Production Functions

Nelli Valmari

Etla

January 2023

Agenda for this week

1. Introduction to production functions
2. An endogeneity problem and traditional solutions
3. Olley and Pakes (1996, Econometrica)
4. Levinsohn and Petrin (2003, Review of Economic Studies)
5. Ackerberg, Caves and Frazer (2015, Econometrica)
6. Other issues (and solutions)

Reading list

Compulsory:

- Olley and Pakes (1996, Econometrica)
- Levinsohn and Petrin (2003, Review of Economic Studies)
- Ackerberg, Caves and Frazer (2015, Econometrica)

Recommended:

- Ackerberg, Benkard, Berry and Pakes (2006, Handbook of Econometrics), Chapter 63, Section 2: Production Functions

Further reading:

- Jan De Loecker and Chad Syverson (2021). "Chapter 3 - An industrial organization perspective on productivity". Handbook of Industrial Organization, Volume 4, 141-223

Introduction to production functions

- A production function relates inputs of a producer (e.g., firm or establishment) to its output
- Inputs: capital, labour, materials, ...
- Productivity of a producer in converting inputs into outputs is unobservable to the econometrician
- Productivity may depend on factors such as management quality and environmental conditions, may be observable or unobservable to the econometrician
- In Cobb-Douglas form:

$$Y_i = A_i K_i^{\beta_K} L_i^{\beta_L}$$

Introduction to production functions

- Production functions underlie the supply side and are one of the basic components explaining market outcomes
- Production functions are estimated to learn about
  - Effects of inputs on output (β_K, β_L)
  - Returns to scale $(\beta_K + \beta_L)$
  - Productivity (A_i) and determinants of productivity and productivity differences between firms
  - Effects of changes in the operating environment, e.g., technical change and changes in competition and regulation
  - Efficiency of resource allocation across firms
  - Growth of firms and industries
  - Price cost markups and market power

Interesting facts about productivity estimates

1. There is enormous dispersion in productivity levels across producers, even within narrowly defined industries. For example, 90-10 percentile productivity ratios are typically estimated to be about 2:1 (North American manufacturing industries) and higher (developing or emerging economies, some service industries)
2. Producers' productivity levels, and therefore productivity differences between producers, are persistent over time
3. Productivity is correlated with "profitability", producer's size, growth, survival probability, low output prices of homogeneous products, and employees' wage level

(De Loecker and Syverson, 2021)

2. Endogeneity problem

- When choosing inputs, the firm's decision-maker observes A_i at least partly
- Decompose A_i into three components:

$$A_i = \beta_0 + \omega_i + \eta_i$$

  - β_0 is an industry-specific term
  - ω_i is observable or predictable to the firm's decision-maker (but not to the econometrician)
  - η_i is unobservable to the firm's decision-maker until the inputs are set; alternative interpretation: measurement error in output
- Consider the choice of L_i to maximise profits; the higher the unobservable $\beta_0 + \omega_i$, the higher the profit-maximising L_i
- Because L_i is endogenous to ω_i, these two are correlated
- If this is not taken into account, the estimate of β_L is biased upwards
- This is referred to as the simultaneity problem (e.g., Marschak and Andrews, 1944)

First traditional solution: instrumental variables

- Use an instrumental variable that is correlated with L_i but not with ω_i or η_i
- If there is perfect competition in the input market, input prices are valid instruments: they affect input demand without affecting the components of productivity
- Even if the assumptions of perfect competition are plausible
    - Input price data is not always available
    - Variation in input prices may not be sufficient for identification
    - Variation in input prices may be due to differences in input quality and be therefore correlated with productivity

Second traditional solution: fixed effects

- Assume that productivity is constant over time, i.e., $\omega_{it} = \omega_i$
- This is a strong assumption

Third traditional solution: First order conditions

- Use first order conditions of profit maximising firms
- If both input and output markets are perfectly competitive, then the output elasticity of a given input equals the input's cost share in revenue (β_L if Cobb-Douglas production technology)
- If competition in the output market is imperfect, one needs to estimate the output elasticity of demand to proceed
- Requires to make the assumption that all inputs are flexible, without any dynamic implications such as adjustment costs; this is a strong assumption on capital

3. Olley and Pakes: The Dynamics of Productivity in the Telecommunications Equipment Industry (1996, Econometrica)

OP

- "Control function approach" or "proxy variable approach"
- Empirical context: measure the impact of deregulation and the breakup of AT&T on productivity of telecommunications equipment producers
- Solve the simultaneity problem (our focus) and another endogeneity issue, referred to as the selection problem (we leave this for later)
- OP's insight for controlling for the unobservable ω_{it}: Firms' investment decisions reveal information about the underlying productivity level of the firm
- Make three key assumptions: assumptions on timing and dynamic nature of inputs, scalar unobservable assumption, and monotonicity assumption; we discuss these shortly

OP

- Firms make production choices to maximise the present discounted value of their current and future profits
- The dynamic optimisation problem doesn't have to be solved to control for the simultaneity issue
- Instead, consider the following production function that is central to the static profit maxmisation problem (OP has also age as one the variables, which I exclude for simplicity):

$$y_{it} = \beta_0 + \beta_K k_{it} + \beta_L l_{it} + \omega_{it} + \eta_{it}$$

OP: Assumptions on timing and dynamic nature of inputs

- Assume that ω_{it} follows an exogenous first-order Markov process, with I_{it} denoting the firm's information set at time t:

$$p\left(\omega_{it+1} | I_{it}\right) = p\left(\omega_{it+1} | \omega_{it}\right)$$

- Note that this is an econometric assumption on unobservables and, at the same time, an economic assumption on how firms learn about their productivity

OP: Assumptions on timing and dynamic nature of inputs

- Given the assumption of a first-order Markov process, conditional expected productivity can be written as a function current productivity:

$$E\left[\omega_{it+1} | I_{it}\right] = g\left(\omega_{it}\right)$$

- Realised ω_{it+1} can be decomposed into conditional expected productivity and an unpredictable component:

$$\omega_{it+1} = g\left(\omega_{it}\right) + \xi_{it+1}$$

where

$$E\left[\xi_{it+1} | I_{it}\right] = 0$$

OP: Assumptions on timing and dynamic nature of inputs

- Choice of l_{jt}
  - after observing ω_{it}
  - no dynamic implications such as adjustment costs
  - In short, labour is a variable and non-dynamic input
- Capital is accumulated as a function of investment, accumulated capital, and its depreciation
  - Investment in capital, i_{it}, is chosen after observing ω_{it}, to maximise the present discounted value of future profits, but it becomes "effective" only in the following period
  - Accumulated capital depreciates at rate δ

$$K_{it} = \delta K_{it-1} + i_{it-1}$$

  - In short, capital is a fixed and dynamic input

OP: Investment choice as a proxy variable

- OP use the investment demand function

$$i_{it} = f_{It}\left(k_{it}, \omega_{it}\right)$$

to invert out the unobservable ω_{it}:

$$\omega_{it} = h_t\left(k_{it}, i_{it}\right)$$

- Investment decisions may be dependent on time-specific factors such as output demand and input prices, which do not vary across producers
- The inversion requires making two more key assumptions: strict monotonicity and scalar unobservable

OP: Strict monotonicity assumption

- i_t is strictly monotonic in ω_{it}
- Intuitive because the marginal product of capital increases in ω_{it}

OP: Scalar unobservable assumption

- i_t has only one unobservable variable: ω_{it}
- For example, no unobservable input price variation across producers; as an exception, the cost of labour may vary across firms if it's not correlated across time
- No measurement or optimisation error in i_{it}
- This is not a light assumption

OP: First stage

- Substitute the inverted investment demand function in the production function:

$$y_{it} = \beta_0 + \beta_K k_{it} + \beta_L l_{it} + h_t(k_{it}, i_{it}) + \eta_{it}$$

- $h_t(k_{it}, i_{it})$ can be treated nonparametrically, as a polynomial of k_{it}, i_{it}:

$$\begin{aligned} y_{it} &= \beta_0 + \beta_K k_{it} + \beta_L l_{it} \\ &+ \boxed{\gamma_{0t} + \gamma_{1t} k_{it} + \gamma_{2t} i_{it} + \gamma_{3t} k_{it}^2 + \gamma_{4t} i_{it}^2 + \gamma_{5t} k_{it} i_{it}} + \eta_{it} \end{aligned}$$

- An OLS regression wouldn't identify all the parameters because $\beta_K k_{it}$ and $\gamma_{1t} k_{it}$ are collinear
- But β_L can be identified because it's not one of the state variables in i_t

OP: First stage

- Combine $\beta_0 + \beta_K k_{it}$ and $h_t(k_{it}, i_{it})$ into one term, denote it by $\phi_t(k_{it}, i_{it})$
- Write $\phi_t(k_{it}, i_{it})$ as a high order polynomial, and estimate the following semiparametric function:

$$y_{it} = \beta_l l_{it} + \phi_t(k_{it}, i_{it}) + \eta_{it}$$

- This way $i_t()$ - probably a complicated function - doesn't have to be considered, and estimation is computationally easier
- This comes at the cost of the monotonicity and scalar unobservable assumptions, and the assumption of labour being a nondynamic variable

OP: Intuition for the second stage

- From the first stage we have an estimate of $\phi_t\left(k_{it}, i_{it}\right)$, i.e., an estimate of $\beta_0 + \beta_K k_{it} + h_t\left(k_{it}, i_{it}\right)$
- How to separate $\beta_K k_{it}$ from $\beta_0 + h_t\left(k_{it}, i_{it}\right)$?
- Recall that the assumption of a first-order Markov process in productivity implies that $h_t\left(k_{it}, i_{it}\right)$ can be decomposed as follows:

$$\omega_{it} = g\left(\omega_{it-1}\right) + \xi_{it}$$

  where

$$E\left[\xi_{it} | I_{it}\right] = 0$$

- Recall also that the assumption on the timing of the choice of K_{it} implies that

$$E\left[\xi_{it} | K_{it}\right] = 0$$

OP: Intuition for the second stage

- Suppose that we knew β_0 and β_K; in that case, given the first stage estimate of $\phi_t\left(k_{it}, i_{it}\right)$, we could compute ω_{it}, or $h_t\left(k_{it}, i_{it}\right)$, as:

$$h_t\left(k_{it}, i_{it}\right) = \phi_t\left(k_{it}, i_{it}\right) - \beta_0 - \beta_K k_{it}$$

- Given $h_t\left(k_{it}, i_{it}\right)$, the nonparametric $g\left(\omega_{it-1}\right)$ can be estimated to obtain ξ_{it}
- If β_0 and β_K are corrected, as we just assumed, then $E\left[\xi_{it} | K_{it}\right] = 0$

OP: Intuition for the second stage

- In practice, for the estimate of $\phi_t(k_{it}, i_{it})$ and a given guess on β_0 and β_K, to obtain $\widehat{\xi}_{it}(\beta_K)$:
  - compute the implied $\widehat{\omega}_{it}$,
  - take lags $\widehat{\omega}_{it-1}$, and
  - regress $\widehat{\omega}_{it}$ on a polynomial of $\widehat{\omega}_{it-1}$
- k_{jt} is a valid instrument because it has been chosen at $t-1$, so β_K is identified using the following moment condition:

$$E\left[\xi_{it}\ k_{it}\right] = 0$$

- Find β_K by minimising (nonlinear GMM):

$$\frac{1}{N}\frac{1}{T}\sum_i\sum_t \widehat{\xi}_{it}\left(\beta_k\right)k_{it} = 0$$

- This is one way to execute the second stage of OP

OP: Second stage

- As β_L is estimated in the first stage, the second stage estimation equation is written as:

$$\begin{aligned} y_{it} - \beta_L l_{it} &= \beta_0 + \beta_K k_{it} + g\left(\omega_{it-1}\right) + \xi_{it} + \eta_{it} \\ &= \beta_0 + \beta_K k_{it} + g\left(\phi_t\left(k_{it-1}, i_{it-1}\right) - \beta_0 - \beta_K k_{it-1}\right) \\ &\quad + \xi_{it} + \eta_{it} \end{aligned}$$

where $\phi_t\left(k_{it-1}, i_{it-1}\right)$ is the lag of $\phi_t\left(k_{it}, i_{it}\right)$ that was estimated in the first stage

- The moment condition to identify β_K is:

$$E\left[\left(\xi_{it} + \eta_{it}\right) \ k_{it}\right] = 0$$

- With sample analogue:

$$\frac{1}{N}\frac{1}{T}\sum_i \sum_t \left(\widehat{\xi_{it} + \eta_{it}}\left(\beta_k\right)\right) k_{it} = 0$$

A one-step estimator of OP

- Wooldridge (2004) develops a one-step estimator of OP
- Deriving standard errors is more straigtforward
- The one-step estimator may be more efficient than the two-step OP estimator

Adjusting OP

- Can allow for dynamic (but variable) labour input
- In this case l_{jt} becomes a state variable and enters the investment function:

$$i_{it} = f_{It}\left(k_{it}, l_{it}, \omega_{it}\right)$$

  which is inverted to control for ω_{it}
- Extend the monotonicity assumption: ω_{it} does not vary across producers with a given k_{it}, i_{it}, and l_{it}
- β_L can be identified only in the second stage
- l_{it-1} is valid because it is correlated with l_{it} but uncorrelated with ξ_{it}

4. Levinsohn and Petrin: Estimating Production Functions Using Inputs to Control for Unobservables (2003, Review of Economic Studies)

LP

- Concern about the monotonicity assumption of OP: the assumption implies that ω_{it} does not vary across producers with a given k_{it} and i_{it}
  - Investments may be "lumpy"
  - Investment decisions may not depend on all kinds of productivity shocks
- Another concern about observations with zero investments
  - In many datasets, especially for developing countries, a considerable share of observations have zero investment
  - Can assume weak monotonicity, which implies disgarding observations with zero investment (like OP do) but, depending on the industry, that may be inefficient use of data

LP

- LP suggest using a different control variable, which is more likely to be strictly monotonic in ω_{it}, like materials
- The production function is now

$$y_{it} = \beta_0 + \beta_K k_{it} + \beta_L l_{it} + \beta_M m_{it} + \omega_{it} + \eta_{it}$$

where materials are a variable, non-dynamic input (like labour)

- The demand for materials is

$$m_{it} = f_{Mt}\left(k_{it}, \omega_{it}\right)$$

which, like in OP, is inverted for $h_t\left(k_{it}, m_{it}\right)$ and substituted into the production function:

$$y_{it} = \beta_0 + \beta_K k_{it} + \beta_L l_{it} + \beta_M m_{it} + h_t\left(k_{it}, m_{it}\right) + \eta_{it}$$

LP

- Again, like in OP, the parameters are estimated in two stages
- In the first stage, estimate the following equation to identify β_L:

$$y_{it} = \beta_L l_{it} + \phi_t \left(k_{it}, m_{it} \right) + \eta_{it}$$

where $\phi_t \left(k_{it}, m_{it} \right) = \beta_0 + \beta_K k_{it} + h_t \left(k_{it}, m_{it} \right)$, and $h_t \left(k_{it}, m_{it} \right)$ is estimated nonparametrically

- In the second stage, estimate β_K and β_M; add one more moment condition with m_{it-1} as the instrument

5. Ackerberg, Caves and Frazer: Identification Properties of Recent Production Function Estimators (2015, Econometrica)

ACF

- Concern about identification problems in OP and LP: does l_{it} vary independently of the nonparametric function of k_{jt} and m_{it}, i.e., is there a source of identifying variation for l_{it}?
- Examine OP and LP and define the conditions under which β_L is defined in the first stage
- Suggest different assumptions on timing of input decisions

ACF: concern about LP

- To understand the concern, assume that l_{jt} is chosen like m_{jt}, i.e., l_{jt} is variable and non-dynamic input:

$$
\begin{aligned}
m_{jt} &= f_{Mt}\left(k_{it}, \omega_{it}\right) \\
l_{jt} &= f_{Lt}\left(k_{it}, \omega_{it}\right)
\end{aligned}
$$

- Substitute the inversion for productivity, $h_t\left(k_{it}, m_{it}\right)$ (like in LP) into the demand function for labour:

$$
l_{jt} = f_{Lt}\left(k_{it}, h_t\left(k_{it}, m_{it}\right)\right)
$$

- l_{it} is deterministic of the other inputs, k_{jt} and m_{it}
- Recall the first stage estimation equation of LP (where $\phi_t\left(k_{it}, m_{it}\right) = \beta_0 + \beta_K k_{it} + h_t\left(k_{it}, m_{it}\right)$):

$$
y_{it} = \beta_L l_{it} + \phi_t\left(k_{it}, m_{it}\right) + \eta_{it}
$$

- So l_{jt} is functionally dependent on $\phi_t\left(k_{it}, m_{it}\right)$; there is no identifying variation for l_{it} for β_L to be identified in the first stage

Solution 1

- ACF come up with two alternative data generating processes for labour that would allow for identifying variation, i.e., something that affects labour demand that is independent of $\phi_t(k_{it}, m_{it})$.
- Assume optimisation error in l_{it} - independent of k_{it} and ω_{it} - but not in m_{it} (to satisfy the assumption of scalar unobservability)
- In this case there is no functional dependence, and β_L can be identified

Solution 2

- Assume a different timing
  - first, after observing ω_{it}, choose m_{it}
  - then, observe a producer-specific shock to the price of labour - independent across time (to satisfy the assumption of scalar unobservability) and other variables
  - after that, set l_{it}
- ACF discuss that these two solutions are not too general

ACF: concern about OP

- ACF make a similar argument for collinearity between l_{jt} and $\phi_t\left(k_{it}, i_{it}\right)$, i.e., for OP where investment is used as the proxy variable
- The collinearity issue is avoided in OP if l_{it} is chosen before ω_{it} is realised, or with incomplete information of ω_{it}
  - there is variation for l_{it} that is independent of ω_{it} and k_{it}
  - even though i_{it} is chosen after l_{it}, it doesn't depend on l_{it}, which is a static input

ACF: an alternative estimator

- Both β_L and β_K are identified in the second stage
- Assume that m_{jt} is conditional on l_{it} (instead of unconditional like in LP), i.e, m_{it} is chosen at the same time or after l_{it} is chosen:

$$m_{it} = f_{Mt}\left(k_{it}, \omega_{it}, l_{it}\right)$$

- Again, invert f_{Mt} and substitute in the production function to replace for the unobservable productivity:

$$y_{it} = \beta_0 + \beta_K k_{it} + \beta_L l_{it} + h_t\left(k_{it}, m_{it}, l_{it}\right) + \eta_{it}$$

- None of the parameters can be identified with this first stage estimation equation
- But, in the first stage, we can separate η_{it} from the composite term:

$$\widehat{\phi}_{it} = \beta_0 + \beta_K k_{it} + \beta_L l_{it} + h_t\left(k_{it}, m_{it}, l_{it}\right)$$

ACF: an alternative estimator

- In the 2nd stage, given the first stage estimate of $\widehat{\phi}_{it}$ and some given values of β_0, β_K, and β_L, to obtain $\widehat{\xi}_{it}(\beta_K, \beta_L)$:
  - compute the implied $\widehat{\omega}_{it}$,
  - take lags $\widehat{\omega}_{it-1}$, and
  - regress $\widehat{\omega}_{it}$ on a polynomial of $\widehat{\omega}_{it-1}$
- Identify β_K and β_L using the following moment conditions:

$$
\begin{aligned}
E\left[\xi_{it}\ k_{it}\right] &= 0 \\
E\left[\xi_{it}\ l_{it-1}\right] &= 0
\end{aligned}
$$

ACF: an alternative estimator

- Find β_K and β_L by minimising (nonlinear GMM):

$$\begin{aligned}
\frac{1}{N}\frac{1}{T}\sum_i\sum_t \widehat{\xi}_{it}\left(\beta_K\right)k_{it} &= 0 \\
\frac{1}{N}\frac{1}{T}\sum_i\sum_t \widehat{\xi}_{it}\left(\beta_K\right)l_{it-1} &= 0
\end{aligned}$$

- If the assumption of fixed labour input is plausible, instrument l_{it-1}can be replaced by l_{it}

Gandhi, Navarro and Rivers (2020, Journal of Political Economy)

- As ACF point out, their estimation method is proposed for estimating value added production functions, not gross output production functions where the proxy variable (materials) is one of the inputs
- If ACF was to be used for identification of a gross output production function, with two variable inputs of which one as a "proxy" variable, in order to identify β_L one would need identifying variation
  - for example, in adjustment costs of labour, such as hiring or firing costs, or
  - firm-specific input shocks to labour, but not to materials (to satisfy the scalar unobservable assumption)
  - ACF allows for these assumptions, and GNR suggest that they are also needed
- In practice, additional sources of variation may not be available
- GNR suggest using first-order conditions for variable inputs

6. Other issues

Selection

- There is another endogeneity issue
- Firms make their entry/exit decisions based on ω_{it} and fixed inputs like capital
- Conditional on being active and in the data set, ω_{it} and capital are negatively correlated
- The intuition is as follows: even if ω_{it} is low, it may be profitable to continue if the capital stock is large; and vice versa, if ω_{it} is high, it may be profitable to continue even if the capital stock is small
- If this is not taken into account, the estimate of β_K is biased downwards
- In practice, the selection problem is treated less often than the simultaneity problem
- Olley and Pakes account for selection by estimating and conditioning on the probability of surviving
- Also LP and ACF can be extended for that

Unobservable output prices and imperfect competition

- Most often production functions are estimated using sales revenue as the measure of output because physical output quantities are unobservable
- Suppose that firms' output doesn't increase in proportion to their inputs; is this due to decreasing returns to scale, or downward sloping demand? If downward sloping demand is not observable, the scalar unobservable assumption cannot be satisfied
- Klette and Griliches (1996): when the data has been generated under imperfect competition (or idiosyncratic demand shifts), output prices and inputs are correlated, and consequently production functions parameters are estimated with omitted price variable bias
- Traditional solutions: deflating output by industry-level price indices or product-level prices
- Preferred solutions: adding a demand system and demand shifters

Unobservable output prices and imperfect competition

- When using the proxy or control variable approach:
  - If there is imperfect competition or heterogeneous demand in the output market, the input demand function has to be written taking account of this, in order to satisfy the scalar unobservable assumption (Doraszelski and Jaumandreu, 2021)
  - In settings with oligopoly, the control function should include also competitors' productivity and state variables

Unobservable output prices and imperfect competition

- Foster Haltiwanger and Syverson (2008) compare physical output productivity and revenue productivity
    - Physical productivity and output prices are negatively correlated
    - Physical productivity and revenue productivity are correlated
    - Young producers set their output prices lower than older producers

Unobservable differences in output quality

- Even if physical output quantities are observable, how to deal with differences in output quality, or otherwise heterogeneous products?

Unobservable input prices

- Intermediate inputs and materials are usually measured in expenditures
- This is not a problem as long as there are no input price differences between producers
- Grieco, Li and Zhang (2016): ignoring input price dispersion (similarly to output price dispersion) using deflated expenditures leads to biased production function estimates
- Use first-order conditions of the firm's profit maximisation problem

Unobservable differences in input quality

- Labour input is often measured in labour hours or the number of full time equivalent employees
- Labour input is likely to be heterogeneous, both within and across producers
- Some studies use expenditures on labour instead of labour hours to control for quality
- Fox and Smeets (2011): wage bill explains as much productivity dispersion as human capital measures

Measurement error in inputs

- Measurement error in capital: capital stock is often estimated assuming the rate of depreciation, and the time when an investment becomes "effective"
- Intangible capital may not be measured at all

Endogenous productivity

- So far we have assumed that productivity evolves as an exogenous process
- But profit-maximising firms have incentives to raise their productivity by, for example, investing in R&D
- As an alternative to building a stock of knowledge, productivity can be estimated as endogenous to R&D input
- For example, Doraszelski and Jaumandreu (2013, Review of Economic Studies) endogenise the productivity process to R&D investments: $\omega_{it} = g\left(\omega_{it-1}, r\&d_{it-1}\right)$
  - R&D can have a heterogenous productivity effect due to nonlinearity with attained productivity, and due to uncertainty in the contribution of R&D to productivity

Multiproduct firms

- A typical, often implicit assumption: each firms produces all of its output with a single-product production technology
- But many firms are multiproduct firms, especially those involved in international trade
- Assuming a single-product production technology may be problematic if
  - production functions are product-specific, and returns to scale take place at the product-level
  - β_K, β_L, and β_M vary across products
  - ω_{it} is product-specific
  - there are economies of scope (this may be one of the reasons why multiproduct firms exist)
  - production technologies are joint due to joint inputs

Multiproduct firms

- Output may be observable at the product-firm or product-plant level but input allocation is observable only at the firm- or plant-level
- At the same time, the unobservable input allocation depends on the unobservable production technology that is to be estimated - including productivity - and perhaps even on output demand
- Multiproduct firms' production and cost functions were studied first in the 1960's and 1970's (e.g., Mundlak (1963, 1964))
- More recently, researchers have again started paying attention to multiproduct firms, perhaps due to product-producer level output data that has become available

Recent solutions for multiproduct firms

- De Loecker, Goldberg, Khandelwal and Pavcnik (2016, Econometrica): Use data on single-product firms to estimate production function parameters, assume that productivity is firm-specific, then solve for the product-level input allocation
- Orr (2022, Journal of Political Economy): Production function parameters do not vary across products but productivity may be product-firm specific, estimate output demand to solve for the input allocations
- Valmari (forthcoming, Review of Economic Studies): Production function parameters and productivity may vary across products, use output demand estimates to solve for product-level inputs
- Dhyne, Petrin, Smeets, Warzynski (2022, working paper): Estimate a transformation function that relates product-level output to firm-level inputs and other products the firm produces; do not have to solve for the unobservable input allocation

Functional forms etc.

- Many production function estimates are for Cobb-Douglas (with unitary elasticity of substitution across inputs) but there are more flexible functional forms (e.g., translog, constant elasticity of substitution)
- Productivity may not be Hicks-neutral, as in Cobb-Douglas; instead, productivity may be factor-biased or factor-augmenting
    - For example, if automatisation and robotisation change labour demand, they set off a factor-biased technological change
- Production function coefficients may not be common across producers
- Production function parameters are likely to change over time

As final remarks

- Productivity is the variation in output that cannot be explained by observable inputs, i.e., it's a residual (or as Abramovitz (1956) puts it: "Productivity is a measure of our ignorance")
- But estimated productivity levels correlate with important outcomes such as entry/exit decisions, participation in international trade, growth, product prices, etc
- So even though productivity is just a "residual" it's an interesting component of the production function

As final remarks

- Whether an estimation method, say, OP or ACF, is appropriate depends on the data generating process of the industry that you study
- When studying the effect of "X" on firm productivity, do include the "X" in the production function, as opposed to including "X" in the analysis only after estimating the production functions
- Researchers have made substantial progress in the productivity literature but interesting questions remain, such as: Why are the measured productivity distributions so wide?