14.662 Spring 2017

Bonus Lecture Notes: Technological Change, Organizational Structure and Wage Structure

David Autor, MIT and NBER

February 11, 2017

Introduction

This bonus note discusses several topics that more deeply probe the relationship between technology, skill demands, organizational structure, and wage setting. In prior semesters, I've cover these topics in lecture, but I'm omitting them this spring to make room for new topics. (To paraphrase Nathan Hale, I only regret that I have but 13 lectures to give to my 14.662 spring 2017 students. Hale's sentiment was far less exalted, of course.) These topics are relevant for *all* lunch-bucket labor economists in training. They won't appear on the 14.662 final exam, but the Labor General exam is of course 'no holds barred.'

1 Why is Information Technology Associated with Greater Worker Discretion and Teamwork (i.e., less 'Taylorism')?

The models discussed so far operate at a very high level. They posit a dramatic change in the organization of work but do not tell a detailed story about how specifically work organization has changed nor why this change has been especially demanding of skilled labor. It would be helpful to have a substantive theory of the nature of recent technological or organizational change that didn't ultimately boil down to the assumed sign on a derivative or the magnitude of some (ultimately unmeasurable) elasticity.

Caroli and Van Reenan study—in a reduced form manner—the impact of work reorganization on skill demands. They provide evidence from British and French establishments that measures of organizational change—such as decentralization of authority—have a strong predictive power for the demand for skills at the establishment level, even after controlling for other determinants of the demand for skills, such as computers. Bresnahan, Brynjolfsson and Hitt (2000) provide U.S. based evidence in a similar vein. Neither of these studies get a great deal further than documenting correlations between reorganization and skill input, but the correlations are striking.

There is also a variety of theoretical and empirical work that offers explicit hypotheses on the link between computerization and changes in the content of work and organization of production.

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 reading list 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, factories have extensively

computerized the process of machine setup for new runs and the coordination of production as products move between machining phases. 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 suggests that product as well as process innovations may contribute to the evolution of skill demands.

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, and enhance quality. In this setting, a good decision can have enormous value added and a bad decision can destroy large amounts of capital.

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 – as well as by Bresnahan, Brynjolfsson 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 these new practices 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. Now in-person handling of checks was reserved for a handful of '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 still performed 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 inspection time 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 micro-foundation for this idea is not given.

2 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. They offer an alternative theory in which specialization comes from a careful optimization problem that balances adaptation, coordination and specialization, and doesn't always increase with larger markets or cheaper communication.

2.1 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 sautéed 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 (that is, there are no within-person coordination costs). In the limit, one can push coordination costs to zero by having all tasks performed by a single worker.

But the downside to bundling of tasks is the loss of 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. One can maximize the gains to specialization by having a separate person do *each* task. But coordination suffers, leading to lower quality or higher coordination costs to achieve the same quality. 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.

2.2 Task structure of production

Production can be thought of as taking place on an $n \times n$ grid of tasks:

a^{11}	a^{12}	a^{13}	...	a^{1n}
a^{21}	a^{22}	a^{13}	...	a^{2n}
a^{31}	a^{12}	a^{33}	...	a^{3n}
...	...	...	...	a^{4n}
a^{n1}	a^{n2}	a^{n3}	...	a^{nn}

The diagonal elements of this grid are the *primary* actions. The off-diagonal actions are the coordinating complementary actions.

Workers are assigned one or more *rows* of tasks, with the number of rows t assigned to each worker determined in equilibrium. For each row that a worker is assigned, it is her responsibility to select both the primary and coordinating actions for the tasks in that row. For simplicity, assume that task assignments are symmetric, so that all workers perform the same number of tasks.

Although workers are responsible for rows, output is a function of the elements of a column. In particular, the complementary actions in each row must be coordinated with the actions in their respective columns *not* the actions in their row. This means that workers responsible for different rows must coordinate with one another to choose tasks optimally within a column.

Before discussing how tasks are chosen, it's useful to understand the sequencing of the model, which goes as follows:

1. The firm determines the number of task per agent, t.
2. Local circumstances θ^i for $i = 1, 2, ..., n$ are realized and observed by employee(s) in charge of task i.
3. Workers communicate the realizations of local information, and each attempt at communication is successful with probability p. Workers conveying this information cannot determine whether or not communication was successful.
4. For each row i, the employee in charge of i chooses actions a^{ij}, where $j \in \{1, 2, ..., n\}$, to maximize the objective function given his information.
5. Profits are realized given the realization of local information, the success of communication, and the chosen values of all tasks.

Now let's consider how actions are chosen.

2.3 Adaptation and coordination

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.

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 $\{a^{i1}, a^{i2}, ..., a^{in}\}$ that are complementary to the primary actions of tasks $j \neq \dot{i}$. 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}. Notice (as above) that worker in charge of task i must coordinate the complementary actions with the tasks in rows $j \neq i$.

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

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

where the parameters $\phi \geq 0$ and $\beta \geq 0$ 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.

2.4 Specialization

Let $T(i)$ equal the set of tasks bundled with worker 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 $T(i) = t$ for all workers (and the total number of workers is n/t).

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:

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

Put more simply $h(0) \geq 0$ and $h_t(\cdot) > 0$. Bundling tasks raises per unit costs of execution by sacrificing gains from specialization. The parameter α reflects the returns to specialization. When α is larger, the gains from specialization are larger because the costs of non-specialization are higher.

2.5 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.

2.6 Profits

Profits of a firm are given by:

$$\begin{aligned}\pi &= -\sum_{i=1}^{n} C^i\left(a^{1i}, a^{2i}, ..., a^{ni}, t|\theta^i\right) \\ &= -\sum_{i=1}^{n}\left[\phi\left(a^{ii}-\theta^i\right)^2 + \sum_{j\notin T(i)}^{n}\beta\left(a^{ji}-a^{ii}\right)^2 + h\left(t,\alpha\right)\right]\end{aligned}$$

You can think of this profit function as a loss function that the firm is attempting to minimize. Hence it might be more intuitive to write $\pi^* = p\times(y+\pi)$, where p is the market price of a perfectly produced unit of output and π is the reduction in quality incurred by imperfections in adaptation and coordination during production.

2.7 Optimal choice of actions

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

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

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:

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

The covariance between local circumstances and the primary action are:

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

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

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

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

Given t, expected profits are:

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

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

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

2.8 Some results

2.8.1 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.

2.8.2 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 this would increase bundling (reduce specialization). In this model, the answer depends on the importance of adaptation ϕ. An increase in task interdependence affects two margins:

1. Holding worker flexibility/adaptation constant, 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:

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

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.

2.8.3 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.

2.8.4 Other comparative statics

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

3 Competing Organizational Forms and Skill Supply: Acemoglu 1999 *AER*

Acemoglu *AER* '99 considers a search model in which firms can use two types of technology: one that uses only high skill workers and the second that uses both high and low skill. The former technology has higher TFP conditional on having a high skill workforce. But if high skill workers are relatively rare, they will have a low arrival rate in the search model. This makes it costly to open vacancies using the strongly skill-dependent technology. Hence, firms will primarily use the mixed technology. However, an exogenous increase in skill supply may induce firms to open skill-dependent jobs. This will raise wages of high skill workers (since they are more productive in the skill-dependent sector) but lower wages of low-skill workers (under the assumption that high and low skill workers are q-complements under the mixed technology).

This model is analytically complex, which may be a virtue if you are a search theorist, otherwise probably not. Acemoglu's 2002 *JEL* paper offers a much simpler alternative that has the same flavor. There is a scarce supply of some factor K, which could be capital or entrepreneurial talent, for example. Skilled workers work with the production function

$$Y_h = A_h^{\alpha} K_h^{1-\alpha} H^{\alpha},$$

and unskilled workers work with the production function

$$Y_l = A_l^{\alpha} K_l^{1-\alpha} L^{\alpha},$$

where Y_h and Y_l are perfect substitutes and $K_l + K_h = K$. Observe: if both production functions are in use, the return to capital K must be equated across H and L sectors.

Now consider a factor-augmenting technological change, specifically an increase in α_h. This raises productivity in the H sector, which increases the return to K in the H sector (since K and H are q-complements). To equilibrate the return to capital across sectors, K must flow from the L sector to the H sector. Given the q-complementarity between K and L, wages of L workers must fall.

Hence, this model generates falling wages of low-skilled workers. This model also implies a rise in the return to K. So, if K is elastically supplied—that is, the stock of K rises until returns fall back to the prior level—the real wages of low-skill workers do not fall.

The closely related paper by Beaudry and Green in *AER* in 2003 ("Wages and Employment in the United States and Germany: What Explains the Differences?") features one critical difference with the model above. In B&R the rising *supply* of skilled workers causes work to reallocate from the low-skill to high-skill sector, ultimately leading to capital starvation in the low-skill sector. Thus, while in the model above, traditional factor-biased technical change leads to rising demand for high skilled workers whereas in B&G, supply itself is sufficient. In both models, it is the inelastic supply of capital that allows low skilled wages to fall. [B&G is actually a traditional Hecksher-Ohlin two-skill, two-sector model, with one unusual wrinkle: both sectors produce an identical good. Thus, changes in factor supplies or prices affect which 'technique' is used to produce the final good, and this in turn affects sectoral allocation, wages and capital demands.]