Case Study: How substitutable are capital and labour?

Estimating the CES with evidence from computerization and AI

How do we know how readily machines can substitute for workers? A computer may take over a calculation, help an engineer design a component, or require a technician to keep it running. These examples suggest different relationships between machines and people. However, economic theory tells us that these relationships show up in different reactions of quantities and prices. So, to quantify the substitutability between capital and labour, economists compare the quantities of inputs used in production with their prices and payments. Ultimately, we want to find a way to *estimate* the elasticity of substitution.

Start with the familiar capital–labour CES:

$$Y = z\left[(A_K K)^{(\sigma-1)/\sigma} + (A_L L)^{(\sigma-1)/\sigma}\right]^{\sigma/(\sigma-1)}. \tag{1}$$

Here K and L are capital and labour services, A_K and A_L are their productivity terms, and z is common productivity. The elasticity of substitution is $\sigma > 0$. Each CES input has coefficient one, with constant scale factors absorbed into productivity terms and the units of effective inputs. The empirical studies estimate fuller specifications with input weights; the simplified forms here explain the substitution relationships.

Let w be the wage, r the rental payment per unit of capital services, and p the output price. Competitive input choice makes w/r equal to the ratio of marginal products. In particular, the first-order conditions for cost minimization imply the following relative demand and relative-payment relationships:

$$\log\frac{K}{L} = \sigma\log\frac{w}{r} + (\sigma-1)\log\frac{A_K}{A_L}, \tag{2}$$

$$\log\frac{rK}{wL} = \frac{\sigma-1}{\sigma}\log\frac{A_K K}{A_L L}. \tag{3}$$

At fixed A_K/A_L, the slope of $\log(K/L)$ against $\log(w/r)$ is σ.

This gives an empirical strategy, but also a problem. The strategy is to collect observations on capital and labour quantities, K_t and L_t, and their prices, r_t and w_t, for different years t. We can then run the regression

$$\underbrace{\log(K_t/L_t)}_{\text{dependent variable}} = b_0 + b_1 \times \underbrace{\log(w_t/r_t)}_{\text{ind. variable}} + u_t. \tag{4}$$

Here b_0 is the intercept, b_1 is the slope, and u_t captures unobserved influences and measurement errors. Suppose that technology and all other determinants of relative input demand are constant, so quantities adjust only in response to changing relative prices. Comparing this regression with equation (2) gives

$$b_0 = (\sigma-1)\log\frac{A_K}{A_L}, \qquad b_1 = \sigma.$$

With accurately measured quantities and prices, observations then lie on a single relative-demand line. The estimated slope gives the estimated elasticity:

$$\widehat{b}_1 = \widehat{\sigma}.$$

The hats indicate estimates obtained from the data. In this ideal case, a steeper fitted line means that firms change their capital–labour ratio more strongly when the relative price changes.

Applying this strategy to real data is more difficult. Measuring the capital stock requires combining different assets, accounting for depreciation and improvements in quality, and determining how much productive service the stock supplies. Investment spending alone does not measure K_t, and a machine's purchase price is different from the rental price r_t of its services. Measuring labour input is also more involved than counting workers: hours differ, and workers have different education, experience, and skills. These types of labour may interact differently with capital, so combining them into one L_t can hide the substitution relationships we want to estimate.

The technological variables A_L and A_K are also changing over time and are not directly observed. If their ratio changes, observations can move between relative-demand curves as well as along them (a change in A_K/A_L shifts the relative-demand curve). An improvement in equipment or workers' productivity may therefore change input quantities even without a change in relative prices. A fitted line through historical observations does not automatically measure the CES elasticity. This is the standard problem economists face when they try to estimate demand curves from historical data.

1 What computerization asks us to explain

An academic article by Krusell, Ohanian, Ríos-Rull, and Violante (2000), *Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis*, published in *Econometrica*, tackles this estimation problem. The authors seek to estimate the elasticities of substitution between capital equipment and different types of labour, and to understand how equipment accumulation affects wage inequality. They begin by examining annual U.S. data on input prices, quantities, and income shares before estimating their production function.

Exhibit 1 reproduces the four figures with which they introduce the evidence. The article classifies workers with at least 16 years of schooling as graduates ("skilled" in the original figures), and all other workers as non-graduates ("unskilled"). These are education groups, with differences in abilities and experience within each group.

Exhibit 1. Equipment prices, labour supply, wages, and income shares

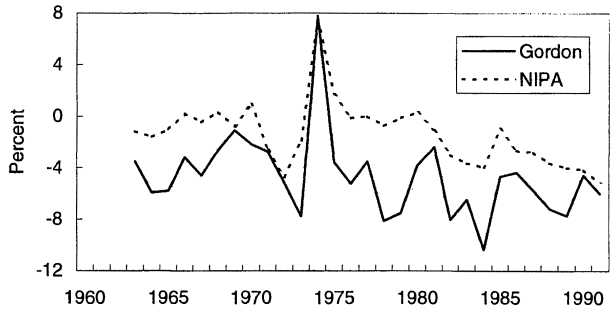


Figure 1. Two measures of changes in capital equipment prices: Gordon's vs. The NIPA's (%).

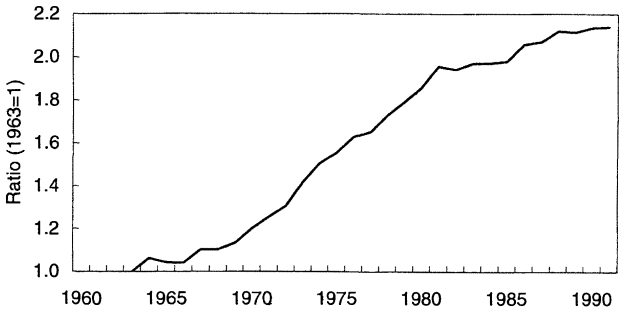


Figure 2. The labor input ratio: Skilled vs. unskilled hours worked (normalized with 1963=1).

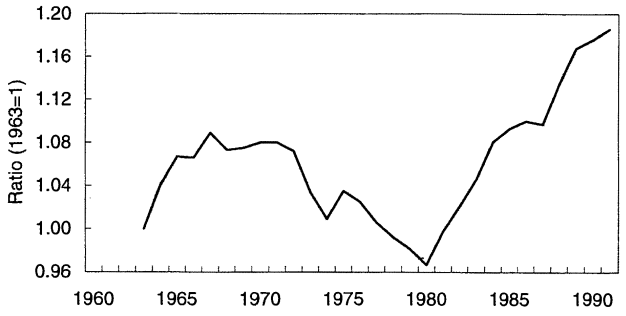


Figure 3. The skill premium: Skilled vs. unskilled wages per hour (normalized with 1963=1).

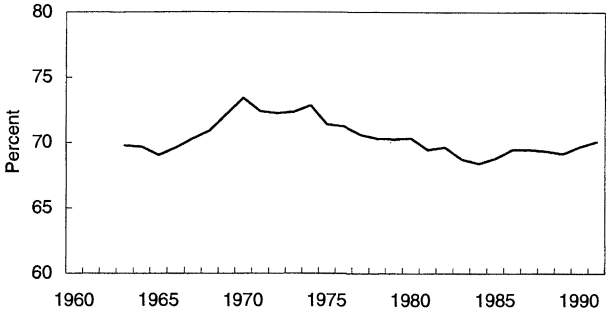


Figure 4. Labor's share of aggregate income (%).

Source: Krusell, Ohanian, Ríos-Rull, and Violante (2000), Figures 1–4, p. 1032. United States, 1963–1991 in the plotted panels; the article's estimation sample extends through 1992. The panels are described below.

Upper left: the price of equipment. This panel compares annual percentage changes in equipment prices relative to consumption prices. The solid line uses Economist's Robert Gordon quality-adjusted equipment-price series, extended by the authors; the dotted line uses the official National Income and Product Accounts (NIPA) series. Most changes are negative, and the quality-adjusted series generally shows larger declines: equipment became cheaper per unit of useful performance. The difference between the lines illustrates why measuring improvements in machines matters. The vertical axis reports price *changes*; estimating the opening regression would also require the rental cost of capital services.

Upper right: the relative supply of graduate labour. This panel uses the Current Population Survey (CPS) to construct the ratio of graduate to non-graduate labour input from hours worked. The authors aggregate worker types using fixed weights to account for changes in workforce composition. The ratio is normalized to one in 1963 and rises to more than two by the end of the plotted period. Thus graduate labour became much more abundant relative to non-graduate labour; this is a relative-input index, rather than a count of graduates or their share of employment.

Lower left: the graduate wage premium. Also constructed from CPS data, this panel shows graduates' hourly wage relative to non-graduates' hourly wage, with its 1963 value normalized to one. The

normalization sets the initial index to one regardless of the actual initial wage gap, this way the index is highlighting the changes of the premium over time and not its actual level. The premium rises modestly in the 1960s, falls through much of the 1970s, and increases sharply after 1980. The last pattern is especially striking alongside the upper-right panel: graduates became more abundant while their relative wage increased.

Lower right: labour's share of aggregate income. This panel uses aggregate income totals to measure wages, salaries, and benefits as a percentage of labour plus capital income. The article's capital-income measure includes depreciation, corporate profits, net interest, and rental income. The share fluctuates around 70% without a sustained trend over this historical window. It describes the division of income between all labour and capital, whereas the wage-premium panel describes differences within labour. The aggregate share therefore adds information that a comparison of graduate and non-graduate wages alone cannot provide.

Together, the panels pose the article's main question: why did the graduate wage premium rise after 1980 despite the growing relative supply of graduates? With an unchanged relative-demand curve, a larger relative supply would tend to reduce the premium. The authors ask whether the expanding stock of increasingly affordable equipment could instead raise demand for graduate labour relative to other labour. To assess that explanation, they need estimates of how differently equipment interacts with the two groups, and a model that can also account for aggregate income shares.

2 Allowing equipment to interact differently with two labour groups

One chief concern in measuring the elasticity of substitution is that capital can interact differently with different types of labour. Equipment may perform activities previously done by some workers while increasing the productivity of others. Combining all workers into one labour input can conceal these differences, so a single capital–labour elasticity may miss an important part of how equipment affects relative wages.

The most crucial distinction for this analysis is between university graduates and other workers. Separating these groups allows the researchers to ask how readily equipment substitutes for each type of labour, and whether it is more complementary with graduates' skills. This distinction connects the estimation problem to the changing graduate wage premium in Exhibit 1.

To represent these relationships, distinguish structures K_S, such as buildings, from equipment K_E, such as machines and computers. Let H denote hours worked by university graduates and L hours worked by other workers. Their productivities per hour are A_H and A_L, giving efficiency-adjusted inputs $A_H H$ and $A_L L$. Equipment is measured in quality-adjusted units, and z denotes productivity common to all inputs.

The Nested CES

First combine equipment with graduate labour in a CES composite C, then combine C with non-graduate labour:

$$C = \left[K_E^{(\sigma_I - 1)/\sigma_I} + (A_H H)^{(\sigma_I - 1)/\sigma_I} \right]^{\sigma_I/(\sigma_I - 1)}, \tag{5}$$

$$Y = z K_S^a \left(\left[(A_L L)^{(\sigma_O - 1)/\sigma_O} + C^{(\sigma_O - 1)/\sigma_O} \right]^{\sigma_O/(\sigma_O - 1)} \right)^{1-a}. \tag{6}$$

Here $0 < a < 1$ is the structures exponent, and $1 - a$ is the exponent on the outer CES composite.
The inner elasticity σ_I concerns equipment and graduate labour; the outer elasticity σ_O concerns non-graduate labour and C.
Capital–skill complementarity means $\sigma_O > \sigma_I$: substitution is easier at the outer margin than between equipment and graduates inside C.

Equipment and the wage premium

Define equipment's payment share *within* C:

$$s_E = \frac{K_E^{(\sigma_I - 1)/\sigma_I}}{K_E^{(\sigma_I - 1)/\sigma_I} + (A_H H)^{(\sigma_I - 1)/\sigma_I}}, \qquad 0 < s_E < 1.$$

The competitive wage premium satisfies

$$\frac{w_H}{w_L} = \frac{A_H}{A_L} \left(\frac{A_H H}{A_L L} \right)^{-1/\sigma_O} \left(\frac{C}{A_H H} \right)^{1/\sigma_I - 1/\sigma_O}. \tag{7}$$

At fixed labour quantities, efficiencies, structures, and other technology,

$$\frac{\partial \log(w_H/w_L)}{\partial \log K_E} = s_E \left(\frac{1}{\sigma_I} - \frac{1}{\sigma_O} \right). \tag{8}$$

For finite changes, with A_H, A_L and both elasticities fixed,

$$\Delta \log \frac{w_H}{w_L} = -\frac{1}{\sigma_O} \Delta \log \frac{H}{L} + \left(\frac{1}{\sigma_I} - \frac{1}{\sigma_O} \right) \Delta \log \frac{C}{A_H H}. \tag{9}$$

Thus the relative-supply contribution and the equipment-composite contribution can push the premium in opposite directions.

3 Taking the production function to the data

What would a researcher actually measure? The original study uses annual U.S. observations for 1963–1992, combining national accounts with survey data on workers. Several variables require construction before an elasticity can be estimated.

Capital follows a simple perpetual-inventory construction:

$$K_t = (1 - \delta)K_{t-1} + \frac{I_t^{\text{nominal}}}{P_t^K},$$

where P_t^K is the investment deflator and δ is depreciation. A better machine can supply more effective services even if its sticker price is unchanged. Failing to adjust for quality can therefore understate real investment and distort the measured capital stock. An initial stock and a depreciation assumption are also needed.

The remaining problem is that labour quantities may respond to unobserved productivity changes. For example, an improvement in graduates' productivity can change both their hours and their wages. Using observed hours directly can then mix the effect of input quantities with the effect of productivity. An *instrument* is a variable that helps predict an input, such as labour hours, while being unrelated to the unobserved shocks in the equation we want to estimate. It lets us use the part of input variation predicted by that variable to separate the input's effect from those shocks. Krusell and coauthors therefore instrument the two *labour quantities*, H_t and L_t. We can explain their estimation procedure in three steps.

Step 1. Regress labour hours on the instruments. First run a separate regression for each labour group:

$$H_t = d_{H,0} + d_{H,1}Z_t + v_{H,t}, \tag{10}$$

$$L_t = d_{L,0} + d_{L,1}Z_t + v_{L,t}. \tag{11}$$

Here Z_t stands for the collection of instruments; $d_{H,1}Z_t$ and $d_{L,1}Z_t$ summarize their contributions, with a separate coefficient on each instrument in each regression. The article's Appendix (p. 1051) lists:

- Current and lagged stocks of equipment and structures.
- The lagged relative price of equipment.
- The lagged U.S. Commerce Department composite index of business-cycle indicators.
- A time trend, with a constant included in each regression.

The instruments must predict labour hours and be unrelated to the relevant current unobserved shocks. The authors treat current capital stocks as exogenous to those shocks. After estimating the first regressions, calculate their fitted values:

$$\widehat{H}_t = \widehat{d}_{H,0} + \widehat{d}_{H,1}Z_t, \qquad \widehat{L}_t = \widehat{d}_{L,0} + \widehat{d}_{L,1}Z_t.$$

These are the hours predicted by the instruments; the remaining first-regression residuals are left out of the predicted inputs.

Step 2. Try parameter values and construct the effective inputs. Next choose trial values for σ_I and σ_O, together with a distribution for the unobserved productivity terms. In our notation, the article's

benchmark assumptions can be written as

$$\log A_{j,t} = \mu_j + \eta_{j,t}, \qquad \eta_{j,t} \sim N(0, \tau^2), \qquad j = H, L. \tag{12}$$

Here μ_j is a constant baseline log productivity and $\eta_{j,t}$ is a random shock with zero mean and variance τ^2. The shocks are independent across years and labour groups in the benchmark, with the same variance for both groups. There is no systematic productivity trend, but productivity can fluctuate from year to year. One baseline productivity level is fixed to set the units; the remaining baseline and the shock variance are among the parameters to be chosen.

Draw possible values of $A_{H,t}$ and $A_{L,t}$ from this distribution; let s identify one such draw. Combine those productivity values with the fitted hours from Step 1. In particular, calculate

$$\widehat{C}_t^{(s)} = \left[K_{E,t}^{(\sigma_I - 1)/\sigma_I} + \left(A_{H,t}^{(s)} \widehat{H}_t \right)^{(\sigma_I - 1)/\sigma_I} \right]^{\sigma_I/(\sigma_I - 1)}. \tag{13}$$

The inner elasticity σ_I determines how equipment and graduate labour combine in $\widehat{C}_t^{(s)}$. The outer elasticity σ_O determines how that composite interacts with $A_{L,t}^{(s)} \widehat{L}_t$ in the production function. Repeating the draws produces many possible predictions for the same observed capital stocks and fitted labour hours.

Step 3. Choose parameters to match the observed income relationships. Use these constructed inputs to predict the division of income and compare those predictions with the data. For example, multiplying equation (7) by H/L and taking logs gives the following wage-bill relationship in our simplified classroom CES:

$$\log \frac{w_{H,t} H_t}{w_{L,t} L_t} \;=\; \left(1 - \frac{1}{\sigma_O}\right) \log \frac{A_{H,t}^{(s)} \widehat{H}_t}{A_{L,t}^{(s)} \widehat{L}_t} \;+\; \left(\frac{1}{\sigma_I} - \frac{1}{\sigma_O}\right) \log \frac{\widehat{C}_t^{(s)}}{A_{H,t}^{(s)} \widehat{H}_t} \;+\; \varepsilon_t^{(s)}. \tag{14}$$

The left-hand side is the observed graduate/non-graduate *wage-bill ratio* in logs, using actual wages and hours. The right-hand side uses fitted hours and a possible realization of productivity; $\varepsilon_t^{(s)}$ records the difference between that prediction and the observed outcome. The wage-bill ratio is also the ratio of the two groups' income shares, so it tells us how labour income is divided between them.

For each set of trial parameters, compare the model's predicted averages and variability across the draws with the observed income relationships. Change the elasticities, the productivity baseline, the shock variance, and the other production parameters, then repeat the calculations in Step 2. The chosen values are those that make the model fit best; the selected elasticities are $\widehat{\sigma}_I$ and $\widehat{\sigma}_O$. The first-step fitted hours stay fixed while these production parameters are varied.

The article fits its fuller weighted wage-bill relationship jointly with the total labour income share and equality of expected investment returns on equipment and structures. Return equality accounts for depreciation and expected equipment-price changes; it does not require equal rental payments per unit of the two assets. Our displayed equation illustrates how the inputs enter; the paper's full fit also

accounts for its additional parameters and disturbances. The instruments address the response of hours to unobserved shocks, while the assumed productivity distribution allows the model to account for productivity that is not directly measured. The results remain conditional on those assumptions.

Exhibit 2. Selected parameter estimates from the historical study

Parameter	Estimate	Meaning in the model
σ_I	0.67	Elasticity between equipment and graduate labour in the inner composite.
σ_O	1.67	Elasticity between non-graduate labour and the equipment–graduate composite.
a	0.117	Structures exponent in the Cobb–Douglas part of production.
τ^2	0.043	Common variance of the shocks to log productivity in equation (12).

Source: Krusell, Ohanian, Ríos-Rull, and Violante (2000), Table I, p. 1040, and Table II, p. 1041; annual U.S. data, 1963–1992. Elasticities are already expressed in the case's notation. Our τ^2 corresponds to the paper's η_ω^2. These selected estimates come from the paper's fuller weighted specification.

The estimated inner elasticity of 0.67 indicates that equipment and graduate labour are relatively difficult to substitute for one another. The outer elasticity of 1.67 indicates greater scope for substitution between non-graduate labour and the composite. The important comparison is $\sigma_O > \sigma_I$: equipment is more complementary with graduate labour, giving equipment accumulation a way to raise the graduate wage premium. These estimates describe aggregate production relationships, so they do not tell us which individual tasks machines replace.

4 The new measurement problem with AI

Can the same approach tell us whether AI substitutes for graduate labour? The extension introduces AI capital K_{AI} separately from ordinary equipment. Its restricted production structure first combines $A_H H$ and K_E into C, then combines C with AI. Using unit coefficients as above:

$$B = \left[C^{(\sigma_{AI}-1)/\sigma_{AI}} + K_{AI}^{(\sigma_{AI}-1)/\sigma_{AI}} \right]^{\sigma_{AI}/(\sigma_{AI}-1)}. \tag{15}$$

The production function becomes

$$Y = z K_S^a \left(\left[(A_L L)^{(\sigma_O-1)/\sigma_O} + B^{(\sigma_O-1)/\sigma_O} \right]^{\sigma_O/(\sigma_O-1)} \right)^{1-a}. \tag{16}$$

The reported estimates come from a separately fitted weighted specification; its productivity terms and input normalizations need not match the historical model. For the classroom example, we retain the constant-returns structure and the notation a and $1-a$ used above.

The new elasticity σ_{AI} concerns *AI and the entire equipment–graduate composite* C. It is not directly an

elasticity between AI and an individual graduate worker. If $\sigma_{AI} < 1$, the two inputs are complements in this nest; if $\sigma_{AI} > 1$, they are more readily substitutable. For $\sigma_{AI} < 1$, the fixed complementary composite limits that gain as K_{AI} becomes arbitrarily large because college-educated workers and equipment capital act as weak links in production.

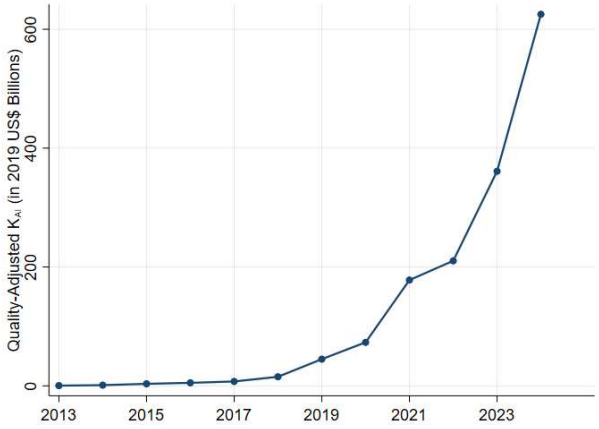
Taking the extension to the data now requires updated measures of both the price of AI capital and the level of its productive stock. Rapid improvements in hardware mean that the same expenditure can buy more effective computing capacity, so the price measure must account for quality changes. Exhibit 3 presents Luduvice and Pinheiro's quality-adjusted hardware-price index and their constructed AI capital stock, showing how the price and quantity of this new input have changed.

Exhibit 3. Faster AI hardware and a constructed capital stock

(a) Quality-adjusted hardware price



(b) Constructed AI capital stock

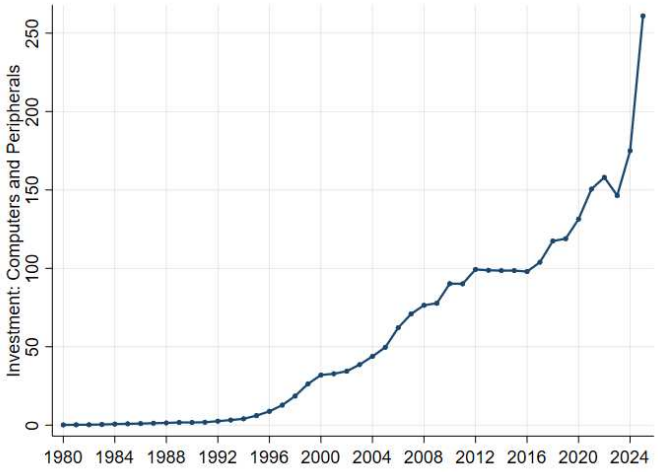


Measurement estimates, 2013–2024. Panel (a) uses accelerator release prices and hardware characteristics to estimate a quality-adjusted price index, with 2019 equal to one. The reported decline from 2013 to 2024 is approximately 97%. Panel (b) accumulates an investment measure deflated by that index, with an initial stock of zero in 2012 and assumed annual depreciation of 15.6%; its units are billions of 2019 U.S. dollars. Both panels are constructed estimates, not a count of AI tasks or jobs. The paper reports uncertainty bands separately.

The price index comes from a regression comparing prices while controlling for processor performance, memory, and power requirements. The constructed stock uses AI-company investment data that include financing transactions, so it also depends on how those transactions map into installed productive assets. It is not a direct national-accounts measure of AI hardware purchases.

The estimation also requires a stock measure that can represent K_{AI} alongside the other model inputs. Because accelerator data cover only a short period, the estimation uses a stock constructed from the broader BEA Computers & Peripherals investment series as its AI proxy. Exhibit 4 shows the investment series underlying that stock, connecting the measurement discussion to the capital input actually used to fit the model.

Exhibit 4. The broader asset category used for estimation



National-accounts evidence. U.S. real investment in Computers & Peripherals, 1980–2024, in billions of dollars as reported in the source figure. This is an investment flow, not the capital stock. The figure does not specify its reference-dollar year.

That category includes computing used for purposes other than AI. From the baseline AI start year, 2013, the corresponding assets are removed from ordinary equipment to avoid counting them twice. The quality-adjusted AI price series also helps connect the AI return relationships to observed price changes. The full estimation sample is 1963–2024, with 60 usable annual observations; the two AI-specific return conditions begin in 2013. Those long historical data do not supply 60 years of observations on modern AI.

Using these capital-price and stock measures together with labour and income data, the authors fit the extended production model. Exhibit 5 reports the resulting elasticities, allowing us to compare the fitted substitution margins with the historical benchmark and assess the relationship between AI and the equipment–graduate composite.

Exhibit 5. Estimated elasticities at the different CES margins

CES margin	Original study	AI extension
Equipment and graduate labour: σ_I	0.67	0.24
Non-graduate labour and the outer composite: σ_O	1.67	2.70
AI and the equipment–graduate composite: σ_{AI}	—	0.89

Estimates from the production models. Original: annual U.S. data, 1963–1992. Extension: annual U.S. data, 1963–2024, with the shorter AI relationships described above. Elasticities have already been calculated. They refer to different margins and come from different specifications and samples. The estimates are conditional on the input definitions, production nests, and assumptions about unobserved technology; they are not direct experimental measurements.

5 Questions

Question 1. When does a slope measure substitution? (20 marks)

(a) In equation (2), hold relative productivity fixed and consider different w/r. For this diagram, set $A_K = A_L$ and normalize both initial ratios to one. Draw $\log(K/L)$ against $\log(w/r)$ for a low and a high elasticity. Explain what the slopes measure. [8 marks]

(b) A researcher fits a line to annual observations on those two ratios and calls its slope the elasticity. Explain why changes in relative productivity can invalidate this interpretation. Name the price and quantity measures the researcher needs, and explain why a computer purchase-price index and nominal investment expenditure are insufficient. [12 marks]

Question 2. What does capital–skill complementarity imply? (25 marks)

(a) Use the historical estimates in Exhibit 2 and equation (8) to show what an increase in K_E does to the wage premium when labour quantities, efficiencies, structures, and other technology are fixed. Explain the mechanism and whether it requires $\sigma_O > 1 > \sigma_I$. [12 marks]

(b) Now allow H/L to rise while efficiencies and parameters remain fixed. State the exact condition in equation (9) for the premium to rise. Use Exhibit 1 to explain why observing rising graduate supply and cheaper equipment does not by itself identify the elasticities or establish that equipment caused the premium's increase. Does the approximately stable total labour share prove that both nests are Cobb–Douglas? [13 marks]

Question 3. How is the model disciplined by real data? (20 marks)

(a) Explain why the original study fits the wage-bill ratio, total labour income share, and expected capital-return condition jointly. What different information does each provide? Explain why a good fit to these relationships is not the same as a controlled experiment changing equipment alone. [12 marks]

(b) A colleague proposes regressing the wage premium on relative graduate hours and interpreting the slope as a supply effect. Explain how an unobserved improvement in graduates' productivity could affect both hours and wages. Why might omitting it distort that interpretation? [8 marks]

Question 4. What exactly is being measured as AI capital? (20 marks)

(a) Explain why quality adjustment matters for the price and stock in Exhibit 3, even if a new processor has the same sticker price as an older one. Identify one further assumption needed to turn investment into a productive capital stock. Does the stock's rapid growth establish whether AI and workers are substitutes? [10 marks]

(b) A commentator says that the elasticity in Exhibit 5 was estimated by regressing graduate wages on the accelerator-based stock in Exhibit 3. Evaluate this description using Exhibit 4 and the estimation account. Explain the purpose of removing the proxy assets from ordinary equipment, and identify one remaining limitation of the proxy. [10 marks]

Question 5. What can we conclude from the AI estimate? (15 marks)

(a) Evaluate the claim: "An elasticity of 0.89 proves that AI cannot replace any graduate task and that every graduate's employment will rise." Identify the inputs to which the estimate applies, and state what follows for B when K_{AI} rises at fixed C and other technology. [9 marks]

(b) Explain why interpreting the estimated AI elasticity depends on the proxy and the assumptions used to construct the capital stock. Would the historical estimate settle the elasticity for a future AI technology or a different production structure? [6 marks]

References

This case study is based on models, results, and exhibits from the following papers. The questions and simplified classroom specifications were prepared for this course.

- Krusell, Per, Lee E. Ohanian, José-Víctor Ríos-Rull, and Giovanni L. Violante (2000). *Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis*. *Econometrica* 68(5), 1029–1053.
- Luduvice, André Victor D., and Roberto Pinheiro (2026). *AI-Augmented Capital-Skill Complementarity*. August version (course copy).