Multiple sectors in an input-output network

Properties of production functions: a review

General properties of constant-returns-to-scale production. First, consider any constant-returns-to-scale production function with multiple inputs, $F(x_1, \ldots, x_n)$.[1] If we log-differentiate this production function, we get

$$d\log F = \frac{\partial \log F}{\partial \log x_1} d\log x_1 + \ldots + \frac{\partial \log F}{\partial \log x_n} d\log x_n$$

which we can also write, using the notation $F_i \equiv \partial F / \partial x_i$, as

$$\frac{dF}{F} = \frac{F_1 x_1}{F}\frac{dx_1}{x_1} + \ldots + \frac{F_n x_n}{F}\frac{dx_n}{x_n} \tag{1}$$

which says that the percent change in output F equals the sum of percent changes dx_i/x_i in each input i, weighted by their respective "shares" $F_i x_i / F$, which are also the elasticities of F with respect to x_i. These shares sum to one by what is often called Euler's theorem, which states that a constant-returns-to-scale function satisfies $F = F_1 x_1 + \ldots + F_n x_n$.[2]

Further, if we suppose that the inputs $x_1, \ldots, x_n$ are chosen given prices $p_1, \ldots, p_n$ in a cost-minimizing way to produce some quantity y, i.e. that they solve the problem

$$\begin{aligned} &\min_{x_1,\ldots,x_n} p_1 x_1 + \ldots + p_n x_n \\ &\text{s.t. } y = F(x_1, \ldots, x_n) \end{aligned} \tag{2}$$

then if we let ψ be the Lagange multiplier on the constraint, which is the marginal cost of producing an additional unit of output, we have the optimality conditions $p_i = \psi F_i$ for each i. It follows that the share of total costs spent on any input i is $p_i x_i / \sum_j p_j x_j = \psi F_i x_i / \sum_j \psi F_j x_j = F_i x_i / F$, i.e. that it is what we are already calling the "share" of i.

Note that given prices p_i, both the marginal cost ψ and the shares $F_i x_i / F$ will be invariant to the level y of output that is demanded. Hence ψ is really a constant cost per unit, not just marginal cost; ψy is the total cost of producing y. This is due to constant returns to scale: any production plan $\{x_i\}$ that works for one y can be scaled up or down to $\{\alpha x_i\}$ to produce αy.

Finally, given the cost-minimization problem (2), we can ask what happens to the total costs of producing y if there is some shock to *prices*. It follows immediately from the envelope theorem, applied to (2) that the change in the objective, total cost, will be $dp_1 x_1 + \ldots + dp_n x_n$. The percent change in total cost will then

[1] In the background, we'll assume that F is strictly concave in every direction except the one where all x_i are increased proportionally, i.e. that $d\mathbf{x}'\mathbf{F_{xx}}d\mathbf{x} < 0$ except when $d\mathbf{x}$ is proportional to $\mathbf{x}$, where it equals zero. This ensures a unique optimum production plan, and other nice behavior.

[2] We can derive this by writing $\alpha F(x_1, \ldots, x_n) = F(\alpha x_1, \ldots, \alpha x_n)$, and then differentiating both sides with respect to α around 1. It is also visible from (1) itself: if all the inputs increase by the same infinitesimal amount $dx_i / x_i = d\epsilon$, then by constant returns to scale we should have $dF/F = d\epsilon$ on the left too, and cancelling these out we get $1 = \frac{F_1 x_1}{F} + \ldots + \frac{F_n x_n}{F}$.

be

$$\frac{d\psi}{\psi} = \frac{dp_1 x_1 + \ldots + dp_n x_n}{p_1 x_1 + \ldots + p_n x_n} = \frac{\psi F \cdot \left(\frac{F_1 x_1}{F}\frac{dp_1}{p_1} + \ldots + \frac{F_n x_n}{F}\frac{dp_n}{p_n}\right)}{\psi F}$$
$$= \frac{F_1 x_1}{F}\frac{dp_1}{p_1} + \ldots + \frac{F_n x_n}{F}\frac{dp_n}{p_n} \tag{3}$$

i.e. it is the share-weighted sum of percent changes in the price of each input![3] This is a very nice, dual counterpart to (1), which states that the percent change in output is the share-weighted sum of percent changes in input. Here, (3) says that the percent change in output *cost* is the share-weighted sum of percent changes in input *prices*.

(Sometimes this result is stated in terms of output prices rather than cost, which is true as long as there is a constant markup. But for now we aren't specifying output prices at all, just solving a cost minimization problem.)

Properties of constant-elasticity-of-substitution (CES) functions. The above results held for any constant-returns-to-scale function. Now let's specify F to a constant-elasticity-of-substitution CES function, of the form

$$F(x_1, \ldots, x_n) = \left(\sum a_i x_i^{\frac{\theta-1}{\theta}}\right)^{\frac{\theta}{\theta-1}} \tag{4}$$

where a_i are nonnegative constants and $\theta > 0$ is the elasticity of substitution. Differentiating F with respect to x_i, we obtain

$$F_i = \frac{\theta}{\theta-1}\left(\sum a_i x_i^{\frac{\theta-1}{\theta}}\right)^{\frac{1}{\theta-1}} \cdot \frac{\theta-1}{\theta} a_i x_i^{-\frac{1}{\theta}}$$
$$= F^{\frac{1}{\theta}} a_i x_i^{-\frac{1}{\theta}} = a_i \left(\frac{x_i}{F}\right)^{-\frac{1}{\theta}} \tag{5}$$

Equating this with $p_i/\psi = F_i$ and rearranging gives

$$\frac{x_i}{F} = a_i^{\theta}\left(\frac{p_i}{\psi}\right)^{-\theta} \tag{6}$$

i.e. that the ratio of input i to total output has a constant elasticity of $-\theta$ with respect to the ratio of input price p_i to output cost ψ. We can also multiply both sides by $F_i = p_i/\psi$ to obtain

$$\frac{F_i x_i}{F} = a_i^{\theta}\left(\frac{p_i}{\psi}\right)^{1-\theta} \tag{7}$$

which states that the input cost share $F_i x_i/F$ of i has a constant elasticity of $1-\theta$ with respect to p_i/ψ. Note that this implies that the cost share of i is increasing in its relative price in the "complements" case $\theta < 1$, and decreasing in its relative price in the "substitutes" case $\theta > 1$; it is constant in the Cobb-Douglas case $\theta = 1$.

[3] Note that the first equality uses $p_1 x_1 + \ldots + p_n x_n = \lambda y = \lambda F$ for the denominator, while multiplying the numerator by $\lambda F_i / p_i = 1$.

Adding a productivity shifter. Suppose further that there is some Hicks-neutral productivity parameter A that scales the production function, so that the production function is now

$$F(x_1, \ldots, x_n; A) = A \left(\sum a_i x_i^{\frac{\theta-1}{\theta}} \right)^{\frac{\theta}{\theta-1}}$$

We then have $F_i = A a_i \left(\frac{x_i}{F/A} \right)^{-\frac{1}{\theta}}$, which equating with $p_i/\psi = F_i$ gives

$$\frac{x_i}{F/A} = a_i^\theta \left(\frac{p_i}{A\psi} \right)^{-\theta} \tag{8}$$

and similarly

$$\frac{F_i x_i}{F} = a_i^\theta \left(\frac{p_i}{A\psi} \right)^{1-\theta} \tag{9}$$

where (8) and (9) are immediate generalizations of (6) and (7).

We note that (9) is particularly nice, since ψ clearly varies inversely with A, so that $A\psi$ is unaffected by changes in A. Combining (9) with (3), we have to first order that

$$d \log \left(\frac{F_i x_i}{F} \right) = (1-\theta) \left(d \log p_i - \sum_j \frac{F_j x_j}{F} d \log p_j \right) \tag{10}$$

i.e. the log change in input share of i equals $(1-\theta)$ times the difference between the log change in input price i and the share-weighted average log change in input price, with A having no effect.

Sectoral production networks: the perfect competition case[4]

We now define a static multisector economy where there are N different production sectors. Households have an endowment L of labor, which will be the unique primary factor of production in the economy, and maximize some CES aggregate

$$C = \mathcal{C}(c_1, \ldots, c_N) \tag{11}$$

of the N goods, subject to a budget constraint $\sum_i p_i c_i = wL$ that equates consumption expenditures with wage earnings.

Each production sector $i \in \{1, \ldots, N\}$ is perfectly competitive and has the technology

$$y_i = A_i f_i(L_i, \{x_{ij}\}_{j \in N}) \tag{12}$$

where A_i is a Hicks-neutral productivity shock and f_i is a CES aggregate with elasticity of substitution θ^i, which takes as inputs labor L_i and other goods j (which can include i itself), which are purchased as intermediates x_{ij}.[5] Due to perfect competition, the price p_i of sector-i output is equated with the constant cost of production given the technology (12).

[4]From now on, the notes will roughly follow section 3 onward of Baqaee and Rubbo (2023), and draw heavily from their narration and notation. For simplicity, however, I impose CES from the beginning, and start with perfect competition rather than markups. I also use c_i to denote consumption and y_i to denote production rather than their convention, which is y_j for consumption and x_j for production—there is some justification for their convention, but I find it hard to remember and it's not needed for our simple purposes.

[5]Note that in (12) we use the shorthand $j \in N$ for $j \in \{1, \ldots, N\}$.

Market clearing in each sector requires that $y_i = c_i + \sum_j x_{ji}$, i.e. that production of a good equals the sum of its use in consumption and intermediate inputs.

Some useful concepts: input-output matrix, Leontief inverse, Domar weights. It will be notationally convenient to treat both consumption and labor as "sectors", where with some abuse of notation consumption is sector C (given index 0) and labor is sector L (given index $N+1$). With this convention, we have $N+2$ sectors: consumption C, the regular production sectors $1, \ldots, N$, and the labor sector L. We will extend the notation naturally for each sector (e.g., for the labor sector, we will have $y_L = L$ and $p_L = w$), and we will sometimes use $\mathcal{U}$ to denote the set of all $N+2$ sectors.

Given some equilibrium of the economy, we define the ***input-output matrix*** Ω as the $(N+2) \times (N+2)$ matrix

$$\Omega_{ij} = \frac{p_j x_{ij}}{p_i y_i} \tag{13}$$

giving sector i's expenditures on j as a share of sector j's sales. Since (i) no sector has consumption as an input, (ii) labor is just an endowment and has no inputs, and (iii) consumption does not have labor as an input, the input-output matrix will take the special form

$$\Omega = \left[\begin{array}{c|ccc|c} 0 & \Omega_{C1} & \cdots & \Omega_{CN} & 0 \\ \hline 0 & \Omega_{11} & \cdots & \Omega_{1N} & \Omega_{1L} \\ 0 & \vdots & \ddots & \vdots & \vdots \\ 0 & \Omega_{N1} & \cdots & \Omega_{NN} & \Omega_{NL} \\ \hline 0 & 0 & \cdots & 0 & 0 \end{array}\right] \tag{14}$$

that has zeros reflecting (i)–(iii).[6]

We then define the ***Leontief inverse*** as $\Psi \equiv (I - \Omega)^{-1}$, where I is the identity. This satisfies the relationship

$$\Psi = (I - \Omega)^{-1} = I + \Omega + \Omega^2 + \ldots \tag{15}$$

and sums "direct" (captured in I and Ω) and "indirect" (captured in higher-order $\Omega^2, \Omega^3, \ldots$) exposures through the production network.

In a sense, Ψ_{ij} measures the fraction of a dollar spent on i that is ultimately spent on j, either directly or indirectly. First, the identity reflects the fact that the spending is on i in the first place. Then, holding the input-output matrix fixed, spending on i generates a fraction Ω_{ij} of spending on each intermediate good j. But those intermediate goods themselves require intermediate goods to produce, leading to additional "second-round" spending measured by $[\Omega^2]_{ij}$, and so on.

We define the ***Domar weight*** of sector i as the ratio of its sales to total consumption spending in the economy:

$$\lambda_i = \frac{p_i y_i}{\sum_{j=1}^N p_j c_j} = \frac{p_i y_i}{p_C C} \tag{16}$$

[6] Of course, the input-output matrix Ω is itself an endogenous object, and the exogenous structural parameters in the model are the coefficients inside each CES production function (i.e. the a_i that we saw earlier in (4)). We will sidestep explicitly dealing with these coefficients, which turns out to be unnecessary. Instead, we will use Ω itself (which has the advantage that it is in principle observable in the data) as a description of a baseline economy, and then obtain formulas in terms of Ω showing how prices and shares locally shift around this economy in response to shocks.

Given this, it turns out that the Leontief inverse Ψ can be written as

$$\Psi = \left[\begin{array}{c|ccc|c} 1 & \lambda_1 & \cdots & \lambda_N & 1 \\ \hline 0 & \Psi_{11} & \cdots & \Psi_{1N} & 1 \\ 0 & \vdots & \ddots & \vdots & \vdots \\ 0 & \Psi_{N1} & \cdots & \Psi_{NN} & 1 \\ \hline 0 & 0 & \cdots & 0 & 1 \end{array}\right] \tag{17}$$

It is worth remarking on two aspects of (17), beyond the straightforward points that the first column is all zeros below the first row (since consumption is not used as an input, no sector except consumption itself "ultimately spends" on consumption) and the last row is all zeros except the last column (since labor uses no inputs, it doesn't "ultimately spend" on anything except itself).

First, the first row of Ψ contains the Domar weights. For all sectors $i \neq C$, market clearing requires that $y_i = c_i + \sum_{j=1}^N x_{ji}$, which can be rewritten as $p_i y_i = p_i c_i + \sum_{j=1}^N p_j y_j \frac{p_i x_{ji}}{p_j y_j} = \sum_{j \in \mathcal{U}} p_j y_j \Omega_{ji}$. Dividing both sides by total consumption, this can be written as $\lambda_i = \sum_{j \in \mathcal{U}} \lambda_j \Omega_{ji}$. For consumption, on the other hand, we have $\lambda_C = 1$ by definition, while $\Omega_{jC} = 0$ for all j, so we can write this for all sectors as $\lambda_i = 1_{i=C} + \sum_{j \in \mathcal{U}} \lambda_j \Omega_{ji}$, or in vector form as $\boldsymbol{\lambda}' = \mathbf{e}_C' + \boldsymbol{\lambda}'\Omega$, where $\mathbf{e}_C$ is the vector with 1 in the 0th entry (corresponding to consumption) and 0s elsewhere. This can be rearranged as $\boldsymbol{\lambda}' = \mathbf{e}_C'(I - \Omega)^{-1} = \mathbf{e}_C'\Psi$, and $\mathbf{e}_C'\Psi$ is just the first row of Ψ.

Formalism aside, this is quite intuitive: consumption is the ultimate source of demand in the economy, so total spending in a sector i, relative to consumption spending, is just given by the Leontief inverse Ψ_{Ci}.

Second (and simpler), the final column of Ψ, corresponding to the share ultimately spent on labor, is all 1s. This is because the only factor of production in this economy is labor, so every dollar spent on any good is ultimately spent, 1-for-1, on labor.

Formally, we can see this by writing $\Psi = I + \Omega\Psi$. Then, if we write the last column as $\Psi\mathbf{e}_L$ (where $\mathbf{e}_L$ is defined analogously to $\mathbf{e}_C$ above, a vector with all 0s except a 1 at the final index for labor), this becomes $\Psi\mathbf{e}_L = \mathbf{e}_L + \Omega\Psi\mathbf{e}_L$, an equation for $\Psi\mathbf{e}_L$. It is clear that $\Psi\mathbf{e}_L = \mathbf{1}$ is a solution to this, since $\Omega\mathbf{1} = \mathbf{1} - \mathbf{e}_L$: every row of Ω sums to 1 except the L row, which sums to zero.

Some simple examples. We'll illustrate some canonical examples of this input-output framework.

First, a ***horizontal economy*** is one where all production sectors produce their goods from labor, without any input-output linkages. Here, we have $\Omega_{iL} = 1$ and $\Omega_{ij} = 0$ for all $i, j \in \{1, \ldots, N\}$. Domar weights λ_i are given simply by the share Ω_{Ci} of each sector in consumption. We also have $\Psi_{ij} = \Omega_{ij}$ except when $i = C$ and $j = L$.

Second, a ***simple roundabout economy*** is one where there is a single production sector 1, which buys some share Ω_{11} of intermediate inputs from itself, with the rest coming from labor: $\Omega_{1L} = 1 - \Omega_{11}$. It follows, for instance, that $\lambda_1 = \Psi_{C1} = \Psi_{11} = 1 + \Omega_{11} + \Omega_{11}^2 + \ldots = \frac{1}{1-\Omega_{11}}$. This was popularized by Basu (1995) as perhaps the simplest way to introduce intermediate inputs.

Third, a ***vertical economy*** is one where each production sector j except N uses the next production sector's output as its only input; sector N uses labor as its only input, and the consumer only consumes sector 1's output. Hence we have $\Omega_{ij} = 1$ if and only if $j = i + 1$, and $\Omega_{ij} = 0$ otherwise.

Note that we have already dealt with a horizontal economy (which we called a "mixture" of different sectors) and a simple roundabout economy in our discussion of pricesetting models.

Price determination. With these concepts, we can now solve for prices, which equal costs. From our earlier result (3), we know that absent technological change, the log change in cost of any sector is given by the cost share-weighted sum of the change in its input prices. Further, here we have a TFP shifter A_i; cost will naturally be inversely proportional to A_i. Hence for any $i \in \{0, \dots, N\}$, we have

$$d \log p_i = -d \log A_i + \sum_{j \in \mathcal{U}} \Omega_{ij} d \log p_j \tag{18}$$

For the labor sector $i = L$, we have $d \log p_i = d \log w$ by definition, and no inputs or productivity shocks. Hence we can extend (18) to the $i = L$ case by adding a $d \log w \cdot 1_{i=L}$ term, and then write as a vector equation

$$d \log \mathbf{p} = -d \log \mathbf{A} + \Omega d \log \mathbf{p} + d \log w \cdot \mathbf{e}_L \tag{19}$$

Moving $\Omega d \log \mathbf{p}$ to the left and then multiplying both sides by $(I - \Omega)^{-1} = \Psi$, we get

$$\begin{aligned} d \log \mathbf{p} &= -\Psi d \log \mathbf{A} + d \log w \cdot \mathbf{\Psi e}_L \\ &= -\Psi d \log \mathbf{A} + d \log w \cdot \mathbf{1} \end{aligned} \tag{20}$$

where we use $\mathbf{\Psi e}_L = \mathbf{1}$ from earlier.

(20) shows the usefulness of the Leontief inverse Ψ. An productivity gain in sector j lowers prices in sector i, relative to wages, in proportion to Ψ_{ij}, which says how much spending on sector i involves ultimately spending (directly or indirectly) on sector j.

Note that if we select the row of (20) corresponding to the consumption sector, then using the fact from earlier that $\Psi_{Cj} = \lambda_j$ (i.e. the first row of the Leontief inverse contains Domar weights), we obtain $d \log p_C = -\sum_{j=1}^{N} \lambda_j d \log A_j + d \log w$, which can be rewritten as

$$d \log \left(\frac{C}{L} \right) = d \log \left(\frac{w}{p_C} \right) = \sum_{j=1}^{N} \lambda_j d \log A_j \tag{21}$$

where we additionally use the household budget constraint $p_C C = wL$.

(21) states that the change in real wages, or equivalently final productivity in consumption goods over labor input, is given by the sum of changes in log productivity in each sector, weighted by their Domar weights. Note that if we choose the aggregate consumption good as the numeraire, setting $p_C \equiv 1$, then (21) can easily be combined with (20) to fully characterize the change in sectoral prices relative to this numeraire.

Equation (21) is an instance of a more general result known as ***Hulten's theorem***, which states that starting from an undistorted economy, if there is a shock to sector-level productivities, then the resulting change in log total factor productivity is equal to the Domar-weighted sum of sectoral log productivity shocks. (Here, we just have one "final output", consumption, and one factor, labor.)

Hulten's theorem suggests that if we just want to know the *aggregate* first-order implications of sector-level productivity shocks, we don't really need to know the structure of the input-output network. Instead, we just need the Domar weights λ_j, which in principle can easily be measured in the data. This "sufficient statistic" result means that papers like Acemoglu, Carvalho, Ozdaglar and Tahbaz-Salehi (2012), which study how network structure affects the transmission from sectoral to aggregate volatility—but focus on the first-order aggregate implications—are very interesting but arguably unnecessary. Thanks to (21), if we know the underlying variances (and possibly covariances) of sectoral productivities $d \log A_i$, all we need to know about the network is the resulting λ_i, which we can measure without knowing anything about the network!

Of course, as we expand either our questions or the model, things become more complicated. For one thing, even if it isn't necessary to obtain aggregate GDP, we might be interested in what happens to sector-level output and spending—which, is indeed, where we'll turn next. This requires knowing the full input-output network and the resulting Leontief inverse.

Further, if in response to a large sectoral productivity shock, the Domar weights $\lambda_i = \frac{p_i y_i}{p_C C}$ vary enough, then the first-order approximation provided by Hulten's theorem will be a bad one, and we will need to know second- or higher-order terms to accurately calculate the effect on aggregate productivity. This is the argument of Baqaee and Farhi (2019).[7]

Also, Hulten's theorem only applies starting from an efficient, undistorted economy, where the structure of production is exactly what a planner would choose. If instead we start from an inefficient economy, then there are first-order deviations from Hulten's theorem, arising from the interaction between substitution and preexisting distortions. For instance, if some sector has large markups and thus produces an inefficiently small quantity, any productivity shock that causes substitution toward that sector will lead to a first-order improvement in allocative efficiency. A number of prominent papers study this kind of setting, including Liu (2019), Baqaee and Farhi (2020), and Bigio and La'O (2020). For now, though, we'll look at the undistorted econony, without markups.

Quantity and share determination. Above, we saw that the effect of sectoral productivity changes on prices was relatively straightforward, given directly by the Leontief inverse. Quantities, it turns out, are more complex—but we can build on the results we've already obtained for prices.

We'll focus our derivation on changes in sales *shares*, i.e. Domar weights λ_i, which can ultimately be combined with our results on prices and aggregate production to obtain sector-level quantities.

Earlier, when we showed that the first row of the Leontief inverse Ψ was λ', we derived the equation $\boldsymbol{\lambda}' = \mathbf{e}'_C + \boldsymbol{\lambda}'\Omega$ from market clearing. In response to a shock, we can totally differentiate this equation to obtain

$$
\begin{aligned}
d\boldsymbol{\lambda}' &= \boldsymbol{\lambda}' d\Omega + d\boldsymbol{\lambda}'\Omega \\
d\boldsymbol{\lambda}'(I - \Omega) &= \boldsymbol{\lambda}' d\Omega \\
d\boldsymbol{\lambda}' &= \boldsymbol{\lambda}' d\Omega (I - \Omega)^{-1} = \boldsymbol{\lambda}' d\Omega \Psi
\end{aligned} \tag{22}
$$

Equation (22) says that the change $d\boldsymbol{\lambda}'$ in shares equals the original shares $\boldsymbol{\lambda}'$ times the change $d\Omega$ in intermediate input shares (giving the "direct" change in demand from substitution), times the Leontief inverse Ψ (which then propagates these changes in demand through the input-output network).

Now we solve for $d\Omega$. For an individual sector i whose technology has an elasticity of substitution θ^i,

[7] Note that Hulten's theorem will holds locally around any point, including the initial equilibrium, the equilibrium following a shock, and everything in between. To get significant nonlinearities, we need some mix of the shock being *really big* and the Domar weights changing *a lot*. In practice, my guess is that this does not happen much in the short run—and that when it does, there may be other complications like price being non-allocative and goods being rationed. This certainly does happen if we are looking at longer-run changes, though. It also helps us understand some hypothetical situations. For instance, the Domar weight on the agricultural sector in the US is probably less than 1%, but if agricultural productivity suddenly fell by 90%, we would probably starve—even though Hulten's theorem implies an effect of final consumption of only a few percent ($d \log C/L = 1\% \times d \log .1 \approx -.023$)! Here, the first-order approximation is a really bad one: in a world where basic foodstuffs became that scarce and we were literally starving, food prices would skyrocket and the Domar weight on the agricultural sector would become close to 1.

our earlier result (10) implies that

$$d\log\Omega_{ij} = (1-\theta^i)\left(d\log p_j - \sum_k \Omega_{ik} d\log p_k\right) \tag{23}$$

This can be combined with (22) and our previous solution for prices to fully characterize changes in shares.

Further analytical traction on shares. We can also attempt to go a bit further analytically with (22) and (23). We write:

$$\begin{aligned} d\lambda_i &= [\boldsymbol{\lambda}' d\Omega \Psi]_i \\ &= \sum_k \lambda_k \left(\sum_l d\Omega_{kl}\Psi_{li}\right) \\ &= \sum_k \lambda_k (1-\theta^k)\left(\sum_l \Omega_{kl}(d\log p_l - \sum_{l'}\Omega_{kl'} d\log p_{l'})\Psi_{li}\right) \\ &= \sum_k \lambda_k (1-\theta^k)\mathrm{Cov}_{\Omega_{kl}}(d\log p_l, \Psi_{li}) \end{aligned} \tag{24}$$

where $\mathrm{Cov}_{\Omega_{kl}}(d\log p_l, \Psi_{li})$ denotes the covariance of $d\log p_l$ and Ψ_{li} across all sectors, weighted by Ω_{kl}.

Now, suppose that we are interested specifically in the effect of a shock $d\log A_j$ to productivity in sector j. Then we note from (20) that $d\log p_l = -\Psi_{lj} d\log A_j + d\log w$, where the $d\log w$ is constant across all sectors and will drop out of the covariance. In this case, we can simplify (24) to just

$$\frac{d\lambda_i}{d\log A_j} = \sum_k \lambda_k(\theta^k - 1)\mathrm{Cov}_{\Omega_{kl}}(\Psi_{lj}, \Psi_{li}) \tag{25}$$

a rather nice expression!

Let's unpack (25). This says that to get the sensitivity of the sector-i share to productivity in sector j, we need to sum substitution effects arising in each sector k. These effects are proportional to λ_k, the Domar weight of each sector, and also to $\theta^k - 1$, which measures the extent to which substitution in sector k will change its spending shares. Finally, the key term is $\mathrm{Cov}_{\Omega_{kl}}(\Psi_{lj}, \Psi_{li})$. This gives the covariance among all input sectors l to k (weighted by their input shares) between their exposure to the shocked sector, j, and their exposure to the sector we're interested in, i.

Why does this covariance matter? The key is whether *the input sectors to k that benefit from the productivity shock in j tend to be the ones that also use more of input i.* If this is true, i.e. if the covariance is positive, then this will contribute to a higher share for input i (because we'll substitute more toward sectors that use it) if $\theta^k > 1$, and a lower share if $\theta^k < 1$.

One note that is obvious from (25): if every sector in the network is Cobb-Douglas, then all terms are zero and no shares ever change! This makes solving the model especially easy in this case.

Backing out quantities. The shares, or Domar weights, λ_i tell us the ratio of nominal sales in each sector i to the nominal value of total consumption, $\lambda_i = \frac{p_i y_i}{p_C C}$. We know from (20) what happens to relative prices like p_i/p_C, and Hulten's theorem (21) tells us what happens to C itself, so by combining this information with the change in λ_i, we can back out the change in the actual quantity y_i in each sector if desired.

An application: substitution between capital and labor

The framework and result above are quite general—even more general than we might realize at first! For instance, the model above only has a single primary factor, labor. But we can tweak it to accommodate other questions, such as how the economy substitutes between capital and labor.

Suppose that there are now two primary factors, capital K and labor L, in our input-output economy. Since the economy is efficient, in principle there is some implied (constant returns to scale) production function $C = F(K, L)$ that optimally uses capital and labor to produce the consumption good. We want to know the elasticity of substitution θ of this F. This gives the relationship between the capital-labor ratio and the factors' relative costs

$$\theta = -\frac{d\log(K/L)}{d\log(r/w)}$$

and equivalently the relationship between the ratio of capital and labor *shares* and the factors' relative costs

$$1 - \theta = \frac{d\log(rK/wL)}{d\log(r/w)} \tag{26}$$

θ is thus a very important macro parameter, since it tells us how the relative income shares of capital and labor will change as their relative costs change (e.g. as interest rates decline and make capital cheaper).

How do we squeeze capital into the input-output model we wrote down, where we had only labor? Remember, we're interested in what happens to shares when the ratio r/w of factor costs changes. We can hardwire a sector K that is "produced" entirely from labor with productivity $A_K = w/r$. Then, the Domar weight λ_K will be the capital share of final output, and we can implement (25) in this case with $i = j = K$ to obtain

$$\frac{d\lambda_K}{d\log A_K} = \sum_k \lambda_k(\theta^k - 1)\text{Var}_{\Omega_{kl}}(\Psi_{lK})$$

where the covariance now becomes a variance, because we're looking at the exposure of capital to itself. Here, we can think of Ψ_{lK} as the overall capital share of sector l, including both direct and indirect exposures.

To connect to expression (26), we note that the capital share is $\lambda_K = rK/C$ and what we normally call the "labor share" is the complement of that, so that $rK/wL = \lambda_K/(1-\lambda_K)$.[8] It follows that $d\log(rK/wL) = d\log\lambda_K - d\log(1-\lambda_K) = \frac{d\lambda_K}{\lambda_K} + \frac{d\lambda_K}{1-\lambda_K} = \frac{d\lambda_K}{\lambda_K(1-\lambda_K)}$, so that we can write

$$\begin{aligned} 1 - \theta &= \frac{d\log(rK/wL)}{d\log(r/w)} \\ &= \frac{d\lambda_K/(\lambda_K(1-\lambda_K))}{-d\log A_K} \\ &= \frac{1}{\lambda_K(1-\lambda_K)}\sum_k \lambda_k(1-\theta^k)\text{Var}_{\Omega_{kl}}(\Psi_{lK}) \end{aligned} \tag{27}$$

which is a simple formula to obtain the aggregate elasticity of substitution between capital and labor. We could make the notation a bit more evocative by letting $\alpha \equiv \lambda_K$ denote the aggregate capital share, and $\bar{\alpha}_l \equiv \Psi_{lK}$ denote the *total* capital share (including indirect exposures) of production in sector l, so that (27)

[8] To squeeze capital into the economy, we're saying that it's produced from labor, but that's just a mathematical fiction; we want to ignore that part of labor here.

becomes (note that we change ks and ls to is and js here to avoid confusion with capital and labor!):

$$1-\theta = \frac{1}{\alpha(1-\alpha)} \sum_i \lambda_i (1-\theta^i) \mathrm{Var}_{\Omega_{ij}}(\bar{\alpha}_j) \tag{28}$$

What is the intuition behind (28)? It's that overall, economy-wide substitution between capital and labor involves substitution in many different sectors. The contribution from any given sector k depends on its size λ_i, its elasticity θ^i, and the *variance* of total capital shares among its inputs j—which gives the extent to which it can meaningfully substitute between sectors of different capital intensity.

An example: two-tier economy. Let's consider a simple economy where the consumer chooses between different production sectors, and each production sector substitutes only between capital and labor.

Both the consumption sector and the production sectors will appear in the sum (28). Let's denote production sector i's capital share by α_i. Then the consumption sector term in the sum in (28) is just $(1-\theta^C)\mathrm{Var}_{\lambda_i}(\alpha_i)$ (where we use the fact that since consumption is the only source of production demand, $\Omega_{Ci} = \lambda_i$, and that the Domar weight of consumption is 1).

Further, for a production sector i, there are two inputs: capital, which has input share α_i and capital share 1, and labor, which has input share $1-\alpha_i$ and capital share 0. The variance is therefore one of a Bernoulli random variable with probability α_i of being 1; this is $\alpha_i(1-\alpha_i)$.

Combining these insights, we can specialize (28) to

$$1-\theta = \frac{1}{\alpha(1-\alpha)} \left((1-\theta^C)\mathrm{Var}_{\lambda_i}(\alpha_i) + \sum_{i=1}^{N} \lambda_i (1-\theta^i)\alpha_i(1-\alpha_i) \right) \tag{29}$$

This says that aggregate $1-\theta$ is a *weighted sum* of the consumer-level substitutability $1-\theta^C$ and the individual producer-level substitutabilities $1-\theta^i$.

Indeed, it's possible to show, although we won't work it out in detail here, that the sum of coefficients on $1-\theta^C$ and $1-\theta^i$ in (29) is 1. Hence, we can rearrange (29) to equivalently write

$$\theta = \frac{1}{\alpha(1-\alpha)} \left(\theta^C \mathrm{Var}_{\lambda_i}(\alpha_i) + \sum_{i=1}^{N} \lambda_i \theta^i \alpha_i (1-\alpha_i) \right) \tag{30}$$

This is essentially the result in the well-known and very nice Oberfield and Raval (2021), which applies it using micro-data in the manufacturing sector.

Just to illustrate the concepts, let's come up with a numerical example. Suppose that the aggregate capital share is $\alpha = 1/3$, but that this is actually the combination of one "capital-intensive" sector with $\alpha_1 = 2/3$, and one "pure-labor" sector with $\alpha_2 = 0$, each with a half share of production. Then $\mathrm{Var}_{\lambda_i}(\alpha_i)$ is $1/9$, $\lambda_1\alpha_1(1-\alpha_1) = 1/9$, $\lambda_2\alpha_2(1-\alpha_2) = 0$, and $\alpha(1-\alpha) = \frac{2}{9}$. It follows that (30) becomes

$$\theta = \frac{9}{2}\left(\frac{1}{9}\theta^C + \frac{1}{9}\theta^1\right) = \frac{1}{2}(\theta^C + \theta^1)$$

In short, in this case, the aggregate elasticity of substitution is actually an average of one-half the elasticity of substitution within the "capital-intensive" sector, and one-half the consumer's elasticity of substitution.

What's remarkable about this is that when we think about capital-labor substitutability, we usually think about firms choosing between capital and labor (and this is the focus of most of the empirical literature).

But here, an entirely non-production choice—how the consumer substitutes between sectors of different capital intensities—is *equally important*! If $\theta^C > \theta^1$, then the majority of aggregate capital-labor substitution will actually come from the consumer's decisions, rather than the firm's decisions.

This broadens our view of what we need to look at when we try to measure aggregate elasticities of substitution—it's not just direct substitution between factors in a firm that matter, but the indirect effects of substitution between inputs at other points in the production network, or even (as in this case) substitution at the consumer level.

References

Acemoglu, D., Carvalho, V.M., Ozdaglar, A. and Tahbaz-Salehi, A. (2012). The Network Origins of Aggregate Fluctuations. *Econometrica* 80(5):1977–2016.

Baqaee, D. and Rubbo, E. (2023). Micro Propagation and Macro Aggregation. *Annual Review of Economics* 15:91–123.

Baqaee, D.R. and Farhi, E. (2019). The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten's Theorem. *Econometrica* 87(4):1155–1203.

Baqaee, D.R. and Farhi, E. (2020). Productivity and Misallocation in General Equilibrium. *The Quarterly Journal of Economics* 135(1):105–163.

Basu, S. (1995). Intermediate goods and business cycles: Implications for productivity and welfare. *The American Economic Review* 85(3):512.

Bigio, S. and La'O, J. (2020). Distortions in Production Networks. *The Quarterly Journal of Economics* 135(4):2187–2253.

Liu, E. (2019). Industrial Policies in Production Networks. *The Quarterly Journal of Economics* 134(4):1883–1948.

Oberfield, E. and Raval, D. (2021). Micro Data and Macro Technology. *Econometrica* 89(2):703–732.