February 2005
Johannes Van Biesebroeck

Topic 3: Productivity (2 lectures)

Background

(don't do too much detail here)

- **Econometric approaches**

  Name of the game is to compare two production plans, where one uses more of input 1, but less of input 2. To compare them we have to control for the input substitution the technology allows.

  For the production function

  $$Q_{it} \quad = \quad A_{it}\ F_{(it)}(X_{it}),$$

  all productivity differences are concentrated in the multiplicative factor A_{it}, which differs between firms and changes over time. Equation (1) contains the different ingredients to calculate productivity and illustrates that productivity is intrinsically a relative concept,

  $$\log \frac{A_{it}}{A_{j\tau}} \quad = \quad \log \frac{Q_{it}}{Q_{j\tau}} - \log \frac{F_{(it)}(X_{it})}{F_{(j\tau)}(X_{j\tau})}.$$

  The calculation of the last term in (1) —the ratio of input aggregators— distinguishes three different methods.

  Two non-parametric approaches: index numbers (exact) — linear programming (flexible approximation to anything)

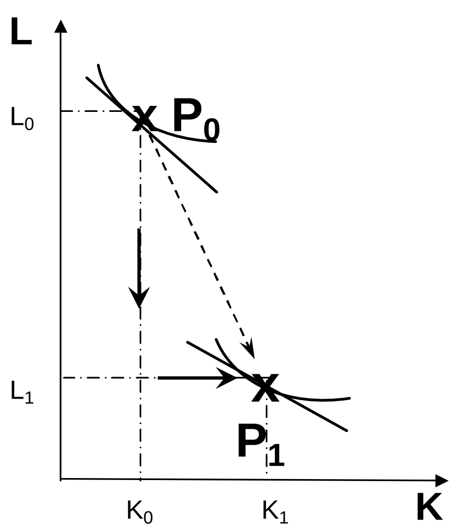

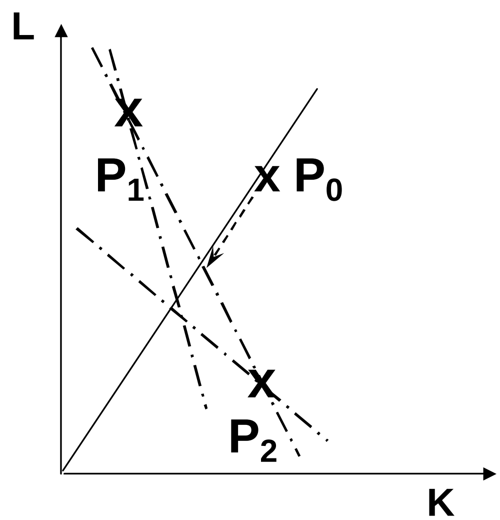

parametric estimation: the slope of the production function or isoquant shows how to control for input-substitution.

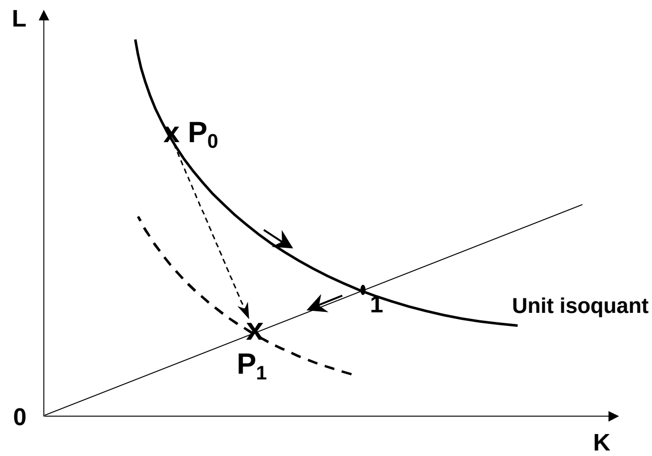

- **Functional forms:**

$$\text{Cobb-Douglas: } Q = AL^{\alpha_L}K^{\alpha_K}$$

many restrictions: homogeneous of degree $\alpha_L + \alpha_K$, unitary elasticities of substitution, constant factor shares, positive amounts of all inputs required, impossible to identify factor bias of productivity growth

$$\text{CES: } Q = A * [\alpha_L L^{-\rho} + \alpha_K K^{-\rho}]^{-\frac{\nu}{\rho}}$$

relaxes the unitary elasticity assumption, still all production factors are substitutes. Elasticity of substitution = ρ and RTS = ν.

$$\begin{aligned}\text{Translog: } \log Q &= \alpha_0 + \alpha_L \log L + \alpha_K \log K + \alpha_t t \\ &+ \frac{1}{2}\beta_{LL}(\log L)^2 + \frac{1}{2}\beta_{KK}(\log K)^2 + \frac{1}{2}\beta_{KL}\log L \log K \\ &+ \beta_{Lt}\log Lt + \beta_{Kt}\log Kt + \frac{1}{2}\beta_{tt}t^2\end{aligned}$$

relaxes some more assumptions—factors can be complements and technical change can be factor-biased— still cannot have zero input for anything.

$$\text{Generalized Leontief: } Q^{1/\nu} = \alpha_L L + \alpha_K K + \beta_{LK}\sqrt{LK}$$

relaxes that one

$$\text{Box-Cox: } Q(\lambda) = \begin{cases} \dfrac{Q^\lambda - 1}{\lambda} & \text{if } \lambda \neq 0 \\ \log Q & \text{if } \lambda = 0 \end{cases}$$

transformed inputs and outputs nest previous two (see Berndt and Khaled (1979)).

- **<u>Estimation:</u>**

  - Marschak and Andrews (1944) were first to point out that input factors were endogenous to the production decision. OLS cannot produce unbiased estimates of production function parameters. Zellner, Kmenta, and Dreze (1966) showed the exact assumptions needed on observability of productivity.
  - Christensen and Greene (1976) show that estimation can be improved by augmenting the cost function with $k-1$ share equations obtained from Shepard's lemma (for k inputs). More generally, duality theory provided a way out by suggestion to estimate cost functions, where factor prices could be assumed to be exogenous, for theory and examples see Fuss and McFadden (1978).
  - However, once imperfect competition is taken into account, output is endogenous as well and unless we are willing to assume constant returns to scale, output is an endogenous variable on the right hand side.
  - Mundlak (1996) points out that the cost function approach does not utilize all available information and is therefore inefficient. Together with the difficulty to calculate a price of capital, most recent application estimate the primal, the production function.
  - Klette and Griliches (1996) show that returns to scale estimates will be biased if the market is not perfectly competitive and deflated sales or value added is used instead of output as left hand side variable in the firm-level production function. Price effects will be misinterpreted as productivity differences. A proposed solution is to augment the model with a demand system and control for price changes. In practice they include aggregate output for the industry in the estimation equation, as firms price setting power is proportional to demand in their CES-demand system.

A functional form for the production function (or any other representation of technology) is assumed. (e.g. Cobb-Douglas)

$$q_{it} = \alpha_0 + \alpha_l l_{it} + \alpha_k k_{it} + \omega_{it} + \epsilon_{it},$$

The endogeneity problem, generated by the simultaneous choice of outputs and inputs, can be solved in three ways:

1. stochastic frontiers: assumptions are made on the parametric distribution of the unobserved productivity difference (ω). (exponential, truncated normal, quadratic in time, ...) The parameters of this distribution are estimated jointly with the parameters of the production function. E.g. $\omega_{it} = \exp(-\eta(t-T))\omega_i$ with $\omega_i \sim N^+(\mu, \sigma^2)$.

2. instrumental variables: GMM estimators are developed in the dynamic panel data literature. Lagged inputs and outputs are weakly endogenous and used as instruments. The GMM-SYS estimator developed by Blundell and Bond (1998) jointly estimates the production function in first differences and in levels. Twice lagged inputs (and output) is used as instrument for the first-differenced equation and lagged differences as instruments for the level equation. (arguments are only econometric, not economic)
   A different approach to generate instruments is provided by Syverson (2001), who exploits market segmentation and explicitly allows for differentiated goods (geographically differentiated).

3. The last method relies on a nonparametric inversion of another decision by the firm to provide information on the unobserved productivity term.

Olley and Pakes (1996)

Olley and Pakes (1996) solve two problems. (Both can be seen from Table 6: replicate in PS)

- Simultaneity bias of input parameters (as both inputs and output is chosen by the firm, conditional on the productivity which is unobservable by the econometrician) is solved by inverting the investment equation nonparametrically. In a first step the coefficient on variable inputs is estimated. In the third step, the capital coefficient is estimated using (implicitly defined) lagged productivity. Alternative approaches include Levinsohn and Petrin (2003), who invert the material input equation instead. In a different context, Van Biesebroeck (2003) inverts the equation governing labor input per shift to control for unobserved capital-productivity differences.
  α_L overestimated. If $\rho(l,k) > 0$ it might lead to underestimation of α_K.

- A second source of (selection) bias occurs when exit from the sample is conditional on productivity falling below a threshold and this threshold is a function of inputs. Capital is an obvious culprit. If capital markets are imperfect and losses in bankruptcy are proportional to the capital stock, this generates one such channel. An alternative possibility is when the profit function is increasing in capital, conditional on productivity. The solution is to include an intermediary step, where the survival probability for each plant at each time is predicted nonparametrically.

α_K **underestimated (because of negative correlation between k and ω.)**

They rely on the firm behavior model in Ericson and Pakes (1995), which Eugene might discuss later, for homogeneous products. The production function is Cobb-Douglas as before, accumulation equations for capital and age are standard, a crucial assumption is the exogenous Markov process for productivity, which is otherwise left unspecified. The distribution of ω_{t+1} conditional on all information known at t is determined by the family of distribution functions

$$F_\omega = \{F(.|\omega),\ \omega \in \Omega\}$$

Bellman equation for an incumbent can be written as

$$V_t(\omega_t, a_t, k_t) = \max\{\Phi, \sup_{i_t \geq 0} \pi_t(\omega_t, a_t, k_t) - c(i_t) + \beta E[V_{t+1}(\omega_{t+1}, a_{t+1}, k_{t+1})|J_t]\}$$

Both profit and value function depend on the three state variables, plus alternative state variables —factor prices and market structure— that don't vary across firms, and is subsumed in the t subscript. The solution to this control problem generates an exit rule and an investment equation:

$$\begin{aligned} \chi_t &= 1 \quad \text{if } \omega_t \geq \underline{\omega}_t(a_t, k_t), \\ & 0 \quad \text{otherwise}, \\ i_t &= i_t(\omega_t, a_t, k_t). \end{aligned}$$

The equations $\underline{\omega}_t(.)$ and $i_t(.)$ are part of the Markov perfect Nash equilibrium. (Threshold is increasing in k_t, alternative model for the same effect is to make $\partial\Phi/\partial k < 1$)

The self-selection generated by the exit behavior implies that $E[\omega_t|a_t, k_t, \omega_{t-1}, \chi_t = 1]$ will be increasing in k_t, leading to a negative bias in the capital coefficient.

Pakes (1994) shows that under certain conditions, $\partial\pi(k, \omega)/\partial k$ is increasing in ω and $\partial c(i, k_1)/\partial i \leq \partial c(i, k_2)/\partial i$ whenever $k_1 \geq k_2$, we know that the investment equation is strictly increasing in ω_t, for each (a_t, k_t) (these conditions guarantee that the total profit function is supermodular in k and ω or in layman's terms that capital and productivity are complements).

Under these conditions, the investment function can be inverted to generate:

$$\omega_t = h_t(i_t, a_t, k_t)$$

Substitute this in the production function generates the **first stage equation**:

$$\begin{aligned} q_{it} &= \alpha_l l_{it} + \phi_t(i_{it}, a_{it}, k_{it}) + \epsilon^1_{it} \qquad \text{where} \\ \phi_t(i_{it}, a_{it}, k_{it}) &= \alpha_0 + \alpha_a a_{it} + \alpha_k k_{it} + h_t(i_t, a_t, k_t). \end{aligned}$$

From this equation we can identify the labor coefficient.

The **second stage** of the model provides an estimate of the survival probability which will be needed later on. We can estimate and predict this probability by the following chain of logic:

$$\begin{aligned} & Pr\{\chi_{t+1} = 1 | \underline{\omega}_{t+1}(k_{t+1}, a_{t+1}), J_t\} \\ = \ & Pr\{\omega_{t+1} \geq \underline{\omega}_{t+1}(k_{t+1}, a_{t+1}) | \underline{\omega}_{t+1}(k_{t+1}, a_{t+1}), \omega_t\} \\ = \ & \psi_t\{\underline{\omega}_{t+1}(k_{t+1}, a_{t+1}), \omega_t\} \qquad \omega \text{ is Markov} \\ = \ & \psi_t(i_t, a_t, k_t) \qquad \text{both } \omega \text{ and } k_{t+1} \text{ and } a_{t+1} \text{ are functions of } (i_t, a_t, k_t) \\ \equiv \ & P_t \end{aligned}$$

A probit regression of the survival dummy on $\psi_t(i_t, a_t, k_t)$ will allow us to construct $\hat{P}_t$.

Finally, the **third stage equation** —where we will identify the coefficients of the state variables a_t and k_t— follows from the following chain:

$$\begin{aligned} & E[y_{t+1} - \alpha_l l_{t+1} | a_{t+1}, k_{t+1}, \chi_{t+1} = 1] \\ &= \alpha_0 + \alpha_a a_{t+1} + \alpha_k k_{t+1} + E[\omega_{t+1} | \omega_t, \chi_{t+1} = 1] \\ &\equiv \alpha_a a_{t+1} + \alpha_k k_{t+1} + g(\underline{\omega}_{t+1}, \omega_t) \qquad \text{where} \\ g(\underline{\omega}_{t+1}, \omega_t) &= g(\psi_t^{-1}(P_t, \phi_t - \alpha_a a_t - \alpha_k k_t), \phi_t - \alpha_a a_t - \alpha_k k_t) \\ &\equiv \xi(P_t, \phi_t - \alpha_a a_t - \alpha_k k_t) \end{aligned}$$

Controlling for exit depends now on two indices ω and $\underline{\omega}_{t+1}$, while more standard selection corrections are all single index problems. (For example, Heckman (1974) makes functional form assumptions on the error in the probability of exit equation, which allows one to get an exact expression for the selection bias—the inverse Mill's ratio—which can then be simply substituted in the equation of interest.) Here we have two indices, but we can express one of them as a function of the second one and the probability of survival (the propensity score). We can invert the second stage and express $\underline{\omega}_{t+1}$ as a nonparametric function of ω_t and P_t. And ω_t can be uncovered from the first stage estimation. It all boils down to

$$\begin{aligned} q_{it+1} - \hat{\alpha}_l l_{it+1} &= \alpha_a a_{it+1} + \alpha_k k_{it+1} + \xi(\hat{\phi}_{it} - \alpha_a a_{it} - \alpha_k k_{it}, \hat{P}_{it}) + \eta_{t+1} + \epsilon^2_{it+1} \\ \text{where} \qquad & \eta_{t+1} = \omega_{t+1} - E[\omega_{t+1} | \omega_t, \chi_{t+1}] \end{aligned}$$

Since capital in use at $t+1$ is known at the beginning of $t+1$, and η_{t+1} is mean independent (by definition) of all variables known at the beginning of $t+1$, η_{t+1} is mean independent of

k_{t+1} and a_{t+1} —making the third stage consistent—, but not of l_{t+1} —necessitating a first preliminary stage, and precluding joint estimation.

The unknown functions $\phi_t(.)$, $\psi_t(.)$, and $\xi(.)$ are approximated by

$$\phi_t(a, b) = \sum_{j=0}^{4-m} \sum_{m=0}^{4} \phi_{mj(t)} a^m b^j$$

and similarly for ψ_t and $\xi(.)$. If the sample period is sufficiently long enough, the coefficients of these approximations ($\phi_{mj(t)}$) can be made time-variant.[1]

They use the estimated production function to illustrate that aggregate productivity gains from deregulation in the telecommunications equipment industry are mainly driven by relocation effects and not plant-level productivity growth. The industry becomes more productive, not because individual plants advance, but because output is relocated from below to above averagely productive plants. How and why is explored in the second problem set.

$$\begin{aligned} p_t &= \sum_{i}^{N} s_{it} p_{it} \\ &= \sum_{i} (\overline{s}_t + \Delta s_{it})(\overline{p}_t + \Delta p_{it}) \\ &= \underbrace{N\overline{s}_t}_{1} \overline{p}_t + \sum_{i} \Delta s_{it} \Delta p_{it} \end{aligned}$$

Levinsohn & Petrin (2003)

All identical, except for the use of the material demand instead of the investment equation. This is also shown to be a monotonically increasing, hence invertible, function of ω, conditional on the state variables.

$$m_t = m_t(\omega_t, a_t, k_t)$$

Reasonable, as material demand is less lumpy and less prone to zero, such that the strict monotonicity in ω is more likely to hold, hence material could be a superior proxy.

[1] An alternative nonparametric estimation technique are kernel regressions.

Ackerberg & Caves (& Frazer) (2004)

Implicit timing assumptions might wreak havoc with the methodologies, especially LP.

In the LP production function first stage, we estimate by substituting the productivity term by the inversion of the material demand equation:

$$\begin{aligned} y_t &= \beta_L l_t + \beta_M m_t + \beta_K k_t + \omega_t + \epsilon \\ &= \beta_L l_t + \beta_M m_t + \beta_K k_t + m_t^{-1}(m_t, k_t) + \epsilon \\ &= \beta_L l_t + \phi_t(m_t, k_t) + \epsilon \end{aligned}$$

The question is what is exogenously moving around the labor demand, such that it is not perfectly collinear with the nonparametric controls for the effect of productivity (an unknown function of m and k.).

Levinsohn-Petrin (show 1-3 side by side with OP equivalent)

- In the case that l and m are chosen simultaneously, there also must exist a labor demand equation that depends on the same state variables, but not m.

  $$l_t = l_t(\omega_t, k_t)$$

  If l and m are really chosen simultaneously, the same state variables will enter this decision and l will be collinear with the included terms.

  $$l_t = l_t(\omega_t, k_t) = l_t(m_t^{-1}(m_t, k_t), k_t) = l_t'(m_t, k_t)$$

  One theoretical possibility for identification in this case is if there is some random optimization error in firms choices of l_t. Such optimization error will move lit independently of the nonparametric function, identifying β_L. However, note that this form of identification is quite tenuous. First, the extent of identification is purely a function of the extent of this optimization error. Second, note the diametrically opposed assumptions - while one needs to assume there is enough optimization error in l_t to identify β_L, one simultaneously needs to assume exactly no optimization error in m_t. If there were optimization error in m_t, the inversion would not be valid.

- If l is chosen prior to m, the material demand equation will depend on it

  $$m_t = m_t(\omega_t, k_t, l_t)$$

  Inverting this to get an expression for ω will necessitate the inclusion of polynomials in l in the production function and again β_L will not be identified.

- Finally, when l is chosen after m there is some hope, but it really depends on what happens with ω (or input prices) between the time m and l are chosen. If nothing happens, it is as if they are chosen simultaneously.

  If ω changes than l will be identified in the production function as it varies with the innovation in ω. Unfortunately, in this case the inversion of the material demand equation does not adequately capture the total unobserved productivity (as we had to assume productivity evolves between m is chosen and production takes place), and its failure to do so will be highly correlated with the variation in the data that identifies β_L. As such, the labor coefficient we estimate will be biased (strongly) by unobserved productivity (i.e. overestimated).

  What is needed is some variation other than ω that determines labor demand, for example the wage rate. Moreover, if we follow OP or LP and make the inversion time specific to control for variations in market structure that influence material demand, the shocks to the wage have to be firm-specific as the component that is common to firms (i.e. the time series variation) will be picked up by the time varying non-parametric inversion, i.e. time-varying coefficients in the function m_t.

  Finally, these price shocks cannot have persistence, otherwise the next periods material m_{t+1} will depend on it and we won't have a good inversion for productivity anymore next period.

Note that, in practice, one probably would not observe this collinearity problem. It is very likely that estimation would produce estimates of β_L. Our point is that unless one believes one of the above two stories, the extent to which the β_L coefficient is identified is the extent to which the LP model is misspecified.

Olley-Pakes

As before, we need something to move around l independent of i.

- Same two (implausible) "stories" that "rescue" LP also work here

- In addition, one can tell a potentially more believable story here if l is chosen before i and ω moves around between the two choices. That way the inversion for i will still recover the full productivity term that is relevant for the production function, while in contrast with LP, the i decision does not depend on l. The reason is that all investment at time t is only relevant for production at time $t + 1$ and only variables that have a lasting impact till next period will matter—l_t is not one of them. In LP, m_t is relevant for production at time t so its demand will depend on l_t.

Alternatives to get around endogeneity problem

Use instruments: e.g. Syverson (2004)

Estimate cost function instead: McElroy (1987)

Use factor prices as explanatory variables, only works under CRS, otherwise need to include output on RHS. To identify all parameters, researchers have often estimated $N-1$ factor share equations jointly with the cost function (give translog example). However, one has to make sure that the model is internally consistent. Theory (Shepard's lemma) states that input demands are directly linked to the cost function: $x_i = \partial C/\partial w_i$.

McElroy (1987) argues that additively adding random error terms to both the cost function and share equations is not consistent with any interpretation of the errors as structural errors (or requires very specific relationships between the errors if they are measurement error). If the relationships we estimate are supposed to be behavioral equations of firms, they should be derived from a consistent underlying model. She shows that this will put (testable) restrictions on the form of error terms.

The following representations of the technology are internally consistent:

$$\begin{aligned}
\text{production function} \quad & q = (x_1 - \epsilon_1, ..., x_n - \epsilon_n; \theta) \\
\text{cost function} \quad & C(q, w, \epsilon) = c(q, w) + \sum_{j=1}^{N} w_j \epsilon_j \\
\text{input demands} \quad & x_i = c_i(q, w) + \epsilon_i \\
\text{share equation} \quad & S_i = \underbrace{\frac{w_i c_i(q, w)}{c(q, w)}}_{s_i(q,w)} + \underbrace{\frac{w_i \epsilon_i - s_i(q, w) \sum_j w_j \epsilon_j}{c(q, w) + \sum_j w_j \epsilon_j}}_{v_i}
\end{aligned}$$

The errors can be interpreted as firm-specific parameters unobserved by the econometrician, technical inefficiencies, or even measurement error.

Advantages of this approach are:

- consistency
- it doesn't matter which representation is estimated
- helps in specification by revealing the interdependence btwn the errors in different eqts,
- interpretation is straightforward: errors are unobserved firm-parameters,
- shares will automatically lie between 0 and 1.

Output measurement: Klette and Griliches (1996)

They show that returns to scale estimates will be biased if the market is not perfectly competitive and deflated sales or value added is used instead of output as left hand side variable in the firm-level production function. Price effects will be misinterpreted as productivity differences. A proposed solution is to augment the model with a demand system and control for price changes.

Production function estimation: The model uses $\log Q$, while we only observe $\log R = \log(QP) = \log Q + \log P$.

The real production function is $q = X\alpha_0 + u^q$, while we estimate $r = X\alpha + u^r$ giving $\hat{\alpha} = (X'X)^{-1}X'r$.

Asymptotically this gives

$$plim_{N\to\infty}(\hat{\alpha}) = \alpha_0 + plim(X'X)^{-1}X'p + plim(X'X)^{-1}X'u^q$$

The last term might be nonzero, called the transmission bias of productivity shocks (see next week). The middle term might be nonzero too, it is the OLS coefficient in the auxiliary regression $p = X\delta + u^p$.

Define RTS $\epsilon = \sum_i \alpha_i$, then

$$plim_{N\to\infty}(\hat{\epsilon} - \epsilon) = plim(\sum_i \hat{\alpha}_i) - \sum_i \alpha_i = \sum_i \delta_i$$

Generally, RTS downwards biased:

- Idiosyncratic changes in factor prices: high cost, high price, low share, low inputs $\rightarrow \delta_i \leq 0$
- Idiosyncratic productivity shocks: high prod., low price and high share $\rightarrow \delta_i \leq 0$
- Demand shocks: IRS (DRS), input growth, lower (higher) price $\rightarrow \delta_i \leq 0 (\geq 0)$

Similarly for cost function estimation.

True cost is $c = W\gamma_0 + \beta_0 q + u^c \equiv W\gamma_0 + \beta_0(q - r) + \beta_0 r + u^c \equiv Z\lambda_0 - \beta_0 p + u^c$, estimated cost is $c = W\gamma + \beta r + \bar{u}^c \equiv Z\lambda + \bar{u}^c$.

As before we can write

$$plim_{N\to\infty}(\hat{\lambda}) = \lambda_0 + \beta_0\ plim(Z'Z)^{-1}Z'p + plim(Z'Z)^{-1}Z'u^c$$

The middle term can be interpreted as the coefficient in the auxiliary regression $p = Z\bar{\delta} + \bar{u}^p = W\bar{\delta}_w + r\bar{\delta}_r + \bar{u}^p$. With elastic demand $\bar{\delta}_r$ will be negative.

As we can write

$$plim_{N\to\infty}(\hat{\beta}) = \beta_0(1-\bar{\delta}_r)$$

RTS, the inverse of $\hat{\beta}$ will be estimated biased downwards.

Solution: deflated sales is: $r_{it} = q_{it} + (p_{it} - p_{It})$

Express price difference (in growth terms) as a function of relative price changes using an industry demand system (CES in this case, i.e. sales are proportional to a firms market share and this proportionality is constant over time):

$$\begin{aligned} Q_{it} &= Q_{It}(\frac{P_{it}}{P_{It}})^{\eta} e^{\epsilon_{it}^D} \\ q_{it} &= q_{It} + \eta(p_{it} - p_{It}) + u_{it}^D \\ r_{it} - (p_{it} - p_{It}) &= q_{It} + \eta(p_{it} - p_{It}) + u_{it}^D \end{aligned}$$

(with $u_{it}^D = \epsilon_{it}^D - \epsilon^D it - 1$)
(with $q_{it} = \log(Q_{it}/Q_{it-1})$)

Rearranging the last equation we find an expression for the unobserved price difference:

$$p_{it} - p_{It} = \frac{1}{1+\eta}(r_{it} - q_{It} - u_{it}^D)$$

(Note that with the demand equation, we could also derive the exact bias in the simple regression without q_{It}, i.e. the δ's)

The last step is to substitute this into the definition for deflated sales $r_{it} = q_{it}+(p_{it}-p_{It})$ and also substitute the production function, which gives (after rearranging again) an expression of r_{it} in terms of observables.

$$\begin{aligned} r_{it} &= \overbrace{\beta_l l_{it} + \beta_k k_{it} + u_{it}^S} + \overbrace{\frac{1}{1+\eta}(r_{it} - q_{It} - u_{it}^D)} \\ (1 - \frac{1}{1+\eta})r_{it} &= \frac{\eta+1}{\eta}(\beta_l l_{it} + \beta_k k_{it}) - \frac{1}{\eta} q_{It} + u_{it}' \end{aligned}$$

Relative to the regular production function,

- it multiplies all input coefficients with $\frac{\eta+1}{\eta}$
- includes q_{It} in the regression.
- Can recover the parameters of interest α_0 and standard errors with Δ-method.

(in the cost function framework, q_{It} would also end up in the factor shares)

References

Berndt, E. R. and M. S. Khaled (1979). Parametric Productivity Measurement and Choice among Flexible Functional Forms. *Journal of Political Economy 87*(6), 1220–45.

Christensen and Greene (1976, August). Economies of Scale in U.S. Electric Power Generation. *Journal of Political Economy 84*(4), 655–676.

Ericson, R. and A. Pakes (1995, January). Markov Perfect Industry Dynamics: A Framework for Empirical Work. *Review of Economic Studies 62*(1), 53–85.

Fuss, M. and D. McFadden (1978). *Production Economics: A Dual Approach to Theory and Applications*. Contributions to Economic Analysis, 111. New York: North-Holland.

Klette, T. J. and Z. Griliches (1996). The Inconsistency of Common Scale Estimators When Output Prices Are Unobserved and Endogenous. *Journal of Applied Econometrics 11*(4), 343–361.

Levinsohn, J. and A. Petrin (2003, April). Estimating Production Functions Using Inputs to Control for Unobservables. *Review of Economic Studies 70*(2), 317–342.

Marschak, J. and W. H. Andrews (1944). Random Simultaneous Equations and the Theory of Production. *Econometrica 12*, 143–205.

McElroy, M. B. (1987). Additive General Error Models for Production, Cost, and Derived Demand or Share Systems. *Journal of Political Economy 95*(4), 737–757.

Mundlak, Y. (1996, March). Production Function Estimation: Reviving the Primal. *Econometrica 64*(2), 431–38.

Olley, G. S. and A. Pakes (1996). The Dynamics of Productivity in the Telecommunications Equipment Industry. *Econometrica 64*(6), 1263–97.

Syverson, C. W. (2001). *Market Structure and Productivity*. Ph.d. dissertation, University of Maryland.

Syverson, C. W. (2004, December). Market Structure and Productivity: A Concrete Example. *Journal of Political Economy 112*(5).

Van Biesebroeck, J. (2003, January). Productivity Dynamics with Technology Choice: An Application to Automobile Assembly. *Review of Economic Studies 70*(1), 167–198.

Zellner, A., J. Kmenta, and J. H. Dreze (1966, October). Specification and Estimation of Cobb-Douglas Production Function Models. *Econometrica 34*, 784–795.