On the Distribution of Income and Worker Assignment under Intrafirm Spillovers, with an Application to Ideas and Networks

Gilles Saint-Paul

Université des Sciences Sociales de Toulouse, Universitat Pompeu Fabra, Centre for Economic Policy Research, and Institut zur Zukunft der Arbeit

I study the earnings structure and the equilibrium assignment of workers when workers exert intrafirm spillovers on each other. I allow for arbitrary spillovers provided that output depends on some aggregate index of workers' skill. Despite the possibility of increasing returns to skills, equilibrium typically exists. I show that equilibrium will typically be segregated and that the skill space can be partitioned into a set of segments and any firm hires from only one segment. Next, I apply the model to analyze the effect of information technology on segmentation and the distribution of income. There are two types of human capital, productivity and creativity, that is, the ability to produce ideas that may be duplicated over a network. Under plausible assumptions, inequality rises and then falls when network size increases, and the poorest workers cannot lose. I also analyze the impact of an improvement in worker quality and of an increased international mobility of ideas.

I. Introduction

Over the last 10 years there has been a renewed interest in problems associated with the distribution of income. This interest has been mo-

I thank Roland Bénabou, Fernando Alvarez, Adriana Kugler, Rafael Repullo, Sherwin Rosen, Masako Ueda, an anonymous referee, participants at the National Bureau of Economic Research 1998 Summer Institute workshop on inequality, the Conference on Technology, Regulation, and Employment (Madrid, 1999), as well as seminar participants at Universitat Pompeu Fabra and Université des Sciences Sociales de Toulouse for helpful comments and suggestions.

[*Journal of Political Economy,* 2001, vol. 109, no. 1]

tivated by the upward trend in inequality that has taken place in the United States and other countries since the early 1970s and has led to a spurt of empirical and theoretical research.[1] This research has studied the role of such diverse factors as the role of international trade, computers, educational achievements, and stratification in firms and neighborhoods.

The present paper is a theoretical contribution to that literature, which was originally motivated by an attempt to understand the effect of information technology on the distribution of income. It is mainly concerned with how technology and stratification interact with each other in determining the distribution of income. I study the earnings structure and the equilibrium assignment of workers to firms in a model in which workers exert intrafirm spillovers on each other. I allow for such spillovers to be somewhat arbitrary, provided that the number of workers within firms is fixed—an assumption that can be relaxed to some extent—and output is a function of some aggregate index of all the workers' skills in the firm. Because of the spillover there are nonconstant marginal returns to skills, and I allow returns to be either increasing or decreasing. Despite the possibility of increasing returns to skills, equilibrium typically exists because the fixed number of workers within each firm prevents them from increasing profits by simply replicating themselves.

I show that equilibrium will typically be segregated, in the sense that the skill space can be partitioned into a set of segments and any firm hires from only one segment. The number of segments depends on the shape of the curve relating total output to aggregate skills, with convex portions acting as a force toward segregation and concave ones as a force toward unification. The wage schedule is always convex in skills, regardless of the shape of spillovers, implying that market forces, by allocating people of similar skills into the same firms, magnify differences in skills and tend to skew the distribution of income to the right. Wages can be interpreted as the difference between the worker's marginal product and a shadow cost of bodies that captures the spillover exerted by the worker on other workers within the firm. This cost is positive if there are increasing returns so that firms are willing to pay a negative wage for some categories of workers.

Next, I develop an application of the model that yields some insights into how information technology affects the distribution of income. There are two extreme ways in which one can think about that problem. At one extreme is the standard view, which ignores job characteristics

[1] See Bound and Johnson (1992), Katz and Murphy (1992), Levy and Murnane (1992), and Juhn, Murphy, and Pierce (1993). Some of the theoretical literature on segmentation and inequality that this burst of empirical research has motivated is referred to later.

and treats each labor type as homogeneous and divisible. According to that view, computers are substitutes for unskilled workers, so that their introduction is quite similar to a rise in the supply of unskilled workers or trade with a country abundantly endowed in that factor.[2] At the other extreme there is the "superstar," or "winner-takes-all," view, according to which information technology allows the best people to spread their talent over an ever larger share of the market, at the expense of less able workers who may find themselves displaced into low-paying jobs, despite being only marginally less talented. This trend toward a society in which all the rewards go to the best has been analyzed in the book by Frank and Cook (1995), who provide numerous examples from sports, performing arts, and top executives. The question, however, is whether such effects are pervasive throughout the labor market or limited to football players and divas. Rosen (1996), in his review of the book, argues that these phenomena are confined to those segments of the market in which a product can be duplicated at no cost (the media) or being the best is the essence of the product (sports competition).

The model I develop takes an intermediate view. I think of wages as the return to two types of human capital. One is productivity, that is, how many units of output one can produce; the other is "creativity," that is, the ability to produce ideas that may be duplicated over a segment of the economy, what I call a "network." In my model, such ideas increase the output of all other workers in the same network; this is the nature of the spillover. I assume that firms and networks coincide, so that the beneficial effect of ideas will increase profitability and, through competition, eventually increase wages for the most creative workers.[3] There is a winner-takes-all dimension to creativity because only the best idea is applied. I model progress in information technology as an increase in the size of the network over which ideas can be spread. It is therefore a force that increases the return to the most creative workers. But, contrary to what happens in, for example, sports, best performers are complementary to other workers since a good idea increases the productivity of all participants in the network. This is a force that mitigates the inegalitarian effects of information technology. Furthermore, as networks get larger, the most creative workers end up competing with each other in the same networks, which reduces the return to creativity.

Under plausible assumptions, I show that inequality rises and then

[2] See Krueger (1993) and DiNardo and Pischke (1997) for the empirical aspect of that debate.

[3] Thus I abstract from the problems of imperfect appropriability of intellectual contributions, as emphasized by the literature on endogenous growth (see Lucas 1988; Romer 1990; Grossman and Helpman 1991; Aghion and Howitt 1998). A Coasian argument would lead to the conclusion that the boundaries of the firms will endogenously adjust so as to internalize the externality; therefore, in equilibrium, firms will coincide with networks.

falls when network size increases. My conclusions are therefore far less pessimistic regarding the impact of information technology on inequality than the extreme superstar approach. I also show that the poorest workers cannot lose from an improvement in information technology and analyze the impact of an improvement in worker quality on income distribution and worker segmentation.

The present paper is related to several strands of literature. One is the literature on job assignment, whose central claim is that an earnings function that relates wages to worker characteristics is insufficient for predictive purposes because it ignores the assignment of workers to jobs (see Sattinger [1993] for a survey). That is, while the standard neoclassical theory of labor demand relates output to homogeneous labor inputs, the assignment approach explicitly recognizes, as this paper does, that both jobs and workers matter. In my approach, the distribution of income depends not only on the distribution of skills but also on the way in which workers are endogenously grouped together in the labor market; a change in the underlying distribution of skills will also affect the way workers are grouped. However, most of the assignment literature assumes that workers are assigned to occupations or jobs whose differences are specified exogenously, whereas in my case, jobs are symmetrical ex ante but end up being different because one's productivity depends on one's colleagues' skills, which differ across firms in equilibrium.

A related line of research has also pointed out that worker characteristics cannot be aggregated into homogeneous factor inputs because they are embodied in workers (see Mandelbrot 1962; Rosen 1983; Heckman and Scheinkman 1987). A similar property arises in my model, where an increase in a firm's aggregate skill can be obtained only by hiring a different mix of worker types.

Another literature is the one about "clubs," ghettoes, and segregation, whose archetypal example is Becker's (1973) theory of marriage. Segregation arises when complementarities are prominent, so that people want to end up in the same club or firm as the "best" ones. This is often the case in problems of local public finance or peer group effects (as in Bénabou [1993, 1996], Fernandez and Rogerson [1996], Epple and Romano [1998], or Rioux [1999]). As in that literature, strong spillovers (which in this case show up in the form of convexity rather than complementarity) lead to segregation. The present paper, however, is more concerned with the determination of earnings and establishes general results regardless of the shape of the spillover.

Finally, the present paper is also related to the literature on superstars, following Rosen (1981), who, in a seminal paper, established that the large income for superstars results from their ability to cover a larger market. The present paper differs from Rosen's superstars model in that spillovers play a key role in determining the distribution of income, whereas in Rosen's paper the key mechanism came from the ability of talented individuals to cover a larger market, much in the fashion in which, under imperfect competition, firms with lower costs have more customers.[4] Kremer's (1993) O-ring production function has more in common with the present paper. In both papers, firms consist of a fixed number of workers,[5] and workers exert spillovers over each other. Here, however, I allow for a much more general pattern of spillovers. Both the O-ring production function and the ideas/networks application that I develop are special cases of that general analysis. In fact, in my case it is the best workers who exert positive spillovers, whereas in Kremer's model it is the worst performer who exerts negative spillovers.[6]

Despite the generality of my results, there is one aspect of that literature that is not captured by my model, namely, the hierarchical one. Many authors, such as Calvo and Wellisz (1979), Rosen (1982), or, more recently, Garicano (1998), have developed models that predict that more talented people are more likely to be "above" in hierarchies and that span of control is increasing with talent. This is clearly not captured in my model, where tasks are modeled as symmetrical; as I briefly argue, to capture these aspects would require a production function in which, in addition to the mean, higher moments of the intrafirm distribution of skills would enter.

The paper is organized as follows: In Section II, I establish my basic general results regarding existence and efficiency of equilibrium, as well as the pattern of segregation and the shape of the wage distribution. In Section III, I show how this framework can be applied to ideas and information technology. In Section IV, I perform comparative statics exercises, looking at the impact on the distribution of income and the pattern of segregation of an increase in average worker quality and improvement in information technology, which I model as an increase in network size. These exercises are motivated by the above-mentioned debate on the rise of inequality. Section V briefly discusses potential extensions of the model, and Section VI contains concluding comments.

[4] Murphy, Shleifer, and Vishny (1991) have applied Rosen's ideas to understand the interaction between economic growth and the allocation of talent across different activities. This is a very different route from the one pursued here, although the two kinds of considerations could be merged.

[5] Or, more fundamentally, the number of workers enters the production function separately from the aggregate labor input.

[6] Furthermore, the correlation between productivity and creativity plays a key role, in my model, in shaping the structure of returns to scale; skill has only one dimension in the O-ring case.

II. A Model of the Labor Market under Intrafirm Spillovers

A. *Basic Setup and Definition of Equilibrium*

Let us consider an economy populated by a continuum of workers who differ by their skill level y, which is distributed over the interval $I = [y_0, y_1]$ with a measure μ, which has full support. The total mass of workers is normalized to one. Firms freely enter the market and consist of a mass s of workers. Therefore, there will be a mass $1/s$ of firms in equilibrium.[7]

A firm's total output is given by $a(\bar{y})$, where $\bar{y}$ is the average skill level of its workers. Therefore, I assume that skills can be defined in such a way that the firm's output depends only on the aggregate skill level within the firm.[8]

In the absence of spillovers, the firm's output is simply the sum of the contributions of each individual worker, which does not depend on his colleagues' characteristics. Consequently, by redefining the skill level as productivity, we get that the $a(\cdot)$ function is then linear. If, on the other hand, there are intrafirm spillovers, then the $a(\cdot)$ function will be nonlinear, which is the case of interest. Workers may exert negative spillovers on each other if there is some fixed factor necessary for running the firm and if they are competing for that fixed factor.[9] In that case, the $a(\cdot)$ function will have decreasing returns. Or they may exert positive spillovers on each other because they cooperate in teams (as in Kremer [1993]) or because of the role of ideas, a phenomenon we shall study below. In that case, each worker's productivity depends positively on other workers' skills, and $a(\cdot)$ will exhibit increasing returns. If, for example, each worker's productivity is the product of his skill and a spillover effect defined as an increasing function of the firm's average skill $f(\bar{y})$, then total output is given by $a(\bar{y}) = s\bar{y}f(\bar{y})$ and exhibits increasing returns.

The production function exhibits increasing returns to skills if $a'(y)y > a(y)$. This condition is different from the condition for increasing *marginal* returns to skills, which is given by $a''(\cdot) > 0$. In the application below, the first condition is always satisfied and the second may not hold. For the time being I do not make any assumption about the

[7] More generally, the assumption of a fixed number of employees within each firm could be relaxed by specifying a production function that depends on both the number of workers and their aggregate skill level. This is discussed below in Sec. V.

[8] Such production functions are also considered by Sattinger (1980) and do not really entail a loss of generality if skills are defined as a multidimensional vector. However, this multidimensionality is associated with problems; see the discussion at the end of Sec. II*B*.

[9] For example, one may need a fixed level of capital k to operate the firm, and if output is given by $f(k, s\bar{y})$, with f having constant returns to scale, then $a(\bar{y}) = f(k, s\bar{y})$ will have decreasing returns.

returns to skills, just assume that $a(\cdot)$ is continuous and twice differentiable.

An equilibrium is characterized by a wage schedule $w(y)$, which tells us how much a worker of skill y will be paid, and by an assignment of workers to firms such that all workers are assigned and firms maximize profits, and no potential entrant could make strictly positive profits. The assignment can be represented by a mapping η, which tells us the skill level of worker i in firm k. The definition below uses such a representation to precisely state the conditions for an equilibrium.

DEFINITION 1. Let M be the set of mappings from $[0, s]$ to I. Let m be the Lebesgue measure on $[0, s] \times [0, 1/s]$. An equilibrium consists of (i) a mapping η from $[0, s] \times [0, 1/s]$ to I, $(i, k) \to y_{ik}$ (worker assignment); and (ii) a mapping w from I to R, $y \to w(y)$ (the wage schedule) such that (a) for all $k \in [0, 1/s]$, $a(\bar{y}_k) = \int_0^s w(y_{ik})di$, where $\bar{y}_k = (\int_0^s y_{ik}di)/s$ is firm k's average skill level; (b) for all $\nu \in M$: $i \to \hat{y}_i$, $a((\int_0^s \hat{y}_i di)/s) \leq \int_0^s w(\hat{y}_i)di$; and ($c$) for all $S \subset I$, $\mu(S) = m(\eta^{-1}(S))$.

The first property tells us that because of the free-entry condition, existing firms make zero profits. The second condition says that potential entrants cannot make positive profits, otherwise they would indeed enter. The third condition is the full-employment condition for any type of worker, which says that the distribution of y implied by the assignment of workers matches the actual one.

B. The Structure of Earnings

The next task is to characterize equilibrium and prove its existence. To do so we shall proceed in several steps. The next proposition tells us how to recover the wage schedule given the initial configuration of firms.

PROPOSITION 1. The equilibrium wage schedule is such that (i)

$$w(y) \geq \frac{a(\bar{y}_k)}{s} + \frac{a'(\bar{y}_k)(y - \bar{y}_k)}{s},$$

for all $y \in I$, for all $k \in [0, 1/s]$; (ii)

$$w(y) = \frac{a(\bar{y}_k)}{s} + \frac{a'(\bar{y}_k)(y - \bar{y}_k)}{s}, \tag{1}$$

if firm k employs some workers of type y;[10] and (iii)

$$w(\bar{y}_k) = \frac{a(\bar{y}_k)}{s}.$$

[10] Formally, this means that $k \in P_2(\eta^{-1}(y))$, where P_2 is the projection operator from $[0, s] \times [0, 1/s]$ to $[0, 1/s]$.

Proof. To compute their desired optimal employment structure, firms maximize

$$\max_{\omega(\cdot)} a\left(\frac{\int_I \omega(y) y d\mu}{s}\right) - \int_I \omega(y) w(y) d\mu$$

subject to

$$\omega(y) \geq 0 \quad \forall y \in I$$

and

$$\int_I \omega(y) d\mu = s. \tag{2}$$

That is, they elect how much of each type they want to hire, as defined by a density $\omega(y)$, subject to the constraint that their total size is s. A type y is employed in that firm if and only if $\omega(y) > 0$.

The first-order condition with respect to any $\omega(y)$ is therefore

$$\frac{a'(\bar{y}_k) y}{s} - w(y) \leq \lambda_k, \tag{3}$$

where λ_k is the Lagrange multiplier[11] associated with (2), and equality holds if $\omega(y) > 0$. Integrating both sides over all employees of firm k and making use of the zero profit condition, we see that

$$\lambda_k = \frac{a'(\bar{y}_k)\bar{y}_k - a(\bar{y}_k)}{s}.$$

Substituting into (3) completes the proof of parts i and ii. To prove part iii, note that part i applied to $y = \bar{y}_k$ implies $w(\bar{y}_k) \geq a(\bar{y}_k)/s$. Assume that part iii does not hold; then it implies that $w(\bar{y}_k) > a(\bar{y}_k)/s$ and that firm k does not employ any worker of type $y = \bar{y}_k$. But then, any firm that employs workers of type $\bar{y}_k$ could strictly increase its profits by replacing them by the same mix of workers as those employed in firm

[11] One cannot readily apply the Kuhn-Tucker theorem because $a(\cdot)$ may not be concave. However, consider a type y employed by firm k and a replacement of a small mass of that type by another type y'. For the firm to be at its optimum, such a change must not increase profits, which is equivalent to

$$a'(\bar{y}_k)\frac{y'}{s} - w(y') \leq a'(\bar{y}_k)\frac{y}{s} - w(y).$$

If y' is also employed in positive quantity by the firm, then the converse must also hold, implying that the quantity $a'(\bar{y}_k)(y/s) - w(y)$ is the same for all workers employed by the firm. Denoting that quantity by λ_k and calling it the Lagrange multiplier, we get back to (3).

k. This would not affect its average productivity and would yield an average labor cost equal to

$$\frac{a(\bar{y}_k)}{s} + \frac{a'(\bar{y}_k)(\bar{y}_k - \bar{y}_k)}{s} = \frac{a(\bar{y}_k)}{s} < w(\bar{y}_k).$$

Consequently, no firm would employ type $\bar{y}_k$ workers, which contradicts the equilibrium assumption. Thus part iii must hold. Q.E.D.

COROLLARY. The wage schedule $w(y)$ is convex.

Proof. Proposition 1 implies that $w(\cdot)$ is the maximum of a set of linear functions. Q.E.D.

The right-hand side of (1) is firm i's willingness to pay for any type y, that is, the marginal product of an individual of type y if he works in firm k. It must be equal to its wage if the firm employs that type and lower than its wage if it does not. Equation (1) implies that wages paid by firm k can be decomposed as $w(y) = [a'(\bar{y}_k)\, y/s] - \lambda_k$. This quantity consists of two terms, reflecting the fact that this individual contributes to the firm's output but also occupies a job. The term $a'(\bar{y}_k)\, y/s$ is the marginal contribution of his *skill* to firm k's output. It is the only term that would remain under constant returns to scale. The second term, $\lambda_k = [a'(\bar{y}_k)\bar{y}_k - a(\bar{y}_k)]/s$, is the shadow price of a job in firm k, that is, the price that anyone has to pay to be a member of that firm. It is positive if there are increasing returns and negative if there are decreasing returns. It measures the average spillover exerted by individuals on each other in firm k. That is, if worker i's contribution is valued when the average productivity of skill in the firm is used, it will be measured as $a(\bar{y}_k)y_i/s\bar{y}_k$. A unit fall in another individual's skill level would then reduce worker i's contribution by $[a'(\bar{y}_k)\bar{y}_k - a(\bar{y}_k)]y_i/(s\bar{y}_k)^2$. This quantity therefore measures the marginal spillover exerted by other employees on worker i's contribution. Its average value in firm k is clearly equal to $\lambda_k/\bar{y}_k$.[12]

One would also get equation (1) if the total size of the firm could vary and output was given by $a(\bar{y}) = sb(\bar{y})$. It is then straightforward to check that the right-hand side of (1) is the marginal product of a type y worker in firm k. The first term, $a(\bar{y}_k)/s = b(\bar{y})$, is the contribution of the associated increase in employment; the second term, $a'(\bar{y}_k)(y - \bar{y}_k)/s = b'(\bar{y}_k)(y - \bar{y}_k)$, represents the effect of that extra individual on firm k's average worker quality. This term implies that individuals who are better than the firm's average are remunerated because their presence tends to increase that average, whereas those who have a skill lower than the firm's average are taxed.

[12] The term λ_k also measures the effect of the presence of an individual in firm k on the size of the cake to be divided between the remaining workers, if that individual gets the marginal product of his skill.

An important aspect of equation (1) is that under increasing returns, there exist worker types such that the willingness to pay for those workers by firm k is *negative*. This arises for y low enough relative to $\bar{y}_k$. It means that the worker's contribution to output is not enough to compensate for the negative spillover he exerts on his colleagues. As in the work by Akerlof (1969, 1981), this arises because employing somebody uses a scarce resource—here, a job slot in firm k. Contrary to Akerlof's work, however, this scarcity holds only at the firm's level. Free entry of firms will guarantee that all workers are employed.

The nonlinearity of the wage schedule implies that despite the fact that an individual firm's output depends only on its workers' aggregate skill, an individual's wage is not equal to the product of his skill level and a unique market price for skills. The reason is that firms are constrained in the number of people they can hire and can vary their average skill level only by hiring a different mix of worker types. That is, the labor market is a market for individuals, not for homogeneous worker characteristics that could be employed independently of people. A similar result may hold in the literature on the aggregation of worker characteristics (Mandelbrot 1962; Rosen 1983; Heckman and Scheinkman 1987), but for a slightly different reason; in these models, different worker characteristics are bundled together in individuals, and the structure of labor supply may be such that some sectors entirely specialize in hiring certain types of individuals, so that these sectors' marginal willingness to pay for a given characteristic will be disconnected from that of other sectors that specialize in different worker types. As the next subsection makes clear, such specialization is also present here, and it accounts for the nonlinearity of the wage schedule.

C. *Segregation and Workers' Assignment*

The next issue is, What does the equilibrium look like in terms of the equilibrium distribution of firms' average skill levels and in terms of the assignment of workers to these firms? Proposition 2 characterizes the equilibrium.

Proposition 2. Equilibrium can be characterized as follows: There exists a partition of I into adjacent intervals $I_\alpha \subset I$, indexed by $\alpha \in A$, such that the following conditions hold:

i. For any firm k there exists $\alpha(k) \in A$ such that all workers of firm k come from I_α: for all $k \in [0, 1/s]$, there exists α, $\eta([0, s], k) \subset I_\alpha$.
ii. The wage schedule $w(y)$ is linear over I_α: for all $\alpha \in A$, for all $y \in I_\alpha$, $w(y) = \omega_\alpha + \delta_\alpha y$.

iii. For any firm k recruiting in I_α, the wage schedule is tangent to the average output schedule at $y = \bar{y}_k$:

$$\delta_{\alpha(k)} = \frac{a'(\bar{y}_k)}{s};$$

$$\omega_{\alpha(k)} = \frac{a(\bar{y}_k) - a'(\bar{y}_k)\bar{y}_k}{s}.$$

iv. The wage schedule is convex.
v. The wage schedule $w(y) \geq a(y)/s$ for all $y \in I$.
vi. The mean of y over I_α is a convex combination of the mean skill level of the firms that recruit in I_α.
vii. Supply equals demand; that is, condition c in definition 1 holds.

Proof. We must prove that these conditions are necessary and sufficient for equilibrium. Let us prove that they are sufficient first.

Over the interval in which a firm hires people, wages are linear. Aggregating them using parts ii and iii implies that a firm's average labor cost is $a(\bar{y}_k)/s$. Thus existing firms make zero profit.

As the wage schedule is convex (part iv), an entrant with average skill $\hat{y}$ cannot do better than hiring workers who all have the same skill $\hat{y}$. But then his profits are given by $a(\hat{y}) - sw(\hat{y})$, which by part v cannot be positive. Thus properties a and b in definition 1 hold. Property c holds because of part vii.

This proves sufficiency.

Let us now prove that these conditions are necessary. Part iv must hold because of the corollary to proposition 1. Part vii must hold because of definition 1. If part v were violated for some $\hat{y}$, an entrant hiring s workers of type $\hat{y}$ would make strictly positive profits, violating condition b in definition 1.

Next, consider two existing firms, k and l, with different means, $\bar{y}_k < \bar{y}_l$. We shall prove that unless they offer the same wage schedule, the most skilled employee of firm k, of skill $y_k^+ \geq \bar{y}_k$, must be less skilled than the least skilled employee of firm l, of skill $y_l^- \leq \bar{y}_l$. To see this, note that proposition 1 implies

$$sw(y_l^-) = a(\bar{y}_l) + a'(\bar{y}_l)(y_l^- - \bar{y}_l) \geq a(\bar{y}_k) + a'(\bar{y}_k)(y_l^- - \bar{y}_k) \tag{4}$$

and

$$sw(y_k^+) = a(\bar{y}_k) + a'(\bar{y}_k)(y_k^+ - \bar{y}_k) \geq a(\bar{y}_l) + a'(\bar{y}_l)(y_k^+ - \bar{y}_l). \tag{5}$$

Taking the difference between the left-hand side of (4) and the right-hand side of (5) and making use of these inequalities, we get

$$[a'(\bar{y}_l) - a'(\bar{y}_k)] \cdot (y_l^- - y_k^+) \geq 0.$$

Thus either $y_l^- \geq y_k^+$ or $a'(\bar{y}_l) \leq a'(\bar{y}_k)$. But we have, again from proposition 1,

$$sw(\bar{y}_l) = a(\bar{y}_l) \geq a(\bar{y}_k) + a'(\bar{y}_k)(\bar{y}_l - \bar{y}_k),$$

$$sw(\bar{y}_k) = a(\bar{y}_k) \geq a(\bar{y}_l) + a'(\bar{y}_l)(\bar{y}_k - \bar{y}_l).$$

These two inequalities imply that

$$a'(\bar{y}_k) \leq \frac{a(\bar{y}_l) - a(\bar{y}_k)}{\bar{y}_l - \bar{y}_k} \leq a'(\bar{y}_l).$$

Thus we can have $y_l^- < y_k^+$ only if

$$a'(\bar{y}_k) = \frac{a(\bar{y}_l) - a(\bar{y}_k)}{\bar{y}_l - \bar{y}_k} = a'(\bar{y}_l) = \delta$$

exactly. In that case the two firms offer the same wage schedule and can recruit within the same cluster (what is called a *dual cluster*). Furthermore, all firms with a mean skill level between $\bar{y}_k$ and $\bar{y}_l$ will also offer the same wage schedule. Let p index such a firm. Assume first that $a(\bar{y}_p) > a(\bar{y}_k) + \delta(\bar{y}_p - \bar{y}_k)$. Then by hiring workers of types $\bar{y}_k$ and $\bar{y}_l$ in proportions $(\bar{y}_l - \bar{y}_p)/(\bar{y}_l - \bar{y}_k)$ and $(\bar{y}_p - \bar{y}_k)/(\bar{y}_l - \bar{y}_k)$, respectively, the firm gets a total labor cost equal to $a(\bar{y}_k) + \delta(\bar{y}_p - \bar{y}_k)$, thus obtaining strictly positive profits, which contradicts the equilibrium assumption. Next, note that $a(\bar{y}_p) < a(\bar{y}_k) + a'(\bar{y}_k)(\bar{y}_p - \bar{y}_k)$ cannot hold since it contradicts part i in proposition 1, since by part iii we have that $a(\bar{y}_p)/s = w(\bar{y}_p)$. Thus we necessarily have

$$a(\bar{y}_p) = a(\bar{y}_k) + \delta(\bar{y}_p - \bar{y}_k).$$

Finally, note that if $\delta \neq a'(\bar{y}_p)$, there exists some y in the neighborhood of $\bar{y}_p$ such that $a(y) > a(\bar{y}_k) + \delta(\bar{y}_p - \bar{y}_k)$. The argument above can then be used to show that an entrant with an average skill level equal to y would make strictly positive profits. Consequently, one must also have $a'(\bar{y}_p) = \delta$. This, along with proposition 1 and the preceding equality, implies that firm p must offer the same wage schedule as firms k and l.

In summary, either two firms offer the same wage schedule, in which case so do all firms in between, or they must obey the "sorting property": the most skilled person employed by firm k cannot be more skilled than the least skilled person employed by firm l.

The preceding results imply that if we order firms by increasing values of $\bar{y}_k$, they can be grouped in clusters that hire within the same adjacent intervals. To construct these intervals, group all firms into subsets of "consecutive" firms offering the same wage schedule. For any such subset α, define $I_\alpha = \bigcup_{k\in\alpha}\eta([0, s], k)$. As the distribution of types has full sup-

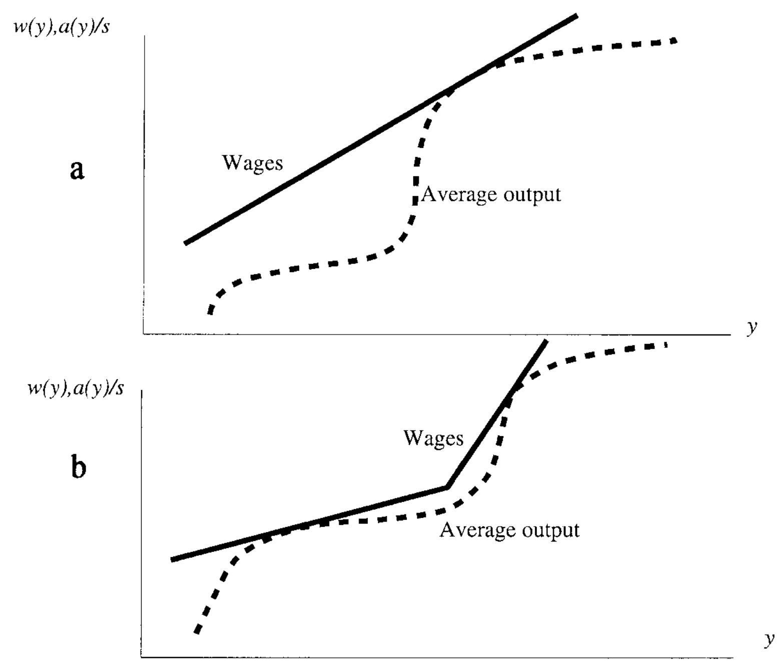

FIG. 1.—*a,* A single cluster covering two humps. *b,* Two unitary zones recruiting from adjacent intervals. *c,* Hypersegregated zone followed by a unitary one. *d,* Two unitary zones separated by a hypersegregated zone. *e,* Dual cluster.

port, these sets cover the whole interval I. Furthermore, to match the sorting property above, they must be intervals.

Within each I_α, the wage schedule is linear since all firms offer the same one and tangent to each firm's mean value of y by construction. This proves parts i, ii, and iii. To prove part vi, note that, in equilibrium, firms recruiting in I_α must hire the whole supply of workers in I_α. Therefore, the average skill level in that interval can be computed by adding $\bar{y}_k$ over all firms k that recruit in I_α and dividing by s times the total mass of such firms. Consequently, part vi must hold. Q.E.D.

Proposition 2 is very useful to understand the structure of equilibria. The typical structure is illustrated in figure 1, which shows us a map of how the economy organizes itself into various zones with very different properties. The average output schedule, which gives us $a(y)/s$ as a function of the firm's average skill level y and is represented by the dotted curve, will in general have convex parts and concave parts. In the concave zones, or "humps," there are local decreasing returns to average human capital, which create a force for agglomeration of different peo-

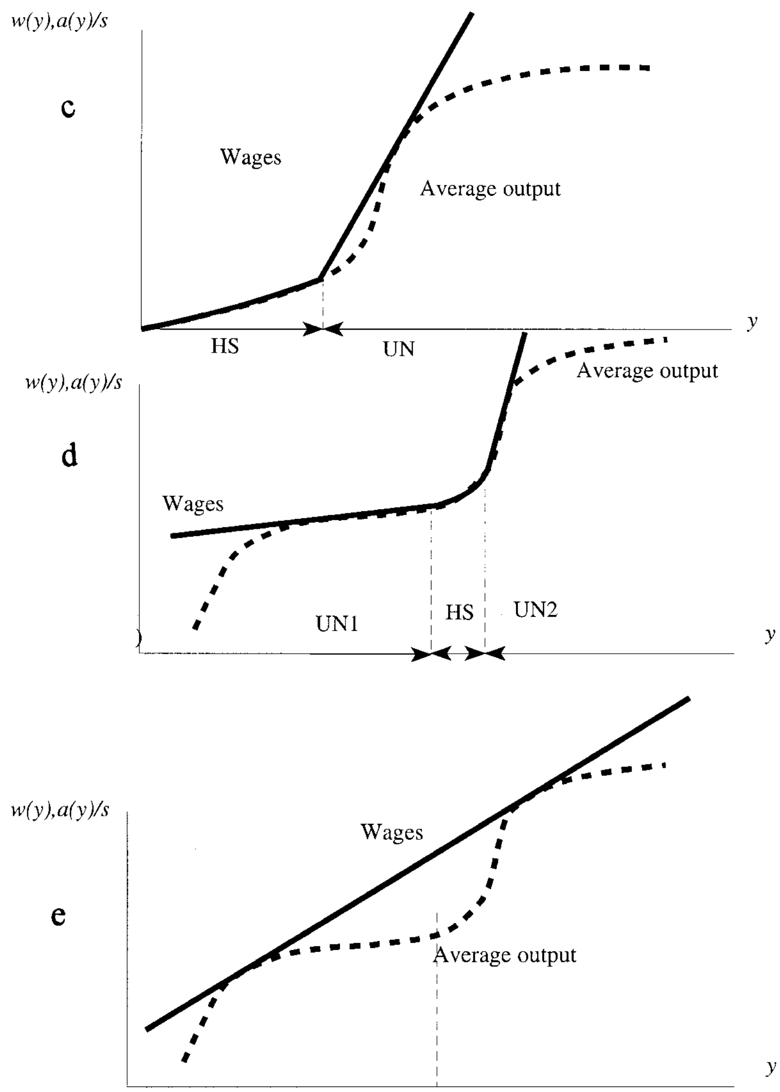

FIG. 1.—*Continued*

ple within the same firm. In the convex zones, there are local increasing returns, which create a force for segregation, that is, matching of people with similar levels of human capital within the same firm.

As can be seen in figure 1, equilibrium is characterized by a map of segregated clusters. The relevant intervals are determined by the points at which the wage schedule $w(y)$, represented by the plain curve, has a kink. Because of convexity, wages are always more sensitive to skill when one moves up the distribution of income. Both that property and the segregation property are quite general and do not depend on the convexity of $a(\cdot)$ throughout. If $a(\cdot)$ were concave, there would be a single cluster and a linear wage schedule, but typically it is enough for it to be convex over some interval to generate segregation and increased steepness of the wage schedule.

The following taxonomy emerges.

1. In many cases there will arise "unitary zones," that is, intervals of values of y in which all people are hired in a single type of firms, with heterogeneous workers but a common average skill level. By single *type* I mean that firms recruiting in a given interval all have the same mean skill level, although they clearly can achieve that mean using different workers, provided that they hire only from that cluster. These firms are segregated in the sense that two firms with different mean skill levels hire from different intervals, but within each interval there is complete agglomeration in the same firm types. Each interval roughly corresponds to a hump or a concave zone. Indeed, parts iii and v in proposition 2 imply that, for any firm that employs several types of workers, the average output schedule is locally concave at that firm's mean skill level.[13] The wage schedule paid by the firm is tangent to that hump at precisely the mean of the distribution of human capital over the corresponding interval. Since in equilibrium all workers in that interval must be hired by firms of the same type, their common mean skill level must be equal to the population mean of that cluster. The mass of firms of a given type is then determined so as to ensure that all workers in the corresponding interval are employed. In figure 1*a,* I have represented a unitary zone covering two humps, and in figure 1*b,* I have two consecutive unitary zones corresponding to two consecutive humps. Therefore, not all humps generate a separate cluster since workers in a hump may be absorbed in firms corresponding to a bigger hump.[14] The number of different unitary zones is not necessarily equal to the number of

[13] This property can also be directly proved using the second-order condition of the firm's optimization problem discussed in proposition 1.

[14] This will occur if these workers are not too far away or too numerous, so their participation in the cluster is compatible with the location of that interval's mean skill level in the zone of the other hump.

humps; it all depends on the underlying distribution of workers. However, the number of unitary zones cannot exceed the number of humps.

2. There may be a "hypersegregated," or "assortative," zone in which there is a continuum of firm types each hiring a single type of workers. In this zone, where the output per capita schedule is necessarily convex, the wage schedule exactly matches the output per capita schedule. As illustrated in figures 1*c* and 1*d,* this typically happens when the wage schedules of two consecutive humps cut the average output schedule before cutting each other. In terms of proposition 2, in a hypersegregated zone, each skill level defines a cluster, which is reduced to one point. The density of workers with a skill level between y and $y + dy$, $d\mu(y)$, is employed by a density $d\mu(y)/s$ of firms hiring only type y workers.

3. Finally, there may exist "dual clusters," where two different types of firm hire from the same pool of workers, offering the same wages. This is illustrated in figure 1*e.* The interval of workers hired in a dual cluster roughly covers the two humps corresponding to each firm type. This corresponds to the case in which the interval I_α contains more than one type of firm. In that case, the two firms may not be segregated: some people employed in the high-productivity firm may be less skilled than some people employed in the low-productivity firm. There are many possible equilibrium distributions of workers between the two firms; the only constraints are that everybody is employed and the mean skill of each type of firm corresponds to the tangency point. While in principle there may be more than two different types of firm in I_α, generically this will never arise.

Now that we understand the structure of equilibria, let us proceed and analyze its efficiency and existence.

D. *Efficiency and Existence*

One interesting question is, Is the equilibrium efficient? Intuitively, the answer should be yes. There is no market failure, and the spillovers exerted by people on each other are entirely internalized by firms, and *in fine,* reflected in the wage structure. The next proposition proves that this is indeed the case.

Proposition 3. An allocation is an equilibrium if and only if it maximizes total output.

Proof. First, let us take an allocation that maximizes total output and prove that it may be supported by an equilibrium. Let us normalize s to one to save on notation. Thus, in equilibrium, there is a continuum of firms of total mass equal to one.

Let k, l be two types of firms in that allocation, with corresponding mean skill levels $\bar{y}_k$, $\bar{y}_l$. Let μ_k (respectively μ_l) be the measure representing the distribution of people working in type k (l) firms.

Consider the following reallocation of output: take a distribution of people of total infinitesimal mass ϵ and mean y out of type k firms and allocate it to type l firms in such a way that all type k (l) firms keep having the same common mean.

Then the new mass of type k firms is $\mu_k(I) - \epsilon$ and their new mean skill level is $\bar{y}_k + [\epsilon/\mu_k(I)](\bar{y}_k - y)$. Similarly, the new mass of type l firms is $\mu_l(I) + \epsilon$ and their new mean skill level is $\bar{y}_l + [\epsilon/\mu_l(I)](y - \bar{y}_l)$.

The contribution of these two types of firms to total output is therefore

$$[\mu_k(I) - \epsilon]a\left(\bar{y}_k + \frac{\epsilon}{\mu_k(I)}(\bar{y}_k - y)\right) + [\mu_l(I) + \epsilon]a\left(\bar{y}_l + \frac{\epsilon}{\mu_l(I)}(y - \bar{y}_l)\right).$$

This cannot exceed the original contribution, equal to $\mu_k(I)a(\bar{y}_k) + \mu_l(I)a(\bar{y}_l)$. Using a first-order Taylor expansion, we see that for ϵ small enough this is equivalent to

$$a(\bar{y}_l) + (y - \bar{y}_l)a'(\bar{y}_l) \leq a(\bar{y}_k) + (y - \bar{y}_k)a'(\bar{y}_k). \tag{6}$$

This must hold for any pair (k, l) and any y employed by some type k firm.

Another option is to create a mass ϵ of firms with average skill level exactly equal to y. The contribution of total output to these firms and type k firms must then be equal to

$$[\mu_k(I) - \epsilon]a\left(\bar{y}_k + \frac{\epsilon}{\mu_k(I)}(\bar{y}_k - y)\right) + \epsilon a(y).$$

This cannot exceed the original contribution of type k firms, that is, $\mu_k(I)a(\bar{y}_k)$. Taking again a first-order Taylor expansion, we see that this is equivalent to

$$a(y) \leq a(\bar{y}_k) + (y - \bar{y}_k)a'(\bar{y}_k). \tag{7}$$

It is then not difficult to see that when we take $w(y) = \max_k a(\bar{y}_k) + (y - \bar{y}_k)a'(\bar{y}_k)$, this allocation supports an equilibrium. Equation (6) and the definition of $w(y)$ guarantee that parts i–iv hold in proposition 2, and part v holds because of (7) and parts vi–vii hold by construction.

Conversely, consider an equilibrium. Let $w(y)$ be the wage schedule associated with this equilibrium. Because of free entry, total output Y must be equal to total wages. Consider any other allocation. Let $\tilde{y}_k$ be the mean skill level of firm k in that allocation, $k \in [0, 1]$. Let $d\mu_k$ be the distribution of workers hired by firm k. Then free entry implies $a(\tilde{y}_k) \leq \int w(y)d\mu_k$. Aggregating over firms yields

$$\int_0^1 a(\tilde{y}_k)dk \leq \int \int w(y)d\mu_k dk = \int w(y)d\mu = Y.$$

Thus total output cannot exceed the equilibrium one. Q.E.D.

Since total output is bounded, there will always exist an allocation that maximizes it. Therefore, there always exists an equilibrium, and it is typically "unique."[15]

In particular, there may well be increasing returns to skills over some range, and this does not prevent an equilibrium from existing. Why? The main difference between that model and a standard production function with increasing returns to *scale* is that it is impossible for a firm to take advantage of increasing returns by simply replicating itself. The size of a firm is fixed by assumption, so we have increasing returns to worker quality but not to scale. This has important consequences for equilibrium. Under economies of scale, replication always leads to an increase in profits because revenues increase more than proportionally to size, whereas costs increase only proportionally. Hence an equilibrium with zero profits cannot exist since a deviator, by increasing its size, always gets positive profits. Here, to increase worker quality, we need to hire a different *mix* of people, which bids up the wages of the most skilled, thus defeating, in equilibrium, the firm's attempt to increase its profits. In other words, while in the standard case increasing returns make equilibrium incompatible with perfect competition, this is not the case here.

E. *Examples and Counterexamples*

The framework above applies to a variety of setups, including the one we shall see in the next section. One example is Kremer's (1993) O-ring production function, where a firm's output is given by $\prod_{i=1}^{n} h_i^{\epsilon/n}$, where each firm has n members and h_i is the human capital of individual i, defined as the probability that the corresponding task is performed properly. This can be rewritten as $\exp(\epsilon\bar{y})$, where an individual i's "skill level" is defined as $y_i = \ln h_i$, and $\bar{y}$ is the firm's average skill level. Hence it corresponds to the case in which $a(\cdot)$ is exponential. The average output schedule is then clearly convex throughout the whole interval of values of y, meaning that the equilibrium is always hypersegregated (fig. 2). Firms employ workers of only one type, and there is a continuum of firms indexed by the worker type that they hire, with a density of firms proportional to the density of workers. The wage of any worker type is simply equal to average output in the firms that employ this type of workers.

By contrast, the model above cannot be applied when one cannot write output as a function of a single aggregate of the firm's employees'

[15] In the sense that the clusters and the distribution of firms' average skill levels are uniquely determined.

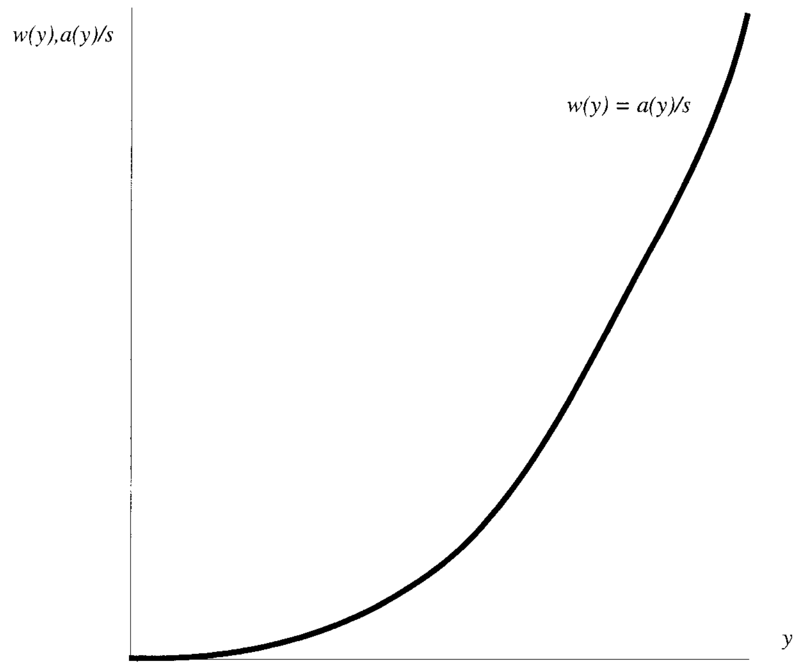

FIG. 2.—Convex case (O-ring)

skill levels. This will be the case whenever there exist asymmetries across workers within the firm, for example, if they perform different tasks. This is the case in the important class of hierarchical models in which the typical prediction is that the most talented workers will be assigned to supervisory tasks, with the span of control increasing with talent (see Calvo and Wellisz 1979; Rosen 1982). In this case, total output will depend not only on the average skill level within the firm but also on some higher moments. Consider the case, for example, in which a fraction θ of the firm's employees are assigned to supervisory tasks, the remaining workers being assigned to production tasks. Assume that output is proportional to $\bar{y}_s\bar{y}_p$, where $\bar{y}_s$ is the average skill level of supervisors and $\bar{y}_p$ the average skill level of production workers. Then, if $\theta < \frac{1}{2}$ and if the firm hires from a uniform distribution over $[\bar{y}_k - \sigma, \bar{y}_k + \sigma]$, one can show that if the best workers are allocated to supervisory tasks, output can be written as $\bar{y} + \sigma(1 - 2\theta)\bar{y}_k - \sigma^2\theta(1 - \theta)$. Thus it is now a function of both average skill and its standard deviation, $a(\bar{y}, \sigma)$.

III. Application: Ideas, Networks, and Information Technologies

Let us now apply the framework above to the issue that was mentioned in the Introduction, namely, the effect of advances in information technology on the distribution of income. We shall first derive the properties satisfied by the $a(\cdot)$ function corresponding to that specific application and then, in the next section, study how the distribution of income changes when the parameter that characterizes the efficiency of information technology shifts; we shall also do another comparative statics exercise, looking at the impact of the supply of skills.

A. Basic Setup

As above, the economy is populated by a continuum of agents, whose skill is given by a number $y \in I = [y_0, y_1]$. The distribution of y is still given by measure μ and the total mass of workers still normalized to one. An agent with skill y has a physical productivity equal to y and at the same time an ability to have ideas h (referred to as "creativity"). Assume that $h = cy + b$. Therefore, productivity and creativity are perfectly correlated.

As above, each firm hires a mass s of workers. Firms are two things. First, they are firms; that is, they hire people and sell output, maximizing profits. Second, they are networks; that is, they are the space over which informational spillovers take place. An idea can be applied at no cost to all the workers of the firm. Firms and networks coincide, so that spillovers are internal to the firm. As already mentioned, this coincidence is what we would expect to arise in equilibrium following a Coasian argument (Coase 1995). It might be interesting, however, to allow for the frontiers of firms to differ from those of networks.

Hence, the size of the firm s is also the range over which ideas can be spread at no cost. Improvements in information technology can therefore be represented by an increase in s: the larger the network, the larger the number of workers who can benefit from a given idea.

The timing of events is as follows:

1. Firms freely enter the market.
2. The labor market operates and hiring takes place. Firms maximize their expected profit. This yields an equilibrium wage schedule $w(y)$. Note that wages are set prior to production.
3. A finite number N of randomly drawn workers have an idea. Each firm makes use of its workers' ideas to improve production. Production takes place and wages are paid. Assume that N is large, so that each firm considers that among its employees there will be exactly $n = sN$ workers who will have an idea.

I now describe the idea process and the production process. Each idea is represented by a random variable $z \in [0, 1]$. It is drawn from a cumulated distribution function $Q(z, h)$, where h is the creativity of the worker who has the idea. More creative people have better ideas. To represent that, we assume that as h increases, the distribution of z changes in such a way that the new distribution dominates the old one in the first-order stochastic dominance sense. That is, assume $\partial Q/\partial h < 0$.

A very convenient specification is $Q(z, h) = F(z)^h$. It guarantees that an increase in h shifts Q downward while preserving the boundary conditions $Q(1, h) = 1$ and $Q(0, h) = 0$. Furthermore, it has very convenient analytical properties.

Then, if an idea z is applied to a given network, total output in that network is given by

$$u(z)\int_0^s y_{ik}di = u(z)s\bar{y}_k,$$

where u is an increasing function of z; as above, y_{ik} denotes the skill level of worker i in firm k; and $\bar{y}_k$ is the average productivity of network k's workers. Thus the production function is linear in individual productivity, and any employee's idea can be applied by all workers in the network, so that it acts as a multiplicative shift to the production function. Below I shall correctly refer to $u(z)$ as the firm's *total factor productivity.*

Assume that ideas are not cumulative; that is, any network will apply the best idea among its employees. An idea is useless even if it is only marginally worse than another one that can be applied over the same network; productivities add up so that if a worker's productivity is marginally lower than another's, so are their marginal products.

The number of ideas in a network is n. The cumulative probability distribution of the best idea is then, conditional on the set of workers who have ideas,

$$\Phi_N(z, \{h_i\}) = \prod_{i=1}^{n} Q(z, h_i). \tag{8}$$

Using our specification, we can see that this is equal to $F(z)^{\Sigma_i h_i}$. When the law of large numbers is applied, this is approximately equal to $F(z)^{n\bar{h}_k}$, where $\bar{h}_k$ is the firm's average creativity.[16] Therefore, the distribution of the best idea depends only on the firm's average creativity,

[16] This is clearly an approximation since the realization of $\sum_i h_i$ will in general differ from $n\bar{h}_k$.

irrespective of how it is distributed across workers. This makes it possible to apply the previous section's results.

B. *The Output Schedule*

To do so, all we have to do is to compute how a firm's expected output depends on its workers' average quality. We get

$$Eu(z)s\bar{y}_k = \left[\int_{z=0}^{z=1} u(z)d(F(z)^{n\bar{h}_k})\right]s\bar{y}_k$$

$$= \left[u(1) - \int_0^1 u'(z)F(z)^{n\bar{h}_k}dz\right]s\bar{y}_k,$$

where we have integrated by parts. It can be shown that the right-hand side is *concave* in $\bar{h}_k$: the marginal return to creativity is decreasing, reflecting the assumption that ideas are bounded. An infinitely creative person is simply sure to have the best idea, but that best idea boosts output by only a finite amount.

The following proposition then characterizes how expected output depends on average skill.

PROPOSITION 4. A network k's expected output can be written as

$$E_k(Y) = \phi(s\bar{h}_k)s\bar{y}_k, \tag{9}$$

where $\phi(\cdot)$ is increasing and concave.

Proof. We have

$$\phi(s\bar{h}_k) = u(\bar{z}) - \int u'(z)F(z)^{n\bar{h}_k}dz.$$

As $n = sN$, this is a function of $s\bar{h}_k$. Taking derivatives, we get

$$\phi' = -N\int u'(z)F(z)^{n\bar{h}_k}\ln F(z)dz > 0 \tag{10}$$

and

$$\phi'' = -N^2\int u'(z)F(z)^{n\bar{h}_k}[\ln F(z)]^2dz < 0. \tag{11}$$

Q.E.D.

Because h is a linear function of y, expected output is a function of the firm's average skill only. That is,

$$E_k(Y) = \phi(s(c\bar{y}_k + b))s\bar{y}_k = a(\bar{y}_k).$$

This equation defines total output as a function of average worker quality so that the previous section's analysis applies. If creativity and productivity are positively correlated ($c > 0$), then this function exhibits increasing returns to worker quality: $ya'(y) > a(y)$, although marginal returns can be either decreasing or increasing. An increase in productivity by a factor λ increases output by more than λ, because these people will also have better ideas. In the less intuitive case in which $c < 0$, there are decreasing returns.

Note that $a(\cdot)$ would be concave if people differed in their creativity h but not in their productivity y. That is, skill is now redefined as creativity only, and total output is now equal to $a(h) = \phi(h)$, where ϕ is concave, as we have seen. This concavity reflects the fact that our production process is in some sense the opposite of the O-ring one. In the O-ring case, output is determined by the worst performance, which leads to a convex dependence in average worker quality.[17] Here it is determined by the best idea, which leads to a concave dependence of total factor productivity on average worker quality. Therefore, if y is the same for all workers, according to proposition 2, there will be a single type of network in equilibrium, and this type's average skill level will be equal to the population average, as in figure 3. Thus, despite the fact that it is only the best idea that is being applied, the positive spillovers associated with the spreading of ideas, to the extent that ideas are bounded, lead to the integration of heterogeneous people into the same networks.

Things are more complicated, however, in the more interesting case in which creativity is correlated with productivity. In such a case, the total output function is given by $a(y) = \phi(s(cy + b))sy$ and may be either concave or convex. In principle, thus, many configurations may arise, including dual clusters. However, a plausible configuration is an S-shape form for $a(y)$, that is, convexity for low values of y and concavity for high values of y. This would be the case, for example, if there were only two possible ideas, a good one, z_1, and a bad one, z_0. Let P be the probability of having a bad idea. Then expected output, as a function of average productivity y, would be equal to

$$a(y) = sy\{u(z_1) - P^{s(cy+b)/N}[u(z_1) - u(z_0)]\}. \tag{12}$$

It is straightforward to check that this is first convex and then concave as y rises. More generally, average output can be written as an integral, over various ideas, of expressions such as (12). So it is plausible that it

[17] To get back to the O-ring case, just change the sign of $u'(z)$ in eqq. (10) and (11): if u is decreasing in z, it is the worst idea that we are picking up, not the best one. In that case it is clear that ϕ is convex rather than concave—implying that convexity is a property inherent in the worst performer's dominance and does not depend on the specific functional form chosen by Kremer. This property and the reverse one discussed in the text are clearly related to the convexity properties of rank-order statistics.

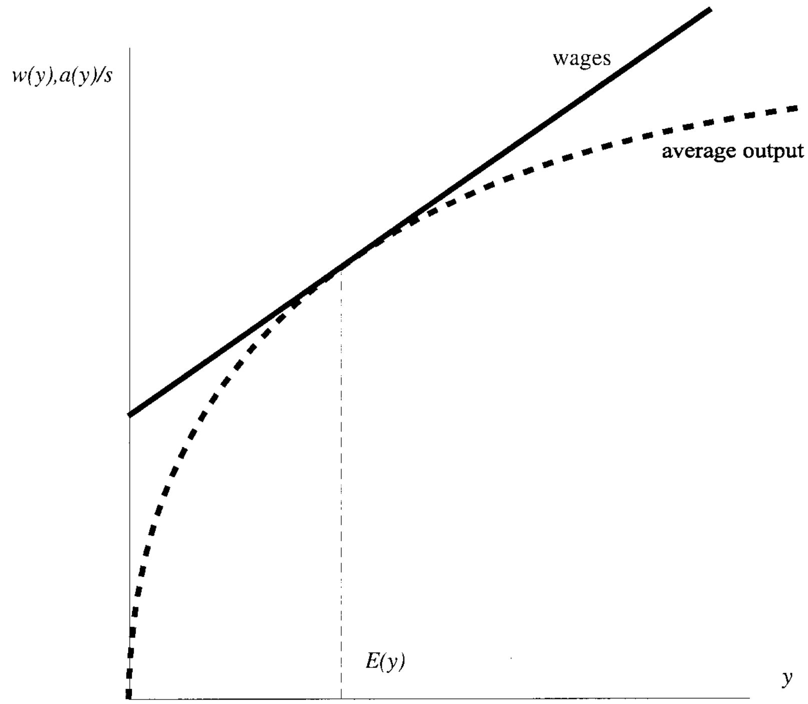

FIG. 3.—Concave case (same productivities)

will have the same S-shape. This pattern reflects the fact that at low levels of worker quality, the best idea is very sensitive to the average skill of the workers employed by the firm. By contrast, at high levels of worker quality, one is almost certain to get the best idea, so that expected output increases almost linearly with average worker quality. In between, there is a zone in which the elasticity of average output to skills falls as one quickly reaches the zone in which the idea potential is almost exhausted.

As illustrated in figure 4, in equilibrium the labor market typically segregates itself into a hypersegregated zone of low-skill workers and a unitary network type of high-skill workers.[18] Hence, people at the bottom of the income distribution work in homogeneous firms employing similarly skilled workers, and those at the top work in diversified networks in which there is, in some sense, a complementarity between the most

[18] If the line tangent to the average output schedule at the population average of y intersects the vertical axis above the average output schedule, then there is a single cluster that also covers the zone in which the average output schedule is convex.

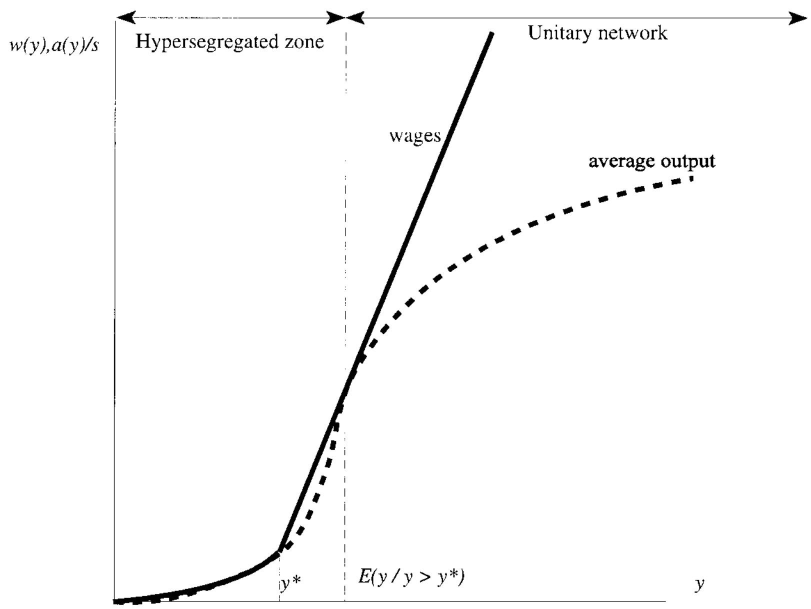

FIG. 4.—S-shaped case

able workers, whose value is in the ideas they have, and the least able ones, whose value is mostly in their contribution to output.

As the nonlinearity in production comes entirely from the existence of ideas, it is interesting to reinterpret the wage equation in the following way. Since $a(y) = \phi(s(ay + b))sy$, we can apply proposition 1 to compute wages in firm k, getting

$$w(y) = \phi(s\bar{h}_k)y + \phi'(s\bar{h}_k)s\bar{y}_k(h - \bar{h}_k). \tag{13}$$

The interpretation I gave of equation (1) still holds. But there is another, interesting, interpretation that can be given to (13). The first term, $\phi(s\bar{h}_k)y$, is the worker's return to his productivity, which is simply equal to his physical product in his firm, with his intellectual contributions ignored. The second term, $\phi'(s\bar{h}_k)s\bar{y}_k(h - \bar{h}_k)$, is a "bonus" paid for creativity. Because of the zero profit condition, this bonus has to average to zero, and we can see that it is proportional to the *deviation* between this individual's creativity and average creativity in the firm. Thus people earn more (less) than their physical marginal product, depending on whether they are more (less) creative than their firm's average. This way, low-creativity people are penalized for the fact that they occupy a job that might be held instead by somebody equally productive but with better ideas. Hence, the wage structure rewards *absolute* productivity but *relative* creativity; the basis of comparison is average creativity in the same network.

IV. The Impact of Technical Change on Segregation and Income Distribution

In this section I shall discuss how technical change—broadly defined—affects the distribution of income in the model. Assume that $\phi(\cdot)$ is such that the average output schedule $a(\cdot)$ has the S-shape studied in Section III, so that worker assignment is summarized by figure 4. More specifically, assume that $\phi(\cdot)$ has the convenient exponential form

$$\phi(x) = A(1 - \beta e^{-\alpha x}). \tag{14}$$

This specification is exact in the special case in which there are only two possible ideas. It yields the desired S-shape for $a(\cdot)$.

A. *An Improvement in the Quality of the Workforce*

Let us start with the simplest experiment, namely, a change in the underlying distribution of skills. The average output schedule is independent of that distribution and therefore does not shift. Equilibrium changes, however, because the initial distribution of network types no longer matches the distribution of agents. Figure 5 illustrates what happens when skills improve. The mean skill level of the initial unitary network increases beyond its original value. To restore equilibrium, unitary networks must have a greater average skill level, which, given the local concavity of the average output schedule, implies that they hire a more diversified set of people. Thus the networks that hire the top of the distribution of income broaden the range of skills they are hiring from. As ideas are "cheaper," they are more willing to hire low-quality people. The least skilled, however, remain in hypersegregated networks and gain nothing from the change, since they do not cooperate with skilled workers in production. Next, there exists a range of people who were previously confined to hypersegregated networks and can now join a unitary network. They enjoy a large wage gain. So do people previously employed in the unitary network but relatively unskilled. Finally, the most skilled experience a wage drop as the value of creativity is reduced. Overall, society is less segregated and more equal, but there remains a mass of poor workers who do not benefit from the change because segregation prevents them from benefiting from the greater flow of

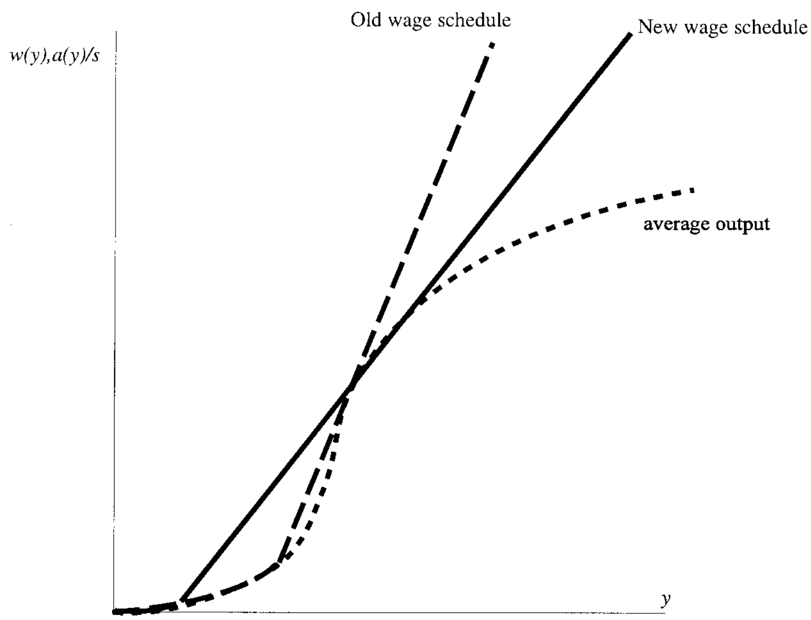

FIG. 5.—Increase in the quality of the workforce

good ideas. Indeed, the gap between that group of people and the average income of the rest of the population has increased.[19]

B. An Increase in Network Size s

An increase in network size s means that the range of spillovers is greater. Formally, it is somewhat similar to an improvement in the quality of the workforce since in both cases a network recruiting any given set of workers will have better ideas. However, the analysis is slightly more complex here, as the average output schedule shifts.

We can interpret such an increase in s as an improvement in information technology. It captures the fact that such improvements increase

[19] The model could be extended to take into account the impact of international trade on inequality. The previous empirical literature has dismissed it as an explanation for the rise in inequality because relative prices have not moved in the direction predicted by that hypothesis (see, e.g., Lawrence and Slaughter 1993). One may speculate that if ideas as well as goods are allowed to cross frontiers, the establishment of transnational networks could generate a positive effect of international trade on inequality in the country endowed with high skills, even though goods prices would not change. The exercise would be similar to the reverse of what has just been discussed. See Saint-Paul (1999) for an informal discussion.

the span over which ideas (more generally, intangible objects) can be spread. On the other hand, this clearly misses some other aspects of information technology (such as improvement in computational ability or transfer of tasks to computers). Hence, the model allows us to insulate one particular effect, but some caution is needed in interpreting the results.

The following proposition tells us that the response of inequality to network size is typically hump-shaped.[20]

Proposition 5. Assume $y_0 > 0$. Let INEQ be any inequality measure, expressed as

$$\text{INEQ} = 1 - \psi^{-1}\left(\int \psi\left(\frac{w(y)}{\bar{a}}\right) d\mu\right),$$

where $\psi(\cdot)$ is a concave, increasing function and $\bar{a}$ is average output per worker (also equal to the average wage). Then there exist s_0 and s_1 such that (i)

$$s \leq s_0 \Rightarrow \frac{\partial \text{INEQ}}{\partial s} > 0$$

and (ii)

$$s \geq s_1 \Rightarrow \frac{\partial \text{INEQ}}{\partial s} < 0.$$

Therefore, inequality rises with s for s small enough but falls for s large enough.

Proof. See the Appendix.

Therefore, an increase in s does not necessarily increase inequality. This is so only below a certain level of technology. This nonmonotonicity is the result of conflicting effects. On the one hand, the most creative workers can spread their ideas over a larger network of economic activity when s increases. This tends to increase their wages relative to others. On the other hand, when networks get larger, the ideas of a given, highly creative worker are less valuable because it is more likely that somebody else in the network would have had an idea almost as good. That is, past some large network sizes, superstars end up competing against each other, which eventually pushes down their wages relative to others. The inegalitarian effects of network size are further mitigated by the fact that in a unitary network, the productivity of the least skilled is enhanced by the application of the most talented workers' ideas.

For s small, all or most of the economy is hypersegregated, and an

[20] That is, I show that an increase in s increases inequality at low values of s and reduces it at high values of s. However, several modes may exist in the middle.

increase in s raises output by much more in high-skill firms than in low-skill ones. That is, the scale effects of network size dominate, whereas the other effects are weak; most workers do not benefit from spillovers from more skilled workers, and skilled workers are not exposed to competition from slightly less creative workers within the same network. Consequently, inequality increases with s. At higher values of s, however, most people work in unitary networks, where the least skilled benefit from the ideas of the most skilled and the most creative types compete against each other. The inequality-reducing effects of network size then dominate.

Proposition 5 characterizes what happens to aggregate measures of inequality. What happens to inequality *within* the unitary network? As it is a network with high-skill workers, proposition 5 may not imply that inequality increases within that network. More precisely, the fact that a unitary network may arise only at a mean firm skill level such that $a(y)$ is locally concave puts restrictions on the effect of s. In fact, one can show that if $b < 0$, there exists a zone in which the slope of the wage schedule offered by the unitary network may increase. In this case, the model predicts that intrafirm inequality (as measured by that slope) may increase over some range. By contrast, if $b > 0$, this cannot happen.[21] The unitary network is then always in the zone in which it has enough ideas so that an increase in network size lowers the return to skills.

Figure 6 summarizes these results. On the horizontal axis there is the average skill level of a firm y. On the vertical axis there is network size s. The first frontier, *HH*, is the frontier between hypersegregated and unitary firms.[22] The second one, *II*, is the frontier between the zone in which local inequality, as measured by $w'(\bar{y})$, increases with s and the zone in which it is reduced.[23] If $b < 0$, *II* is above *HH*, so that inequality may increase within a unitary network as s rises. If $b > 0$, *II* is below *HH*, so this cannot happen.

Another interesting question is whether an increase in network size could actually make some people poorer. Here, this can happen only at the top of the distribution of income, not at the bottom. Thus, despite the possible increase in the steepness of the wage schedule offered by

[21] Let us prove this claim. The unitary network has a $\bar{y}_k$ such that $a(\cdot)$ is locally concave or, equivalently, $2\phi' + cs\bar{y}_k\phi'' < 0$, which may be rewritten as $\bar{y}_k > 2/\alpha cs$. Computing the expression $(2c\bar{y}_k + b) - \alpha(c\bar{y}_k + b)cs\bar{y}_k$ at $\bar{y}_k = 2/\alpha cs$, we simply get $-b$. So, if $b > 0$, $\partial w'(y;\ s)/\partial s < 0$ over the whole range of values of $\bar{y}_k$, where the unitary network may arise; if $b < 0$, for some distributions μ, one may see a unitary network with a value of $\bar{y}_k$ such that $\partial w'(y;\ s)/\partial s > 0$.

[22] It is defined by $s = 2/\alpha cy$. Note 21 implies that above that locus, the average output schedule is locally concave at y, a necessary condition for a firm with average skill y to define a unitary zone. Conversely, if that condition holds, we can always pick up a distribution μ such that the unitary zone has a mean precisely equal to y.

[23] This frontier is defined by $s = (2cy + b)/[\alpha(cy + b)cy]$.

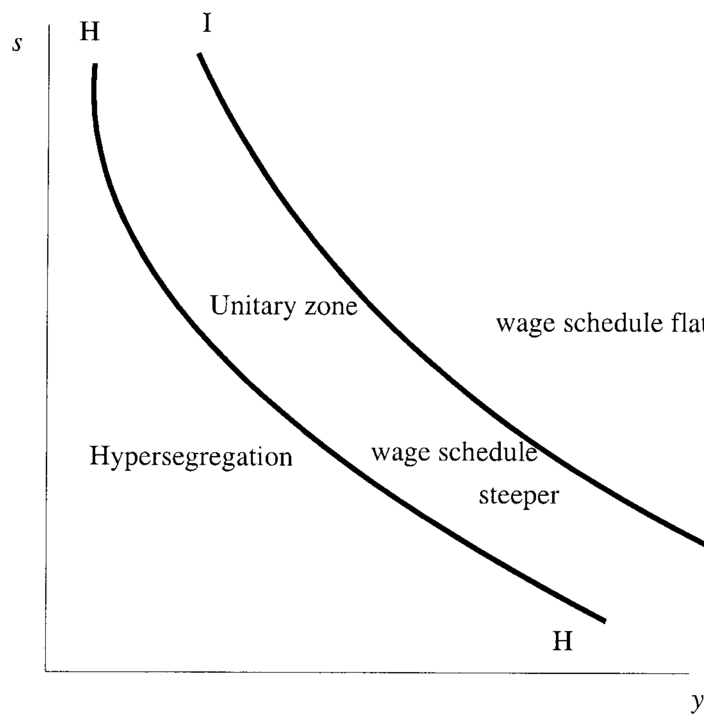

FIG. 6.—Effect of an increase in s on the slope of the wage schedule in the unitary zone ($b < 0$).

the unitary network, technical progress in the form of increased network size cannot impoverish the poorest. This is summarized in the following "anti-Marxist" result.

PROPOSITION 6. *No impoverishment.*—Remember that y_0 is the lowest value of y. Then $\partial w(y_0)/\partial s > 0$.

Proof. If y is in a hypersegregated zone, wages in a firm employing type y workers are just equal to output per capita, which obviously increases with s since these people, when more numerous, will share better ideas. If y_0 is included in the unitary network, then this network must cover the whole interval of skill levels. Then its mean skill level is equal to the population one, $\bar{y}$, which is unaffected by an increase in s. Next, we can compute

$$\frac{\partial w(y)}{\partial s} = y\bar{h}_k\phi'(s\bar{h}_k) + \phi'(s\bar{h}_k)\bar{y}_k(h - \bar{h}_k) + \phi''(s\bar{h}_k)s\bar{y}_k\bar{h}_k(h - \bar{h}_k).$$

Because the average output schedule must be concave at $y = \bar{y}_k$, we must have

$$2\phi'(s\bar{h}_k) + sa\bar{y}_k\phi''(s\bar{h}_k) \leq 0,$$

implying, as long as $h < \bar{h}_k$,

$$\frac{\partial w(y)}{\partial s} \geq [h\bar{y}_k + \bar{h}_k(\bar{y}_k - y)]\phi' > 0.$$

Q.E.D.

If one is willing to interpret an increase in s as an improvement in information technology, the results above are at variance with some popular pessimism, which holds that it is inegalitarian and harmful to the poorest. Our results suggest that there are inequality-reducing effects as well and that these effects are likely to prevail eventually. They highlight the difference between a "normal" labor market, where people cooperate to create wealth, and a market such as the one for "superstars," where they crowd each other out in competition for a fixed prize.

V. Extensions

In this section I briefly discuss how relaxing some of the assumptions of the models may affect the results. The first two extensions are concerned with the specific ideas/networks application, and the last ones deal with the more general framework.

A. More than One Good

The interpretation of the results in terms of income distribution may be slightly more complex if there is more than one good in the economy. Above we considered that people produced a single homogeneous good. We were therefore looking at changes in the distribution of wages expressed in terms of that good (call it good A) rather than in terms of welfare. If there are other goods and if these goods are less intensive in information technology, then an improvement in information technology will lower the relative price of good A.[24] It is now possible for those workers who gain little when their wage is expressed in terms of good A to actually lose in welfare terms. Consider, for example, an improvement in the economy's average skill level. I have shown that if the least skilled workers were in a hypersegregated zone, their wages—in terms of good A—did not change. If the relative price of good A falls, then their real wage will actually fall, and reallocation of these workers to other sectors may only partially mitigate that. The same thing may

[24] See Cohen and Saint-Paul (1994) for an analysis of the implications of asymmetrical technical change on wages and employment dynamics.

happen in the case of an increase in s if the rise in the poorest's wage in terms of good A is small relative to the fall in the price of that good.

B. Unbounded Ideas

The results established in Section IV depend on the assumption that output is S-shaped in average skill, which came from the concavity of $\phi(\cdot)$. This occurred because ideas were bounded. However, if ideas are unbounded, the output schedule may be convex throughout.[25] The zone in which inequality falls with network size may then disappear.

C. Endogenous Firm Size

My results clearly depend on the assumption of a fixed firm size; however, the framework can be extended to a variable size. Going back to the general model, assume that total output is a function of both the number of workers and their average skill. Let us denote such a function by $a(\bar{y}, s)$. Then one has to add to the analysis the determination of optimal firm size. One can check by extending the proof of proposition 1 that the corresponding first-order condition is

$$\frac{\partial a}{\partial s}(\bar{y}_k, s) = \frac{a(\bar{y}_k, s)}{s} - \frac{1}{s}\bar{y}\frac{\partial a}{\partial \bar{y}}(\bar{y}_k, s). \tag{15}$$

This equation determines the firm's optimal size as a function of its average level of human capital. Note that the rest of the analysis is unchanged; in particular, the wage schedule offered by a firm still obeys (1), so that the basic results are still likely to hold, if one allows for variations in firm size.

The right-hand side of (15) is nothing but $-\lambda_k$, the opposite of the shadow price of bodies. Hence, if size is determined optimally, the shadow price of bodies will be equal to the marginal cost of increasing the size of the firm, $-\partial a(\bar{y}_k, s)/\partial s$.[26]

Equation (15) implies that firms that have the same type, that is, the same average skill level, will also have the same size, but that size will in general differ across firms. It sounds reasonable to think that firms with a greater average skill level will be larger; this is what happens in

[25] One may formalize unbounded ideas by assuming that the $u(\cdot)$ function has a vertical asymptote at $z = 1$. In that case, integration by parts is no longer possible, implying that one cannot prove that $\phi(\cdot)$ is concave.

[26] Under increasing returns to worker quality, the firm will end up in a zone in which $\partial a/\partial s < 0$, meaning that an increase in the number of employees will reduce net output. For this to occur, there must exist a cost of increasing the size of the firm, say a setup or investment cost, which ends up increasing more than the direct contribution of the extra employees to output.

Kremer (1993) and is in accordance with the often-found result that more talented people have a greater span of control (as in Rosen [1982]). However, it all depends on cross derivatives and on whether positive spillovers are stronger in firms with more talented people. If this is true, then the right-hand side of (15) will become more negative as $\bar{y}_k$ gets larger, implying that these firms will have a bigger marginal cost of size and therefore will typically[27] be larger.

In the case of the idea/networks model, size can be made variable by deducting a convex setup cost from average output, so that we would have

$$a(\bar{y},\ s) \equiv \phi(s(c\bar{y} + b))s\bar{y} - \frac{\gamma s^2}{2}.$$

One can then easily check that the first-order condition determining optimal firm size is equivalent to

$$\frac{\gamma s}{2} = \phi(s(c\bar{y} + b))\bar{y} + 2sc\bar{y}^2\phi'(s(c\bar{y} + b)).$$

It is very likely that this defines a positive relationship between s and $\bar{y}$, but if ϕ'' is negative enough, this may not happen.

D. Endogenous Firm Capital

The model can also be extended by introducing endogenous firm capital. To the extent that capital and skills are complementary, in equilibrium the high-skill firm will work with more capital. This makes it more likely that the average output schedule is convex, thus reinforcing the scope for segregation and typically increasing inequality. Assume, for example, that total firm output is $m(k)a(\bar{y})$, where $m(\cdot)$ is increasing and concave and k is the firm's capital. Let r be the cost of capital. Then a firm with average skill $\bar{y}$ will set k such that $m'(k)a(\bar{y}) = r$. This defines k as an increasing function of $\bar{y}$, $k(\bar{y})$. Firms with a higher $\bar{y}$ will then have a higher capital/labor ratio. This tends to increase the returns to skill. That is, the slope of the average output (net of capital costs) schedule is

$$\frac{d\{[m(k)a(\bar{y}) - rk]/s\}}{dy} = \frac{m(k)a'(\bar{y})}{s}.$$

Its second derivative is

[27] By "typically" I actually mean provided that the cross derivative $\partial^2 a/\partial\bar{y}_k\partial s$ is not too negative.

$$\frac{m(k)a''(\bar{y}) + m'(k)a'(\bar{y})(dk/d\bar{y})}{s},$$

which is greater than if capital were the same across firms. Formally, the analysis is the same as in Section II, with the output function replaced by its reduced form net of capital costs: $m(k(\bar{y}))a(\bar{y}) - rk(\bar{y})$.

VI. Conclusion

The framework developed here is flexible and general enough to be applicable to the analysis of people's assignment and segregation in a variety of settings. Potential examples include schools, neighborhoods and ghettoes, and social networks.

On the other hand, as any model, it has some limitations. In particular, as we have seen, in order to capture the hierarchical aspects of the firm's organization, one would need a more flexible specification of spillovers that would make use of higher moments of the intrafirm distribution of skills.

The application to information technology has yielded a variety of insights; let me insist on two of them, which I believe are most relevant. First, increases in network size have both positive and negative effects on inequality. Under reasonable assumptions, one finds that inequality typically rises and then falls as network size increases. This is somewhat consistent with the observation that after years of an upward trend, inequality seems to be falling again in the 1990s. Second, in many cases the bottom of the distribution of income was locked into a "hypersegregated" zone; that is, they were not interacting with high-skill workers in production. This prevented them from benefiting from an increase in the supply of skilled workers, and they were likely to benefit less than other workers from an improvement in information technology. Their real wage could even fall if one considered a multigood world.

Appendix A

Proof of Proposition 5

The derivative of the measure of inequality with respect to s has the same sign as

$$-\int \psi'\left(\frac{w(y)}{\bar{a}}\right)\left[\frac{\partial}{\partial s}w(y;\ s)\,\bar{a}(s) - w(y;\ s)\,\frac{\partial \bar{a}(s)}{\partial s}\right]d\mu, \tag{A1}$$

where the dependence in s has been made explicit for the sake of clarity. The term in brackets defines a function $H(y;\ s)$, which, since $\bar{a} = \int w(y)d\mu$, satisfies

$$\int H(y)d\mu = 0. \tag{A2}$$

The average output schedule is defined by (9), and ϕ is defined by (14). One can readily see that the second derivative of the average output schedule has the same sign as $2 - \alpha scy$. Therefore, it is convex if $s < 2/\alpha c y_1$, concave if $s > 2/\alpha c y_0$, and S-shaped in the middle. As s increases, the interval in which it is concave gradually increases from an empty set to the whole $[y_0, y_1]$ interval.

If $s < 2/\alpha c y_1$, the whole economy is hypersegregated. Equation (13) then implies that wages are given by

$$w(y;\ s) = Ay[1 - \beta e^{-\alpha s(cy+b)}],$$

implying

$$\frac{\partial}{\partial s} w(y;\ s) = A\alpha\beta y(cy + b)e^{-\alpha s(cy+b)}.$$

Consequently, $H(y)$ has the same sign as

$$\bar{a}(s)\beta\alpha(cy + b)e^{-\alpha s(cy+b)} - \left[\frac{\partial}{\partial s}\bar{a}(s)\right][1 - \beta e^{-\alpha s(cy+b)}].$$

Inspection of this formula reveals that, as both $\bar{a}(s)$ and $\partial\bar{a}(s)/\partial s$ are bounded away from zero and infinity as s goes to zero, for s small enough, say smaller than some $\tilde{s}$, it is increasing with y over $[y_0, y_1]$ (the cy term dominates). Consequently, since (A2) must hold, $H(y;\ s)$ is first negative and then positive when y increases. Since $\psi'(w(y)/a)$ is falling with y, this in turn implies that the expression in (A1) is positive. This proves claim i in proposition 5: for s smaller than $\min(\tilde{s},\ 2/\alpha c y_1)$, inequality increases with s.

If $s > 2/\alpha c y_0$, the average output schedule is concave throughout, and there is a single type of firm, which hires workers whose average skill is equal to $\bar{y}$, the economywide average skill level. To show that inequality falls with s, I simply show that the slope of the wage schedule falls with s for s large enough. Given that $\bar{a}(s)$ unambiguously increases with s, these two facts imply that in that zone, INEQ unambiguously falls with s.

Equations (13) and (14) imply that in that zone, wages are given by

$$w(y) = A[1 - \beta e^{-\alpha s(c\bar{y}+b)}]y + A\alpha s\bar{y}c(y - \bar{y})\beta e^{-\alpha s(c\bar{y}+b)}.$$

The slope of the wage schedule therefore is

$$w'(y) = A[1 - \beta e^{-\alpha s(c\bar{y}+b)}] + A\alpha s\bar{y}c\beta e^{-\alpha s(c\bar{y}+b)}.$$

When s increases, this slope increases by an amount equal to

$$\frac{\partial}{\partial s} w'(y) = A\beta\alpha(2c\bar{y} + b)e^{-\alpha s(c\bar{y}+b)} - A\alpha^2 sc\bar{y}\beta(c\bar{y} + b)e^{-\alpha s(c\bar{y}+b)}.$$

This expression is negative for s larger than $\hat{s} = (2c\bar{y} + b)/\alpha c\bar{y}(c\bar{y} + b)$. Consequently, INEQ falls with s for $s \geq \max(2/\alpha c y_0,\ \hat{s})$. Q.E.D.

If $y_0 = 0$, there is always a convex portion in the average output schedule. The last part of the proof is no longer valid, strictly speaking. However, if μ has no mass in the neighborhood of $y = 0$, then the proof above can be extended, since for s large enough the convex zone accounts for an arbitrarily small part of the inequality measure. A similar line of reasoning can extend the proof to the case in which y_1 is infinite.

References

Aghion, Philippe, and Howitt, Peter. *Endogenous Growth Theory.* Cambridge, Mass.: MIT Press, 1998.

Akerlof, George A. "Structural Unemployment in a Neoclassical Framework." *J.P.E.* 77 (May/June 1969): 399–407.

———. "Jobs as Dam Sites." *Rev. Econ. Studies* 48 (January 1981): 37–49.

Becker, Gary S. "A Theory of Marriage: Part I." *J.P.E.* 81 (July/August 1973): 813–46.

Bénabou, Roland. "Workings of a City: Location, Education, and Production." *Q.J.E.* 108 (August 1993): 619–52.

———. "Heterogeneity, Stratification, and Growth: Macroeconomic Implications of Community Structure and School Finance." *A.E.R.* 86 (June 1996): 584–609.

Bound, John, and Johnson, George. "Changes in the Structure of Wages in the 1980's: An Evaluation of Alternative Explanations." *A.E.R.* 82 (June 1992): 371–92.

Calvo, Guillermo A., and Wellisz, Stanislaw. "Hierarchy, Ability, and Income Distribution." *J.P.E.* 87, no. 5, pt. 1 (October 1979): 991–1010.

Coase, Ronald H. "The Nature of the Firm." In *The Legacy of Ronald Coase in Economic Analysis,* vol. 1, edited by Steven G. Medema. Aldershot, U.K.: Elgar, 1995.

Cohen, Daniel, and Saint-Paul, Gilles. "Uneven Technical Progress and Job Destruction." Working paper. London: Centre Econ. Policy Res., 1994.

DiNardo, John E., and Pischke, Jorn-Steffen. "The Returns to Computer Use Revisited: Have Pencils Changed the Wage Structure Too?" *Q.J.E.* 112 (February 1997): 291–303.

Epple, Dennis, and Romano, Richard E. "Competition between Private and Public Schools, Vouchers, and Peer-Group Effects." *A.E.R.* 88 (March 1998): 33–62.

Fernandez, Raquel, and Rogerson, Richard. "Income Distribution, Communities, and the Quality of Public Education." *Q.J.E.* 111 (February 1996): 135–64.

Frank, Robert H., and Cook, Philip J. *The Winner-Take-All Society.* New York: Simon and Schuster, 1995.

Garicano, Luis. "Knowledge-Based Hierarchies." Ph. D. dissertation, Univ. Chicago, 1998.

Grossman, Gene M., and Helpman, Elhanan. *Innovation and Growth in the Global Economy.* Cambridge, Mass.: MIT Press, 1991.

Heckman, James J., and Scheinkman, José A. "The Importance of Bundling in a Gorman-Lancaster Model of Earnings." *Rev. Econ. Studies* 54 (April 1987): 243–55.

Juhn, Chinhui; Murphy, Kevin M.; and Pierce, Brooks. "Wage Inequality and the Rise in Returns to Skill." *J.P.E.* 101 (June 1993): 410–42.

Katz, Lawrence F., and Murphy, Kevin M. "Changes in Relative Wages, 1963–1987: Supply and Demand Factors." *Q.J.E.* 107 (February 1992): 35–78.

Kremer, Michael. "The O-Ring Theory of Economic Development." *Q.J.E.* 108 (August 1993): 551–75.

Krueger, Alan B. "How Computers Have Changed the Wage Structure: Evidence from Microdata, 1984–1989." *Q.J.E.* 108 (February 1993): 33–60.

Lawrence, Robert Z., and Slaughter, Matthew J. "International Trade and American Wages in the 1980s: Giant Sucking Sound or Small Hiccup?" *Brookings Papers Econ. Activity: Microeconomics* (1993), pp. 161–210.

Levy, Frank, and Murnane, Richard J. "U.S. Earnings Levels and Earnings Inequality: A Review of Recent Trends and Proposed Explanations." *J. Econ. Literature* 30 (September 1992): 1333–81.

Lucas, Robert E., Jr. "On the Mechanics of Economic Development." *J. Monetary Econ.* 22 (July 1988): 3–42.

Mandelbrot, Benoit. "Paretian Distributions and Income Maximization." *Q.J.E.* 76 (February 1962): 57–85.

Murphy, Kevin M.; Shleifer, Andrei; and Vishny, Robert W. "The Allocation of Talent: Implications for Growth." *Q.J.E.* 106 (May 1991): 503–30

Rioux, Laurence. "L'importance des qualifications dans les nouvelles organisations: quelles conséquences pour le marché du travail?" Ph. D. dissertation, Ecole des Hautes Etudes en Sciences Sociales, 1999.

Romer, Paul M. "Endogenous Technological Change." *J.P.E.* 98, no. 5, pt. 2 (October 1990): S71–S102.

Rosen, Sherwin. "The Economics of Superstars." *A.E.R.* 71 (December 1981): 845–58.

———. "Authority, Control, and the Distribution of Earnings." *Bell J. Econ.* 13 (Autumn 1982): 311–23.

———. "A Note on Aggregation of Skills and Labor Quality." *J. Human Resources* 18 (Summer 1983): 425–31

———. Review of *The Winner-Take-All Society,* by Robert H. Frank and Philip J. Cook. *J. Econ. Literature* 34 (March 1996): 133–35.

Saint-Paul, Gilles. "On the Distribution of Income and Worker Assignment under Intra-firm Spillovers, with an Application to Ideas and Networks." Working paper. Madrid: Univ. Pompeu Fabra, 1999.

Sattinger, Michael. *Capital and the Distribution of Labor Earnings.* Amsterdam: North-Holland, 1980.

———. "Assignment Models of the Distribution of Earnings." *J. Econ. Literature* 31 (June 1993): 831–80.