The Cost Share Approach to Production Functions

A New Semi-parametric Method for Estimating Production Functions with Noisy Data on Input Expenditures and Revenue

Giulio Gottardo

University of Oxford and Pembroke College

Motivation

- Applications for production functions:
  - input elasticities for counterfactuals
  - total factor productivity (efficiency and misallocation)
  - price-cost markups (market power)

Motivation

- Applications for production functions:
  - input elasticities for counterfactuals
  - total factor productivity (efficiency and misallocation)
  - price-cost markups (market power)
- Challenges with standard data:
  - inputs are correlated with unobserved productivity
  - noisy data, measurement of output and capital costs
  - with market power: inputs are correlated with unobserved prices

Literature Gap

- Most PF estimation methods are for competitive output markets (Olley and Pakes, 1996; Blundell and Bond, 2000; Levinsohn and Petrin, 2003; Ackerberg et al., 2015; Gandhi et al., 2020)

Literature Gap

- Most PF estimation methods are for competitive output markets (Olley and Pakes, 1996; Blundell and Bond, 2000; Levinsohn and Petrin, 2003; Ackerberg et al., 2015; Gandhi et al., 2020)
- With market power regressing revenue on inputs is tricky

$$\underbrace{r_{it}}_{\text{log revenue}} = \underbrace{q_{it}}_{\text{log output}} + \underbrace{p_{it}}_{\text{log price}} = \underbrace{q(\omega_{it}, l_{it}, k_{it})}_{\text{productivity and inputs}} + \underbrace{p(q(\omega_{it}, l_{it}, k_{it}), \Upsilon_{it})}_{\text{also (but not only) productivity and inputs!}}$$

  - revenue cannot proxy for output → price bias (Klette and Griliches, 1996; Bond et al., 2021). how to even measure physical output?
  - input demand inversion of proxy methods may fail (Doraszelski and Jaumandreu, 2023; Biondi, 2024)
  - current fixes rely on strong assumptions on demand and/or production (Raval, 2023; Ackerberg and De Loecker, 2024)
  - also: strong assumptions on the dynamic process of productivity ω_{it} are needed

Literature Gap

- Most PF estimation methods are for competitive output markets (Olley and Pakes, 1996; Blundell and Bond, 2000; Levinsohn and Petrin, 2003; Ackerberg et al., 2015; Gandhi et al., 2020)
- With market power regressing revenue on inputs is tricky

$$\underbrace{r_{it}}_{\text{log revenue}} = \underbrace{q_{it}}_{\text{log output}} + \underbrace{p_{it}}_{\text{log price}} = \underbrace{q(\omega_{it}, l_{it}, k_{it})}_{\text{productivity and inputs}} + \underbrace{p(q(\omega_{it}, l_{it}, k_{it}), \Upsilon_{it})}_{\text{also (but not only) productivity and inputs!}}$$

  - revenue cannot proxy for output → price bias (Klette and Griliches, 1996; Bond et al., 2021). how to even measure physical output?
  - input demand inversion of proxy methods may fail (Doraszelski and Jaumandreu, 2023; Biondi, 2024)
  - current fixes rely on strong assumptions on demand and/or production (Raval, 2023; Ackerberg and De Loecker, 2024)
  - also: strong assumptions on the dynamic process of productivity ω_{it} are needed
- What can we learn with data on revenue and inputs alone?

Methodological Contribution

- A new estimator of output elasticities, markups, and revenue productivity

[1] This can be relaxed with additional data.

Methodological Contribution

- A new estimator of output elasticities, markups, and revenue productivity
  - based on *cost* shares $\rightarrow$ no need for firm-level output prices
  - no assumption on the functional form of demand, the dynamic process of productivity and/or price
  - new method for removing noise from data on the cost of capital/inputs with frictions

[1] This can be relaxed with additional data.

Methodological Contribution

- A new estimator of output elasticities, markups, and revenue productivity
  - based on *cost* shares → no need for firm-level output prices
  - no assumption on the functional form of demand, the dynamic process of productivity and/or price
  - new method for removing noise from data on the cost of capital/inputs with frictions
- It requires some of the standard assumptions of the literature
  - profit maximizing firms
  - competitive input markets[1]
  - at least one flexible input
  - Hicks-neutral technology
  - fixed returns to scale

[1] This can be relaxed with additional data.

Summary of Findings

- The new approach is useful for recovering assumption-light distributions $\rightarrow$ who are the superstar firms of the US stock market?

Summary of Findings

- The new approach is useful for recovering assumption-light distributions → who are the superstar firms of the US stock market?

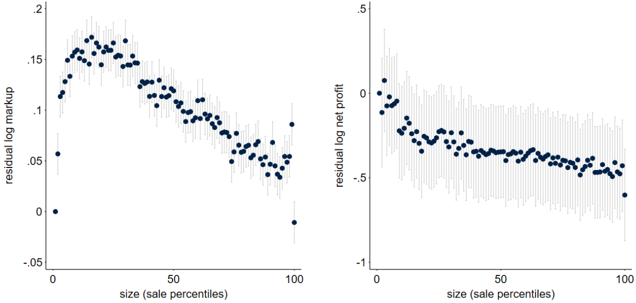

Figure: Log markups vs size (left panel) and log net profits vs size (right panel). Everything is net of year-by-sector-by-exchange fixed effects. 95% CIs in grey.

- They are smaller, more sunk cost-intensive firms. The top 1% by size does not record abnormal net profits

Outline

- Intro
- Theory
- The New Estimator
- The Superstars of Compustat
- Conclusion

Flexible Inputs

- Flexible input definition: the firm is a price taker in its market, faces no frictions, and there are no dynamic implications

Flexible Inputs

- Flexible input definition: the firm is a price taker in its market, faces no frictions, and there are no dynamic implications
- Profit maximizing firm i at time t chooses flexible input X_{it} s.t.:

$$\underbrace{MRX_{it}}_{\text{marginal revenue of X}} = \underbrace{P_{it}^{X}}_{\text{cost of X, exogenous}}$$

  - irrespective of functional forms and assumptions on other inputs ▸ Explicit Dynamic Model

Flexible Inputs

- Flexible input definition: the firm is a price taker in its market, faces no frictions, and there are no dynamic implications
- Profit maximizing firm i at time t chooses flexible input X_{it} s.t.:

$$\underbrace{MRX_{it}}_{\text{marginal revenue of X}} = \underbrace{P^X_{it}}_{\text{cost of X, exogenous}}$$

  - irrespective of functional forms and assumptions on other inputs Explicit Dynamic Model

- A bit of manipulation delivers the well-known result of De Loecker and Warzynski (2012)

$$\alpha^{R,X}_{it} = \frac{\theta^{Q,X}_{it}}{\mu_{it}}$$

  - $\alpha^{R,X}_{it} := \frac{P^X_{it} X_{it}}{P_{it} Q_{it}}$ is the revenue share of the flexible input
  - $\theta^{Q,X}_{it} := \frac{\partial Q_{it}}{\partial X_{it}} \frac{X_{it}}{Q_{it}}$ is its output elasticity
  - $\mu_{it} = \left(\frac{\partial P_{it}}{\partial Q_{it}} \frac{Q_{it}}{P_{it}} + 1 \right)^{-1}$ is the price-marginal cost markup

Inputs with Frictions

- Capital (dynamic + adjustment frictions): the "price" of capital and its economic cost differ. I need the *economic cost* of capital

Inputs with Frictions

- Capital (dynamic + adjustment frictions): the "price" of capital and its economic cost differ. I need the *economic cost* of capital
- Regardless of frictions, profit maximization still requires

$$\underbrace{MRK_{it}}_{\text{marginal revenue of } K_{it}} = \underbrace{\rho_{it}}_{\text{cost of } K_{it}}$$

  - ρ_{it} is endogenous, as it captures all dynamics and frictions ▸ Explicit Dynamic Model Derivation

Inputs with Frictions

- Capital (dynamic + adjustment frictions): the "price" of capital and its economic cost differ. I need the *economic cost* of capital
- Regardless of frictions, profit maximization still requires

$$\underbrace{MRK_{it}}_{\text{marginal revenue of } K_{it}} = \underbrace{\rho_{it}}_{\text{cost of } K_{it}}$$

  - ρ_{it} is endogenous, as it captures all dynamics and frictions ▸ Explicit Dynamic Model Derivation
- Same manipulation as before

$$\alpha_{it}^{R,K} = \frac{\theta_{it}^{Q,K}}{\mu_{it}}$$

  - it is the same as the flexible input with appropriately defined ρ_{it}

Returns to Scale and Cost Shares

- Fix the returns to scale $\theta_{it}^{Q,X} + \theta_{it}^{Q,K} = 1$[2]

[2] I could fix any number. 1 is just plausible and common

Returns to Scale and Cost Shares

- Fix the returns to scale $\theta_{it}^{Q,X} + \theta_{it}^{Q,K} = 1$[2]
- Combining all previous results

$$\begin{aligned} \mu_{it} &= \frac{P_{it} Q_{it}}{P_{it}^{X} X_{it} + \rho_{it} K_{it}} \\ \alpha_{it}^{C,X} &:= \frac{P_{it}^{X} X_{it}}{P_{it}^{X} X_{it} + \rho_{it} K_{it}} = \theta_{it}^{Q,X} \\ \alpha_{it}^{C,K} &:= \frac{\rho_{it} K_{it}}{P_{it}^{X} X_{it} + \rho_{it} K_{it}} = \theta_{it}^{Q,K} \end{aligned}$$

- Equilibrium output elasticities equal *cost* shares, and the markup is the ratio of revenue to (economic) total cost ▸ general case

[2] I could fix any number. 1 is just plausible and common

Returns to Scale and Cost Shares

- Fix the returns to scale $\theta_{it}^{Q,X} + \theta_{it}^{Q,K} = 1$[2]
- Combining all previous results

$$\mu_{it} = \frac{P_{it} Q_{it}}{P_{it}^{X} X_{it} + \rho_{it} K_{it}}$$

$$\alpha_{it}^{C,X} := \frac{P_{it}^{X} X_{it}}{P_{it}^{X} X_{it} + \rho_{it} K_{it}} = \theta_{it}^{Q,X}$$

$$\alpha_{it}^{C,K} := \frac{\rho_{it} K_{it}}{P_{it}^{X} X_{it} + \rho_{it} K_{it}} = \theta_{it}^{Q,K}$$

- Equilibrium output elasticities equal *cost* shares, and the markup is the ratio of revenue to (economic) total cost general case
- Main problem in applications: data on ρ_{it} is too noisy...

[2] I could fix any number. 1 is just plausible and common

Outline

- Intro
- Theory
- The New Estimator
- The Superstars of Compustat
- Conclusion

Data on the Cost of Capital is Noisy

- The rental rate ρ_{it} is not directly observed. It is usually calculated from data
  - e.g. CAPM for listed companies. interest expenses/debt for private ones

Data on the Cost of Capital is Noisy

- The rental rate ρ_{it} is not directly observed. It is usually calculated from data
  - e.g. CAPM for listed companies. interest expenses/debt for private ones
- Say we observe $\tilde{\rho}_{it}$:

$$\tilde{\rho}_{it} = \rho_{it} + \varepsilon_{it}^{\rho}$$

  - with $\mathrm{E}[\varepsilon_{it}^{\rho}] = \mathrm{E}[\varepsilon_{it}^{\rho}|x_{it}, k_{it}] = 0$
  - this is the case when the cost of capital is from the empirical CAPM derivation
- Specifying this in logs is more sensible in some settings. All results apply regardless

Data on the Cost of Capital is Noisy

- The rental rate ρ_{it} is not directly observed. It is usually calculated from data
  - e.g. CAPM for listed companies. interest expenses/debt for private ones
- Say we observe $\tilde{\rho}_{it}$:

$$\tilde{\rho}_{it} = \rho_{it} + \varepsilon^{\rho}_{it}$$

  - with $\mathrm{E}[\varepsilon^{\rho}_{it}] = \mathrm{E}[\varepsilon^{\rho}_{it} | x_{it}, k_{it}] = 0$
  - this is the case when the cost of capital is from the empirical CAPM derivation
- Specifying this in logs is more sensible in some settings. All results apply regardless
- If data on the cost of capital is noisy, simple cost shares are *badly* biased estimators of output elasticities

First Stage: Purging the Cost of Capital

- Idea: purge measurement error from data on the cost of capital

First Stage: Purging the Cost of Capital

- Idea: purge measurement error from data on the cost of capital
- Say $\tilde{\rho}_{it} = \rho_{it} + \varepsilon^{\rho}_{it}$, and combine the FOCs of X_{it} and K_{it}

$$\tilde{\rho}_{it} = \underbrace{\frac{\theta^{Q,K}_{it} \cdot P^X_{it} X_{it}}{\theta^{Q,X}_{it} \cdot K_{it}}}_{\text{equilibrium } \rho_{it}} + \varepsilon^{\rho}_{it} = \rho(x_{it}, k_{it}, p^X_{it}) + \varepsilon^{\rho}_{it}$$

  - The true rental rate is a function of input levels and prices alone because it equals marginal revenue

First Stage: Purging the Cost of Capital

- Idea: purge measurement error from data on the cost of capital
- Say $\tilde{\rho}_{it} = \rho_{it} + \varepsilon^{\rho}_{it}$, and combine the FOCs of X_{it} and K_{it}

$$\tilde{\rho}_{it} = \underbrace{\frac{\theta^{Q,K}_{it} \cdot P^X_{it} X_{it}}{\theta^{Q,X}_{it} \cdot K_{it}}}_{\text{equilibrium } \rho_{it}} + \varepsilon^{\rho}_{it} = \rho(x_{it}, k_{it}, p^X_{it}) + \varepsilon^{\rho}_{it}$$

  - The true rental rate is a function of input levels and prices alone because it equals marginal revenue
- Implementation: non-parametric regression of $\tilde{\rho}_{it}$ on $(x_{it}, k_{it}, p^X_{it})$

$$\hat{\rho}_{it} = \tilde{\rho}_{it} - \hat{\varepsilon}^{\rho}_{it}$$

First Stage: Purging the Cost of Capital

- Idea: purge measurement error from data on the cost of capital
- Say $\tilde{\rho}_{it} = \rho_{it} + \varepsilon^{\rho}_{it}$, and combine the FOCs of X_{it} and K_{it}

$$\tilde{\rho}_{it} = \underbrace{\frac{\theta^{Q,K}_{it} \cdot P^X_{it} X_{it}}{\theta^{Q,X}_{it} \cdot K_{it}}}_{\text{equilibrium } \rho_{it}} + \varepsilon^{\rho}_{it} = \rho(x_{it}, k_{it}, p^X_{it}) + \varepsilon^{\rho}_{it}$$

  - The true rental rate is a function of input levels and prices alone because it equals marginal revenue
- Implementation: non-parametric regression of $\tilde{\rho}_{it}$ on $(x_{it}, k_{it}, p^X_{it})$

$$\hat{\rho}_{it} = \tilde{\rho}_{it} - \hat{\varepsilon}^{\rho}_{it}$$

- This works regardless of the source of frictions and data noise
  - note: $E[\hat{\rho}_{it}] = E[\tilde{\rho}_{it}]$. Choosing whether to purge in levels or logs matters for the average

Second Stage: Output Elasticities

- Calculate purged cost shares $\hat{\alpha}^{C,X} = \frac{P_{it}^X X_{it}}{P_{it}^X X_{it} + \hat{\rho}_{it} K_{it}}$ and $\hat{\alpha}^{C,K} = \frac{\hat{\rho}_{it} K_{it}}{P_{it}^X X_{it} + \hat{\rho}_{it} K_{it}}$

Second Stage: Output Elasticities

- Calculate purged cost shares $\hat{\alpha}^{C,X} = \frac{P^X_{it} X_{it}}{P^X_{it} X_{it} + \hat{\rho}_{it} K_{it}}$ and $\hat{\alpha}^{C,K} = \frac{\hat{\rho}_{it} K_{it}}{P^X_{it} X_{it} + \hat{\rho}_{it} K_{it}}$
- True cost shares and output elasticities must coincide, but more can be learned

$$\hat{\alpha}^{C,X}_{it} = \underbrace{\theta^{Q,X}(x_{it}, k_{it})}_{\text{part. derivative of the p.f.}} + u_{it}$$

  - $u_{it} = \hat{\alpha}^{C,X}_{it} - \alpha^{C,X}_{it}$ would be correlated with (x_{it}, k_{it}) without purging

Second Stage: Output Elasticities

- Calculate purged cost shares $\hat{\alpha}^{C,X} = \frac{P^X_{it} X_{it}}{P^X_{it} X_{it} + \hat{\rho}_{it} K_{it}}$ and $\hat{\alpha}^{C,K} = \frac{\hat{\rho}_{it} K_{it}}{P^X_{it} X_{it} + \hat{\rho}_{it} K_{it}}$
- True cost shares and output elasticities must coincide, but more can be learned

$$\hat{\alpha}^{C,X}_{it} = \underbrace{\theta^{Q,X}(x_{it}, k_{it})}_{\text{part. derivative of the p.f.}} + u_{it}$$

  - $u_{it} = \hat{\alpha}^{C,X}_{it} - \alpha^{C,X}_{it}$ would be correlated with (x_{it}, k_{it}) without purging
- Implementation: non-parametric regression of $\hat{\alpha}^{C,X}_{it}$ and $\hat{\alpha}^{C,K}_{it}$ on input levels
  - result: $\hat{\theta}^{Q,X}_{it} = \mathrm{E}[\theta^{Q,X}_{it} | x_{it}, k_{it}]$ and $\hat{\theta}^{Q,K}_{it} = \mathrm{E}[\theta^{Q,K}_{it} | x_{it}, k_{it}]$ (partial derivatives of the p.f.)

Markups and Revenue Productivity

- Markup (De Loecker and Warzynski, 2012)

$$\hat{\mu}_{it} = \frac{\hat{\theta}_{it}^{Q,X}}{\tilde{\alpha}_{it}^{R,X}}$$

Markups and Revenue Productivity

- Markup (De Loecker and Warzynski, 2012)

$$\hat{\mu}_{it} = \frac{\hat{\theta}_{it}^{Q,X}}{\tilde{\alpha}_{it}^{R,X}}$$

- Revenue productivity in logs $\omega_{it}^{R} := \omega_{it} + p_{it}$ (TFPQ + log price):

$$\hat{\omega}_{it}^{R} = \tilde{r}_{it} - \int_{0}^{x_{it}} \hat{\theta}^{Q,X}(x, k_{it})dx - \int_{0}^{k_{it}} \hat{\theta}^{Q,K}(0, k)dk$$

Two-Stage Estimator Recap

Stage 1: Purge noise from cost of capital

- Regress noisy $\tilde{\rho}_{it}$ on (x_{it}, k_{it}) non-parametrically:

$$\tilde{\rho}_{it} = \rho(x_{it}, k_{it}) + \varepsilon^{\rho}_{it}.$$

- Obtain purged $\hat{\rho}_{it}$ and re-compute cost shares

Stage 2: Semi-parametric elasticity estimation

- Use cleaned shares $\hat{\alpha}^{C,X}_{it}$, and $\hat{\alpha}^{C,K}_{it}$ to fit:

$$\hat{\alpha}^{C,X}_{it} = \theta^{Q,X}(x_{it}, k_{it}) + u_{it}$$

- Use non-parametric output elasticity estimates to calculate markups and revenue TFP
- Monte Carlo simulations for the skeptics Monte Carlo

Taking Stock

- Doraszelski and Jaumandreu (2023): it is not possible to remove measurement error in revenue without strong assumptions on demand → markups and revenue TFP are noisy

Taking Stock

- Doraszelski and Jaumandreu (2023): it is not possible to remove measurement error in revenue without strong assumptions on demand → markups and revenue TFP are noisy
- Measurement error in input levels cannot be purged → but can use IVs (Collard-Wexler and De Loecker, 2024)

Taking Stock

- Doraszelski and Jaumandreu (2023): it is not possible to remove measurement error in revenue without strong assumptions on demand → markups and revenue TFP are noisy
- Measurement error in input levels cannot be purged → but can use IVs (Collard-Wexler and De Loecker, 2024)
- Market power in input markets → external supply elasticity/markdown estimates can be incorporated into the purged cost shares adding markdowns

Taking Stock

- Doraszelski and Jaumandreu (2023): it is not possible to remove measurement error in revenue without strong assumptions on demand → markups and revenue TFP are noisy
- Measurement error in input levels cannot be purged → but can use IVs (Collard-Wexler and De Loecker, 2024)
- Market power in input markets → external supply elasticity/markdown estimates can be incorporated into the purged cost shares ▸ adding markdowns
- Different returns to scale → the assumption can be changed arbitrarily, and tested with data on Q_{it} (Syverson, 2004) ▸ changing RTS

Taking Stock

- Doraszelski and Jaumandreu (2023): it is not possible to remove measurement error in revenue without strong assumptions on demand → markups and revenue TFP are noisy
- Measurement error in input levels cannot be purged → but can use IVs (Collard-Wexler and De Loecker, 2024)
- Market power in input markets → external supply elasticity/markdown estimates can be incorporated into the purged cost shares ▸ adding markdowns
- Different returns to scale → the assumption can be changed arbitrarily, and tested with data on Q_{it} (Syverson, 2004) ▸ changing RTS
- Advantages:
  - no assumption on the functional form of demand
  - no assumption on the functional form of output elasticities
  - no assumption on the dynamic process of productivity and/or price
  - no stance on the measurement of physical output

Outline

- Intro
- Theory
- The New Estimator
- The Superstars of Compustat
- Conclusion

Compustat Data

- US listed firms, 1962–2024
  - sales
  - COGS (cost of goods sold), the flexible input
  - PPEGT (property, plant, and equipment), tangible capital
  - SGA (selling, general, and administrative expenses), spending on overhead
  - intan (intangible capital), I construct the "true" intangible capital stock following Chiavari and Goraya (2025)
- External sources for the cost of capital: FRED (inflation, inflation expectations, treasury yields), WRDS (CAPM-based betas) Data Appendix

Estimation

- Constant returns to scale + flexible COGS
- Proceed by sector-year
- First stage: regress $\tilde{\rho}_{it}$ on a 2nd order polynomial of log COGS and PPEGT
- Second stage: regress purged cost shares on 1st order polynomials of log COGS and PPEGT
- Calculate markups and revenue TFP
- I treat SGA/intangible capital as sunk costs - in line with most of the literature
  - only in the paper: treat intangible capital as a production input

Results: COGS Elasticities, Markups, Revenue TFPR

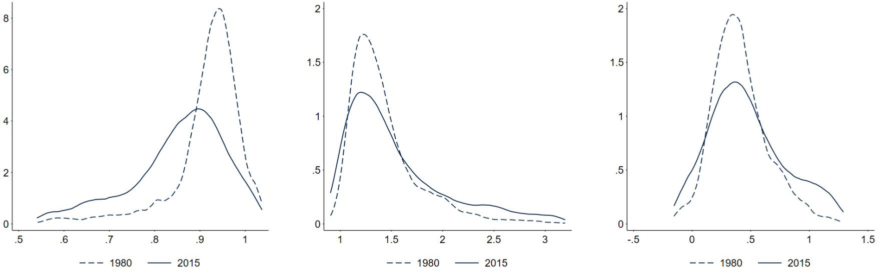

Figure: Distribution of the output elasticity of COGS (left), markups (middle), log TFPR (right): 1980 (dash) vs 2015 (solid)

- Cobb Douglas assumption is increasingly wrong over time
- Log markup and log productivity dispersion increase 6x - especially their right tails
  TFPR variance decomposition
  - note that TFPR is in logs, the markup is in levels

Question: Who are the Superstars?

- Increased markup/TFPR dispersion means that the top firms are doing better than average more than they used to
  - the 'superstars' of Autor et al. (2020)

Question: Who are the Superstars?

- Increased markup/TFPR dispersion means that the top firms are doing better than average more than they used to
  - the 'superstars' of Autor et al. (2020)
- My method imposes no functional form assumptions on demand, the production function, and productivity $\rightarrow$ I can make comparisons across firms

Question: Who are the Superstars?

- Increased markup/TFPR dispersion means that the top firms are doing better than average more than they used to
  - the 'superstars' of Autor et al. (2020)
- My method imposes no functional form assumptions on demand, the production function, and productivity $\rightarrow$ I can make comparisons across firms
- Who are the superstars?
  - do larger firms have higher markups/productivity?
  - is spending on overhead/intangible stock correlated with higher markups/productivity?
  - markups vs net accounting profits

Larger Firms Spend Less on Sunk Costs

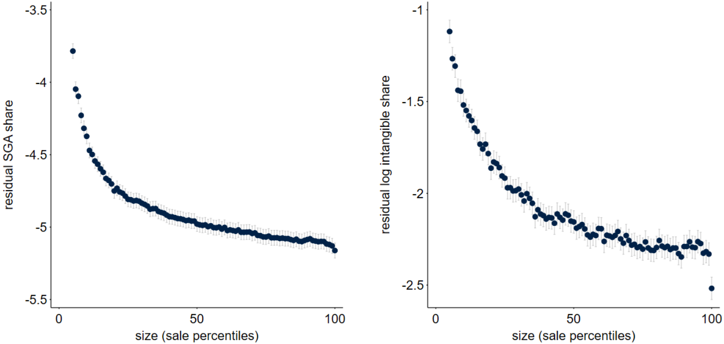

Figure: SGA expenditure share of sales (left panel) and intangible share of sales (right panel) vs size (sales). Everything is net of year-by-sector-by-exchange fixed effects. 95% CIs in grey.

- Smaller firms spend relatively more on overhead and have relatively larger intangible capital stocks (very robust to definition of the latter)
  - consistent with the 'sunk cost' interpretation

Sunk Costs “Require” Higher Markups

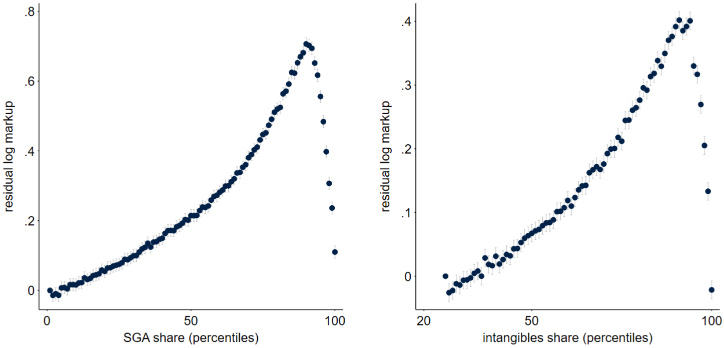

Figure: Log markups vs SGA share of sales (left panel) and log markups vs intangible share of sales (right panel). Everything is net of year-by-sector-by-exchange fixed effects. 95% CIs in grey.

- Sunk cost-intensive firms must charge higher murkups to cover non-production sunk costs

Smaller Firms Have Higher Markups, but Not Net Profits

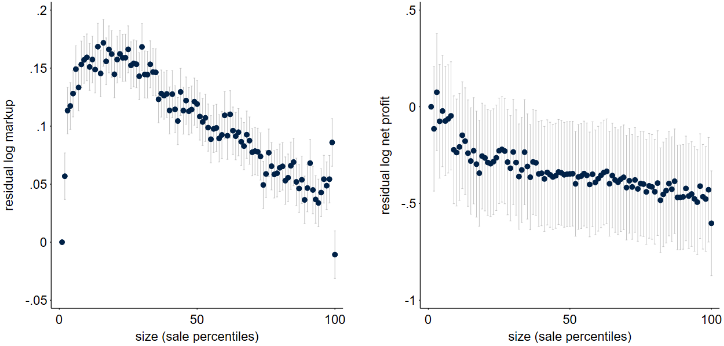

Figure: Log markups vs size (left panel) and log net profits vs size (right panel). Everything is net of year-by-sector-by-exchange fixed effects. 95% CIs in grey.

- Price-marginal cost markups are decreasing is size because of sunk cost-intensive small firms
- Net accounting profits (which include SGA in total cost) are not

Revenue Productivity Displays The Same Patterns

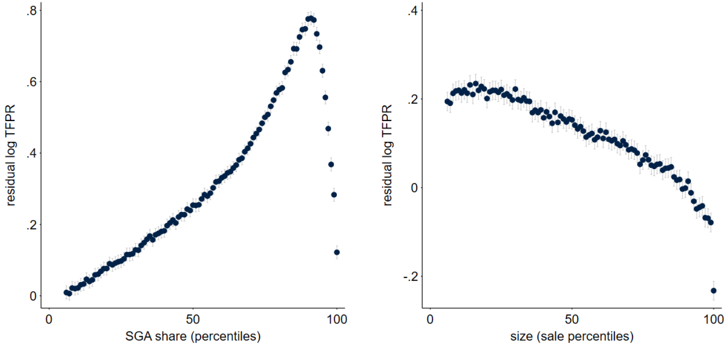

Figure: Log TFPR vs SGA share of sales (left panel) and TFPR vs size (right panel). Everything is net of year-by-sector-by-exchange fixed effects. 95% CIs in grey.

- Most of the variance in TFPR is accounted for by markups → all previous results apply to TFPR too

Outline

Conclusion

- **Methodological contribution:** A new two-stage estimator for production functions
  - Can be applied to any firm data
  - Purges noise in capital costs
  - Recovers heterogeneous output elasticities, markups, and revenue productivity with minimal demand, productivity process, output measurement, and functional form assumptions

Conclusion

- **Methodological contribution:** A new two-stage estimator for production functions
  - Can be applied to any firm data
  - Purges noise in capital costs
  - Recovers heterogeneous output elasticities, markups, and revenue productivity with minimal demand, productivity process, output measurement, and functional form assumptions
- **Compustat application:** who are the superstar firms?
  - Markup and revenue productivity dispersion has increased 6x since the 1980s
  - Firms with higher spending on overhead/larger intangible capital stocks have higher markups and revenue productivity
  - These are the smaller firms - consistently with a sunk costs story. The largest firms spend less on overhead, and thus have lower markups and revenue productivity
  - no significant relationship between net profits and size

References I

Ackerberg, D. A., Caves, K. and Frazer, G. (2015), 'Identification properties of recent production function estimators', *Econometrica* **83**(6), 2411–2451.

Ackerberg, D. A. and De Loecker, J. (2024), 'Production function identiÖcation under imperfect competition', *Working paper* .

Autor, D., Dorn, D., Katz, L. F., Patterson, C. and Van Reenen, J. (2020), 'The Fall of the Labor Share and the Rise of Superstar Firms*', *The Quarterly Journal of Economics* **135**(2), 645–709.

Biondi, F. (2024), 'Firm productivity and derived factor demand under variable markups', *Working paper* .

Blundell, R. and Bond, S. R. (2000), 'Gmm estimation with persistent panel data: An application to production functions', *Econometric Reviews* **19**, 321–340.

Bond, S., Hashemi, A., Kaplan, G. and Żoch, P. (2021), 'Some unpleasant markup arithmetic: Production function elasticities and their estimation from production data', *Journal of Monetary Economics* **121**(C), 1–14.

Chiavari, A. and Goraya, S. (2025), 'The rise of intangible capital and the macroeconomic implications', *Working paper* .

Collard-Wexler, A. and De Loecker, J. (2024), 'Productivity and capital measurement error', *NBER Working Paper* .

References II

De Loecker, J. and Warzynski, F. (2012), 'Markups and firm-level export status', *American Economic Review* **102**(6), 2437–2471.

Doraszelski, U. and Jaumandreu, J. (2023), 'Reexamining the de loecker warzynski (2012) method for estimating markups', *CEPR Discussion Paper No. DP16027* .

Gandhi, A., Navarro, S. and Rivers, D. A. (2020), 'On the identification of gross output production functions', *Journal of Political Economy* **128**(8), 2973 – 3016.

Klette, T. J. and Griliches, Z. (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 Petrin, A. (2003), 'Estimating Production Functions Using Inputs to Control for Unobservables', *The Review of Economic Studies* **70**(2), 317–341.

Olley, S. G. and Pakes, A. (1996), 'The dynamics of productivity in the telecommunications equipment industry', *Econometrica* **64**(6), 1263–1297.

Raval, D. (2023), 'A flexible cost share approach to markup estimation', *Economics Letters* **230**, 111262.

Syverson, C. (2004), 'Market structure and productivity: A concrete example', *Journal of Political Economy* **112**(6), 1181–1222.

Firm Problem

- Firm i at time t maximizes its NPV by choosing flexible input X_{it} and investment I_{it}[3]
- Standard Bellman equation:

$$V(K_{it}, \Omega_{it}, \Upsilon_{it}) = \max_{X_{it}, I_{it}} P(Q_{it}, \Upsilon_{it})Q_{it} - X_{it} - \Phi_i(K_{it}, I_{it}) + \beta_i E_t[V(K_{it+1}, \Omega_{it+1}, \Upsilon_{it+1})]$$

$$\text{s.t. } Q_{it} = \Omega_{it} \cdot F(X_{it}, K_{it})$$

$$K_{it+1} = (1 - \delta_i)K_{it} + I_{it}.$$

  - Υ_{it} is a vector that captures residual demand heterogeneity, including differences in market structure/strategic interactions
  - $\Phi_i(\cdot)$ captures payments to capital and investment frictions. Discounting β_i and depreciation δ_i are firm specific
  - Ω_{it} is Hicks-neutral quantity total factor productivity (TFPQ), and $F(\cdot)$ the production function
  - the flexible input X_{it} is chosen as the *numeraire* wlog Back

[3]This is a standard, compact, and yet general model. All my results extend to more inputs and different input choice frictions.

The Cost of Capital: Derivation

- Envelope condition of the dynamic problem $\rightarrow$ equilibrium rental rate ρ_{it}

$$\rho_{it} := \underbrace{\frac{\partial(P_{it}Q_{it})}{\partial K_{it}}}_{MRK_{it}} = \underbrace{\frac{\partial \phi_{it}}{\partial K_{it}}}_{\text{direct payments to } K} + \underbrace{\frac{\partial V_{it}}{\partial K_{it}}}_{\text{marginal value of } K \text{ today}} - \underbrace{\beta_i \cdot (1-\delta_i) \cdot \mathsf{E}_t \left[\frac{\partial V_{it+1}}{\partial K_{it+1}} \right]}_{\text{current marginal value of } K \text{ tomorrow}}$$

- This is always true, even for a flexible input (the last two terms would be zero) Back

Cost Shares: General Case

- Say that there is a third input, L_{it} which faces no adjustment frictions but is monopsonistic
  - i.e. it faces a markdown $\nu_{it} > 1$ (ratio of MRL_{it} and W_{it}, its price)
- Also, let the returns to scale be $\theta_{it}^{Q,X} + \theta_{it}^{Q,K} + \theta_{it}^{Q,L} = \mathcal{T}$
- The markup and elasticity expressions become

$$
\begin{aligned}
\mu_{it} &= \mathcal{T} \frac{P_{it} Q_{it}}{X_{it} + \rho_{it} K_{it} + \nu_{it} W_{it} L_{it}} \\
\alpha_{it}^{C,X} &:= \mathcal{T} \frac{X_{it}}{X_{it} + \rho_{it} K_{it} + \nu_{it} W_{it} L_{it}} = \theta_{it}^{Q,X} \\
\alpha_{it}^{C,K} &:= \mathcal{T} \frac{\rho_{it} K_{it}}{X_{it} + \rho_{it} K_{it} + \nu_{it} W_{it} L_{it}} = \theta_{it}^{Q,K} \\
\alpha_{it}^{C,L} &:= \mathcal{T} \frac{\nu_{it} W_{it} L_{it}}{X_{it} + \rho_{it} K_{it} + \nu_{it} W_{it} L_{it}} = \theta_{it}^{Q,L}
\end{aligned}
$$

- i.e. the markdown must be incorporated into "actual" total cost, and everything is scaled by the returns to scale parameter Back

Testing the Returns to Scale

- With data on Q_{it}, the returns to scale (RTS) can be tested (Syverson, 2004)
- Say the true RTS parameter is $\mathcal{T} \neq 1 \rightarrow$ second-stage estimates under constant RTS are 'wrongly scaled'

$$\theta_{it}^{Q,X} = \mathcal{T} \cdot \hat{\theta}_{it}^{Q,X} \quad \text{and} \quad \theta_{it}^{Q,K} = \mathcal{T} \cdot \hat{\theta}_{it}^{Q,K}$$

- Plug these into the production function:

$$q_{it} = \mathcal{T} \cdot \left(\int_0^{x_{it}} \hat{\theta}^{Q,X}(x, k_{it})\, \mathrm{d}x + \int_0^{k_{it}} \hat{\theta}^{Q,K}(0, k)\, \mathrm{d}k \right) + \omega_{it}^{Q}$$

- The inside of the bracket can be treated as data. With instruments $\mathbf{z}_{it}$ s.t. $\mathrm{E}[\omega_{it}^{Q} \mathbf{z}_{it}] = 0$, $\mathcal{T}$ is identified
- $\hat{\mathcal{T}}$ can be used to rescale all estimates Back

Noise from CAPM

- CAPM (for listed companies): each company pays a different risk premium to its investors

$$\underbrace{R_{it} - R_{ft}}_{\text{real return of the stock}} = \beta_i \underbrace{(R_{mt} - R_{ft})}_{\text{equity premium}} + \alpha_i + \varepsilon_{it}$$

  - R_{it} is the return on firm i's stock, R_{ft} is the return on the safe asset, R_{mt} is the return of the market portfolio, and β_i is a measure of the stock's risk/volatility

- The CAPM is usually estimated with OLS from daily/weekly/monthly stock market data. It follows that

$$\hat{\beta}_i = \beta_i + \frac{\sum_1^T (R_{mt} - R_{ft})\varepsilon_{it}}{\sum_1^T (R_{mt} - R_{ft})^2}$$

  - the expected value of the second term is zero, and it approaches zero as $T \to \infty$

- It follows that $\hat{\beta}_i = \beta_i + \varepsilon_i^\beta$ with $\text{E}[\varepsilon_i^\beta] = 0$ for finite T
- If we use $\hat{\beta}_i$ to calculate the firm-level cost of capital (with depreciation δ)

$$\tilde{\rho}_{it} = R_{ft} + \hat{\beta}_i (R_{mt} - R_{ft}) + \delta = \rho_{it} + \varepsilon_i^\beta (R_{mt} - R_{ft}) = \rho_{it} + \varepsilon_{it}^\rho$$

  - taking the equity premium as fixed implies $\text{E}[\varepsilon_{it}^\rho] = 0$ Back

Monte Carlo Setting

- Same as theory. Translog production function + linear demand (unforgiving heterogeneity)
- The researcher observes (noisy) revenue, noisy cost-of-capital, (noisy) capital stock, and expenditure on the flexible input
- Benchmark: 'naive' accounting (no purging of $\tilde{\rho}_{it}$, no regressions)
- I look at average estimates and correlations between estimates and true values

Monte Carlo: Output Elasticity, Error in Capital Cost

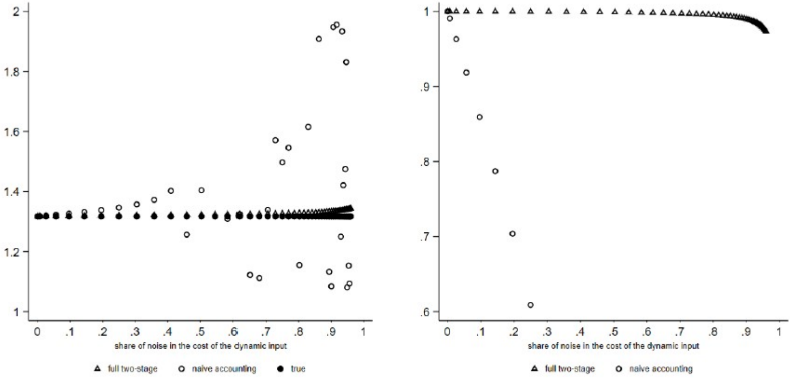

Figure: Sensitivity to error in the cost of capital: average $\hat{\theta}_{it}^{Q,X}$ (left panel), and corr($\hat{\theta}_{it}^{Q,X}$, $\theta_{it}^{Q,X}$) (right panel)

- Two-stage essentially nails $\theta_{it}^{Q,X}$ even when $\tilde{\rho}_{it}$ is $>$ 90% noise

Monte Carlo: Output Elasticity, Error in Capital Stock

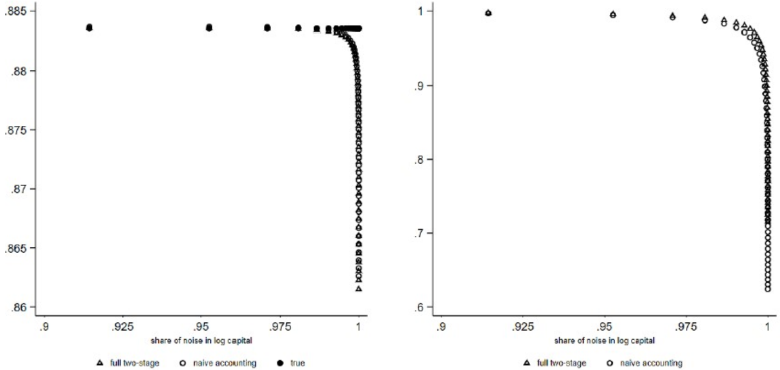

Figure: Sensitivity to error in the capital stock: average $\hat{\theta}_{it}^{Q,X}$ (left panel), and corr($\hat{\theta}_{it}^{Q,X}$, $\theta_{it}^{Q,X}$) (right panel)

- Issues with noise in the capital stock kick in with extremely poor measurement (look at the horizontal axis!)
  - the capital elasticity/share is plausibly low - and the variance of K_{it} is too low ▸ Back

Data Appendix

- Cost of capital: $\tilde{\rho}_{it} = \mathrm{E}_t[\pi_{it}^{10Y}] + \hat{\beta}_{it} \cdot 0.07 + 0.05$
  - $\hat{\beta}_{it}$ is an OLS estimate of the empirical CAPM by WRDS
  - $\mathrm{E}_t[\pi_t^{10Y}]$ is the expected real return on the 10-year Treasury
  - $+0.07$ is the average equity premium, and $+0.05$ is depreciation
- Intangible capital: combine accounts and perpetual inventory method (Chiavari and Goraya, 2025)
  - $K_{it}^{I,ACC} = \text{intan}_{it} + \text{amortization of intan}_{it} - \text{goodwill}_{it}$
  - $K_{it}^{I,PIM} = (1 - 0.26) \cdot K_{it-1}^{I,PIM} + \text{SGA}_{it}$, with $K_{i0}^{I,PIM} = 0$
  - $K_{it}^{I} = K_{it}^{I,ACC} + K_{it}^{I,PIM}$
- Cleaning: everything is deflated with the CPI, the 1st and 99th percentiles of all revenue shares are dropped Back

Results: Revenue TFP Variance Decomposition

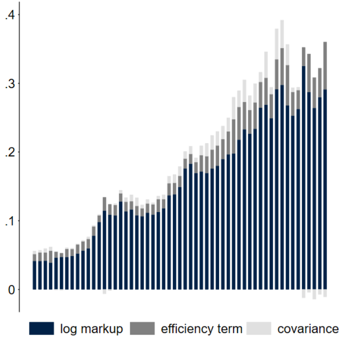

Figure: TFPR variance decomposition 1962-2024: markups vs efficiency term

- $\omega_{it}^{R} := \underbrace{\omega_{it}}_{\text{log quantity TFP}} + \underbrace{p_{it}}_{\text{log price}} = \underbrace{\log \mu_{it}}_{\text{markup}} + \underbrace{(\log \lambda_{it} + \omega_{it})}_{\text{NOT quantity TFP}}$
- TFPR dispersion is largely due to markup dispersion, which increased ▸ Back