

MIT 14.662 Graduate Labor Economics II Spring 2009
Lecture Note 4: Market Structure, Organizational
Structure and Wage Structure

David Autor
MIT and NBER

April 16, 2009

1 MARKET STRUCTURE, ORGANIZATIONAL STRUCTURE AND WAGE STRUCTURE

There is much more important material to cover about aggregate production functions than we can possibly accommodate in 14.662. For example, where do aggregate production functions come from? We've been using them rather blithely this semester, but have not so far attempted to justify whether there is any justification for their use (other than their analytic tractability). You may ask yourself: how much do I want to know about aggregate production functions? The answer is: a lot. Because they are an incredibly important part of both labor and macroeconomics, and indeed stand at the intersection of the two. Some recommended papers are: Acemoglu (2003) "Labor and Capital Augmenting Technical Change" in the *JEEA* and Caselli and Coleman (2006) "The World Technology Frontier" in the *AER*, and Jones 2005 in the *QJE*.

2 WHY IS INFORMATION TECHNOLOGY ASSOCIATED WITH GREATER WORKER DISCRETION AND TEAMWORK (I.E., LESS 'TAYLORISM')?

Many authors have noted that problem solving teams and other 'lean production techniques' are often paired with information technologies, suggesting that newer Human Resource Practices (HRM) and IT are complements. Why would this be so? Some recent careful case studies are in my view among the richest material in this literature (which suggests to me that theory would benefit if theorists would leave their desks more often). Ann Bartel, Casey Ichniowski, and Kathryn Shaw have performed detailed case studies of changing production technology in three industries: steel, medical devices, and valves. The 2007 QJE paper on your syllabus is specifically about the valve industry.

My interpretation of their main findings is as follows:

1. Information technology has removed much of the mechanized, repetitive, rote components of production. Many repetitive tasks are now performed by machines that coordinate and monitor production. In the BIS valve study, computer-guided industrial lasers inspect completed valves for specification compliance with precision measured in microns. Previously, this time-consuming inspection step was done by hand. Similarly, the process of machine setup for new runs and the coordination of production as products move between machining phases have also been extensively computerized. This automation is feasible because computerized machining tools are far more flexible than older forms of capital. Much of the cumbersome reconfiguration of the assembly line needed for different

products and batches is now handled directly by the machinery. All of these production changes are *process improvements*.

2. One consequence of the increased flexibility of the *process* is a change in the set of *products* produced. The valve firms studied by BIS, which are those that have continued producing in the U.S., have increasingly moved out of commodity production and into ‘mass customization.’ They exploit the flexibility of the new capital to do shorter runs of more complex products. This in turn requires greater flexibility on the part of the workforce. But notice that commodity valve manufacturing, which is increasingly performed overseas, may not be undergoing similar changes in skill demand. [A 2005 paper by Chong Xiang of Purdue in the *ReStat* (“New Goods and Rising Skill Premium: An Empirical Investigation”) presents detailed evidence that *new goods* are increasingly skill intensive. This is valuable, representative evidence on the importance of product (rather than process) innovations for changing skill demands. I neglected to put this paper on the syllabus.]
3. Workers are increasingly required to use abstract reasoning to manage production. Whereas workers used to spend much of their time in contact with physical materials, much of the ‘work’ is now performed from the ‘control pulpit’ of highly automated factories where the key task is to monitor multiple assembly lines via large displays. The challenge is to be alert to problems as they arise and make subtle adjustments that improve efficiency, reduce error, enhance quality. In this setting, a good decision can have enormous value added and a bad decision can destroy large amounts of output.
4. Production work is frequently reorganized into teams where workers have responsibility for quality control and are required to solve problems and develop better ways to organize production. As observed by BIS, and Bresnahan, Bryjolffson and Hitt (2002), and Caroli and van Reenan (2001) problem solving teams and other ‘lean production techniques’ are often paired with information technologies, suggesting that these Human newer Resource Practices (HRM) and IT are complements. (More on why this might be below...)

These observations are drawn from the manufacturing sector, which provides a declining share of employment in almost all advanced countries. How relevant are they to the service

sector? Similar case studies for the service sector are not in abundance. One by Autor, Levy, Murnane (2002, ILRR) on the computerization of the back office check-processing function of a bank provides descriptive evidence that is consistent with the observations in Bartel, Ichniowski and Shaw (though the examples are not nearly as dramatic as those in manufacturing). In the bank studied by ALM, the introduction of Optical Character Recognition and electronic imaging of paper checks reduced the rote component of check processing for the ‘exceptions processors,’ reducing time spent on routine ‘paper chase’ activities and increasing time spent on problem solving and account management. Notably, for the ‘check preparers’ who perform the physical handling of the checks, there was little net change in skill demands—but there was a dramatic reduction in employment.

BIS specifically investigate four hypotheses:

1. New IT-enhanced machines improve production process efficiency. Setup time, run time, and inspection time fall after new IT-enhanced equipment in these stages is adopted.
2. New IT promotes product customization and innovation. New 3D-CAD technologies should directly affect the plant’s capabilities of designing more customized valves, while other technologies that reduce setup time would also promote customization.
3. IT adoption may increase (or decrease) skill demand.
4. IT adoption may require new HRM practices.

The theoretical foundation for these hypotheses is BIS’ observation that IT reduces setup time, which is otherwise particularly costly for customized, small batch jobs. This cost reduction differentially reduces the price of producing customized relative to commodity products. What happens when IT prices fall:

1. Firms purchase more IT
2. production efficiency rises, setup time, run time and inspectime time fall.
3. Firms make a strategic move towards producing more customized products. This is due to fall in setup times.

4. Optimal skill demand changes, but the direction is ambiguous. If setup is the most skill-intensive task (as seems likely), when setup time falls, skill demand falls. But the move to more setup-intensive products exerts a countervailing force. Third, IT-based machinery increasingly displaces routine tasks, thereby raising skill content of labor input. This goes in the direction of increasing skill requirements.
5. Finally, new HRM practices may complement use of higher skill levels or new machinery, though a microfoundation for this idea is not given.

3 TECHNOLOGY AND THE ORGANIZATION OF PRODUCTION

A 2007 *JPE* paper by Wouter Dessein and Tanos Santos (“Adaptive Organizations”) offers an ingenious hypothesis on the subtle linkages between information (or communication) technology and work organization.

Quoting Dessein-Santos (p. 17):

Adam Smith’s famous observation that the “the division of labor is limited by the extent of the market,” has been challenged by both the management literature as well as economists such as Becker and Murphy (1992). These two strands of the literature have emphasized that perhaps, more importantly, specialization is mainly constrained by the need to coordinate activities. In particular, a straightforward comparative static prediction in Becker and Murphy (1992 [QJE]) is that as coordination costs increase, one should see less specialization.”

The ‘Smithian’ view would suggest that the expansion of markets via population growth, international trade, etc. would increase specialization by enlarging the extent of the market. The Becker-Murphy view would also suggest that the falling price of information technology would increase specialization by lowering communications/coordination costs. Yet, we seem to see a world where jobs are becoming broader and workers are increasingly asked to exercise discretion. What is going on?

One possibility is that casual empiricism is simply wrong: work *is* becoming more compartmentalized, perhaps not in advanced countries but in the world at large (though why *not* in advanced countries?). A second possibility is that the theory is wrong. The Dessein-Santos’ paper suggests the latter.

3.1 ADAPTATION VERSUS COORDINATION

Dessein-Santos consider a model of organizational design in which there are a number of tasks, n , to be performed. Performance of each task is optimized through two actions: (1) adaptation and (2) coordination.

Adaptation is the adjustment that should be made to a particular task in response to current conditions (which are assumed not entirely forecastable). For example, when cooking a Japanese dinner, the optimal quantity of water to add to the rice cooker depends in part on the current humidity and temperature as well as the season when the rice was harvested relative to the cooking date.

Coordination is the adjustment of other complementary tasks to the adaptations made on the primary task. For example, adding water to the rice cooker might mean that the meal will be delayed, in which case the vegetables should be sauteed several minutes later so they do not turn to mush.

One way to accomplish both objectives simultaneously (adaptation *and* coordination) is to assign one person to perform both tasks (*bundle*). The authors assume that *coordination* is costless (or very low cost) if one person performs all tasks.

But the downside to bundling of tasks is that we lose the gains to *specialization*. Productive efficiency at each task is probably greatest if one chef cooks the rice and the other the vegetables (picture a large restaurant). The authors therefore assume that the more tasks done by one worker, the lower her productivity at each.

The fact that there are gains to specialization suggests that there is a value of *unbundling* tasks. But in the cooking example, it's easy to see the problem that may arise. If tasks are highly specialized, there is a significant risk of coordination failures—information may be lost in coordinating among the workers performing complementary tasks. If the rice chef adjusts the water perfectly but neglects to inform the vegetable chef (or the information is lost with some probability) then we'll end up with better rice but perhaps a worse meal.

There is therefore a trade-off between allowing discretion to maximize adaptation (which comes at the cost of greater coordination failures) and reducing discretion to improve coordination (which comes at the cost of reduced adaptation). One can also bundle more tasks into a single job, which improves both adaptation and coordination. But this reduces gains from specialization.

Dessein-Santos map out two general overarching organizational strategies:

1. ‘Ex ante coordination’—The goal here is to reduce coordination failures by *limiting discretion*. So, the rice chef does not adapt the recipe by much to current conditions. This means the rice is perhaps of lower quality but it is well timed with the vegetables. In this case, it is clear that a firm will also want to subdivide tasks finely to take advantage of gains to specialization since there is little need to coordinate (i.e., the coordination is built-in since there is no discretion). This the so-called ‘rigid’ organization. In this type of organization, there will be little investment in communication (meetings, teams, etc.) because such communication is not needed. One could call this production mode ‘Fordism.’
2. ‘Ex post coordination’—Here, the firm allows workers discretion to adapt to local conditions but then must rely on high quality *ex post* communications to coordinate among adapted tasks. The firm will therefore want to invest in high quality communications to make ex post coordination work. Notice also that firms may want to use *broad* job assignments in this model. The reason is that if communication is costly but workers can still coordinate costlessly *within* their groups of tasks, it is valuable to both have high quality communications and to reduce the degree of specialization so that these communications are not too costly in practice (e.g., not too many meetings). One might label this organizational model as ‘flexible production.’

Now consider what happens when communications costs fall (i.e., due to computerization). Firms will find ex post coordination cheaper (i.e., communications among adapted tasks). This may increase the value of giving workers discretion to adapt their tasks (less coordination failure ex post). But if so, it will also increase the value of bundling tasks so that workers can costlessly coordinate within tasks they are adapting. Thus, broad task assignments and intensive ‘horizontal communications’ are complements under some circumstances.

Josh Angrist has often observed that the principle effect of computerization on his work activities is that he now does his own typing (whereas secretary used to do this for him). According to Dessein-Santos, this is efficient!

3.2 MODEL BASICS: COORDINATION, SPECIALIZATION, AND ADAPTATION

3.2.1 ADAPTATION

Task i consists of undertaking a primary action a^{ii} , whose effectiveness depends on how well it is adapted to the local environment. Adaptation calls for use of *local information*, which pertains only to this task and can only be observed by the worker assigned to the task.

The local information is a random variable θ^i with mean $\hat{\theta}^i$ and common variance σ_θ^2 . To achieve perfect adaptation, the primary action a^{ii} should be equal to θ^i . The realization of θ^i is independent across tasks.

3.2.2 COORDINATION

To ensure that task i is coordinated with all tasks $j \neq i$, the employee in charge of task i must perform a sequence of $n - 1$ actions $\{a^{i1}, a^{i2}, \dots, a^{in}\}$ that are complementary to the primary actions of tasks $j \neq i$. To achieve perfect coordination between tasks i and j , action a^{ij} of task i should be set equal to the primary action a^{jj} .

So, if the organization consists of two tasks, then adaptation and coordination losses amount to:

$$\phi \left[(a^{11} - \theta^1)^2 + (a^{22} - \theta^2)^2 \right] + \beta \left[(a^{12} - a^{22})^2 + (a^{21} - a^{11})^2 \right],$$

where the parameters ϕ and β determine the importance of adaptation and coordination respectively. Here, a^{11} and a^{22} are the primary tasks, and actions a^{12} and a^{21} are the subordinate tasks for each primary action respectively.

3.2.3 TASK ASSIGNMENT

Let $T(i)$ equal the set of tasks bundled with i . For simplicity, assume that organizational design is symmetric, so that all workers in an organization have an identical number of tasks t assigned to them, so that $T(i) = n/t$ for all workers.

Task variety is costly in the ‘Smithian’ sense in that forgoes the gains from specialization. Specifically, the labor cost of carrying out task i , denoted by $h(t, \alpha)$, is increasing in the level of task bundling t and in α . Thus $h(\cdot)$ is a per-task cost, with:

$$h(\bar{t}) - h(\underline{t}) \geq 0, \text{ and } h_\alpha(\bar{t}) - h_\alpha(\underline{t}) \geq 0, \text{ and } \bar{t} > \underline{t}.$$

Bundling tasks therefore raises per unit costs of execution.

3.2.4 COMMUNICATION

To improve coordination, workers can communicate the realization of the local information θ^i prior to the actual implementation of the actions. If tasks i and j are assigned to different employees, then with probability p the value of θ^i will be communicated successfully and with probability $(1 - p)$, it will be pure noise. Thus, p is a measure of the quality of the communication channel. It is assumed that agents who receive communication know whether communication was successful or pure noise. Agents sending this information do not know whether their communication was successful.

3.2.5 PROFITS

Profits of an organization are given by

$$- \sum_{i=1}^n C^i (a^{1i}, a^{2i}, \dots, a^{ni}, t | \theta^i),$$

where

$$C^i (a^{1i}, a^{2i}, \dots, a^{ni}, t | \theta^i) = \phi (a^{ii} - \theta^i)^2 + \sum_{j \notin T(i)} \beta (a^{ji} - a^{ii})^2 + h(t, \alpha)$$

3.2.6 TIMING

1. Organization determines the number of task per agent, t .
2. Local circumstances θ^i for $i = 1, 2, \dots, n$ are realized and observed by employee(s) in charge of task i .
3. Workers communicate the realizations of local information, and each is successful with probability p .
4. For each task i , the employee in charge of that task chooses actions a^{ij} to maximize the objective function given his information.

3.2.7 OPTIMAL CHOICE OF ACTIONS

One can show that employees optimally choose the following primary actions:

$$a^{ii}(\tau) = \hat{\theta}^i + \left[\frac{\phi}{\phi + \beta(n-t)(1-p)} \right] (\theta^i - \hat{\theta}^i).$$

Notice that the degree of adaption is increasing in ϕ , decreasing in β , increasing in the quality of communication p , and increasing in task bundling t .

Complementary actions are chosen as:

$$a^{ji}(t) = \begin{cases} a^{ii(t)} & \text{when task } j \text{ learns } \theta^i \\ \hat{\theta}^i & \text{when task } j \text{ does not learn } \theta^i \end{cases} .$$

The covariance between local circumstances and the primary action are:

$$\sigma_{a\theta}(t) = Cov[a^{ii}(t), \theta^i] = \left[\frac{\phi}{\phi + \beta(n-t)(1-p)} \right] \sigma_{\theta}^2.$$

Thus, $\sigma_{a\theta}(t)$ characterizes how strictly employees adhere to the ex ante strategy ($\hat{\theta}^i$) versus tailoring their actions to local circumstances. Notice that $\sigma_{a\theta}(t)$ is increasing in the *variability* of local circumstances (in addition to the other comparative statics above).

A key observation is that $\sigma_{a\theta}(t)$ is increasing in task bundling:

$$\sigma_{a\theta}(\bar{t}) > \sigma_{a\theta}(\underline{t})$$

Given t , expected profits are:

$$\Pi(t) = n\phi[\sigma_{a\theta}(t) - \sigma_{\theta}^2] - nh(t, \alpha),$$

$$t^* = \arg \max_{t \in T} \Pi(t),$$

as a function of $\phi, \alpha, \sigma_{\theta}^2, \beta$ and p .

3.3 SOME RESULTS

3.3.1 SPECIALIZATION

Task specialization is decreasing in the importance of adaption, ϕ , and the variance of local circumstances σ_{θ}^2 , but increasing in the returns to specialization, α .

3.3.2 THE RELATIONSHIP BETWEEN COORDINATION COSTS AND SPECIALIZATION

The Becker-Murphy 1992 *QJE* article would suggest that as coordination costs fall, specialization increases. In the current paper, coordination costs operate through β (holding communication technology fixed). When β is higher, tasks are more interdependent, so coordination costs are higher.

How does a rise in β affect task bundling? Becker-Murphy suggests that it would reduce it. In this model, the answer depends on the importance of adaptation ϕ . An increase in task interdependence affects two margins:

1. For given ϕ , a rise in β makes it more important to improve coordination, which leads to increased bundling.
2. But a rise in β may also spur the organization to reduce employee flexibility and become less adaptive. This reduces the need for task bundling. If β is very large, an employer may eliminate all flexibility, in which case it is optimal to fully specialize $t^* = 1$. Conversely, if β is very small so that tasks are virtually independent, then coordination can be ignored and so it is also optimal to specialize.

Accordingly, the relationship between t^* and β is non-monotone. For $h(n, \alpha) < \infty$ and given α , there exists a unique $\hat{\phi}$ such that:

$$\lim_{\beta \rightarrow \infty} t^* = \begin{cases} n & \text{if } \phi > \hat{\phi} \\ 1 & \text{if } \phi < \hat{\phi} \end{cases}$$

Moreover, when $\phi < \hat{\phi}$, t^* is non-monotone in β , increasing and then decreasing. Initially, a rise in β leads to more bundling to improve coordination. As β gets larger, firms simply reduce adaption and make jobs very narrow. This resolves coordination problems, though also makes firms less adaptive.

3.3.3 QUALITY OF COMMUNICATION

Intuition would suggest that one would see more specialization as communication costs $(1 - p)$ fall since coordination gets easier. But organizations may also take advantage of improved communications to become more adaptive. In this case, by bundling, the organization increases its adaptiveness to local environmental variables. This increased adaptation is complemented by improved communication because it becomes more important to coordinate the other tasks when the primary action is allowed to vary more.

So, there is again a trade-off:

1. For given level of employee flexibility, $\sigma_{a\theta}(t)$, an improvement in communications makes it easier to coordinate, reducing the need for bundling.

2. But as coordination through communication improves, firms will want to increase employee adaptability. This favors additional task bundling.

When communication channels are very poor, task bundling often increases as communication technology improves. Firms go from being extremely inflexible to somewhat more adaptive. In this range, communications and bundling are complements.

But when communication technology gets sufficiently good, bundling looks less and less attractive since it forgoes ‘Smithian’ gains from specialization but saves little in coordination errors (since the quality of communications is good). In this range, communications and bundling are substitutes.

3.3.4 OTHER COMPARATIVE STATICS

There are many interesting extensions in the paper, which I leave for your further reading pleasure. An important choice variable that the paper studies is the optimal level of p when improvements in p can be purchased at a positive cost.

3.4 ROSEN: THE ECONOMICS OF SUPERSTARS

The 1981 *AER* article by Sherwin Rosen, “The Economics of Superstars,” is often cited as a prescient explanation for the rise in returns to skills that was to occur in the years immediately following its publication. The Superstars hypothesis probably may or may not help to explain the phenomenon of rising inequality. Nevertheless, the paper offers a fascinating insight that has considerable currency as an explanation for why wages of CEOs, entertainers and athletes are incomparably higher than for other occupations.

Certain services appear to be characterized by a demand structure that places considerable weight on quality versus quantity. For example, a patient would rather have one really good heart operation than two mediocre ones; an opera fan would rather see one Pavoratti concert than ten mediocre opera performances; no amount of Applebee’s food amounts to a meal at Oleana. Hence, if there is widespread agreement on *who* is talented among the providers of such services, one can imagine that talented service providers could earn considerable rents.

There is also a natural check on the possibility that one ‘talent’ will control an entire market: congestion. A single heart surgeon cannot perform all of the surgeries; a great cook can only

make so many dinners in an evening. It may be these congestion externalities that prevent single talents from controlling entire markets.

Now, imagine a technological improvement that reduces congestion, allowing holders of talent to more efficiently serve larger markets. Examples: television, recording technology, the Internet, etc. These types of technological advances may increase the rents accruing to the most talented individuals by increasing their market share (possibly to a global scale). Some professions in which this is likely to be important: all performing artists and entertainers; managers of large corporations (CEOs); athletes; fund managers; possibly some academics. Is it actually hard to make a very long list.

The Rosen model almost surely helps to explain why athletes and movie-stars earn such enormous salaries, and why these salaries have risen as mass communications and entertainment have improved. The 2008 *AER* paper by Terviö explores this idea rigorously.

4 WHAT ARE CEO'S WORTH (AND WHY)?

CEO pay is a divisive topic on which there has been an enormous amount of research. While there is room for debate about whether this topic receives more attention than it merits, it still offers 'teachable moments.' One such moment is the 2008 *AER* by Terviö. This paper studies CEO pay and productivity through the lens of an assignment model. Assignment models were introduced by Nobel Laureate Jan Tinbergen in the 1950s, and they were 'popularized' (okay, made less obscure) by Sattinger's widely cited 1993 *JEL* survey. The defining feature of assignment models is the presence of indivisibilities among factors of production. In a situation where the amounts of two matching factors cannot be shifted across different units of production, factors are *not* paid their marginal products in the standard sense. As we will see, it's possible (indeed) likely for the wage paid to a given worker to rise or fall due to changes in the distribution of ability of other workers *without* any change in the productivity of that worker or in the value of his output.

Assignment models have defied widespread use, perhaps in part because they are typically rather non-intuitive and analytically difficult. The paper by Terviö presents a beautiful exposition of the intellectual foundation of assignment models by reformulating the basic assignment model using distributional ranks rather than actual distribution functions. This formulation is completely natural because assignment models are inherently ordinal rather than cardinal. The

paper’s substantive conclusions on CEO pay and its relationship to the distribution of CEO talent, are both insightful and intuitive.

4.1 BASIC SETUP

There are three assumptions that simplify the model:

1. One-dimensional quality of factors
2. Continuity of quality distributions
3. Complementarity between the qualities of factors

The first two assumptions are simply for analytical convenience. The third assumption is essential to the conceptual framework. The paper further assumes, risk neutrality and symmetric information (which implies an absence of agency problems).

Worker quality is denoted as a . Firm quality is denoted by b , which will be referred to as firm size. There is a unit mass of workers and firms with continuous, finite support without gaps. (The unit mass should be thought of as a normalization of all units that are active in equilibrium.)

The production function is continuous and strictly increasing in both its arguments. The key assumption is that this function has a positive cross-partial between a and b , so these inputs are complements. Accordingly, efficiency requires positive assortative matching.

As is explained in the paper, one can write the production function *without loss of generality* as:

$$Y(a, b) = a \cdot b.$$

The notation in the paper is actually an important contribution. Order abilities by quantile so that $a[i]$ is the ability of the i^{th} quantile individual, and $a'[i] > 0$. Denoting the distribution function by F_a , the profile of a is defined by

$$a[i] = a \leftrightarrow F_a(a) = i.$$

4.2 EQUILIBRIUM CONDITIONS

There are two types of constraints on the problem: a sorting constraint, which requires that no worker/firm pair wish to rematch; and an incentive compatibility constraint, which requires

that all workers and firms earn at least their outside income:

$$\begin{aligned}
Y(a[i], b[i]) - w[i] &\geq Y(a[i], b[j]) - w_j \quad \forall i, j \in [0, 1] && SC(i, j) \\
Y(a[i], b[i]) - w[i] &\geq \pi^0 \quad \forall i \in [0, 1] && PC b[i] \\
w[i] &> w^0 \quad \forall i \in [0, 1] && PC a[i]
\end{aligned}$$

It is assumed that π^0 and w^0 are the same for all units, though it would be sufficient to assume that outside opportunities increase slower along the profile than do equilibrium incomes. The lowest active pair breaks even:

$$Y(a[0], b[0]) = \pi^0 + w^0.$$

While there are in theory $2n!$ sorting constraints, where n is the number of units (if there were a finite number, which formally there are not), most constraints are redundant since for $i \geq j \geq k$, $SC(i, j) + SC(j, k)$ implies $SC(i, k)$. The binding constraints are therefore: (1) the marginal sorting constraints that keep firms from wanting to hire the next best individual; and (2) the participation constraints of the lowest types.

Ordering the binding sorting constraints, one can write:

$$\frac{Y(a[i], b[i]) - Y(a[i - \varepsilon], b[i])}{\varepsilon} \geq \frac{w[i] - w[i - \varepsilon]}{\varepsilon}.$$

This equation becomes an equality as $\varepsilon \rightarrow 0$, and yields the slope of the wage profile (since the limit of the above expression is the partial derivative):

$$w'[i] = Y_a(a[i], b[i]) a'[i],$$

where Y_a is the partial derivative. This gives us the slope. To get the full profile, we need to integrate over the production function and add in the binding participation constraint:

$$w[i] = w^0 + \int_0^i Y_a(a[j], b[j]) a'[j] d[j].$$

Similarly, the profile of profits must satisfy:

$$\begin{aligned}
\pi'[i] &= Y_b(a[i], b[i]) b'[i] \\
\pi[i] &= \pi^0 + \int_0^i Y_b(a[j], b[j]) b'[j] d[j].
\end{aligned}$$

These conditions also imply that $y = \pi + w$ at each firm.

This is an extremely tight set of constraints on the problem. It implies that the wages and profits of each factor depend at quantile i depend on the full profile of factors from quantile 0 to $i - \varepsilon$ (but not on the profile above i).

A number of useful observations follow (and you should read the paper, which is chock full of astute insights, and even more so in the footnotes):

1. There is no bargaining in this model because of the continuity of the distribution of both factors. If there was a jump at some point in the profile of one factor, then all of the surplus would go to the factor with the jump because the other side is perfectly competitive.
2. Payments to factors are only affected by the quality of those *below* them in the ranking. This is because the binding constraint on each worker or firm is that the quality/price of the worker/firm just below it in the distribution.
3. The unobserved productivity characteristics a and b are essentially ordinal. Any increasing transformation of the scale of measurement for a factor's quality, combined with the inverse change in the functional form of the production function changes nothing of substance in the model. In this sense $Y(a[i], b[i]) = a[i] \times b[i]$ is a general functional form (so long as we are assuming complementarity).
4. Quoting Tervio, "It would be incorrect to say that factors earn their marginal productivity by the usual definition of marginal productivity, because the increase in output if the individual of ability $a[i]$ were to increase in ability is proportional to $b[i]$. But if she were to increase in ability, then, in equilibrium, she would also move up in the ranking and be matched with a higher b — and other individuals would have to move down and experience a decrease in productivity. This means that individuals in fact do receive their marginal product, once the margin is defined correctly. As ability cannot conceivably be extracted from one individual and poured into another, the relevant margin here is whether an individual will participate in the industry or not—and if not, then the effect of the resulting rearrangement of remaining individuals is part of the marginal product."

4.3 COMPARATIVE STATICS

Consider a change by which the production function Y is multiplied by some constant G with no change in distributions of a and b . Uniform productivity growth is mathematically equivalent to changing the units of measurement for output. This gives rise to a useful lemma:

Scaling lemma: If $Y_t(a, b) = GY(a, b)$, $w_t^0 = Gw^0$ and $\pi_t^0 = G\pi$, then $w_t[i] = Gw_t[i]$ and $\pi_t[i] = G\pi_t[i]$ for all $i \in [0, 1]$.

Figure 1 summarizes much of the intuition of the model. In the strictly multiplicative case, the output accruing from matching a worker of ability a and a firm of ability b is $a \times b$, which is simply a rectangle in a Cartesian graph. This graph is drawn as follows. Let $a = \varphi(b)$, defined by $a[F_b(b) = \{(a, b) \text{ st. } F_a(a) = F_b(b)\}]$ with slope

$$\varphi'(b) = a'[F_b(b)] f_b(b) = \left. \frac{a'[i]}{b'[i]} \right|_{i=F_b(b)}.$$

$\varphi(b)$ is strictly increasing in b , and the slope is given by the relative steepness of a and b at each quantile i .

Be sure to understand Figure 1, since essentially all of the intuition of the assignment model is contained in this figure.

4.4 APPLYING THIS MODEL TO CEOs

There are numerous conceptual and practical issues to confront before bringing this model to the data. I will discuss these assumptions and results only briefly.

The issues:

1. The surplus created by the CEO-firm interaction is unobserved. The market value of the firm is affected by the current CEO and by expectations of future productivity (which also depend on future CEOs)—thus firm size, an outcome, cannot be treated as the b variable.
2. A second issue is that part of the market value of the firm will surely reflect the value of capital that can be readily transferred/resold among firms. This capital is *not* indivisible and so is not part of the surplus represented in $Y(\cdot)$.

3. A third issue is that the current market value of a firm depends on the quality of the current CEO and the quality of past and (in expectation) future CEOs. Thus, one cannot necessarily infer the current CEO's quality purely from contemporaneous data.
4. Productivity tends to grow over time, and the expectation of growth further affects current market value.
5. The distribution of CEO ability and latent (exogenously determined) firm size (not market value) can change over time.
6. Outside options may shift.

Thus, the subsequent analysis rests on a number of strong assumptions that necessarily make the exercise speculative (though perhaps still informative if one believes the foundational assumptions of the assignment model):

1. The distribution of $a [i]$ and $b [i]$ are time invariant.
2. Productivity grows deterministically at rate g at all firms (that's why the scaling lemma is needed)
3. The value of outside options grow at rate g . Tervio uses values between 0.2 and 0.025.
4. The discount rate is constant. Tervio uses values between 0.08 and 0.05.
5. The impact of past *and future* CEO quality on current firm performance decays at a constant rate $\alpha_{\tau+1} = \alpha_{\tau}\lambda / (1 + \lambda)$. λ determines the decay rate. With $\lambda \rightarrow \infty$, only the current CEO affects contemporaneous earnings. Tervio uses values between ∞ and 0.1
6. Since adjustable capital must earn the market rate of return, it must be subtracted from Y . To determine the contribution of adjustable capital, Tervio assume's that the gross surplus has constant elasticity θ with respect to adjustable capital. Tervio sets the share of adjustable capital in Y at values between 0 and 0.8.

These assumptions are sufficient to pin down the contribution of firm size and ability to output up to an additive constant (one constant each for the two factors). The reason the constant is indeterminate is that one cannot infer the contribution of firm size and ability to

surplus at the smallest firm in the sample since this depends on the full unobserved distributions of size and ability below this firm. However, these constants drop out when estimating counterfactuals of the form ‘how much would surplus rise or fall *relative to* the current baseline if the distribution of CEO ability were modified in various ways?’

Tervio explores three counterfactuals for CEO ability: all firms are paired with a CEO of the ability of the lowest ranked CEO; all firms are paired with the median CEO; all firms are paired with the highest ranked CEO. One can also perform a similar counterfactual for firm size. See Tables 1 and 2.

Several points:

1. The contribution of firm size to CEO rents is much larger than the contribution of CEO ability. That is, in the logic of this model, the high levels of CEO pay at the top are mostly due to the exogenous component of firm scale.
2. The distribution of CEO ability *relative to* the lowest ability CEO can be inferred from the differential equations that describe the equilibrium outcomes. See Figure 5. A key conclusion, and one that should be intuitive, is that the distribution of CEO ability is *not* wide. It is estimated that the CEO at the top ranked firm is only 5 to 8 percent more able than the CEO at the bottom ranked firm.
3. To see the role of firm size in CEO pay, consider figure 4 where all firms are replaced with the lowest, median or highest ranked firm while the CEO ability distribution is held constant. Panel B is perhaps most informative. If all firms were the median size, the top CEO would earn 12.5 million less and the bottom CEO would earn 2.5 million more.

4.5 CONCLUSIONS

You do not have to be convinced that this model or empirical exercise fully (or even partly) explains the level or trends of CEO pay to find much of value in this paper. As a labor economist, you should add the basic workings of assignment models to your intellectual toolkit. And this is best paper I’ve seen for acquiring those tools. As I will discuss in class, the assignment model can be used to interpret the Rosen 1981 *AER* superstar’s model, though Rosen’s original model is not constructed in the form of an assignment model.

I will not spend class time on the related paper by Gabaix and Landier (*QJE* 2008) and the critique of Gabaix-Landier by Gordon and Dew-Becker (“Unresolved Issues in the Rise of American Inequality”, *Brookings*). This is an informative debate.

5 TERVIÖ: MEDIOCRITY IN TALENT MARKETS

A paper by Marko Terviö proposes an alternative model that casts the same facts in quite a different light. The point of departure between the Rosen and Terviö articles is in the mechanism by which talent is discovered. Rosen takes it as given that talent is *known* to the market. Given full information, the wages paid to talent are likely to be efficient. Terviö’s paper takes a step back by asking: how does talent become known? Terviö observes that to be ‘discovered,’ one must first ‘audition.’ This observation seems plausible for many talent-occupations. In the Terviö model, knowledge about worker quality is a joint output of production. (That is, to find out if a worker would make a good actor, he needs to appear in a movie.) Economist Richard Caves refers to this *ex ante* uncertainty about talent as the “Nobody Knows” property – one cannot evaluate talent without putting it to use. Now, add a second key assumption: there are positive production costs. That is, workers cannot audition ‘for free’ because some other scarce is used for the audition. This resource could be capital or, more realistically, it could be the time of skilled evaluators (talent agents, moviegoers, academics, passengers on an airplane, etc.). More generally, consumers cannot become informed about the entire array of talents potentially available in the market because time and attention devoted to discovery is scarce. To evaluate and disseminate information about talent is costly. If these assumptions are correct, there is a positive opportunity cost to auditioning a worker’s talent by using him for production. The cost is that someone else better could fill the slot (therefore making better use of the same scarce resource).

Now assume: 1) learning is symmetric (i.e., my talent is revealed to the public at the same time as it is revealed to the employer – this assumption can be weakened); 2) workers cannot commit to binding wage contracts (i.e., no indentured servitude); 3) workers are capital constrained so that they cannot simply ‘buy’ a job to audition their talents.

What does this structure imply about the operation of the market for talent? The answer is surprising. Rather than fully develop the model – which Terviö does elegantly – I will work through the example in the paper.

5.1 TERVIÖ EXAMPLE

- Consider a competitive industry that combines workers with capital. Firms need one worker and one unit of capital at cost \$4 million to produce output.
- There is free entry by firms.
- There is an unlimited supply of potential workers, all of whom have opportunity cost \$0. Hence, *talent is not scarce*.
- Output (more naturally ‘quality’ rather than ‘quantity’) entirely depends on the talent of the worker.
- The talent distribution is uniform on 0 to 100. $T \sim U [0, 100]$.
- Talent is ex ante unknown to workers and becomes public after 1 period of work.
- Careers are finite and last at most 16 periods.
- The industry output faces a downward-sloping demand curve, and firms take the market price as given.
- Workers are constrained to take a non-negative wage (no borrowing).
- For simplicity, there is no discounting.

5.2 EQUILIBRIUM WITH CONSTRAINED WORKERS

What does the equilibrium of this model look like? Initial results:

- All workers who turn out to be above the population mean of 50 will stay in the industry until they retire. Why? These veterans create more revenue than a novice in expectation, so they can always out-compete them for a job in the industry.
- The market price of output must be such that novice-hiring firms break even. Since novices are *not* scarce, they will always be paid zero (their opportunity cost).
- A potential novice is expected to make 50 units, so an output price of \$80,000 is needed to recover the \$4 million capital cost. At this price, there is no entry or exit.

- The average output of veterans will be 75 (since veterans are retained if their talent exceeds the mean).
- The average output in the industry will be lower, since this output includes 1st period novices. In fact, the average output is 72. Why is this so? In equilibrium, outflows and inflows must be equal. Since one cohort of veterans retires each period, there must be sufficient novices to replace them given quality threshold θ . Hence, $1/\theta$ novices must be hired each period. The overall fraction of novices in the industry in each period will be:

$$\frac{(1/((1 - \theta)))}{(1/(1 - \theta) + 15)} = \frac{1}{1 + 15(1 - \theta)}.$$

With $\theta = 0.5$, the fraction of industry novices is 0.118 and average talent is 72.

- A novice has a fifty-fifty chance of being retained in the industry, in which case his expected wage for the next fifteen periods is $15 \cdot [(75 \times 80,000) - 4,000,000] = \30 million.
- Hence, the lifetime rent associated with being hired into the industry is one-half of this amount, \$15 million (since there is a 50 percent chance of exiting after one period).

5.3 EFFICIENT BENCHMARK

Now we want to ask: what should happen?

Suppose that workers are risk neutral and are not credit constrained. They will be willing to bid to enter the industry up to the value of expected future talent rents.

- Start by assuming that novices offer \$1.5 million to firms for the chance to work (this will turn out to be the equilibrium).
- Then, at output price $\$P$, a novice-hiring firm will in expectation generate $50 \times \$P$ in revenue, and have a net cost of \$2.5 million (since the worker has already paid \$1.5 million for the job and the capital cost is \$4 million).
- For firms to break even, the equilibrium price of output must therefore be \$50K ($\$50K \times 50 = \2.5 million).
- At this price, it is only worthwhile for a novice to stay in the industry (that is, become a veteran) if his talent is greater than or equal to 80 (since $\$80K \times 50 = \4 million). The expected talent of veterans is therefore 90.

- The average per period wage of veterans is $\$90K \times 50 - \$4,000K = \$0.5$ million.
- So, the expected lifetime rent of novices entering the industry is $.2 \times 15 \times .5$ million = \$1.5 million, which is the price of buying a job we had assumed initially.
- Each period, this many new novices must be hired (as a fraction of all employed): $1 / (1 + 15(1 - .8)) = .25$
- The average output in the industry, including novices will be: $.25 \times 50 + .75 \times 90 = 80$

5.4 COMPARISON

Let's compare the characteristics of this equilibrium to the previous one:

	Original	Benchmark
Output price	\$80K	\$50K
Threshold talent	50	80
Average talent	72	80
Top wage	\$4 million	\$1 million
Proportion novices	12%	25%

What is quite striking here is that the original equilibrium has *lower talent and higher wages* than the efficient benchmark. Why? Because the original equilibrium has artificial talent scarcity. Because talent rents accrue to workers not firms, firms have insufficient incentive to create new opportunities for talent revelation (in fact, they have none). So, there is under-experimentation, excess retention of mediocre talent, and inefficiently high rents.

This problem could be solved in either of two ways: (1) the legalization of binding indentured servitude contracts such that a worker's wage is set independent of output and the worker cannot quit once hired (but can be fired in the 1st period) – this used to happen in European Soccer through the transfer fees system; (2) the removal of credit constraints so that workers can bid for jobs. This bidding will generate an efficient market in talent revelation. You might ask yourself whether the markets for CEOs, athletes, movie stars and other celebrity professionals appears to satisfy (1) or (2).

The paper is well worth reading for its formal development of these ideas. The extension of the model to ‘has beens’ is particularly elegant. An interesting point on which to speculate is whether this model is, in some sense, a more general case than the Rosen model; that is, one could potentially nest the Rosen model as a special case of this model in which talent revelation was costless.

To sum up, the key insight suggested by this model is that talent can receive large rents even if that talent is mediocre. The reason is that the (ex-post) realization of the known talent distribution may have artificial scarcity.

It will be interesting to observe whether this paper has an impact on the economic presumption that the market for superstars is in some intrinsic sense efficient.