Graduate Labor Economics

Notes to Accompany Lectures 1 and 2: Static Labor Demand

Brendan M. Price*
Federal Reserve Board

Our first two lectures will lay out the neoclassical theory of static labor demand—which is basically just applied producer theory. Later in the course we'll incorporate frictions instead of assuming frictionless, perfectly competitive labor markets.

Why start here? A practical answer is that many of the empirical papers we'll be discussing in the next few lectures lean heavily on this body of theory. A more compelling answer is that many of the big questions in contemporary labor economics—such as why the college wage premium has grown dramatically over time, whether immigration depresses the wages of native-born workers, whether raising the minimum wage will lower employment, and whether robots will take all of our jobs—are ultimately questions about relative factor demands, and labor demand theory gives us an organizing framework for thinking coherently about factor demands. So let's dive in.[1]

1 One-factor model

We'll start with the simplest possible model, in which a profit-maximizing firm uses a single factor of production ("labor") to produce its output. We'll add capital in a little while.

- Production function: $Y = F(L)$, where F is increasing, smooth, and concave.
  - Usually we'll think of this as the production function for a particular firm, but sometimes it makes sense to regard this as an industry-level (or even aggregate) production function.
  - The one-factor model is often motivated as a short-run production function holding other factors (e.g., capital) fixed: $F(L) \equiv \widetilde{F}(L, \overline{K})$. The presumption here is that labor is easier to adjust in the short run than capital (which can take a long time to install). In practice, there are lots of costs to adjusting labor inputs, but the "short-run" view may be appropriate in some settings.
  - The underlying multi-factor production function may exhibit constant returns to scale, but—holding other factors fixed—the one-factor production function exhibits decreasing returns to scale. The fixed factors act as a "drag" on output as employment expands.
  - Example: coffee shop with a fixed number of espresso machines ($\overline{K}$), deciding how many baristas to employ today. Adding baristas boosts output (customers can be processed faster), but with diminishing returns: eventually all the espresso machines are being

 Email: brendan.m.price@frb.gov. These notes are adapted from UC Davis course ECN 250A, which I taught in Winter 2018 and Spring 2019. Any views or opinions expressed here are my own and do not necessarily represent the views or policies of the Board of Governors of the Federal Reserve System or its staff. This version: 2020/03/21.

[1] This lecture note draws on Hamermesh (1996) and Cahuc et al. (2014). The proof contained in the appendix is adapted from David Card's lecture notes.

used at maximum capacity. More machines can be installed in the long run, but not by 7 a.m. today.

- Assume perfect competition in both factor (input) and product (output) markets.
  - Firm is a price-taker and a wage-taker ("small" in both markets). Equivalently, the firm faces infinitely elastic product demand and infinitely elastic labor supply.
  - Example: a coffee shop in a dense urban area, where consumers can easily find equally good alternative coffee shops and where baristas can easily find equally remunerative alternative jobs.
  - By contrast, the coffee-shop *industry* presumably faces downward-sloping product demand (and maybe upward-sloping labor supply, if it employs a big enough share of low-skill workers).
- Profit maximization problem (PMP):

$$\pi(p, w) = \max_{L} \; pF(L) - wL \tag{1}$$

  - If there are other factors in the background (e.g., capital stock $\overline{K}$), actual profits are really $\pi(p, w) = \max_L \; pF(L) - wL - r\overline{K}$. But if the other factors are fixed, the non-labor factor payments are sunk costs and won't affect the optimal choice of labor, so we can ignore them.
- First-order condition (FOC), assuming an interior solution: $pF'(L^*) = w$.[2]
  - Firms set wages equal to the marginal revenue product of labor (MRPL or simply MPL). Standard logic: keep hiring baristas until you just break even on the marginal barista.
  - Labor demand slopes downward (why?). Together with a flat (firm-level) labor supply curve, this yields a unique solution for optimal labor demand.
  - We often speak of "derived demand for labor". The idea is that the firm chooses Q given product demand, then hires as much L as it needs to produce Q. (Indeed, one can rewrite the maximization problem to solve first for Q and then for L.)
- Comparative statics: How does L^* respond to (i) changes in product price? (ii) changes in productivity (e.g., a proportional increase in $F(L)$)? (iii) a sales tax? (iv) a payroll tax? (v) a tax on espresso machines, in the short run?
- It's easy to weaken our price-taking assumptions:
  - Product market power: if the firm faces downward-sloping demand for its output, as represented by the price function $p(Q)$, then it maximizes $p(F(L))F(L) - wL$, with corresponding FOC

$$p(F(L^*))F'(L^*)\left(1 + \frac{1}{\phi}\right) = w \tag{2}$$

    where $\phi \equiv \frac{\mathrm{d}\log Q}{\mathrm{d}\log p}$ is the price elasticity of product demand.[3] How does monopoly affect employment relative to the competitive benchmark? (Think of this through the lens of derived demand.)

[2] We can ensure an interior optimum by adding the Inada-type conditions $\lim_{L\to 0} F'(L) = \infty$ and $\lim_{L\to\infty} F'(L) = 0$. Since my goal in these lectures is to focus on the main results and insights, I'll be a little casual about technical conditions like these.

[3] I'm being casual with notation here: I really mean $\phi(F(L^*))$, since the elasticity of product demand may vary at different consumption levels. Going forward, I'll sometimes use abbreviated notation like this without comment.

- Factor market power ("monopsony"): suppose the firm faces upward-sloping labor supply but is still constrained to pay everyone the same wage. Then the firm maximizes $pF(L) - w(L)L$:

$$pF'(L^*) = \left(1 + \frac{1}{\varepsilon}\right) w(L^*) \tag{3}$$

  where $\varepsilon \equiv \frac{\mathrm{dlog}\, L}{\mathrm{dlog}\, w}$ is the (firm-level) elasticity of labor supply. In words: MRPL is set equal to the *marginal factor cost* (MFC). Under monopsony, MFC exceeds the wage because—to hire an extra worker—the firm also has to raise wages for the *inframarginal* workers. How do wages and employment compare to the competitive benchmark?

- The one-factor model is easily thought of as representing a spot market for homogeneous workers, but it can accommodate a limited degree of worker heterogeneity.
  - The labor input L is an aggregate of the *efficiency units* of labor supplied by all workers: $L \equiv \sum_i e_i h_i$, where h_i is the number of hours supplied by worker i and e_i is the worker's "efficiency" or "skill". Since workers differ along only a single productive dimension, this is often called a *single-index model* of labor demand.
  - Correspondingly, worker i's earnings are $y_i = w e_i h_i$, where w is the market price for an efficiency unit of labor. Earnings inequality reflects both hours worked and worker efficiency.
  - Key assumption: workers of different skill levels are perfect substitutes in production. Put differently, a barista who's twice as "efficient" as another will produce—and earn—twice as much per hour regardless of the mix of workers employed by the cafe. This is a strong assumption!
  - The same logic applies to hours: two four-hour shifts are equivalent to one eight-hour shift, so all shifts are perfect substitutes. Is this reasonable for a coffee shop? What about a consulting firm?

- Labor market equilibrium is pinned down by the existence of a market-level labor supply curve. If (as we typically assume) the supply curve is (weakly) upward sloping, we'll obtain equilibrium employment and wages (L^*, w^*). Labor supply, of course, is a vast topic in its own right.

2 Two-factor model

At the end of the day, the one-factor model can only say so much. So let's add a second factor, which I'll call "capital". (Later we'll regard the two factors as "low-skill" and "high-skill" workers, with capital implicitly in the background.) The two-factor model will let us answer questions like this: how do changes in the cost of capital affect labor demand? Some of the results below have n-factor generalizations, but we'll stick with two factors for now.

2.1 Constant returns to scale

- Production function: $Y = F(L, K)$, where F is increasing, smooth, and concave in both arguments ($F_L > 0$, $F_K > 0$, $F_{LL} < 0$, $F_{KK} < 0$, where subscripts denote partial derivatives).

- Common assumption: constant returns to scale (CRS). Many of the results below hold more generally, but some rely explicitly on CRS.

- Formally: $F(\lambda L, \lambda K) = \lambda F(L, K)$ for any $\lambda > 0$.[4]
- CRS is often motivated by an idea of "replicating production sites": if a cafe can produce Q lattes using L workers, K espresso machines, and one parcel of land, the owner can open a second identical cafe and produce $2Q$ lattes using twice the inputs.
- Her ability to do this may be limited by scarce factors (competent baristas, her own managerial oversight), but it's probably best to model that through upward-sloping factor supply curves rather than decreasing returns in the technology itself.
- One noteworthy implication of CRS: labor productivity $Y/L = F(1, K/L) \equiv \widetilde{F}(K/L)$ only depends on the capital-labor ratio.

- Mathematically, CRS is equivalent to a super-useful property: *homogeneity of degree one* (HOD[1]).
  - A function $f(\mathbf{x})$ is homogeneous of degree k (HOD[k]) $\iff$ $f(\lambda \mathbf{x}) = \lambda^k f(\mathbf{x})$ for any $\lambda > 0$.
  - Euler's homogeneous function theorem: if $f(\mathbf{x})$ is HOD[k], then $\mathbf{x} \cdot \nabla f(\mathbf{x}) = k f(\mathbf{x})$, where ∇ is the gradient operator. This implies $F(L, K) = F_L L + F_K K$, where F_L and F_K are partial derivatives. You'll want to remember this.
  - Euler's theorem has a nice corollary: if $f(\mathbf{x})$ is HOD[k], then its partial derivatives are HOD[$k-1$]. In particular: if $F(L, K)$ is HOD[1], then F_L and F_K are HOD[0]. Economically, this means that scaling all inputs proportionately doesn't affect marginal factor products: $F_L(\lambda L, \lambda K) = F_L(L, K)$ and $F_K(\lambda L, \lambda K) = F_K(L, K)$ for any $\lambda > 0$.
  - Applying Euler's theorem to F_L, we have $F_{LL} L + F_{LK} K = 0 \implies F_{LK} = -\frac{F_{LL} L}{K} > 0$ under our assumptions. So, concavity + CRS imply $F_{LK} > 0$. Note also that, by Young's theorem, we can switch the order of differentiation so that $F_{LK} = F_{KL}$.

2.2 The primal problem: profit maximization (PMP)

- Firm optimization:
$$\pi(p, w, r) = \max_{L,K} pF(L, K) - wL - rK \tag{4}$$
  where w is the wage and r is the rental rate of capital.
  - Once again assuming perfect competition in product and factor markets.
  - Solving this problem gives us the unconditional factor demands: $L^*(p, w, r), K^*(p, w, r)$.

- But it's usually easier to work with the dual problem of *cost minimization*, for two reasons:
  - It's conceptually useful to split the producer problem into a two-step procedure: choose the optimal factor mix for a given level of output; then choose the optimal output given the cost-minimizing factor mix. As we'll see below, these steps correspond to *substitution effects* and *scale effects*, respectively.
  - The PMP doesn't always have an interior solution. If p is too small relative to w and r, the firm sets $L^* = K^* = 0$ (i.e., it shuts down). If p is too large, the firm sets $L^* = K^* = \infty$, yielding infinite profits. So the unconditional factor demands are poorly behaved for some (p, w, r). More on this anon.

[4] We can define *decreasing* and *increasing* returns to scale in the same fashion. Suppose that $F(\lambda L, \lambda K) = \lambda^\theta F(L, K)$ for some $\theta > 0$. Then $0 < \theta < 1$ represents DRS, $\theta = 1$ represents CRS, and $\theta > 1$ represents IRS.

2.3 The dual problem: cost minimization (CMP)

- Firm optimization:
$$C(w,r,\overline{Y}) = \min_{L,K} wL + rK \text{ s.t. } F(L,K) \geq \overline{Y} \tag{5}$$
where $\overline{Y}$ is a "target" level of output.

- FOCs: $\mu F_L(\overline{L},\overline{K}) = w$, $\mu F_K(\overline{L},\overline{K}) = r$, where μ is the Lagrange multiplier.[5]
  - Conditional factor demands: $\overline{L}(w,r,\overline{Y}), \overline{K}(w,r,\overline{Y})$. These are well-defined.
  - Inputs are chosen so that the marginal product of each factor equals the factor price.
  - Dividing yields $\frac{F_L(\overline{L},\overline{K})}{F_K(\overline{L},\overline{K})} = \frac{w}{r}$: at the optimal factor mix, *the marginal rate of technical substitution (MRTS) equals the factor-price ratio.* The MRTS tells us how much capital would need to increase to offset a one-unit reduction in the labor input, holding output constant.
  - Graphically, the MRTS is the (absolute value of the) slope of the isoquant $\{(L,K) \mid F(L,K) = \overline{Y}\}$, and the factor-price ratio is the slope of the isocost curve $\{(L,K) \mid wL + rK = \overline{C}\}$. The firm chooses the factor mix such that the isocost curve is tangent to the isoquant.
  - A concave production function implies convex isoquants, ensuring a unique point of tangency.

- Cost function: $C(w,r,\overline{Y}) = w\overline{L}(w,r,\overline{Y}) + r\overline{K}(w,r,\overline{Y})$. Several useful properties (any graduate micro textbook should have the proofs):
  - $C(w,r,\overline{Y})$ is increasing in all of its arguments. It is HOD[1] in (w,r): for given output, doubling all factor costs doubles total cost.
  - $C(w,r,\overline{Y})$ is concave in (w,r), i.e. $C_{ww} \leq 0$ and $C_{rr} \leq 0$. Why? Because of substitution. Holding factor mix constant, total cost is linear in both factor prices—but in general, the firm can do better by substituting away from a factor whose price is rising.
  - Shephard's lemma:
$$\overline{L} = \frac{\partial}{\partial w} C(w,r,\overline{Y}) = C_w \quad \text{and} \quad \overline{K} = \frac{\partial}{\partial r} C(w,r,\overline{Y}) = C_r \tag{6}$$
This follows from the envelope theorem: if a firm is using its optimal factor mix, then factor-price changes have first-order effects on costs but only negligible second-order effects. If wages rise by \$1, then to a first approximation total costs rise by \$$L$.

- Under CRS, the cost function is HOD[1] in output: $C(w,r,\lambda\overline{Y}) = \lambda C(w,r,\overline{Y})$. Why?
  - Implies $C(w,r,\overline{Y}) = c(w,r)\overline{Y}$, where $c(w,r) \equiv C(w,r,1)$ is the *unit cost function.*
  - Dividing total factor demands by total output gives us the unit factor demands: $\overline{l}(w,r) \equiv \overline{L}(w,r,1) = \frac{1}{\overline{Y}}\overline{L}(w,r,\overline{Y})$, $\overline{k}(w,r) \equiv \overline{K}(w,r,1) = \frac{1}{\overline{Y}}\overline{K}(w,r,\overline{Y})$. These tell us what mix of labor and capital the firm uses to produce a single unit of output.
  - Note that $c(w,r) = w\overline{l} + r\overline{k}$. $c(w,r)$ inherits all the key properties of $C(w,r,\overline{Y})$.
  - Again using Shephard's lemma, $\overline{l} = c_w(w,r)$ and $\overline{k} = c_r(w,r)$.

[5] Under perfect competition and CRS, μ equals the marginal cost of production at the optimal factor mix, which in turn equals the product price (in equilibrium), so these conditions reduce to $pF_L = w$ and $pF_K = r$.

2.4 Substitution between factors

- Own-price effects: by Shephard's lemma, $\frac{\partial \bar{l}}{\partial w} = \frac{\partial c_w}{\partial w} = c_{ww} \leq 0$. Holding output constant, labor demand is (again) downward sloping. Symmetrically, $\frac{\partial \bar{k}}{\partial r} \leq 0$.

- Cross-price effects: how about $\frac{\partial \bar{l}}{\partial r}$ and $\frac{\partial \bar{k}}{\partial w}$?
  - Heuristically: in the cost-minimization problem, only relative factor prices matter. Increasing r decreases $\frac{w}{r}$, so that the firm uses more labor and less capital to produce each unit of output (picture the isocost and isoquant curves). Therefore $\frac{\partial \bar{l}}{\partial r} \geq 0$ and $\frac{\partial \bar{k}}{\partial w} \geq 0$.[6]

- Key parameter: elasticity of substitution $\sigma = \frac{\mathrm{dlog}(\overline{L}/\overline{K})}{\mathrm{dlog}(r/w)} = \frac{\mathrm{dlog}(\bar{l}/\bar{k})}{\mathrm{dlog}(r/w)} \geq 0$
  - The standard convention is to write r/w (not w/r) in the denominator to make this positive.
  - σ is the percent change in relative factor inputs chosen by an optimizing firm when confronted with a 1% change in factor costs, holding output constant. When σ is large, even a small change in relative factor prices induces a large change in firms' optimal factor mixes. When σ is small, firms have limited ability to substitute between factors, so there is less change in the factor mix.
  - The elasticity of substitution must be (weakly) positive. Why? This fact follows from the partial derivatives we've already signed.

- It can be shown that, under CRS, $\sigma = \frac{CC_{wr}}{C_w C_r} = \frac{c\, c_{wr}}{c_w c_r}$. The proof involves some laborious algebra (and not much economic insight), so I've relegated it to an appendix at the end of these notes. There is a dual result: CRS also implies that $\sigma = \frac{F_L F_K}{F F_{LK}}$.

2.5 The substitution effect

With all of this machinery in place, we can now state a set of classic results: under CRS production and perfect competition, the following "laws of demand" hold:

- $\overline{\eta}_{Lw} = -s_K \sigma \leq 0$, where $\overline{\eta}_{Lw} \equiv \frac{\mathrm{dlog}\, \bar{l}}{\mathrm{dlog}\, w}$ is the constant-output wage elasticity of labor demand and $s_K = \frac{r\bar{k}}{c}$ is the capital share of production costs.[7]
  - This is a *substitution effect*: holding output constant, firms substitute away from labor and towards capital when wages rise. Not surprisingly, the substitution effect is bigger when capital and labor are highly substitutable (σ large).

---

[6] Rigorously: because the unit cost function is HOD[1] in (w, r), we have $c(w, r) = rc(w/r, 1)$. Using Shephard's lemma and the fact that $c_{ww} \leq 0$,

$$\frac{\partial \bar{l}}{\partial r} = \frac{\partial}{\partial r}\frac{\partial}{\partial w}c(w,r) = \frac{\partial}{\partial r}\frac{\partial}{\partial w}rc\left(\frac{w}{r},1\right) = \frac{\partial}{\partial r}c_w\left(\frac{w}{r},1\right) = -\frac{w}{r^2}c_{ww}\left(\frac{w}{r},1\right) \geq 0 \tag{7}$$

Since $\frac{\partial \bar{l}}{\partial r} = c_{wr}$, this argument also shows that $c_{wr} \geq 0$, which I use in a proof later on. We also have $\frac{\partial \bar{l}}{\partial r} = c_{wr} = c_{rw} = \frac{\partial \bar{k}}{\partial w}$. This is a variant of "Slutsky symmetry" from consumer theory.

[7] The proof relies on Euler's theorem. Since $c(w, r)$ is HOD[1] in prices, its partial derivative c_w is HOD[0]. Using Euler's theorem, we can write $c_{ww}w + c_{wr}r = 0 \implies c_{ww} = -\frac{c_{wr}r}{w}$. Using the fact that $s_K\sigma = \frac{r}{\bar{l}}c_{wr}$ (shown in Footnote 8 below),

$$\overline{\eta}_{Lw} = \frac{w}{\bar{l}}\frac{d\bar{l}}{dw} = \frac{w}{\bar{l}}c_{ww} = \frac{w}{\bar{l}}\left(-\frac{c_{wr}r}{w}\right) = -\frac{r}{\bar{l}}c_{wr} = -s_K\sigma \tag{8}$$

- Why does a big capital share augment the substitution effect? Loosely, σ tells us the proportional change in $\overline{L}/\overline{K}$, but not which part ($\overline{L}$ or $\overline{K}$) is changing more. When s_K is large, $\overline{L}/\overline{K}$ will tend to change through a big proportional change in $\overline{L}$ and a smaller one in $\overline{K}$. Conversely, as $s_K \to 0$, we recover the one-factor model and there's no way to substitute at all, so $\overline{\eta}_{Lw} = 0$.

- $\overline{\eta}_{Lr} = s_K\sigma \geq 0$, where $\overline{\eta}_{Lr} \equiv \frac{\mathrm{dlog}\,\bar{l}}{\mathrm{dlog}\,r}$ is the constant-output elasticity of labor demand with respect to the rental price of capital.[8] This is the same substitution effect working in reverse (in fact $\overline{\eta}_{Lr} = -\overline{\eta}_{Lw}$).

But the *substitution effect* is only part of the story. To derive the *unconditional* response of labor demand to changes in factor prices, we need to account for the *scale effect*: an increase in *either* factor price reduces the firm's optimal scale, reducing the demand for *both* factors.

2.6 Unconditional demand and the scale effect

- Given its optimal factor mix (summarized by the cost function), the firm solves a second-stage problem:
$$\pi(p, w, r) = \max_{Y}\; pY - C(w, r, Y) \tag{10}$$

- The profit function has various properties analogous to those of the cost function: it is HOD[1] and convex in (p, w, r), and it satisfies Hotelling's lemma: $Y^* = \frac{\partial \pi}{\partial p}$, $L^* = -\frac{\partial \pi}{\partial w}$, and $K^* = -\frac{\partial \pi}{\partial r}$.

- Under CRS, the PMP becomes
$$\pi(p, w, r) = \max_{Y}\; pY - c(w, r)Y = \max_{Y}\; (p - c(w, r))Y \tag{11}$$

- Trouble in paradise! For a generic vector (p, w, r), either $p > c(w, r)$ so the firm sets $Y^* = \infty$, or $p < c(w, r)$, so the firm sets $Y^* = 0$. Can we ever hope to get an interior solution?
  - Yes! In equilibrium, free entry of firms will drive down p until each firm makes zero profit—which occurs precisely when $p = c(w, r)$. So the knife-edge case is the one that will hold in equilibrium.
  - Under CRS and perfect competition, *firm size is indeterminate*: to see this, observe that if $(L^*, K^*) \in \mathrm{argmax}_{L,K}\, F(L, K) - wL - rK$, then $(\lambda L^*, \lambda K^*) \in \mathrm{argmax}_{L,K}\, F(L, K) - wL - rK$ for any λ. Thus theory is silent as to whether we'll have one firm employing the whole market or every firm employing one worker.[9]
  - If there are many firms, Bertrand competition drives prices down to the breakeven level. If there's only one firm, the *threat* of entry compels it to charge the zero-profit price.

[8] The proof relies on repeated use of Shephard's lemma, plus $\sigma = \frac{c\, c_{wr}}{c_w c_r} \implies c_{wr} = \frac{c_w c_r}{c}\sigma$:

$$\overline{\eta}_{Lr} = \frac{r}{\bar{l}}\frac{d\bar{l}}{dr} = \frac{r}{\bar{l}}c_{wr} = \frac{r}{\bar{l}}\frac{c_w c_r}{c}\sigma = \frac{r}{\bar{l}}\frac{\bar{l}\,\bar{k}}{c}\sigma = \frac{r\bar{k}}{c}\sigma = s_k\sigma \tag{9}$$

[9] This is a useful abstraction but we'll break it later in the course. There are several ways to do so. One is to assume that firms face upward-sloping supply of at least one factor, such as entrepreneurial ability. Another is to assume that firms face downward-sloping product demand, perhaps because each firm produces its own "variety" and consumers like variety. Then we'll have a determinate set of firms (and firm sizes) operating in equilibrium.

- The combination of CRS with perfect competition has another important implication. By Euler's theorem,

$$pF(L,K) - wL - rK = pF(L,K) - pF_L L - pF_K K = p[F(L,K) - F_L L - F_K K] = 0, \quad (12)$$

so that an optimizing firm is just able to recoup its (factor) costs.

  - We're in a zero-profit world: revenue is divided between labor and capital according to the factor shares $s_L = \frac{pF_L L}{pY} = \frac{wL}{C}$ and $s_K = \frac{pF_K K}{pY} = \frac{rK}{C}$.
  - With DRS or imperfect competition, this breaks down: revenues = "pure profits" + factor costs.

- With the second stage in place, we can see where the scale effect comes from. Since $c(w,r)$ is increasing in both arguments, an increase in either w or r must be met by an increase in p if the product market is to clear. But since product demand is downward sloping, an increase in p implies a decrease in Y: output must fall, so that firms will scale back their demands for both labor and capital.

2.7 Putting it all together: the Hicks-Marshall laws of derived demand

- Under CRS, there are simple expressions for the *unconditional* elasticities of labor demand with respect to changes in factor prices. Since firm size (and hence scale effects) is indeterminate under CRS + perfect competition, it's best to think of these as industry-level relationships. The proofs are a bit more involved, so I state these laws without proof.

- (1) $\eta_{Lw} = -s_K\sigma + s_L\phi \leq 0$, where $\eta_{Lw} \equiv \frac{\mathrm{dlog}\, L}{\mathrm{dlog}\, w}$ is the *unconditional* wage elasticity of labor demand and, as before, ϕ is the price elasticity of product demand.

  - The first term is the substitution effect ($\overline{\eta}_{Lw} \leq 0$), which we saw earlier.
  - The second term is the scale effect, which is negative since $\phi \leq 0$. Notice that the scale effect depends on labor's share of production costs: the bigger s_L, the bigger the "cost shock" caused by an increase in wages and the more firms scale back their production. It's also proportional to the elasticity of product demand. If consumers are very price-sensitive, then an increase in wages (which in turn boosts the output price) leads to a drastic reduction in derived demand for labor.
  - The takeaway here is that an increase in wages causes firms to reduce employment for two reasons: substitution towards other factors and a scaling back of overall operations. Thus $|\eta_{Lw}| \geq |\overline{\eta}_{Lw}|$. As in the one-factor model, we have a downward-sloping labor demand curve.

- (2) $\eta_{Lr} = s_K\sigma + s_K\phi$, where $\eta_{LK} \equiv \frac{\mathrm{dlog}\, L}{\mathrm{dlog}\, r}$ is the *unconditional* elasticity of labor demand with respect to the rental price of capital.

  - The first term is the substitution effect ($\overline{\eta}_{LK} \geq 0$).
  - The second term is the scale effect, which is again negative. Now the scale effect depends on the *capital* share of costs (which determines the size of the "cost shock").
  - Given the offsetting scale and substitution effects, the net effect of r on L^* is ambiguous.

2.8 Many factors: notions of complementarity

- The PMP and CMP generalize to an arbitrary number of factors:

$$\pi(p, \mathbf{w}) = \max_{\mathbf{x}\in\mathbb{R}_n^+} pF(\mathbf{x}) - \mathbf{w}\cdot\mathbf{x} \tag{13}$$

$$C(\mathbf{w}, \overline{Y}) = \min_{\mathbf{x}\in\mathbb{R}_n^+} \mathbf{w}\cdot\mathbf{x} \quad \text{s.t. } F(\mathbf{x}) \geq \overline{Y} \tag{14}$$

  The multi-factor profit functions and cost functions have the same properties as before (e.g., Shephard's lemma still holds).

- But now substitution patterns can become much more complicated.

  - It remains true that $\frac{\partial \overline{x}_i}{\partial w_i} \leq 0$ for any factor: own-price constant-output effects and elasticities are negative. The same is also true for unconditional own-price effects. The logic is just as in the two-factor case: if wages go up, then (all else equal) scale and substitution effects both call for a reduction in labor demand.
  - What about cross-price effects? When w_i increases, all we can say is that $\overline{x}_j$ will increase for *at least one* other input $j \neq i$. Demand for some other factors $j' \neq i$ may actually decrease.
  - Example: in the near future, Amazon may begin using unmanned drones to deliver packages. Presumably drones and human delivery people will coexist for a while, so that the production function for delivery services is $F(D, g(L, V))$ where D = drones, L = labor, V = manned vehicles, and $g(\cdot)$ is a subaggregate of delivery people and the vehicles they drive. As drone prices fall, companies like Amazon are likely shift towards drones and away from humans and manned vehicles. By contrast, a fall in wages would encourage firms to substitute away from drones towards *both* workers and manned vehicles, since these are complements in production.

- A variety of terms are used to describe the degree of substitutability/complementarity between pairs of factors in a multi-factor world (Seidman, 1984).

  - Factors i and j are q-complements [respectively, substitutes] if, holding the quantities of all other factors fixed, an increase in the quantity of factor i increases [decreases] the marginal product of factor j. That is, q-complementarity means that $F_{ij} > 0$. (Since $F_{ij} = F_{ji}$, the definition is symmetric.) Lawyers and paralegals are q-complements: when a law firm hires more paralegals, the lawyers can devote more of their time to lucrative specialized tasks, letting the paralegals handle more routine matters. Driverless cars and cab drivers are q-substitutes: as more and more driverless cars appear on the road (holding the supply of cab drivers fixed), they will reduce cab-driver productivity since each cab driver will get fewer customers.
  - Factors i and j are p-substitutes [complements] if an increase in the price of factor i increases the conditional demand for factor j. Drones and trucks are p-substitutes; trucks and truck drivers are p-complements. Cross-price effects on conditional factor demand obey a symmtery condition, so the definition of p-substitutes/complements is also symmetric. (In the two-factor model, the factors are necessarily p-substitutes, but p-complements emerge once we add a third factor.)

- Factors i and j are *gross* substitutes [complements] if an increase in the price of factor i increases the *unconditional* demand for factor j. Not all p-substitutes are gross substitutes. Cahuc et al. (2014) give the example of aggregate labor and aggregate energy. Holding output fixed, an increase in the price of oil may induce firms to substitute away from energy towards labor, in which case labor and energy are p-substitutes. But rising oil prices will also have negative scale effects in many industries; if the scale effect dominates, labor and energy will be gross complements.

3 CES production functions

- A great deal of applied work uses a class of production functions characterized by a *constant elasticity of substitution* (CES). For two factors, a CES production function takes the form

$$F(L,K) = [(A_L L)^\rho + (A_K K)^\rho]^{\frac{1}{\rho}} \quad \text{with } \rho \in (-\infty, 1]$$

where (A_L, A_K) are factor-augmenting technology shifters.[10]

- Taking FOCs in the cost minimization problem, and noting that the Lagrange multiplier equals p in equilibrium, we obtain the usual condition that marginal factor products equal marginal costs:

$$\begin{aligned} pA_L^\rho L^{\rho-1} \left[(A_L L)^\rho + (A_K K)^\rho\right]^{\frac{1-\rho}{\rho}} &= w \\ pA_K^\rho K^{\rho-1} \left[(A_L L)^\rho + (A_K K)^\rho\right]^{\frac{1-\rho}{\rho}} &= r \end{aligned}$$

where I've dropped the overbars in $\overline{L}$ and $\overline{K}$ to unclutter the notation.

- This implies $\frac{w}{r} = \left(\frac{A_L}{A_K}\right)^\rho \left(\frac{L}{K}\right)^{\rho-1}$, or equivalently

$$\log \frac{L}{K} = \frac{\rho}{1-\rho} \log \frac{A_L}{A_K} + \frac{1}{1-\rho} \log \frac{r}{w} \tag{15}$$

- Elasticity of substitution: $\sigma \equiv \frac{\text{dlog}\, L/K}{\text{dlog}\, r/w} = \frac{1}{1-\rho}$. Note that, as the name "CES" implies, σ is a constant (not a function of L or K).[11] We can rewrite the relative skill demands as

$$\log \frac{L}{K} = (1-\sigma) \log \frac{A_L}{A_K} + \sigma \log \frac{r}{w} \tag{16}$$

[10] There are other ways of writing this. For example, we can incorporate a Hicks-neutral technology shifter A, so that

$$F(L,K) = A\left[(\widetilde{A}_L L)^\rho + (\widetilde{A}_K K)^\rho\right]^{\frac{1}{\rho}}.$$

But setting $A_L = A\widetilde{A}_L$ and $A_K = A\widetilde{A}_K$ recovers our original expression, so this is without loss of generality. Another modification is to allow for *factor-replacing* rather than *factor-augmenting* technological change: this is commonly done in the literature on job tasks but is beyond the scope of what we're doing here. See David Autor's lecture notes for MIT course 14.662.

[11] Given this result, CES production functions are often written as

$$F(L,K) = \left[(A_L L)^{\frac{\sigma-1}{\sigma}} + (A_K K)^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}} \quad \text{with } \sigma \geq 0.$$

- A nice (and elegant) feature of the CES is that it nests three economically interesting special cases:

  - $\sigma = 0$: Leontief (i.e., perfect complements). Factors must be used in fixed proportions: a trucking company might employ equal numbers of workers and trucks.[12]
  - $\sigma = \infty$: Linear technology (i.e., perfect substitutes). Firms adopt a "bang-bang" solution of using whichever factor offers a better value.
  - $\sigma = 1$: Cobb-Douglas. This is a knife-edge case, in which each factor receives a fixed share of total factor payments. (Show this!)

  The linear case is immediate: set $\rho = 1$ in the CES function, and you get the linear function $F(L, K) = A_L L + A_K K$. The Cobb-Douglas and Leontief cases are trickier, but they can be derived using L'Hôpital's rule. The derivations are a time-honored problem set ritual, but one without much economic intuition, so I will spare you.

- Now let's do some comparative statics. The CES yields a lot of important insights:

  - It's easy to show (show it!) that the MRPL (and hence also equilibrium wages) is increasing in K: hiring an extra worker boosts output by a greater amount when she has more capital to work alongside. This property is known as q-complementarity, which is important for thinking about how changes in skill supplies should affect the *absolute* wages of different skill groups.
  - Now consider factor-augmenting technological change—let's say an increase in A_K/A_L. Naïvely, we might presume that increasing the relative productivity of capital will surely depress wages relative to capital rentals. For $\sigma > 1$ this indeed holds, but for $\sigma = 1$ there's no change, and for $\sigma < 1$ we actually get the reverse implication. The intuition here is that, for sufficiently small σ, an increase in A_K effectively creates a "surplus" of capital, depressing the rate of return to each efficiency unit of capital to such an extent that the falling price per efficiency unit dominates the increasing efficiency of each machine.

- The CES production function looms large in the literature on wage inequality and skill-biased technical change. You'll see it again when we talk about Katz and Murphy (1992) and Card and Lemieux (2001).

Appendix: Proof that $\sigma = \frac{cc_{wr}}{c_w c_r}$

Proof. By Shephard's lemma, $\bar{l} = c_w$ and $\bar{k} = c_r$. Because $c_w(w, r)$ is HOD[0] (as in consumer theory, only relative prices matter!), we can express c_w—and therefore $\bar{l}$—as a function simply of the factor price ratio:

$$\bar{l} = c_w(w, r) = c_w(w/r, 1) \equiv g(w/r) \tag{17}$$

Likewise, we can write

$$\bar{k} = c_r(w, r) = c_r(w/r, 1) \equiv h(w/r) \tag{18}$$

[12] In the real world, there's often more opportunity to substitute than one might initially suspect. If wages are low enough, firms might assign two drivers to each truck, letting them trade off to minimize time lost due to sleep (or accidents due to fatigue). Or a company might use larger/smaller tractor-trailers, altering K for a given L.

for some functions $g(\cdot)$ and $h(\cdot)$. Differentiating these expressions with respect to w yields

$$\frac{\partial \bar{l}}{\partial w} = c_{ww} = \frac{1}{r} g'\left(\frac{w}{r}\right) \quad \text{and} \quad \frac{\partial \bar{k}}{\partial w} = c_{wr} = \frac{1}{r} h'\left(\frac{w}{r}\right) \tag{19}$$

Next, take the log of the factor ratio to obtain

$$\log \frac{\bar{l}}{\bar{k}} = \log g\left(\frac{w}{r}\right) - \log h\left(\frac{w}{r}\right) \tag{20}$$

Using properties of the logarithm, the elasticity of substitution is

$$\sigma = \frac{\mathrm{dlog}\, \bar{l}/\bar{k}}{\mathrm{dlog}\, r/w} = -\frac{w}{r}\frac{\mathrm{dlog}\, \bar{l}/\bar{k}}{d\left(\frac{w}{r}\right)} = -\frac{w}{r}\left(\frac{g'\left(\frac{w}{r}\right)}{g\left(\frac{w}{r}\right)} - \frac{h'\left(\frac{w}{r}\right)}{h\left(\frac{w}{r}\right)}\right) = -w\left(\frac{c_{ww}}{c_w} - \frac{c_{wr}}{c_r}\right) \tag{21}$$

where I've substituted in terms derived above. Finally, again using the fact that c_w is HOD[0], Euler's theorem implies that

$$c_{ww} w + c_{wr} r = 0 \implies c_{ww} = -\frac{c_{wr} r}{w} \tag{22}$$

Substituting this into Equation (21), we have

$$\sigma = -w\left(-\frac{c_{wr} r}{c_w w} - \frac{c_{wr}}{c_r}\right) = \frac{c_{wr} r}{c_w} + \frac{c_{wr} w}{c_r} = c_{wr}\frac{c_w w + c_r r}{c_w c_r} = \frac{c c_{wr}}{c_w c_r} \tag{23}$$

□

References

Cahuc, P., Carcillo, S., and Zylberberg, A. (2014). *Labor Economics*. MIT Press, Cambridge, MA. Second Edition.

Card, D. and Lemieux, T. (2001). Can Falling Supply Explain the Rising Return to College for Younger Men? A Cohort-Based Analysis. *Quarterly Journal of Economics*, 116(2):705–746.

Hamermesh, D. S. (1996). *Labor Demand*. Princeton University Press, Princeton, NJ.

Katz, L. F. and Murphy, K. M. (1992). Changes in Relative Wages, 1963-1987: Supply and Demand Factors. *Quarterly Journal of Economics*, 107(1):35–78.

Seidman, L. S. (1984). Complements and Substitutes: The Importance of Minding p's and q's. *Southern Economic Journal*, 56(1):183–190.