factor shares in period t to price and quantity changes between periods t and t+1, it should

[6]P-values with the normal likelihood & clustered/clustered bootstrap covariance estimates are .041/.255, while the corresponding p-values for the t-likelihood are .238/.327.

Table 1: Baseline Maximum Likelihood Modelling & Associated Monte Carlos
(dependent variables = panel data on growth of prices and non-own-use input demand)

	multivariate normal errors				multivariate-t errors			
	0 lags	1 lag	2 lags	3 lags	0 lags	1 lag	2 lags	3 lags
v					1.38 (.175/.176)	1.27 (.146/.162)	1.27 (.139/.146)	.125 (.143/.159)
ln L	3298	3155	3040	2932	4711	4572	4397	4220
N	1416	1357	1298	1239	1416	1357	1298	1239
	p - values of joint hypothesis tests: clustered/clustered bootstrapped							
$\mathbf{L}_1 = \mathbf{0}_{2x2}$		.801/.887	.903/.965	.851/.941		.000/.000	.000/.000	.000/.001
$\mathbf{L}_2 = \mathbf{0}_{2x2}$			.000/.005	.000/.000			.000/.065	.000/.067
$\mathbf{L}_3 = \mathbf{0}_{2x2}$				.000/.000				.000/.000
	Monte Carlo probability of rejecting true null regarding $\mathbf{L}_l$: clustered/clustered bootstrapped							
$\mathbf{L}_1$		.98/.96	.97/.89	.98/.97		.73/.04	.71/.04	.71/.03
$\mathbf{L}_2$			.99/.93	1.0/.93			.62/.09	.60/.07
$\mathbf{L}_3$				.99/.93				.65/.04
	Monte Carlo probability of rejecting false null of $\mathbf{L}_l = 0$: clustered/clustered bootstrapped							
$\mathbf{L}_1$		1.0/.99	1.0/1.0	1.0/.98		1.0/1.0	1.0/1.0	1.0/1.0
$\mathbf{L}_2$			1.0/1.0	1.0/1.0			1.0/1.0	1.0/1.0
$\mathbf{L}_3$				.99/.97				1.0/1.0

Notes: Lags indicates number of lags of the dependent variables included in the regression, as in (12) above. lnL = ln likelihood, N = # of observations, v = estimated degrees of freedom. Standard error estimates in () are clustered/clustered bootstrapped, both at the industry level. Bootstrapped coefficient standard errors and covariance matrices calculated using 99 draws. Monte Carlo rejection probabilities over 100 runs using data generating processes based on the t-likelihood point estimates in the last four columns.

come from serial correlation in the residuals. Whitening the residuals using lagged dependent variables as regressors corrects for any such endogeneity. Below results are reported for all lag specifications and throughout I cluster at the industry level, as is done in Table 1, to allow for any unaddressed intertemporal dependence between industry residuals in a given specification.

The estimated degrees of freedom of the multivariate-t in Table 1 are between 1.25 and 1.38, implying extremely heavy tailed data with no higher than a first (integer) moment.[7] This brings into question the accuracy of inference with or without the bootstrap, as higher moments are usually specified as sufficient (albeit not necessary) for consistency. Moreover, with extremely low degrees of freedom the t-likelihood often has multiple local maxima, as for a given estimate of the degrees of freedom $\hat{v}$, different weights $w_i = [\hat{v} + \hat{\boldsymbol{\varepsilon}}_i' \widehat{\boldsymbol{\Sigma}}^{-1} \hat{\boldsymbol{\varepsilon}}_i]^{-1}$ allow multiple fixed point solutions with $\widehat{\boldsymbol{\beta}} = (\mathbf{X}'\mathbf{W}\mathbf{X})^{-1}(\mathbf{X}'\mathbf{W}\mathbf{Y})$, $\widehat{\boldsymbol{\Sigma}} = \hat{\boldsymbol{\varepsilon}}'\mathbf{W}\hat{\boldsymbol{\varepsilon}}/\mathrm{sum}(\mathbf{W})$ & $\hat{\boldsymbol{\varepsilon}} = \mathbf{Y} - \mathbf{X}\widehat{\boldsymbol{\beta}}$.

[7]The t-distribution has moments of order up to its degrees of freedom.

I address these issues using an estimation algorithm whose accuracy is verified in Monte Carlos. In determining the coefficients $\widehat{\boldsymbol{\beta}}$ for a given t-distribution estimate of the degrees of freedom, I set weights initially at one and then iterate least squares recursively to the fixed point, in a manner similar to that used in weighted least squares and GMM estimation (which share the same issue of potential multiplicity).[8] This concentrates the likelihood as a function of the degrees of freedom,[9] which is generally well behaved, as illustrated in the on-line appendix, so that a global maximum is easily found. I use Monte Carlos to show that this algorithm, in combination with the bootstrap, achieves conservative control of size while retaining power, as detailed now.

Panel D of Table 1 above begins by reporting Monte Carlo rejection rates at the .05 level of true nulls equal to the parameters of the data generating process (*dgp*), which is given by the parameters and degrees of freedom of the multivariate-t estimates with the corresponding lag structure in the last columns of the table. 100 data realizations are made per *dgp* and 99 industry clustered draws used for each bootstrap. Rejection rates of the true null using the multivariate normal likelihood, whether bootstrapped or not, are near 1.0 at the .05 nominal level. This stems from a substantial bias in these estimates that arises with normal estimation of a lag structure with such heavy tailed data. In contrast, when estimated using the multivariate-t, although rejection rates using the conventional covariance estimate have large size distortions, the bootstrap brings these down close to nominal level while retaining the same 1.0 probability (power) of rejecting incorrect nulls of 0 for the lag coefficients.

Our interest in this paper, however, lies in accurate inference regarding the mismeasurement parameters γ_j in (9) above. To this end, Table 2 reports Monte Carlo rejection probabilities using data generating processes based upon the t-distribution point estimates of the model in (9). There are 4 computer related and 69 other factor shares, and 4 different lag lengths

[8]This concentration of the likelihood using iteratively reweighted least squares is essentially the application of the EM algorithm (Dempster, Laird & Rubin 1977) to the t-distribution advocated by Liu & Rubin (1995).

[9]The non-linear mismeasurement models are linear conditional on the mismeasurement parameter, so in that case the likelihood is concentrated as a function of that parameter and the degrees of freedom.

(0 … 3), for a total of 73 x 4 = 292 data generating processes. For each I estimate the multivariate-t on the data and then use the estimated parameters as the data generating process, using 50 Monte Carlos for each *dgp* in the case of the four computer related factor shares and 10 in the case of other input shares. The table presents average rejection probabilities for the mismeasurement parameter across these Monte Carlos using: the conventional industry clustered standard error estimate; the bootstrapped standard error estimate (bootstrap-se); the percentiles of the bootstrapped coefficient distribution (bootstrap-c); the percentiles of the bootstrap t-distribution (bootstrap-t); and versions of the bootstrap -se, -c and -t based upon bootstrap sub-sampling. The bootstrap-t is known to provide an asymptotic refinement, i.e. faster convergence to the true distribution, but in finite samples may provide no benefits if the finite sample standard error estimate is inaccurate, as noted by Hall (1992, p. 167). Sub-sampling asymptotically fixes bootstrap failures in some contexts and, with regards to heavy tailed data, Hall & LePage (1996) show that sub-sampling M from N observations, with M/N going to zero, the bootstrap-t asymptotically allows for accurate inference with no more than first moments in the data generating process. I sub-sample by drawing 30 industry clusters from the 59 of the data. Bootstrap p-values are calculated using 99 industry clustered draws, with percentile -c and -t methods using the absolute value of the test statistic and a p-value equal to (G+u)/100, where G is the number of bootstrapped test statistics greater than that of the original sample and u is a draw from the (0,1) uniform distribution.[10] Reported in the table are mean rejection probabilities of true nulls equal to the parameter of the *dgp* and false nulls of zero effects, illustrating aspects

[10]There is a popular misconception that the bootstrap distribution must be approximated with a large number of draws to achieve accurate inference. Hope (1968) noted that with k an integer and m draws from a continuous bootstrap distribution, an exact test (relative to the distribution) at level $\alpha = k/(m+1)$ is achieved when the null is rejected if k-1 or less draws are greater than the sample test statistic. Jockel (1986) showed the same is true for draws from an arbitrary distribution, if $(G+(T+1)*u)/(m+1)$ is less than α, where T is the number of ties with the sample test statistic (the +1 treating the sample test statistic as a tie with itself) and u a draw from the [0,1] uniform distribution. The on-line appendix to Young (2019) shows, rather trivially, that this is true for arbitrary α, i.e. there is no need for $\alpha(m+1)$ to be an integer. In this paper, ties do not occur. Given the high computation costs, I select the smallest number (99) such that integer values of G (0,4, and 9) indicate rejection at the .01, .05 and .1 levels, regardless of the draw u (the sample's tie with itself). For the bootstrap-se, a degrees of freedom adjustment should be made for the sampling distribution of the s.e. estimate, but as the chi-squared distribution is commonly used, whose critical values are very close to those of the squared-t with 98 degrees of freedom, I use the chi-squared as well.

Table 2: Monte Carlo Tests of Unmeasured Productivity Growth
(data generating processes based upon t-likelihood point estimates for 292 models)

	multivariate normal likelihood							multivariate-t likelihood						
	clus-	bootstrapped			sub-sampling			clus-	bootstrapped			sub-sampling		
level	tered	se	c	t	se	c	t	tered	se	c	t	se	c	t
(A) computer related inputs (4 inputs x 4 lag structures x 50 Monte Carlos)														
rejection probabilities of true null of the data generating process														
.01	.080	.000	.000	.041	.004	.004	.035	.141	.010	.004	.095	.003	.001	.076
.05	.141	.014	.010	.079	.014	.009	.064	.201	.033	.024	.174	.013	.006	.170
.10	.223	.037	.036	.110	.029	.021	.100	.264	.066	.070	.247	.030	.028	.226
rejection probabilities of false null of 0 mismeasurement														
.01	.226	.005	.001	.168	.006	.004	.149	.431	.190	.070	.376	.164	.022	.346
.05	.286	.025	.019	.220	.018	.015	.205	.520	.321	.239	.481	.325	.140	.454
.10	.344	.061	.049	.263	.041	.039	.253	.578	.407	.367	.541	.427	.263	.520
coefficient of variation of true null rejection probabilities														
.01	1.3	.	.	1.4	2.1	2.9	1.6	1.0	1.3	2.1	0.9	2.7	4.0	1.2
.05	1.0	1.5	1.5	1.1	1.3	1.9	1.4	0.8	0.9	1.0	0.7	1.5	1.9	0.8
.10	0.8	0.8	0.9	1.0	0.8	1.1	1.0	0.6	0.7	0.7	0.6	1.1	1.2	0.6
coefficient of variation of false null rejection probabilities														
.01	1.5	2.3	200.0	80.9	1.9	2.9	1.7	0.9	1.3	1.9	1.1	1.4	2.7	1.1
.05	1.1	1.2	59.9	68.9	1.2	1.4	1.5	0.7	1.0	1.1	0.7	0.9	1.5	0.8
.10	0.9	0.8	40.3	60.2	0.9	1.1	1.3	0.6	0.8	0.8	0.6	0.8	1.0	0.6
(B) other inputs (69 inputs x 4 lag structures x 10 Monte Carlos)														
rejection probabilities of true null of the data generating process														
.01	.052	.003	.002	.011	.003	.001	.008	.099	.010	.005	.051	.004	.002	.042
.05	.111	.020	.018	.047	.015	.013	.034	.167	.036	.033	.126	.015	.010	.114
.10	.170	.043	.054	.087	.034	.037	.063	.229	.068	.074	.184	.030	.035	.173
rejection probabilities of false null of 0 mismeasurement														
.01	.095	.013	.007	.029	.005	.002	.021	.346	.137	.082	.220	.178	.032	.199
.05	.158	.040	.042	.082	.024	.021	.063	.451	.223	.204	.361	.254	.126	.345
.10	.235	.073	.093	.133	.047	.058	.107	.519	.279	.303	.442	.307	.204	.436
coefficient of variation of true null rejection probabilities														
.01	1.8	5.5	6.7	3.2	5.8	8.3	3.9	1.4	3.2	4.3	1.5	4.9	6.7	1.8
.05	1.1	2.2	2.4	1.7	2.7	2.7	2.0	1.0	1.7	1.8	1.0	2.5	3.1	1.1
.10	0.8	1.5	1.4	1.2	1.9	1.7	1.4	0.8	1.4	1.2	0.8	2.0	1.9	0.8
coefficient of variation of false null rejection probabilities														
.01	1.6	3.5	44.2	24.8	4.2	7.4	3.0	0.9	1.8	2.2	1.1	1.5	3.3	1.1
.05	1.1	2.0	20.5	15.5	2.3	2.4	1.8	0.7	1.4	1.4	0.8	1.2	1.8	0.8
.10	0.8	1.5	14.2	11.5	1.6	1.5	1.3	0.6	1.1	1.1	0.6	1.1	1.4	0.7

Notes: Rejection probabilities calculated across 50 (10) realizations of each of 16 (276) *dgps* in panel A (B). Coefficient of variation is across the average rejection rate in 50 (10) iterations of each *dgp*. Standard error estimates cluster at the industry level, bootstraps use 99 industry clustered draws, and sub-sampling draws 30 of 59 industry clusters with replacement.

of size and power. Also reported is the coefficient of variation of the average rejection rate in 50 draws each of 16 computer-related *dgps* in panel A and 10 draws each of 276 *dgps* for other inputs in panel B. If the tests had true null rejection probabilities equal to nominal level in each and every *dgp*, at the .01, .05 and .1 levels these should be 1.4, 0.8, and 0.4 in panel A and 3.1, 1.4 and 0.9 in panel B.

Two patterns are immediately apparent in the table. First, use of the mis-specified normal likelihood on this heavy tailed data results in low and inconsistent power. While normal likelihood rejection probabilities of the true mismeasurement parameters less than or equal to nominal value can be achieved with the use of the bootstrap, the volatility of the normal coefficient estimates is so great that the false null of zero is also rejected with low average frequency. In addition, both true and false null rejection probabilities vary greatly from one *dgp* to another, as shown by the coefficients of variation. Estimation and inference using the multivariate normal in the presence of extremely heavy tailed data is not a sensible empirical strategy. Results using the familiar normal likelihood are presented below to assure the reader that point estimates for key results are of the same sign and magnitude as found using the multivariate-t, but their bootstrapped standard errors are much larger.

Second, the table indicates that the simplest bootstrap, the -se, provides conservative control of the true null rejection probability, while (in the case of the multivariate-t) retaining power and showing a variability across *dgps* close to what might be expected from an exact test. True null rejection probabilities with the conventional industry clustered covariance estimate are well above nominal value, as are those using the bootstrap-t (for the t likelihood), notwithstanding its asymptotic properties.[11] Again, despite its asymptotic virtues, sub-sampling provides no systematic advantages in finite sample mean rejection probabilities, but, in the case of the t-likelihood, has more variable true null rejection probabilities, indicating less consistent performance across *dgps*. The bootstrap-se and -c have similar performance, although the -c has

[11]This mirrors results found for 2SLS in Young (2022) where, because the standard error estimate is very poor, the bootstrap-t works worse than the -c in finite samples, as was anticipated by Hall (1992).

less power at low p-values, which matters a great deal with Bonferroni-Holm multiple testing adjustments made below. Given these advantages, and as it allows the conventional presentation of results in coefficient (standard error) form, I adopt the bootstrap-se as the baseline approach in tables below.[12] However, for key results I also report p-values using all inference methods described in Table 2, as well as additional details on their Monte Carlos.

To summarize, with heavy tailed data in tests of the mismeasurement parameter the bootstrapped normal likelihood controls size at the expense of power to reject false nulls of 0 effects that is often below nominal level as well. However, using the simplest bootstrap, the bootstrap-se, the t-likelihood achieves conservative control of the true null rejection probability, while retaining power to reject false nulls that is a multiple of the nominal level. These results are achieved using a maximization algorithm that avoids multiplicity of local maxima found by selecting coefficients to minimize the weights on individual observations by, for a given degrees of freedom, setting weights initially at 1 and iteratively reweighting least squares until the weights converge, a procedure analogous to that used in weighted least squared and GMM estimation. This concentrates the likelihood and, as shown in Monte Carlos, despite the lack of higher than first integer moments in the data generating process, with standard errors calculated using the bootstrap-se maintains conservative control of true null rejection probabilities while retaining substantive power.

IV. Results

Table 3 reports estimated mismeasurement rates for computer related inputs using multivariate normal and t likelihoods. The estimating equations follow (9) above, augmented with 0 to 3 lags of the vectors of dependent variables as in (12). Each cell reports the mismeasurement parameter γ_j for a separately estimated model, its bootstrap-se standard error estimate in (), and the delete-one-industry-at-a-time min to max coefficient range in []. With the

[12]The reader should keep in mind, however, that the coefficients may not have second moments. The reported standard error is merely a bootstrapped number which, when multiplied by percentiles of the normal distribution and added and subtracted from the point estimate, conservatively covers the true parameter value.

Table 3: Mismeasurement Associated with Computer Technology
(each cell a separately estimated model)

	normal likelihood				t likelihood			
	no lags	1 lag	2 lags	3 lags	no lags	1 lag	2 lags	3 lags
computer hardware capital	-.36	-.10	-.32	-.38	-.22	-.10	-.11	-.07
	(.56)	(1.1)	(1.4)	(1.1)	(.29)	(.18)	(.19)	(.22)
	[-.43/-.25]	[-.19/.08]	[-.43/.06]	[-.90/.16]	[-.64/-.19]	[-.23/-.08]	[-.28/-.08]	[-.17/-.04]
software capital	.00	.22	.02	-.39	-.30	-.11	-.11	-.10
	(.40)	(.49)	(.62)	(.69)	(.16)	(.18)	(.18)	(.21)
	[-.14/.33]	[-.02/.67]	[-.12/.84]	[-.60/.53]	[-.36/-.26]	[-.17/-.05]	[-.25/-.04]	[-.18/.01]
computer & electronic intermediates	-.49	-.51	-.64	-.48	-.48	-.35	-.44	-.33
	(.22)	(.24)	(.33)	(.21)	(.12)	(.10)	(.13)	(.10)
	[-.51/-.32]	[-.86/-.45]	[-1.1/-.53]	[-.64/-.42]	[-.48/-.46]	[-.46/-.34]	[-.51/-.41]	[-.39/-.30]
computer systems design intermediates	2.8	4.2	3.5	2.2	.16	.30	.18	.17
	(2.2)	(2.8)	(2.7)	(2.2)	(1.1)	(.87)	(.72)	(.63)
	[.61/3.9]	[.96/5.4]	[.84/4.8]	[.67/3.1]	[-.89/.56]	[-.51/.50]	[-.58/.39]	[-.38/.33]

Notes: Mismeasurement parameters as in the systems estimation (9) above augmented with lag vectors as in (12). Bootstrap-se standard errors in () based upon 99 industry-clustered draws; min-max delete-one industry coefficient range in []. 59 industries x 24 years = 1416 observations without lags, reduced by 59 for each lag.

exception of computer systems design intermediates and software capital in some specifications, all point estimates are negative, implying that national income accountants implicitly overestimate computer augmenting technical change. However, aside from the estimates for computer and electronic intermediates, all point estimates are statistically insignificant, with t-stats less than 1 in absolute value or, when somewhat higher, easily reduced below 1 through a change in the number of lags. The delete-one min to max coefficient range for computer capital, software capital and computer systems design is also negative to positive with the normal and/or t likelihood, showing that point estimates, and particularly the positive point estimates of systems design, are not at all robust.

In contrast, the results for computer and electronics intermediates are consistently significant, with t-stats of -2 with the normal likelihood and -4 to -3 with more powerful t-likelihood estimation. Point estimates with the t-likelihood in this case are also quite stable, ranging between -.5 and -.3 across all lag structures and all delete-one industry min-max bounds. Estimates with the normal likelihood are more negative and volatile, with point estimates ranging from -.6 to -.5 and delete-one industry min-max bounds of -1.1 to -.3 across all lag structures.

Throughout the presentation below, normal likelihood results follow those with the t-likelihood, but with larger standard errors and point estimate volatility, in keeping with the lower precision and power of these estimates shown in Monte Carlos. In the interest of transparency, these results are presented, but I view estimates with the multivariate t as much more reliable and focus the discussion on these.

The results for computer & electronics intermediates are not only unique amongst computer related inputs; they are unique across all 73 primary and intermediate inputs tested in this paper. Table 4 reports the results for the remaining five categories of capital, two categories of labour and three broad intermediate input factor income share categories given in the BEA total factor productivity data base. For R&D capital with normal estimation, point estimates are positive with some t-stats as high as 1.8, but these largely vanish, with a positive to negative delete-one industry min-max range, when the t-likelihood is used. Otherwise, no t-stat is much greater than 1 in absolute value and the min-max delete-one industry coefficient range generally includes zero with at least one lag structure. The bottom three rows of the table report results for the three detailed domestic intermediate input industry categories, out of the remaining 57 non-computer categories taken from the input output tables, which have an absolute t-stat greater than 1.9 in some specification. For machinery and management of companies & enterprises intermediates this t-stat appears with normal estimation, but point estimates are near zero with t-likelihood estimation. For amusement, gambling and recreation intermediates, the absurdly large point estimates for these inputs seem most easily explained by the equally large standard errors. Needless to say, with any adjustment for the vast amount of multiple testing underlying this table, all results are statistically insignificant at the .05 or .1 levels.

The results above use bootstrap-se standard errors, as these were found in the Monte Carlos of section III to conservatively control size, retain power and provide the most uniform results across data generating processes patterned on the models tested in these tables. Table 5 focuses in more narrowly on estimates for computer & electronics intermediates, reporting the Monte Carlo simulation rejection rates and p-values in the sample using all inference techniques

Table 4: Mismeasurement Estimates for Non-Computer Inputs

	normal likelihood				t-likelihood			
	no lags	1 lag	2 lags	3 lags	no lags	1 lag	2 lags	3 lags
r & d capital	.40 (.23)	.44 (.24)	.51 (.29)	.34 (.28)	.28 (.23)	.23 (.18)	.20 (.20)	.22 (.17)
	[.14/.50]	[.22/.56]	[.19/.66]	[.05/.46]	[-.12/.41]	[-.06/.32]	[-.10/.35]	[-.05/.31]
communications capital	.62 (.83)	.67 (1.0)	.28 (1.0)	-.13 (.93)	-.30 (.47)	.03 (.34)	-.15 (.38)	-.38 (.36)
	[.27/1.1]	[.37/1.4]	[-.02/.83]	[-.51/.41]	[-.46/.00]	[-.09/.19]	[-.36/.10]	[-.56/-.15]
instruments capital	1.2 (1.6)	2.4 (2.1)	3.1 (2.4)	2.8 (2.2)	-.33 (.90)	.12 (.85)	.21 (.88)	.30 (.82)
	[.10/2.0]	[1.2/3.4]	[1.6/4.3]	[1.6/3.7]	[-.76/-.06]	[-.15/.56]	[-.21/.65]	[.02/.83]
transport equipment	-.07 (.33)	.08 (.52)	.11 (.64)	.18 (.50)	-.04 (.24)	-.16 (.20)	-.01 (.23)	-.06 (.21)
	[-.23/.22]	[-.09/.57]	[-.10/.79]	[.02/.70]	[-.23/.03]	[-.25/.02]	[-.10/.07]	[-.13/.08]
other equipment	.14 (.40)	.38 (.38)	.43 (.48)	.51 (.43)	.12 (.11)	.08 (.09)	.20 (.17)	.16 (.11)
	[-.01/.21]	[.22/.46]	[.25/.56]	[.34/.70]	[.10/.14]	[.05/.11]	[.17/.23]	[.13/.19]
art capital	-.17 (.23)	-.06 (.26)	-.12 (.33)	-.06 (.23)	.02 (.11)	.07 (.11)	.02 (.11)	.03 (.10)
	[-.24/-.08]	[-.20/.05]	[-.32/.03]	[-.23/.04]	[-.04/.13]	[.01/.21]	[-.05/.18]	[-.04/.12]
structures capital	-.08 (.11)	-.11 (.12)	.01 (.11)	.05 (.11)	.06 (.04)	.00 (.03)	.05 (.04)	.05 (.03)
	[-.12/.05]	[-.15/.04]	[-.02/.07]	[.00/.09]	[.04/.07]	[-.01/.01]	[.04/.07]	[.04/.07]
college labour	.21 (.14)	.23 (.14)	.23 (.18)	.09 (.15)	.02 (.07)	.03 (.06)	.03 (.06)	.03 (.04)
	[.12/.24]	[.14/.27]	[.14/.28]	[.02/.13]	[-.03/.04]	[.00/.04]	[-.01/.05]	[.00/.04]
non-college labour	-.03 (.18)	-.12 (.20)	-.20 (.34)	-.17 (.26)	.03 (.05)	.03 (.03)	.01 (.04)	.01 (.04)
	[-.12/.02]	[-.26/-.01]	[-.39/-.02]	[-.32/-.03]	[.00/.05]	[.02/.04]	[-.01/.03]	[-.01/.02]
energy intermediates	.00 (.16)	-.01 (.15)	-.02 (.18)	.08 (.18)	-.04 (.08)	-.04 (.05)	-.05 (.05)	-.05 (.05)
	[-.08/.10]	[-.09/.09]	[-.12/.06]	[-.04/.16]	[-.06/.00]	[-.05/.00]	[-.07/-.03]	[-.07/-.03]
materials intermediates	-.06 (.08)	-.08 (.08)	-.11 (.10)	-.05 (.10)	-.04 (.08)	-.06 (.06)	-.06 (.07)	-.08 (.06)
	[-.11/.04]	[-.13/.01]	[-.15/.00]	[-.12/.04]	[-.05/.04]	[-.07/.00]	[-.08/.00]	[-.09/-.02]
services intermediates	-.01 (.05)	.00 (.05)	-.04 (.05)	-.03 (.05)	-.02 (.02)	.00 (.02)	-.01 (.02)	-.01 (.02)
	[-.03/.01]	[-.02/.02]	[-.06/-.01]	[-.05/.00]	[-.03/-.01]	[.00/.01]	[-.02/-.01]	[-.01/.00]
machinery	1.5 (.92)	1.3 (.82)	1.2 (.80)	2.1 (.88)	-.32 (.48)	-.47 (.32)	-.46 (.54)	.03 (.53)
	[.74/2.2]	[.72/1.9]	[.56/1.7]	[1.4/2.6]	[-.50/-.06]	[-.56/-.25]	[-.58/-.21]	[-.15/.38]
management of companies & enterprises	-.47 (.32)	-.89 (.46)	-1.3 (.65)	-.79 (.41)	-.23 (.17)	-.07 (.13)	-.08 (.66)	-.19 (.14)
	[-.58/-.30]	[-1.0/-.33]	[-1.5/-.51]	[-.95/-.41]	[-.31/-.18]	[-.10/-.03]	[-.11/-.01]	[-.24/-.11]
amusements, gambling & recreation intermediates	-20.4 (29.5)	-21.6 (31.1)	-28.6 (33.0)	-25.3 (32.2)	-14.3 (7.2)	-11.6 (6.4)	-10.3 (5.9)	-10.6 (6.0)
	[-48.6/-12.1]	[-49.0/-15.7]	[-58.2/-21.5]	[-53.9/-19.6]	[-21.5/-12.6]	[-15.9/-9.9]	[-14.3/-8.8]	[-14.7/-9.3]

Notes: as in Table 3.

Table 5: Monte Carlo Rejection Rates & Alternative P-Values for Mismeasurement
(computer & electronics intermediates)

	Monte Carlo rejection probabilities at .01 level								p-values in the sample			
	true nulls = parameter				false nulls = 0 effects							
	0 lags	1 lag	2 lags	3 lags	0 lags	1 lag	2 lags	3 lags	0 lags	1 lag	2 lags	3 lags
					multivariate t – likelihood							
clustered t-stat	.38	.44	.28	.32	1.0	1.0	1.0	1.0	.0000	.0000	.0000	.0000
bootstrap-se	.00	.00	.00	.00	.70	.42	.46	.12	.0000	.0007	.0011	.0015
bootstrap-c	.00	.00	.00	.00	.22	.04	.08	.00	.0100	.0300	.0300	.0300
bootstrap-t	.22	.32	.16	.22	1.0	.98	.98	1.0	.0000	.0100	.0000	.0000
subsampling-se	.00	.00	.00	.00	.34	.14	.10	.00	.0194	.0413	.0149	.0347
subsampling-c	.00	.00	.00	.00	.06	.00	.00	.00	.0200	.1000	.0500	.0600
subsampling-t	.22	.26	.18	.18	1.0	.94	.98	.90	.0000	.0000	.0000	.0000
					multivariate normal likelihood							
clustered t-stat	.24	.26	.28	.20	.86	.76	.80	.72	.0000	.0000	.0000	.0000
bootstrap-se	.00	.00	.00	.00	.00	.00	.00	.02	.0277	.0364	.0517	.0200
bootstrap-c	.00	.00	.00	.00	.00	.00	.00	.00	.0600	.0500	.0600	.0300
bootstrap-t	.12	.12	.18	.12	.74	.64	.56	.52	.0000	.0000	.0100	.0100
subsampling-se	.00	.00	.00	.00	.00	.00	.00	.00	.1219	.0822	.0958	.1041
subsampling-c	.00	.00	.00	.00	.00	.00	.00	.00	.1300	.0600	.1000	.1100
subsampling-t	.08	.14	.14	.14	.70	.56	.56	.48	.0000	.0000	.0200	.0100

Notes: Monte Carlo rejection rates in 50 draws from *dgps* patterned after the multivariate-t sample point estimates, as in Table 2 earlier. Bootstraps use 99 industry-clustered draws from the bootstrap distribution.

examined earlier above. As seen in the table, with the multivariate t the conventional clustered t-stat, the bootstrap-t and the sub-sampling bootstrap-t all have a great deal of power to reject the false null of 0 effects when the *dgp* has the parameter of the point estimate, but also have large size distortions with true null rejection probabilities that are 16 to 44 times the .01 nominal level. In contrast, the -se & -c based techniques are extremely conservative, with true null rejection probabilities near zero. Power for these methods with the t-likelihood declines as lags are added, with the bootstrap-se being the only technique that combines conservative control of size with non-vanishing power. Empirical rejection rates for the normal likelihood show less power, as was the case in the broader Monte Carlos above, with a zero probability of rejecting the false null of zero effects outside of t-stat based methods with size distortions.

Table 5 also reports p-values for computer & electronics intermediates in the sample itself using the different methods. The clustered t-stat and bootstrap-se methods report a p-value based upon the chi^2 distribution, while the percentile -c & -t methods report a p-value equal to the number of greater (G) bootstrap occurrences divided by 100, dropping the (0,1) uniform random variable in the p-value formula (G+u)/100 for clarity. As can be seen, the sample p-values mirror the relative power of the different methods more than their relative size distortions. While the .00 to .01 p-values of t-likelihood t-stat based methods might be attributed to either power or size, the frequent rejections of the null using -se and -c methods are unlikely to be Type-I errors given the zero frequency of these in the Monte Carlos. Instead, the t-likelihood -se & -c p-values follow their simulated relative power, drifting up as lags are added. P-values for the same using the normal likelihood are systematically higher, in keeping with their lower simulated power, but still below .05 in a few instances.

The impact of multiple testing adjustments on the above results is easily seen. The conservative Bonferroni adjustment to control the family wise error rate, i.e. the probability of one or more Type-I errors in a set of tests, multiplies the p-values by the number of tests.[13] By this criterion, we should multiply the computer intermediates p-values for each lag structure by 4 to control error probabilities across the four computer categories tested in Table 3. For the percentile -c & -t bootstrap methods, we need to keep in mind the addition of a random draw from the (0,1) uniform distribution divided by 100. Thus, for example, with 1 lag using the t-likelihood bootstrap-t one would reject the null at the .05 level using a 4 test adjustment with .25 probability, as (.01+.0025)*4 = .05.

With these adjustments in mind, we see that putting aside -t methods with severe excess Type-I rejection probabilities, only the t-likelihood bootstrap-se results remain consistently significant at the .01 level (.012 in the case of 3 lags) with a four test adjustment. Once one moves beyond 0 lags, the bootstrap-c and subsampling-se and -c results are not significant at the .05 or .1 levels with a Bonferroni adjustment.[14] However, to reject the multiple testing null at these levels one needs to be able to reject the single-test null at about a

[13]There are step-down procedures to control the family wise error rate (Holm 1979) or the false discovery rate (Benjamini & Hochberg 1985) having made an initial rejection using the lowest p-value, but they are not relevant as all other individual mismeasurement p-values are already insignificant without adjustment for multiple testing, or easily made so with a change in the number of lags.

[14]For the bootstrap percentile methods, one can follow Westfall-Young (1993) and use the joint distribution of the bootstraps to calculate the distribution of the minimum p-value across tests, as is done in the on-line appendix. While the adjusted p-values are lower than with Bonferroni, it makes no difference to the evaluation of the -c results as insignificant at the.10 level arrived at by multiplying the numbers above by 4. Similarly, while percentile -t methods for the most part reject the null at the .05 level, this again is no different than what is easily seen in Table 5 with a Bonferroni adjustment.

.01 or .02 level which, as the Monte Carlos show, these tests have an evanescent probability of achieving as lags are added when the 0 null is actually false and the mismeasurement parameter equals its estimated value.

The credibility of results also rests on their robustness to changes in the specification. To this end, the tables above use alternative likelihoods, vary the lag structure, and report the min-max delete-one industry coefficient range. Table 6 considers additional specification checks. Row 1 of the table repeats the baseline t-likelihood estimates of Table 3 with both year and industry fixed effects, with rows 2 through 4 reporting results removing year fixed effects, industry fixed effects and both fixed effects. All point estimates are negative across all lag structures and, excepting the removal of year fixed effects with longer lag structures in row 2 alone, retain t-statistics of between -4 and -2 and delete-one industry min to max ranges that remain negative to negative, excluding zero. Results removing fixed effects for the normal likelihood, in rows 5 through 8 of the table, show similar patterns. The ln t-likelihood, given in {}, falls by 400 points with the removal of the 20+ year fixed effects and using these to control for business cycle induced correlations between price, quantity and total factor productivity growth (also mismeasured due to changes in capacity utilization), seems appropriate. However, the table shows that this specification is by and large not an essential determinant of the results.

Row 9 of Table 6 considers whether the statistical significance of the computer input mismeasurement parameter in other regressions spuriously proxies for industry heterogeneity in the elasticity of supply and demand. To this end, it runs a specification that allows for industry heterogeneity (β_i) in the impact of total factor productivity growth on prices and quantities:

$$(13)\quad \begin{aligned} \hat{P}_{it}^{M} &= \beta_i^{P}\big[\hat{A}_{it}^{M} + \gamma_j\big(\theta_{jit} - \Omega_{jit}\big)\big] + \gamma_j\theta_{jit} + \eta_i^{P} + \eta_t^{P} + \varepsilon_{it}^{P} \\ \hat{X}_{it}^{M} &= \beta_i^{X}\big[\hat{A}_{it}^{M} + \gamma_j\big(\theta_{jit} - \Omega_{jit}\big)\big] - \gamma_j\theta_{jit} + \eta_i^{X} + \eta_t^{X} + \varepsilon_{it}^{X}, \end{aligned}$$

augmented, of course, with vector auto-regression lags of the lefthand side variables, as in (12) earlier. As seen in the table, allowing industry specific heterogeneity in the β's has very little effect on the results. Point estimates and the delete-one industry min to max bounds with the t-likelihood across all lag structures range from -.27 to -.53 and are very similar to

Table 6: Specification Checks
(computer & electronics intermediates: each cell a separately estimated model)

	no lags	1 lag	2 lags	3 lags
impact of fixed effects (t-likelihoods)				
(1) baseline t likelihood	-.48 (.12) [-.48/-.46] {4823}	-.35 (.10) [-.46/-.34] {4666}	-.44 (.14) [-.51/-.41] {4488}	-.33 (.10) [-.39/-.30] {4304}
(2) no year fixed effects	-.41 (.11) [-.41/-.28] {4403}	-.29 (.11) [-.31/-.12] {4275}	-.30 (.17) [-.31/.01] {4104}	-.22 (.17) [-.23/.09] {3940}
(3) no industry fixed effects	-.43 (.08) [-.47/-.42] {4579}	-.24 (.07) [-.31/-.22] {4495}	-.26 (.09) [-.35/-.24] {4300}	-.16 (.07) [-.25/-.15] {4149}
(4) no fixed effects	-.41 (.09) [-.49/-.40] {4202}	-.19 (.05) [-.23/-.18] {4141}	-.21 (.06) [-.27/-.19] {3953}	-.14 (.06) [-.22/-.13] {3811}
impact of fixed effects (normal likelihoods)				
(5) baseline normal likelihood	-.49 (.22) [-.51/-.32] {3429}	-.51 (.24) [-.86/-.45] {3285}	-.64 (.33) [-1.1/-.53] {3164}	-.48 (.21) [-.64/-.42] {3044}
(6) no year fixed effects	-.44 (.17) [-.45/-.27] {3151}	-.36 (.20) [-.38/.00] {3022}	-.41 (.33) [-.43/.23] {2918}	-.26 (.45) [-.29/.67] {2804}
(7) no industry fixed effects	-.45 (.09) [-.48/-.40] {3370}	-.45 (.10) [-.51/-.40] {3225}	-.56 (.13) [-.60/-.46] {3096}	-.38 (.08) [-.40/-.34] {2995}
(8) no fixed effects	-.43 (.09) [-.46/-.38] {3099}	-.40 (.08) [-.42/-.37] {2970}	-.48 (.10) [-.50/-.42] {2858}	-.31 (.08) [-.32/-.28] {2758}
heterogeneity in the response of prices and quantities to TFP growth (industry-specific β)				
(9) t-likelihood	-.53 (.13) [-.53/-.51] {5092}	-.46 (.11) [-.51/-.41] {4953}	-.40 (.14) [-.49/-.39] {4740}	-.32 (.14) [-.34/-.27] {4567}
(10) normal likelihood	-.49 (.24) [-.51/-.47] {3878}	-.44 (.23) [-.91/-.40] {3758}	-.49 (.29) [-1.1/-.39] {3590}	-.39 (.30) [-.92/-.32] {3441}
thick tailed distributions based on normal likelihoods with scaled covariance matrices				
(11) industry scaled	-.47 (.15) [-.48/-.28] {5043}	-.36 (.12) [-.37/-.30] {4858}	-.36 (.14) [-.38/-.31] {4645}	-.30 (.13) [-.31/-.14] {4457}
(12) year scaled	-.45 (.26) [-.48/.07] {3838}	-.43 (.20) [-.45/-.17] {3706}	-.51 (.25) [-.55/-.26] {3599}	-.41 (.23) [-.44/.00] {3501}
(13) industry and year scaled	-.47 (.20) [-.48/-.07] {5294}	-.32 (.13) [-.33/-.10] {5137}	-.34 (.15) [-.35/-.10] {4920}	-.29 (.18) [-.31/.04] {4728}
industry level unemployment controls (unemployment level or change)				
(14) *U* level t-likelihood	-.46 (.21) [-.47/-.35] {4962}	-.40 (.15) [-.41/-.38] {4791}	-.45 (.19) [-.46/-.43] {4639}	-.40 (.17) [-.41/-.37] {4459}
(15) *U* change t-likelihood	-.47 (.14) [-.47/-.33] {5047}	-.37 (.12) [-.38/-.32] {4927}	-.39 (.13) [-.40/-.35] {4738}	-.33 (.15) [-.36/-.01] {4554}
(16) *U* level normal likelihood	-.47 (.36) [-.50/.06] {3482}	-.52 (.27) [-.58/-.45] {3352}	-.66 (.37) [-.80/-.55] {3230}	-.48 (.32) [-.50/-.42] {3105}
(17) *U* change normal likelihood	-.49 (.25) [-.51/-.15] {3730}	-.49 (.23) [-.72/-.44] {3605}	-.55 (.30) [-.87/-.47] {3456}	-.34 (.17) [-.39/-.33] {3365}
estimation using industry mean data (no lags)				
(18) t - likelihood	-.46 (.23) [-.49/-.38] {348}	(19) normal likelihood		-.47 (.14) [-.52/-.45] {339}

Notes: Unless otherwise noted, as in Table 3. {} = ln likelihood. Scaled (2x2) covariance matrices are scalar multiples of each other, as described in the text.

those found in the baseline specification in the first row of the table. Point estimates with the normal likelihood, in row 10, all remain negative, as do the min to max delete one bounds, but with larger standard errors and coefficients which shrink toward zero with lags, are no longer significant at the .05 level with lags. Relative to the baseline specifications in rows (1) and (5) of the table, likelihoods rise by hundreds of points and likelihood ratio tests reject the null that all β_i^P and β_i^X are identical with a p-value that is basically 0,[15] but the industry-specific β are estimated with great imprecision and the implied point estimates of elasticities of supply and demand are often negative, so I place preference on the pooled estimates of the baseline t-likelihood specification. Rows 9 and 10 show, however, that heterogeneity in the industry level response of price and quantity to total factor productivity growth is not creating the spurious appearance of mismeasurement.

Rows 11 - 13 of Table 6 consider alternative heavy tailed distributions. The multivariate t is the distribution of a normal vector which on each industry x time draw has its covariance matrix multiplied by the inverse square root of an independently drawn chi-squared variable. This suggests the possibility of creating a heavy-tailed distribution by multiplying the covariance matrix of a multivariate normal vector by only industry, year, or industry & year fixed effects, i.e. one in which the covariance matrix for industry *i* in period *t* equals $\lambda_i * \mathbf{\Sigma}$, $\lambda_t * \mathbf{\Sigma}$, or $\lambda_i * \lambda_t * \mathbf{\Sigma}$, with λ_i and λ_t suitably normalized to equal 1 in some base industry and year. Maximum likelihood estimation of this model, like that of the t-likelihood, easily finds multiple maxima, as coefficient estimates are once again weighted OLS estimates based upon covariance matrices whose relative magnitude is in turn determined by the OLS residuals. I follow the procedure used in estimating the multivariate-t: for each possible value of the mismeasurement parameter computing a ln-likelihood by setting initial weights to 1 and iteratively reweighting until the coefficient and covariance matrix estimates converge to a fixed point.

As seen in Table 6, results using the normal likelihood with industry and year scaled covariance matrices closely resemble those with the multivariate-t. Point estimates range

[15] As there are as many coefficients as clusters used in the construction of covariance estimates, Wald tests cannot be used to assess the significance of differences in these coefficients.

from -.5 to -.3 and, with the exception of year scaling factors combined with 0 or 3 lags, the delete-one industry min-max range is consistently negative and t-statistics are greater than 2 in absolute value. The 58 estimated industry scaling factors in row 11 are particularly important, increasing the unscaled normal likelihood in row 5 by 150 to 200 ln points more than the multivariate-t in row 1, while the 20+ year scaling factors in row 12 are less effective, having ln likelihoods 800 to 1000 points less than the baseline multivariate-t. Estimates for these models, as well as those with industry specific β_i in row 9, are difficult & costly to compute, while bootstrap draws often do not converge,[16] so I am unable to evaluate the accuracy of these procedures using Monte Carlos. For this reason the baseline results focus on the multivariate-t (with just industry and year fixed effects), which models heavy-tailed data with a parsimonious number of parameters (one), remaining tractable and computable while consistently showing likelihoods 1300 to 1400 ln points above the normal, as noted earlier.[17]

Changes in capacity utilization with the business cycle or demand shocks create mismeasurement of TFP growth, which potentially makes this righthand side variable endogenous to quantity and price shocks, biasing the estimates of β^P & β^X and hence, perhaps, the mismeasurement parameter γ_j. While the inclusion of year fixed effects offers a baseline correction for these forces, the remaining rows of Table 6 address this issue further. First, in rows 14 and 16 I add the national unemployment rate interacted with an industry indicator to the righthand side, so that the specification (not counting lags of the dependent variables) becomes

$$(14)\quad \hat{P}_{it}^{M} = \beta^{P}\left[\hat{A}_{it}^{M} + \gamma_j\left(\theta_{jit} - \Omega_{jit}\right)\right] + \gamma_j\theta_{jit} + \delta_i^{P}U_t + \eta_i^{P} + \eta_t^{P} + \varepsilon_{it}^{P}$$

[16]For these models, starting with initial weights of 1, about 1/5 of the time (1/3 in some instances) the iteratively reweighting procedure generates zero residuals for groups of observations, producing near singular covariance matrices whose determinant eventually exceeds machine precision and produces undefined likelihoods. This only happens in about 1 in 100 bootstrap draws with the multivariate-t in the baseline specification. In all tables and Monte Carlos above, bootstrap draws which do not converge are discarded until 99 convergent draws are found. For the rare instances when these issues arise in the calculation of the delete-one industry min-max range, likelihoods that do not converge using the iteratively reweighted concentrated likelihood procedure are simultaneously maximized across all parameters, converging to some local maximum.

[17]In addition to rows 11 – 13 in Table 6 I have tried the multivariate generalized gaussian distribution (Pascal et al 2017), for which the weights are the inverse of the squared Mahalanobis distance alone, without the addition of the t degrees of freedom v. For my heavy tailed data, estimation using this distribution is utterly hopeless, as no matter the algorithm used the weights move to infinity for one group of observations or another.

$$\hat{X}_{it}^{M} = \beta^{X}\left[\hat{A}_{it}^{M} + \gamma_j\left(\theta_{jit} - \Omega_{jit}\right)\right] - \gamma_j\theta_{jit} + \delta_i^{X}U_t + \eta_i^{X} + \eta_t^{X} + \varepsilon_{it}^{X},$$

where U_t is the national unemployment rate and the δ_i are industry specific responses to the business cycle. Rows 15 and 17 use a similar specification, with the change in the unemployment rate, U_t - U_{t-1}, as the regressor. These specifications move beyond common year fixed effects in removing industry specific business cycle variation from the identification of β^P & β^X . As seen in the table, the specification with the change in the unemployment rate fits the data better (with higher ln likelihoods) and moves the normal likelihood point estimates closer to those of the t-likelihood, which however change little from the baseline specification with year fixed effects in row (1). Although standard errors increase, absolute t-stats with the t-likelihood remain greater than 2 or 3 and all delete one industry min-max ranges are negative.

Rows 18 and 19 of Table 6 consider a more radical adjustment, running the analysis using the 24 year industry averages of the variables, which should eliminate any substantive effect from changes in capacity utilization. Naturally, the industry and year fixed effects are dropped, so the specification is

$$\text{(15)} \quad \hat{P}_i^{M} = \beta^{P}\left[\hat{A}_i^{M} + \gamma_j\left(\theta_{ji} - \Omega_{ji}\right)\right] + \gamma_j\theta_{ji} + c^{P} + \varepsilon_i^{P}$$
$$\hat{X}_i^{M} = \beta^{X}\left[\hat{A}_i^{M} + \gamma_j\left(\theta_{ji} - \Omega_{ji}\right)\right] - \gamma_j\theta_{ji} + c^{X} + \varepsilon_i^{X},$$

where c^P and c^X are constants and all variables are industry averages. As shown in the table, the t- and normal likelihood point estimates are all but identical to their baseline specifications, as well as to each other, and despite having only 59 observations, have bootstrapped-se based t-stats of -2 and -3.4, respectively. Shortly below I give evidence that point estimates are largely invariant across an enormous range of elasticities of supply and demand determining β^P & β^X, further confirming that they cannot possibly be heavily influenced by any bias in β^P & β^X due to changes in capacity utilization and endogenous total factor productivity growth.

Table 7 below reports the estimated response of intermediate input prices and demand (β^P and β^X) to total factor productivity growth in the baseline, unemployment change augmented, industry scaled and industry mean models of Table 6. One can infer from these

Table 7: Estimates of β^P and β^X and Implied Demand (σ) & Supply (ρ) Elasticities
(computer & electronics intermediates)

	baseline model				industry scaled covariance matrices	
	t - likelihood		normal likelihood			
	no lags	3 lags	no lags	3 lags	no lags	3 lags
β^P	-.19 (.05) [-.21/-.17]	-.22 (.07) [-.25/-.19]	-.44 (.21) [-.50/-.20]	-.40 (.18) [-.46/-.20]	-.18 (.05) [-.21/-.17]	-.21 (.07) [-.23/-.18]
β^X	.46 (.08) [.43/.50]	.60 (.16) [.50/.65]	.68 (.21) [.60/.76]	.68 (.20) [.58/.77]	.55 (.15) [.50/.64]	.64 (.16) [.59/.72]
σ	2.5 (.81) [2.3/2.7]	2.8 (2.2) [2.5/3.2]	1.5 (1.7) [1.3/3.7]	1.7 (1.5) [1.3/3.6]	3.0 (.91) [2.7/3.3]	3.1 (1.6) [2.9/3.6]
ρ	.57 (.11) [.53/.62]	.77 (.21) [.62/.83]	1.2 (2.1) [.92/1.5]	1.1 (.67) [.93/1.4]	.67 (.20) [.61/.80]	.81 (.24) [.75/.93]
	unemployment change controls				industry means	
	t - likelihood		normal likelihood		t - likelihood	normal likelihood
	no lags	3 lags	no lags	3 lags		
β^P	-.16 (.05) [-.18/-.15]	-.20 (.06) [-.26/-.18]	-.42 (.19) [-.48/-.21]	-.48 (.22) [-.57/-.25]	-.89 (.26) [-.96/-.81]	-.66 (.18) [-.81/-.60]
β^X	.46 (.09) [.42/.52]	.81 (.14) [.54/.84]	.60 (.19) [.51/.68]	.61 (.18) [.52/.70]	1.5 (.55) [1.2/1.6]	1.0 (.47) [.87/1.2]
σ	2.8 (1.0) [2.6/3.1]	4.1 (1.9) [2.1/4.5]	1.4 (1.1) [1.1/3.2]	1.3 (1.1) [1.1/2.7]	1.7 (1.1) [1.4/1.8]	1.5 (.77) [1.3/1.8]
ρ	.55 (.11) [.50/.63]	1.0 (.19) [.70/1.1]	1.0 (.82) [.82/1.3]	1.2 (1.8) [.91/1.6]	14.2 (49.7) [6.6/34.4]	3.8 (13.5) [2.2/6.5]

Notes: Results for models estimated in rows of Table 6 above: baseline (1 & 5), industry scaled (11), unemployment change controls (15 & 17), & industry means (18 & 19). Otherwise, as in Table 3.

the elasticities of relative demand (σ) and supply (ρ), as $\sigma = -\beta^X/\beta^P$ and, under the assumption that total factor productivity growth shifts the supply curve down one for one, $\rho = \beta^X/(1+\beta^P)$. Where 3 lags of the left hand side variables are included in the estimating equation, I report the estimated long run response of price and quantity to total factor productivity growth, equal to $(\mathbf{I}_2\text{-}\mathbf{L}_1\text{-}\mathbf{L}_2\text{-}\mathbf{L}_3)^{-1}\boldsymbol{\beta}$, were the $\mathbf{L}_l$ are the 2x2 matrices of coefficients on lags of order l of the dependent variables and $\mathbf{I}_2$ is the 2x2 identity matrix. The reported elasticities in this case are those implied by these long run responses to total factor productivity growth.

As seen in the table, estimated elasticities of demand using industry x year data range from 1.3 to 4.1, while elasticities of supply range from .6 to 1.2. Point estimates with multivariate normal likelihoods are lower for the elasticity of demand and higher for the elasticity of supply than those with t-likelihoods or with industry scaled covariances, but standard errors are sufficiently large that all estimates aren't meaningfully different from each other. Industry mean data for 1997-2021 find much flatter industry relative supply curves, with a t-likelihood supply elasticity of 14.2 and normal likelihood supply elasticity of 3.8, but with only 59 observations and standard errors of 50 and 14, respectively, these are not meaningfully different from the other results or indeed relatively extreme alternatives.

There are relatively few empirical estimates of the elasticity of relative demand or supply between industry classifications such as are used in this paper. Using the annual BEA input output data and military expenditures as an instrument, Atalay (2017) finds an elasticity of substitution across industries of .2 or less, while Miranda-Pinto (2021) using the same data and instrument finds an elasticity of substitution across industry inputs of 0 in manufacturing and .5 in services. In contrast, Peter and Ruane (2025), using India's trade liberalization as a natural experiment, find an elasticity of substitution across industry groups in that country of 2.5 over a seven year horizon, which is in line with what is found in this paper. As shown by the much higher ln likelihood of the specification with heterogeneity in β^P and β^X in Table 6 earlier, there is considerable cross industry heterogeneity in σ and ρ, and hence variation induced by changes in military expenditures, which only impacts certain industries, may not yield an estimate that is representative of average industry level elasticities. The estimates given in this paper treat total factor productivity growth as a broad instrument and, as shown in Table 6, do not change when adjustments for endogeneity due to business cycle capacity utilization are made. Nevertheless, to further address concerns about the possible impact of bias in the estimation of the demand and supply elasticities, Table 8 estimates the mismeasurement parameter across the whole range of possible elasticities.

Specifically, Table 8 estimates the baseline model (9) for computer and electronics intermediates using the t-likelihood with "known" values of the elasticity of relative demand (σ) and supply (ρ), and by implication known values of β^P and β^X. I consider values of σ and ρ equal to 0, .25, .5, .75, 1, 2, 5, 10, 1000, and 100000, which should run the gamut of possible priors, and estimate the mismeasurement parameter for all possible binary combinations of these,[18] as well as all three lag structures. As can be seen, the estimated mismeasurement parameter remains solidly negative across all possibilities within each lag structure. Moreover, for elasticities of demand less than or equal to .50, the estimates hardly

[18]Except σ=ρ=0, when the supply and demand curves are both vertical and the model is unidentified.

Table 8: Estimated Computer & Electronics Intermediates Mismeasurement Parameters for Different "Known" Values of σ and ρ (t-likelihood)

by ρ: / by σ	0	.25	.50	.75	1	2	5	10	1000	100000
					no lags					
0		-.45	-.45	-.45	-.45#	-.45	-.45	-.45#	-.45#	-.45
.25	-.48#	-.47#	-.47	-.47#	-.47#	-.47	-.47	-.47#	-.47#	-.47
.50	-.48#	-.48#	-.48#	-.47#	-.47#	-.48	-.48	-.48	-.48#	-.48
.75	-.48#	-.49#	-.49#	-.48*	-.48*	-.48#	-.49	-.49#	-.49#	-.49
1	-.48#	-.49#	-.49#	-.49*	-.48*	-.48#	-.49	-.49#	-.50	-.50
2	-.48#	-.49#	-.50#	-.50*	-.50*	-.48#	-.49	-.51#	-.51	-.51
5	-.48#	-.50#	-.51#	-.53*	-.53*	-.49*	-.48#	-.48*	-.48#	-.48#
10	-.48#	-.50#	-.52#	-.54*	-.56*	-.52#	-.48*	-.48*	-.47*	-.47#
1000	-.48#	-.49#	-.52#	-.55*	-.60*	-.54#	-.48*	-.48*	-.46*	-.46#
100000	-.48#	-.49#	-.52#	-.55*	-.60*	-.54#	-.48*	-.48*	-.46*	-.46#
					1 lag					
0		-.35	-.35	-.35	-.35	-.35	-.35	-.35	-.35	-.35
.25	-.44*	-.37	-.35	-.34	-.34	-.35	-.35	-.35	-.36	-.36
.50	-.44*	-.43#	-.37	-.35	-.35	-.34	-.35	-.36	-.36	-.36
.75	-.44*	-.44#	-.41	-.37	-.36	-.35	-.35	-.36	-.36	-.36
1	-.44*	-.44#	-.43#	-.40#	-.37	-.35	-.35	-.36	-.37	-.37
2	-.44*	-.44#	-.44#	-.43*	-.41*	-.36	-.34	-.35	-.36	-.36
5	-.44*	-.44#	-.43*	-.42*	-.41*	-.35#	-.35	-.33	-.34	-.34
10	-.44*	-.43#	-.42*	-.41*	-.39*	-.36#	-.37#	-.33#	-.34#	-.34
1000	-.44*	-.43#	-.41*	-.39*	-.38*	-.40*	-.37*	-.34*	-.35#	-.35#
100000	-.44*	-.43#	-.41*	-.39*	-.38*	-.40*	-.37*	-.34*	-.35#	-.35#
					2 lags					
0		-.40	-.40	-.40	-.40	-.40	-.40	-.40	-.40	-.40
.25	-.45#	-.43	-.41	-.43	-.42	-.41	-.41	-.41	-.41	-.41
.50	-.45#	-.45	-.43	-.43	-.43	-.44	-.43	-.43	-.43	-.43
.75	-.45#	-.47#	-.43	-.43	-.43	-.44	-.45	-.44	-.45	-.45
1	-.45#	-.47#	-.45#	-.43	-.43	-.44	-.46	-.46	-.47	-.47
2	-.45#	-.47#	-.47#	-.47*	-.43#	-.42	-.45	-.48	-.47	-.47
5	-.45#	-.46#	-.46*	-.46*	-.46*	-.40#	-.44	-.45	-.43	-.43
10	-.45#	-.46#	-.46#	-.45*	-.44*	-.40#	-.44#	-.45#	-.42#	-.42#
1000	-.45#	-.45#	-.45#	-.43*	-.41*	-.46*	-.45*	-.44*	-.41#	-.41#
100000	-.45#	-.45#	-.45#	-.43*	-.41*	-.46*	-.45*	-.44*	-.41#	-.41#
					3 lags					
0		-.32	-.32	-.32	-.32	-.32	-.32	-.32	-.32	-.32
.25	-.27	-.33	-.33	-.32	-.31	-.32	-.32	-.32	-.32	-.32
.50	-.27	-.34#	-.33	-.33	-.33	-.32	-.33	-.33	-.33	-.33
.75	-.27	-.34#	-.35	-.34	-.34#	-.34	-.33	-.34	-.35	-.35
1	-.27	-.33#	-.35	-.35#	-.35	-.34	-.34	-.35	-.37	-.37
2	-.27	-.31#	-.35#	-.36*	-.37*	-.36	-.32	-.34	-.37	-.37
5	-.27	-.28	-.31#	-.35*	-.38*	-.37#	-.34	-.30	-.29	-.29
10	-.27	-.27	-.27#	-.33#	-.36*	-.39*	-.33#	-.34	-.33	-.33
1000	-.27	-.26	-.24	-.28#	-.32*	-.36#	-.35#	-.34#	-.32#	-.32#
100000	-.27	-.26	-.24	-.28#	-.32*	-.36#	-.35#	-.34#	-.32#	-.32#

Notes: results based upon (9) augmented with lags as in (12), with the exception that β^P and β^X are not estimated, but rather taken as given by the values implied by σ and ρ. Absolute value of t-statistic (*) greater than 3; (#) greater than 2.

vary at all from those of the baseline results, ranging from -.48 to -.45 with no lags and from -.34 to -.27 with 3 lags across all elasticities of supply. Thus, while the point estimates of the elasticity of substitution underlying the results above are more similar to those of Peter and Ruane (2025) than those of Atalay (2017) and Miranda-Pinto (2021), the estimate of the mismeasurement parameter would be virtually identical if the reader were to impose a value of σ equal to the point estimates of Atalay & Miranda-Pinto and any value of ρ they like.[19]

That said, it is not the case that the estimates of the elasticities & β's do not matter. The reader will note that many of the point estimates in Table 8, when evaluated using the bootstrap-se standard error, do not have absolute t-statistics larger than 3 or even 2 (denoted by * and $^{\#}$, respectively), despite the fact that the baseline estimates in row 1 of Table 6 have absolute t-statistics of 4 or 3. This stems from the significant cross-industry heterogeneity in the β*'s* (and by implication σ's and ρ's) found in the model of rows 9 and 10 of Table 6. When the β's are taken as unknown, the estimated values vary with each bootstrap sample of industries, and this limits the variation in the estimate of the mismeasurement parameter, which stays relatively constant across subsamples. In contrast, when the β's are taken as known and fixed across industry samples, greater variation in the estimate of the mismeasurement parameter emerges. t-statistics of -4.1, -3.4, -3.3, and -3.2 with 0 through 3 lags in the baseline estimates of row 1 of Table 6, which allow β to vary across bootstrap samples, become -2.6, -2.8, -3.3, and -2.7, respectively, if one takes the point estimates of β^P and β^X found in each specification and imposes it on the subsequent bootstrap.[20]

To summarize, the results above show that movements of intermediate input price and non-own-use quantity demanded, net of movements implied by total factor productivity growth in the intermediate input producing industry, vary systematically with the share of computer and electronics inputs in an industry's cost structure. This result is robust to the use of alternative normal or thick tailed likelihoods and multiple lags of the dependent variables.

[19]The on-line appendix presents a similar table using the normal likelihood, finding that the estimate of the mismeasurement parameter is always solidly negative and for $\sigma \leq .5$ ranges from -.46 to -.41 with 0 lags and from -.44 to -.33 with 3 lags.

[20] Similarly, with normal likelihood estimation, t-stats of -2.2, -2.1, -1.9, -2.3 in the baseline specification with 0 through 3 lags become -1.8, -1.4, -1.7, and -1.9, respectively, if the point estimates of the β found in each specification are imposed on the subsequent bootstrap.

It is found with industry and year fixed effects (using within industry variation) and when estimated across industry means alone (using cross industry variation). It is robust to deleting individual industries from the sample, so it is not determined by any industry. It is robust to allowing heterogeneity across intermediate input producing industries in the response of price and quantity demanded to total factor productivity growth; in fact, heterogeneity in the elasticities plays an important role in the bootstrap, as the mismeasurement estimate varies less when the average sample response to total factor productivity growth is allowed to vary across industry samples. It is not driven by any bias due to endogeneity of total factor productivity growth through capacity utilization mismeasurement, as the result is robust to adjustments for business cycle variation and, furthermore, holds for long run industry averages, where capacity utilization is hardly relevant. Finally, the point estimates are stable across the full range of possible elasticity of demand and supply combinations, although, as already noted, bootstrapped significance depends upon acknowledging that the average sample elasticity does vary with the industries considered. Put simply, the greater the input share of computer and electronics inputs, the more price and quantity growth deviates from predicted values based upon measured total factor productivity growth. No such robust relationship is found for any other input.

V. Confirmatory Evidence from TFP Growth Projected on Factor Shares

The variation underlying the preceding results is admittedly difficult to intuit. At the least squares solution which maximizes each likelihood, the "regressors" are the derivatives of the non-linear equations with respect to each parameter:

$$(16)\quad \frac{\partial \hat{P}_{it}^{M}}{\partial \beta^{P}} = \hat{A}_{it}^{M} + \gamma_j\left(\theta_{jit} - \Omega_{jit}\right),\ \ \frac{\partial \hat{P}_{it}^{M}}{\partial \gamma_j} = \beta^{P}\left(\theta_{jit} - \Omega_{jit}\right) + \theta_{jit}$$
$$\frac{\partial \hat{X}_{it}^{M}}{\partial \beta^{X}} = \hat{A}_{it}^{M} + \gamma_j\left(\theta_{jit} - \Omega_{jit}\right),\ \ \frac{\partial \hat{X}_{it}^{M}}{\partial \gamma_j} = \beta^{X}\left(\theta_{jit} - \Omega_{jit}\right) - \theta_{jit}.$$

Thus, at the point estimates, the mismeasurement parameter is determined by the variation of $\beta\left(\theta_{jit} - \Omega_{jit}\right) \pm \theta_{jit}$ with $\hat{P}_{it}^{M}$ & $\hat{X}_{it}^{M}$ that is orthogonal to unobserved true total factor productivity growth $\hat{A}_{it}^{M} + \gamma_j\left(\theta_{jit} - \Omega_{jit}\right)$. Unfortunately, little credibility is likely to be gained by pointing readers to variation that is orthogonal to constructed variation. One might be tempted to gain some insight by regressing $\hat{P}_{it}^{M}$ & $\hat{X}_{it}^{M}$ on $\hat{A}_{it}^{M}$, θ_{jit} and Ω_{jit}, but this is senseless, because the whole point of mismeasurement is that it affects $\hat{A}_{it}^{M}$, so not much can credibly be learnt from variation of factor shares that is orthogonal to measured total factor productivity growth.

We can, instead, look for confirmation in variation in factor shares that is not orthogonal to measured total factor productivity growth. Specifically, consider running measured total factor productivity growth on own and upstream factor shares plus industry & year fixed effects

$$(17)\quad \hat{A}_{it}^{M} = \beta^{\theta}\theta_{jit} + \beta^{\Omega}\Omega_{jit} + \eta_{i}^{A} + \eta_{t}^{A} + \varepsilon_{it}.$$

The mismeasurement model predicts that $\beta^{\theta} = -\beta^{\Omega}(= -\gamma_{j})$, mismeasurement of what users accomplish with computer inputs should result in opposite effects on measured productivity growth in users of users. While proponents of the benefits associated with computer related inputs would expect to find a positive coefficient on own factor shares θ_{jit}, non-zero effects of the opposite sign on the concatenated computer factor shares of upstream industries Ω_{jit} are hard to explain as a benefit of computer use. If the mismeasurement model is true, efficiency can be gained by imposing the constraint $\beta^{\theta} = -\beta^{\Omega}$ and estimating

$$(18)\quad \hat{A}_{it}^{M} = \beta^{\theta-\Omega}\left(\theta_{jit} - \Omega_{jit}\right) + \eta_{i}^{A} + \eta_{t}^{A} + \varepsilon_{it}.$$

Table 9 below reports such estimates for computer capital and intermediate inputs, augmented with 0, 1 or 2 lags of total factor productivity growth on the righthand side.[21] Reported coefficients with lags are the long run effect on TFP growth, i.e. the estimated effect divided by 1 minus the sum of lag coefficients. Also reported are results using industry mean data and just constant terms and factor shares as regressors. Standard errors are based upon the bootstrap-se with industry clustered draws.

As seen in the table, the results for computer & electronics intermediates are consistent with the mismeasurement model, albeit with somewhat more imprecision and higher p-values than found with systems estimation earlier. The coefficients on θ_{jit} are all

[21]Using normal and t-likelihoods a 3rd lag is estimated to have a near zero coefficient of .012 & .007 and bootstrap-se p-values of .807 & .857, respectively, while the normal likelihood finds .05 level bootstrap-se significant 1st and 2nd lag coefficients of -.17 to -.13 in specifications with 1 through 3 lags, while the t-likelihood finds .05 significant 1st lags of -.06 to -.08 in similar specifications.

Table 9: Measured Total Factor Productivity Growth Projected on Computer Input Shares

	normal likelihood				t likelihood			
	means	no lags	1 lag	2 lags	Means	no lags	1 lag	2 lags
				computer hardware capital				
β^{θ}	-.03 (.25)	-.26 (.22)	-.31 (.25)	-.50 (.38)	-.15 (.31)	-.28 (.27)	-.48 (.35)	-.60 (.43)
	[-.17/.11]	[-.30/.02]	[-.36/-.03]	[-.57/-.35]	[-.32/.17]	[-.33/-.04]	[-.55/-.38]	[-.76/-.38]
β^{Ω}	.39 (1.2)	-1.8 (2.5)	-2.9 (2.5)	-3.7 (2.3)	.48 (1.2)	1.7 (1.8)	1.9 (1.7)	.38 (1.5)
	[-.11/1.1]	[-2.5/.54]	[-3.9/-.75]	[-4.9/-2.1]	[-.13/.91]	[.53/2.4]	[.87/2.3]	[-.29/.91]
$\beta^{\theta} = -\beta^{\Omega}$	.746	.401	.185	.068	.755	.405	.383	.885
$\beta^{\theta\text{-}\Omega}$	-.02 (.25)	-.23 (.28)	-.27 (.32)	-.45 (.45)	-.12 (.29)	-.28 (.30)	-.49 (.36)	-.59 (.41)
	[-.17/.11]	[-.29/.06]	[-.34/.05]	[-.52/-.28]	[-.30/.17]	[-.35/-.02]	[-.57/-.40]	[-.75/-.38]
				computer software capital				
β^{θ}	.15 (.17)	.07 (.20)	.13 (.21)	.02 (.27)	.10 (.11)	.04 (.30)	-.08 (.52)	-.05 (.45)
	[.08/.20]	[.02/.15]	[.07/.22]	[-.05/.18]	[.02/.11]	[-.05/.21]	[-.20/.24]	[-.24/.23]
β^{Ω}	-.34 (.63)	.48 (.63)	.13 (.68)	-.14 (.75)	-.27 (.47)	.18 (1.0)	.29 (1.1)	-.16 (1.1)
	[-.52/.10]	[.16/.69]	[-.23/.36]	[-.55/.07]	[-.35/.04]	[-.53/.86]	[-.60/.64]	[-.91/.28]
$\beta^{\theta} = -\beta^{\Omega}$	.772	.321	.666	.862	.681	.799	.773	.791
$\beta^{\theta\text{-}\Omega}$	.13 (.14)	.07 (.21)	.13 (.22)	.02 (.26)	.08 (.10)	.04 (.32)	-.07 (.52)	-.05 (.44)
	[.08/.17]	[.02/.15]	[.06/.22]	[-.05/.17]	[.01/.10]	[-.04/.19]	[-.19/.23]	[-.24/.23]
				computer & electronics intermediates				
β^{θ}	.53 (.23)	.68 (.33)	.67 (.32)	.67 (.31)	.52 (.17)	.56 (.27)	.57 (.25)	.62 (.26)
	[.47/.59]	[.59/.85]	[.58/.84]	[.46/.78]	[.32/.58]	[.51/.68]	[.52/.66]	[.48/.68]
β^{Ω}	-.92 (.61)	-1.4 (.75)	-1.3 (.78)	-1.3 (.72)	-.87 (.47)	-.92 (.54)	-.96 (.52)	-1.1 (.60)
	[-1.2/-.63]	[-1.9/-1.1]	[-1.8/-1.0]	[-1.7/-1.1]	[-1.2/-.44]	[-1.1/-.75]	[-1.2/-.79]	[-1.3/-.86]
$\beta^{\theta} = -\beta^{\Omega}$	.321	.185	.235	.219	.260	.345	.309	.268
$\beta^{\theta\text{-}\Omega}$	.42 (.10)	.36 (.28)	.36 (.27)	.35 (.23)	.42 (.10)	.40 (.26)	.39 (.24)	.43 (.22)
	[.29/.43]	[.35/.55]	[.35/.59]	[.29/.37]	[.25/.43]	[.39/.57]	[.38/.58]	[.35/.44]
				computer systems design intermediates				
β^{θ}	.41 (.28)	-.50 (1.1)	-.76 (.99)	-.87 (.91)	.34 (.32)	.46 (1.3)	.34 (1.3)	.11 (1.2)
	[.32/.54]	[-.77/.27]	[-1.0/-.05]	[-1.2/-.36]	[.15/.51]	[.00/1.6]	[-.15/1.6]	[-.30/1.3]
β^{Ω}	-.19 (.96)	.20 (2.5)	-.44 (2.4)	-.85 (2.8)	-.29 (.72)	-1.9 (2.9)	-2.0 (2.3)	-1.9 (1.9)
	[-.48/.07]	[-1.2/1.1]	[-1.7/.16]	[-2.4/-.18]	[-1.0/-.06]	[-3.1/-.53]	[-3.1/-.95]	[-3.0/-1.2]
$\beta^{\theta} = -\beta^{\Omega}$	.817	.895	.615	.535	.948	.578	.387	.274
$\beta^{\theta\text{-}\Omega}$	.40 (.27)	-.51 (1.1)	-.77 (.99)	-.89 (.89)	.33 (.29)	.46 (1.3)	.38 (1.3)	.14 (1.1)
	[.31/.53]	[-.75/.26]	[-1.0/-.05]	[-1.1/-.35]	[.17/.49]	[.07/1.6]	[-.04/1.7]	[-.24/1.3]

Notes: Estimating equations as in (17) for β^{θ} & β^{Ω} and (18) for $\beta^{\theta\text{-}\Omega}$, with a constant term for industry means specification, or industry & year fixed effects plus indicated number of lags of TFP growth otherwise. 59 observations with industry means, 1416 minus number of lags*59 otherwise. Bootstrap-se standard errors in () based upon 99 industry-clustered draws; min-max delete-one industry coefficient range in []. $\beta^{\theta} = -\beta^{\Omega}$: p-value of the test based upon bootstrap-se covariance matrix & chi^2 distribution.

positive with t-stats of between 2 and 3, the coefficients on Ω_{jit} are all negative with t-stats of between -1.5 and -2, and the delete-one-industry min to max ranges for the two coefficients are positive to positive and negative to negative, respectively, across all likelihoods and lag orders. The restriction $\beta^{\theta} = -\beta^{\Omega}$ is not rejected and point estimates of $\beta^{\theta-\Omega}$ found when the constraint is imposed are very similar to those found for -γ_j earlier above, ranging between .35 and .43 across the normal and t-likelihoods with various lags, and between .25 and .59 across all min to max delete-one-industry ranges. In contrast, for the other computer inputs, while the restriction $\beta^{\theta} = -\beta^{\Omega}$ is also not rejected, we find that the point estimates of β^{θ}& β^{Ω} are of the same sign in a few specifications and these, along with $\beta^{\theta-\Omega}$, vary in sign across the min-max delete-one-industry range. In fact, when these same specifications are run across all 59 input-output intermediate inputs and 14 BEA total factor productivity input shares described earlier above, computer & electronics intermediates are one of only three for which the signs of β^{θ}& β^{Ω} are opposite and, along with $\beta^{\theta-\Omega}$, do not vary across specifications, and the only across all 73 inputs for which the signs of $\beta^{\theta}, \beta^{\Omega}$ & $\beta^{\theta-\Omega}$ do not vary in any delete-one-industry min-max range for any of the eight normal and t-likelihood specifications of Table 9. The computer & electronics intermediates results, while not always statistically significant, are unique in the robustness of their sign (and even their magnitude, as seen in the table) to the likelihood and lag specification and delete-one-industry range.

Table 10 runs some of the same specification checks considered earlier for computer & electronics intermediates. We see that when the regressions are run without industry, year, or industry and year fixed effects the estimated values of $\beta^{\theta-\Omega}$, with the constraint $\beta^{\theta} = -\beta^{\Omega}$ imposed, change little from what is found with both fixed effects in Table 9, with point estimates from .39 to .43 and a delete-one-industry min to max range of .31 to .51 across both likelihoods and three lag structures. Without industry or industry and year fixed effects, the point estimates and min-max delete-one-industry ranges of β^{θ} and β^{Ω} are always positive and negative respectively. However, as in the case of systems estimation earlier, the year fixed effects do play a greater role in the results, as without these the t-stats on $\beta^{\theta-\Omega}$ fall below 2 and the estimated values of β^{Ω} turn positive. As argued earlier, controlling for business cycle mismeasurement and variation with year fixed effects seems reasonable. Additional controls for industry specific business cycle responses through industry specific unemployment controls, in the bottom panels of the table, have no substantive effects on the baseline estimates with year fixed effects, as was the case with systems estimation earlier.

Table 10: Specification Checks: Measured Total Factor Productivity Growth
Projected on Computer & Electronics Intermediate Input Shares

	normal likelihood			t likelihood		
	no lags	1 lag	2 lags	no lags	1 lag	2 lags
			no industry fixed effects			
β^{θ}	.55 (.14)	.55 (.16)	.55 (.17)	.52 (.17)	.52 (.19)	.53 (.20)
	[.53/.60]	[.53/.65]	[.53/.64]	[.50/.63]	[.51/.67]	[.50/.68]
β^{Ω}	-1.0 (.34)	-.97 (.37)	-1.0 (.39)	-.90 (.40)	-.90 (.42)	-.95 (.48)
	[-1.2/-.93]	[-1.2/-.91]	[-1.2/-.89]	[-1.2/-.76]	[-1.2/-.79]	[-1.3/-.70]
$\beta^{\theta} = -\beta^{\Omega}$	.036	.055	.060	.123	.129	.179
$\beta^{\theta\text{-}\Omega}$	.41 (.08)	.42 (.08)	.42 (.09)	.41 (.08)	.42 (.08)	.43 (.10)
	[.31/.42]	[.36/.43]	[.35/.44]	[.34/.43]	[.37/.43]	[.37/.45]
			no year fixed effects			
β^{θ}	.40 (.31)	.39 (.30)	.37 (.31)	.25 (.26)	.19 (.26)	.18 (.26)
	[.33/.49]	[.31/.50]	[.21/.44]	[.16/.40]	[.11/.40]	[.02/.33]
β^{Ω}	-.44 (.58)	-.39 (.55)	-.33 (.54)	.08 (.38)	.34 (.38)	.52 (.54)
	[-.59/-.23]	[-.57/-.15]	[-.53/-.02]	[-.06/.36]	[.19/.60]	[.28/1.1]
$\beta^{\theta} = -\beta^{\Omega}$	.933	.994	.919	.208	.039	.062
$\beta^{\theta\text{-}\Omega}$	.39 (.28)	.39 (.27)	.39 (.27)	.41 (.26)	.42 (.25)	.45 (.22)
	[.38/.49]	[.39/.51]	[.22/.40]	[.41/.45]	[.42/.45]	[.37/.46]
			no year or industry fixed effects			
β^{θ}	.48 (.14)	.49 (.15)	.49 (.15)	.46 (.13)	.45 (.14)	.46 (.16)
	[.40/.50]	[.45/.50]	[.45/.51]	[.42/.48]	[.43/.46]	[.42/.48]
β^{Ω}	-.71 (.32)	-.70 (.33)	-.72 (.33)	-.61 (.31)	-.56 (.32)	-.59 (.37)
	[-.81/-.58]	[-.76/-.61]	[-.80/-.59]	[-.68/-.42]	[-.62/-.40]	[-.66/-.30]
$\beta^{\theta} = -\beta^{\Omega}$	.247	.283	.248	.444	.586	.606
$\beta^{\theta\text{-}\Omega}$	.41 (.08)	.42 (.08)	.42 (.09)	.42 (.08)	.42 (.08)	.43 (.10)
	[.31/.43]	[.35/.43]	[.35/.44]	[.33/.43]	[.36/.44]	[.36/.45]
			both fixed effects with additional industry specific unemployment levels controls			
β^{θ}	.76 (.38)	.77 (.37)	.77 (.36)	.57 (.31)	.57 (.33)	.66 (.33)
	[.66/.96]	[.68/.99]	[.52/.93]	[.38/.62]	[.25/.63]	[.27/.74]
β^{Ω}	-1.6 (.77)	-1.7 (.83)	-1.7 (.79)	-.98 (.62)	-1.0 (.61)	-1.2 (.67)
	[-2.2/-1.4]	[-2.3/-1.4]	[-2.2/-1.4]	[-1.2/-.77]	[-1.2/-.65]	[-1.5/-.89]
$\beta^{\theta} = -\beta^{\Omega}$	.098	.126	.125	.376	.316	.239
$\beta^{\theta\text{-}\Omega}$	.34 (.33)	.34 (.31)	.35 (.27)	.37 (.30)	.38 (.30)	.44 (.26)
	[.33/.53]	[.33/.52]	[.34/.37]	[.27/.38]	[.16/.39]	[.18/.45]
			both fixed effects with additional industry specific unemployment changes controls			
β^{θ}	.64 (.34)	.61 (.33)	.60 (.33)	.49 (.23)	.48 (.24)	.54 (.31)
	[.52/.80]	[.51/.78]	[.26/.71]	[.35/.54]	[.37/.56]	[.27/.59]
β^{Ω}	-1.2 (.79)	-1.2 (.84)	-1.2 (.79)	-.73 (.51)	-.74 (.54)	-.92 (.68)
	[-1.7/-.91]	[-1.7/-.85]	[-1.5/-.87]	[-.91/-.58]	[-1.0/-.57]	[-1.1/-.64]
$\beta^{\theta} = -\beta^{\Omega}$	.266	.345	.332	.538	.537	.460
$\beta^{\theta\text{-}\Omega}$	.35 (.25)	.35 (.23)	.34 (.23)	.37 (.22)	.37 (.21)	40 (.26)
	[.33/.35]	[.34/.38]	[.10/.35]	[.24/.38]	[.29/.38]	[.14/.41]

Notes: Industry specific unemployment controls as in (14) earlier. Otherwise, as in Table 9.

In sum, when measured total factor productivity growth is projected on computer and electronics intermediates expenditure shares, it is robustly positively correlated with own expenditure on those inputs and negatively correlated with upstream expenditure on those inputs. Moreover, when the mismeasurement constraint is imposed on the coefficients of this regression, the combined point estimate is very close to that found in the earlier price and quantity regressions.

VI. Summary

In the above, the coefficient $\beta^{\theta-\Omega}$ found when projecting measured productivity growth $\hat{A}_{it}^{M}$ on own minus concatenated upstream factor shares $\theta_{jit} - \Omega_{jit}$, for j = computer & electronics intermediates, is very close to the negative of the mismeasurement parameter γ_j that adjusts prices and downstream intermediate input demand quantities, $\hat{P}_{it}^{T} = \hat{P}_{it}^{M} - \gamma_j\theta_{jit}$ & $\hat{X}_{it}^{T} = \hat{X}_{it}^{M} + \gamma_j\theta_{jit}$, to make their relation to adjusted productivity growth $\hat{A}_{it}^{T} = \hat{A}_{it}^{M} + \gamma_j(\theta_{jit} - \Omega_{jit})$ consistent across varying levels of use θ_{jit} and use of users Ω_{jit}. Thus, putting each estimate into the equations of the other, if $\beta^{\theta-\Omega}$ adjusted price & quantity ($\hat{P}_{it}^{M} + \beta^{\theta-\Omega}\theta_{jit}$ & $\hat{X}_{it}^{M} - \beta^{\theta-\Omega}\theta_{jit}$) are projected on $\beta^{\theta-\Omega}$ adjusted productivity [$\hat{A}_{it}^{M} - \beta^{\theta-\Omega}(\theta_{jit} - \Omega_{jit})$], the relationship requires no further adjustment to be consistent across levels of θ_{jit} & Ω_{jit}, while if γ_j adjusted productivity growth [$\hat{A}_{it}^{M} + \gamma_j(\theta_{jit} - \Omega_{jit})$] is projected on θ_{jit} & Ω_{jit}, no significant productivity effects are found from the use of computer inputs. Two completely different estimation methodologies, using different dependent variables and identification from completely different variation, give roughly the same answer: productivity growth attributable to computer & electronics intermediates is overestimated by about .3 to .4 per percentage expenditure on those inputs. These results are robust to a broad range of likelihood, specification and sample changes.

Table 11 reports the 1997-2021 mismeasurement adjustments of BEA industry level gross output and value added growth based upon the estimated -.33 mismeasurement rate of computer and electronics intermediates factor augmenting productivity growth found in the baseline specification with 3 lags in Table 3 earlier.[22] As one might expect, the largest

[22]The adjustment of output growth is -.33θ_{jit}, while the adjustment of value added growth is -.33*(θ_{jit}-Ω_{jit}) times the ratio of gross output to value added.

Table 11: Average Annual Mismeasurement Adjustment 1997-2021 by Industry

	measured growth		adjustment	
	gross output	value added	gross output	value added
computer & electronic products	.041	.107	-.051	-.080
other transportation equipment	.011	.017	-.018	-.023
motor vehicles, bodies & trailers & parts	.015	.031	-.010	-.018
electrical equipment, appliances & components	-.008	.005	-.010	-.014
machinery	-.001	.008	-.007	-.009
broadcasting & telecommunications	.041	.051	-.007	-.008
fabricated metal products	-.004	-.002	-.005	-.007
computer systems design & related services	.068	.089	-.005	-.005
⋮	⋮	⋮	⋮	⋮
support activities for mining	.033	.042	.000	.001
farms	.011	.015	.000	.001
mining, except oil & gas	-.011	-.008	.000	.001
pipeline transportation	.008	.049	.000	.001
securities, commodity contracts & investments	.034	.021	.000	.002
apparel, leather & allied products	-.067	-.046	.000	.002
food services & drinking places	.022	.015	.000	.002
water transportation	.000	.011	.000	.006
private sector gross domestic product:		.023		-.003

Notes: Based on -.33 mismeasurement parameter in computer & electronics intermediates with 3 lags in Table 3. Adjustment of output growth = -.33θ_{jit}; adjustment of value added growth = -.33*(θ_{jit}-Ω_{jit}) times the ratio of gross output to value added.

negative gross output and value added adjustments are in durables goods industries that use these inputs directly in production, while positive value added adjustments occur in non-durable goods industries whose upstream suppliers use these inputs. The biggest adjustment is for computer & electronics products, itself a big user of its own inputs, but as has been repeatedly emphasized above, the point estimates do not depend heavily upon any individual industry, including this one. While the adjustments of computer and electronics output may seem large, Aizcorbe & Ripperger-Suhler (2024), specialists within the BEA itself, estimate that chain drift in hedonic price indices has resulted in these overstating the 2011-2020 decline in computer & notebook prices by 6 and 8 percent per annum, respectively.

As noted in Table 11, a -.33 exaggeration of factor augmenting technical change in the use of computer & electronics intermediates implies that annual private sector GDP growth is overestimated by .003, or about ⅛$^{\text{th}}$ of the .023 growth per annum recorded between 1997 and 2021. As private sector total factor productivity growth during this period averages .006 per annum, absent any adjustment for mismeasurement of the growth of capital inputs, this would also imply a 50 percent reduction in estimated total factor productivity growth. Unfortunately, the estimates above provide little guidance as to potential mismeasurement of the growth of capital goods in the BEA productivity accounts, as US equipment investment is increasingly sourced from other countries. According to the input-output tables, while the nominal value of computer & electronics output used in gross fixed investment in equipment rose from $142 billion in 1997 to $208 billion in 2021, during the same period imports of computer & electronics for that purpose rose from $37 billion to $131 billion, so that the share of domestically sourced computers & electronics in investment fell from 74% to 37%. The estimates above are based upon domestically sourced computer & electronics intermediates used in domestic production which, again based upon the input-output tables, have accounted for more than 94% of total computer & electronics intermediates used in domestic production in every year from 1997 to 2021. While they provide insight into mismeasurement associated with domestically produced intermediates, they have little to say regarding mismeasurement of the growth of computer & electronics capital equipment in the BEA productivity accounts, as in recent decades these are driven by imported goods. As noted earlier, point estimates in the on-line appendix of mismeasurement of the growth of primary inputs using a related dark matter methodology indicate an overstatement of computer and software capital growth. But all are statistically indistinguishable from 0 and sensitive to the specification and the removal of individual industries from the sample.

While this paper is not about the methodology of price indices, some perspective on the results above can be found by considering two characteristics of durable goods that impart a downward bias to model matching and hedonic price indices, two common price deflation techniques used in the US.[23] First, as emphasized by Harper (2007), the prices of durable

[23]A third technique used heavily in the evaluation of automobiles and other products with frequent model changes (see Groshen et al 2017) asks manufacturers to identify the cost increment associated with new characteristics and treats that cost increment as quality, i.e. assumes cost per unit quality remains constant. This is problematic, not least because it treats government mandated features such as catalytic converters and fuel

goods reflect the net present value of the flow of quasi-rents from their use. However, as technology progresses durable goods become obsolete because of issues with interoperability and the shadow value of complementary resources. Thus, old computers are regularly scraped not because they wear out but because they lack the capacity needed to interface with more modern software and, most importantly, because the growing capabilities of alternatives raises the shadow value of users' time. Obsolescence means that the same physical product purchased in later years is not the same product from the point of view of consumers, as it has a shorter expected life. This leads hedonic price indices to overstate the value of increasing characteristics, as there is a hidden characteristic (expected service life) that is correlated with quality, and hence overstate price deflation. Similarly, matching the "same" product through time overstates deflation as later versions of the same good embody a shorter stream of expected rents. The use of the flow rental as the price for the flow value of durable goods services would address this issue, but unfortunately, outside of housing, in our national accounts the benefits of durables are measured using their sales prices and not their rentals.

Second, as emphasized by Aizcorbe and Copeland (2007), durable goods are purchased intermittently and have a product life cycle. When initially introduced higher quality models are purchased by consumers who place the highest valuation on quality. As prices decline over time, either due to cost reducing learning by doing or price discrimination, consumers who place a lower valuation on quality purchase them. Model matching price indices will overstate gains from price declines because those who place a low value on quality only gain from the part of the price decline that is below their low initial reservation value, while those with a high reservation value may not gain from price declines at all if they go on to purchase newer high quality vintages. Hedonic price indices will also overstate the value of quality and produce downward biased price trends because when comparing the prices and characteristics of contemporaneously sold vintages there is an omitted variable, the value of quality to those buying the goods, that is positively correlated with goods characteristics.

---

efficiency as providing quality gains to individual consumers equal to their costs, which would make such mandates unnecessary in the first place.

Methods such as model matching and hedonic price indices are well suited to measuring price and quality changes in environments where all consumers literally "consume" all products all of the time. Once one allows that products are durable and purchased intermittently by consumers whose characteristics vary systematically with the product life cycle, it is not hard to see that such techniques could easily produce upward biased estimates of quality-adjusted output growth. Since in the modern era computer intermediates are intimately tied to improvements in the characteristics and quality of goods, these biases could produce the exaggeration of the benefits of computer intermediate use found in this paper.

Bibliography

Acemoglu, Daron, David Autor, David Dorn, Gordon H. Hanson and Brendan Price (2014). "Return of the Solow Paradox? IT, Productivity, and Employment in US Manufacturing." *American Economic Review* 104: 394-399.

Aeppel, Timothy (2015). "Silicon Valley Doesn't Believe U.S. Productivity Is Down." *Wall Street Journal*, July 16.

Aizcorbe, Ana and Adam Copeland (2007). "Intermittent Purchases and Welfare-Based Price Deflators for Durable Goods." Bureau of Economic Analysis, working paper.

Aizcorbe, Ana and Daniel Ripperger-Suhler (2024). "Do Price Deflators for High-Tech Goods Overstate Quality Change?" BEA Working Paper Series, WP2024-5.

Alloway, Tracy (2015). "Goldman: How 'Grand Theft Auto' Explains One of the Biggest Mysteries of the U.S. Economy." *Bloomberg*, May 26.

Atalay, Enghin (2017). "How Important are Sectoral Shocks?" *American Economic Journal: Macroeconomics* 9: 254-280.

Benjamini, Yoav, and Yosef Hochberg (1995). "Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing." *Journal of the Royal Statistical Society, Series B* (*Methodological*) 57: 289-300.

Bils, Mark and Peter J. Klenow (2001). "Quantifying Quality Growth." *American Economic Review* 91: 1006-1030.

Byrne, David M., John G. Fernald and Marshall B. Reinsdorf (2016). "Does the United States have a Productivity Slowdown or a Measurement Problem?" *Brookings Papers on Economic Activity*, Spring 2016: 109-182.

Dempster, A.P, N. M. Laird and D. B. Rubin (1995). "Maximum Likelihood from Incomplete Data via the EM Algorithm." *Journal of the Royal Statistical Society. Series B* (*Methodological*) 1: 1-38.

Feenstra, Robert C. (1995). "Exact Hedonic Price Indexes." *Review of Economics and Statistics* 77: 634-653.

Gordon, Robert J. (1990). The Measurement of Durable Goods Prices. Chicago: University of Chicago Press.

Gordon, Robert J. (2000). "Does the "New Economy" Measure up to the Great Inventions of the Past?" *Journal of Economic Perspectives* 14: 49-74.

Gordon, Robert J. (2009). "Apparel Prices 1914-1993 and the Hulten/Bruegel Paradox." W. Erwin Diewert, John S. Greenlees and Charles R. Hulten, editors. Price Index Concepts and Measurement. University of Chicago Press.

Gordon, Robert J. and Todd VanGoethem (2007). "Downward Bias in the Most Important CPI Component: The Case of Rental Shelter 1914-2003." In Ernst R. Berndt and Charles R. Hulten, editors. Hard to Measure Goods and Services: Essays in Honor of Zvi Griliches. University of Chicago Press.

Groshen, Erica L., Brian C. Moyer, Ana M. Aizcorbe, Ralph Bradley and David M. Friedman (2017). "How Government Statistics Adjust for Potential Biases from Quality Change and New Goods in an Age of Digital Technologies: A View from the Trenches." *Journal of Economic Perspectives* 31: 187-210.

Hall, Peter (1992). The Bootstrap and Edgeworth Expansion. New York: Springer-Verlag, 1992.

Hall, Peter and Raoul LePage (1996). "On Bootstrap Estimation of the Distribution of the Studentized Mean." *Annals of the Institute of Statistical Mathematics* 48: 403-421.

Harper, Michael J. (2007). "Technology and the Theory of Vintage Aggregation." In Ernst R. Berndt and Charles R. Hulten, editors. Hard to Measure Goods and Services: Essays in Honor of Zvi Griliches. University of Chicago Press.

Hobijn, Bart (2002). "On Both Sides of the Quality Bias in Price Indexes." Federal Reserve Bank of New York, working paper.

Holm, Sture (1979). "A Simple Sequentially Rejective Multiple Test Procedure." *Scandinavian Journal of Statistics* 6: 65-70.

Hope, Adery C.A. (1968). "A Simplified Monte Carlo Significance Test Procedure." Journal of the Royal Statistical Society, Series B (Methodological) 30: 582-598.

Jockel, Karl-Heinz (1986). "Finite Sample Properties and Asymptotic Efficiency of Monte Carlo Tests." The Annals of Statistics 14: 336-347.

Jones, Charles I. (1995). "R&D-Based Models of Economic Growth." *Journal of Political Economy* 103: 759-784.

Jones, Charles I. (2002). "Sources of U.S. Economic Growth in a World of Ideas." *American Economic Review* 92: 220-239.

Liu, Chuanhai and Donald B. Rubin (1977). "ML Estimation of the t Distribution using EM and its Extensions, ECM and ECME." *Statistica Sinica* 5: 19-39.

Miranda-Pinto, Jorge (2021). "Production Network Structure, Service Share and Aggregate Volatility." *Review of Economic Dynamics* 39: 146-173.

Oliner, Stephen D. and Daniel E. Sichel (2000). "The Resurgence of Growth in the Late 1990s: Is Information Technology the Story?" *Journal of Economic Perspectives* 14: 3-22.

Pascal, Frederic, Lionel Bombrun, Jean-Yves Tourneret and Yannick Berthoumieu (2013).

Peter, Alessandra and Cian Ruane (2025). "The Aggregate Importance of Intermediate Input Substitutability." NBER Working Paper 31233.

Solow, Robert M. (1987). "We'd better watch out." *New York Times Book Review* (July 12): 36.

Stiroh, Kevin J (2002). "Information Technology and the US Productivity Revival: What Do the Industry Data Say?" *American Economic Review* 92: 1559-1576.

Syverson, Chad (2017). "Challenges to Mismeasurement Explanations for the US productivity Slowdown." *Journal of Economic Perspectives* 31: 165-186

Westfall, Peter H., and S. Stanley Young (1993). Resampling-Based Multiple-Testing: Examples and Methods for p-Value Adjustment. New York: John Wiley & Sons, 1993.

Young, Alwyn (2019). Online appendix for "Channelling Fisher: Randomization Tests and the Statistical Insignificance of Seemingly Significant Experimental Results." *Quarterly Journal of Economics* 134: 557-598.

Young, Alwyn (2022). "Consistency without Inference: Instrumental Variables in Practical Application." *European Economic Review* 147: 104112.