Learning About the Nature of Production from Equilibrium Assignment Patterns

Luis Garicano
London School of Economics

Thomas N. Hubbard
Northwestern University

July 10, 2012

Abstract

This paper exploits empirically a key insight from Lucas (1977) and Rosen (1982): that the organization of production and the distribution of earnings across individuals are jointly determined by the equilibrium assignment of individuals to firms and hierarchical positions. We study how different classes of production functions generate alternative equilibrium assignments. We then use confidential Census data on U.S. law offices to investigate the form that the production function should take to rationalize earnings patterns in legal services. We argue that earnings patterns in this industry are consistent with a production function that is characterized by asymmetric sensitivity to the skill of agents in different organizational positions, complementarity between managers' and workers' skill, and scale effects in individual skill.

1 Introduction

What is the nature of human capital intensive production? How does the nature of human capital intensive production affect the equilibrium organization of human capital intensive industries? This paper investigates these questions.

We exploit a powerful idea from Lucas (1977) and Rosen (1982): that the organization of production and earnings patterns within industries are jointly determined by the same underlying mechanism – the equilibrium assignment of individuals to firms and hierarchical positions. This equilibrium assignment, in turn, reflects the characteristics of the underlying production function.[1] This idea contains an important empirical implication: earnings patterns contain a wealth of information that can allow researchers to better understand the nature of production in an industry, and in turn, the industry's equilibrium organization.

We first discuss several classes of production functions and examine equilibrium assignment patterns under each. To fix ideas, we begin by analyzing equilibrium assignment when individuals' skills are strictly and symmetrically complementary (as in Kremer (1993)) and when individuals' skills are strict substitutes (as discussed in Grossman and Maggi (2000)), and how the analysis changes when individuals' skills are complements but affect production asymmetrically (Kremer and Maskin (2004)); for example, when some tasks are more skill-sensitive than others. We then study production functions with scale effects but in which there is perfect substitution between the quality and quantity of workers' human capital, such as the ones proposed by Lucas and Rosen. Finally, we explore production functions that combine several of these elements: complementarities, asymmetric sensitivity, and scale effects. We label these "hierarchical production functions" because we intend them to capture a wide range of production processes in which production is more sensitive to managers' than workers' skill and where worker skill allows managers to reduce the time they spend per worker, including contexts where managers help, monitor, coordinate, or train workers.[2] This class of production function, as Garicano and Rossi-Hansberg (2006) show, generates equilibrium assignments with the following features: scale effects associated with managerial skill, positive assortative matching, and cross-matching. That is, better managers work with more and better workers, and managers work with workers with dissimilar skill levels. These authors show further how in problem-solving contexts, cross-matching can necessarily involve strong stratification by skill: the least skilled manager is more skilled than the most skilled worker. These assignment patterns gener-

[1] Rosen notes that "the firm cannot be analyzed in isolation from other production units in the economy. Rather, each person must be placed in his proper niche, and the marriage of personnel to positions and to firms must be addressed directly." (322)

[2] This type of production function was first proposed by Garicano (2000), in a context where managers' role is to help subordinates solve problems.

ate distinct earnings and organizational patterns. Scale effects and positive assortative matching together imply that managerial earnings, worker earnings, and the worker/manager ratio should be positively correlated; strong stratification implies that managers earn more than workers, even when comparing managers at firms with low worker/manager ratios to workers at firms with high worker/manager ratios. These patterns differ sharply from those generated by production functions in which individuals' skills are substitutes or are strictly and symmetrically complementary.

We then examine earnings patterns of lawyers in the United States, using data from the 1992 Census of Services. These data contain law-office-level information about revenues, the number of partners, the number of associates, and associate earnings. We use these data to infer how much partners and associates earn at each office. We use earnings patterns to draw inferences about the equilibrium assignment of lawyers to each other, to organizational positions, and to firms. Our evidence indicates that higher-earning partners work with more, and higher-earning, associates, and that this is true both within and across local geographic markets. We also find evidence consistent with strong skill stratification: controlling for their field of specialization, partners in offices with the lowest partner-associate ratios earn more than associates in offices with the highest partner-associate ratios. That is, the least-leveraged partners earn more than associates do, even those associates who work at offices with high partner-associate ratios. These empirical patterns are consistent with the equilibrium assignments generated by hierarchical production functions such as the one we explore, and suggest that asymmetric complementarities and increasing returns to skill are key aspects of the production process in this human-capital-intensive context.

We then consider the implications of this class of production functions with respect to the equilibrium assignment of individuals to markets of different sizes. If increasing returns lead highly-skilled managers to work in the largest markets, the equilibrium assignment patterns depicted above imply a distinctive cross-market pattern in which the probability that an individual works in a large market rises, falls, then rises with their skill. This is because as an individual's skill increases, their comparative advantage changes from being a worker supporting a highly-skilled, highly-leveraged manager (who works in a large market) to being a low-leverage manager who works in a small market. We show evidence that suggests such a pattern: the relationship between individual lawyers' earnings and the size of the local market in which they work is non-monotonic. When looking at local markets in a given size range, earnings distributions among lawyers appear to be bimodal, with the location of the modes in the earnings distributions increasing as one moves from smaller to larger local markets. These spatial patterns are easily rationalized by production functions such as those we propose that involve limited quality-quantity substitution in individuals' human capital, and in which cross-matching obtains in equilibrium.

We see the contribution of the paper as follows. First, we wish to reintroduce the idea that earnings patterns can say a lot about the nature of human capital intensive production and about the underlying reasons for industries' equilibrium organization. This idea has been underexploited, in part because of the lack of data sets that contain not only information about individuals' earnings, but also on their position within their firms' organization and their firms' characteristics.[3] Second, we wish to emphasize the power of combining equilibrium analysis with organizational models. Evidence on who works with whom and in what capacity can be enormously informative, but inferences from such evidence must be based on equilibrium models since such models allow assignments to be based on individuals' comparative rather than absolute advantage. Third, we provide some empirical evidence using industry-wide data from one human-capital-intensive industry, legal services. While our data are not ideal for this purpose – it would be better to have individual-level earnings data – they have several aspects that lend themselves to an analysis of equilibrium assignment: they are industry-wide, and they allow us to connect earnings to consistently-defined organizational positions and to firms. Earnings patterns from these data suggest that positive assortative matching, scale effects associated with managerial skill, and stratification by skill are important elements of equilibrium assignment in this industry. These elements are hallmarks of production processes in which there is limited substitution between the quality and quantity of individuals' human capital, individuals' skills are complements but affect production asymmetrically, and there exist scale economies associated with individuals' skill.

The paper is structured as follows. Section 2 discusses the general existing theoretical results on equilibrium assignment under different assumptions about production, and about scale of operations effects. In section 3 we describe our data and analyze earnings patterns in legal services in light of

[3]It might also reflect an intellectual separation between the fields of labor economics and industrial organization that Rosen (1982) was trying to bridge.

these models.

2 Production and Equilibrium Assignment

2.1 Non-Hierarchical Production Functions and Assignment

The existence of indivisibilities in individuals' characteristics, such as ability, leads to an assignment problem in production – each individual must be matched to some specific other individuals. Consider in particular a production function that depends on the productive abilities of two individuals, $y = f(z_1, z_2)$, where z_i represents the ability (equivalently, "skill" or "talent") of the individual given task or "position" i. A central theme of the literature on equilibrium assignment is that the nature of the interaction in production between individuals' skills determines the equilibrium assignment of individuals in an economy to each other (and thus to productive units) and to positions. This, in turn, shapes the equilibrium organization of production and the distribution of earnings. We discuss here the main existing results from this literature. Although we keep the discussion quite informal and focus on presenting the ideas behind these results,[4] the reader should keep in mind an economy with a continuum of agent types and a type space that is a compact subset of the real line, where there is a continuous probability distribution over types.

The main results concern the relationship between the form of the production function and the form of the economy-wide equilibrium assignment. We distinguish among production functions according to how agents' skills interact – in particular, whether production is supermodular or submodular in agents' abilities – and whether production is symmetrically or asymmetrically sensitive to individuals' abilities.

First, consider production functions that are supermodular in individuals' abilities and where production is symmetrically sensitive to individuals' abilities. One such case is the Leontief function $y = \min(z_1, z_2)$, which would capture situations where production involves two tasks, both of which must be accomplished for production to take place, and more able individuals

[4] See Sattinger (1993) for a good review of the literature on this topic and Legros and Newman (2007) for the formal exposition of a set of general conditions characterizing positive and negative assortative matching in equilibrium.

are more productive at their task. Another such case is the multiplicative production function $y = z_1 z_2$. This production function was used, and its properties investigated, first in the context of marriage by Becker (1981) and later in a production context by Kremer (1993). Supermodularity, combined with symmetric ability sensitivity, carries a strong implication for the equilibrium assignment of individuals to each other: it implies self-matching or segregation. Kremer (1993) obtains the equilibrium assignment of individuals to workgroups in symmetric production functions of this form, and shows that if an economy consists of individuals who differ in their ability, the optimal organization of these individuals will be such that those with similar amounts of ability will work with each other. The most able individual will work with the second most able, the third most able will work with the fourth, and so on. Self-matching is not only optimal, but is also a characteristic of the competitive equilibrium, because individuals' willingness to pay to be paired with an individual with a given talent level is increasing with their own talent. Thus when tasks are symmetric and complementary, an extreme form of segregation should obtain in equilibrium where individuals work only with others like themselves.

Second, consider production functions that are submodular in individuals' abilities, so that individuals' abilities are substitutes. This a characteristic of production processes where only the best idea or most skillful execution matters; other ideas or efforts turn out to be redundant. Suppose, for example, that production requires two individuals, and output takes place if and only if at least one individual knows the solution to a particular problem, and suppose that the probability that individual i knows the solution is z_i.Then output is given by the submodular production function $y = 1-(1-z_1)(1-z_2)$.[5] Submodularity carries an implication for equilibrium sorting opposite to that of supermodularity: negative assortative matching is the outcome; looking across productive units, more able individuals work with less able individuals. (See Legros and Newman (2007) for a precise statement of this result.) In the symmetric case depicted here, an extreme form of mixing obtains in equilibrium in which the highest ability individual is matched with the lowest ability one, the second highest individual is matched with the second lowest, and so on.

[5] A production function like this was first suggested by Sah and Stiglitz (1986) in the context of project screening within a "polyarchy:" a project is approved if at least one division head likes it.

Third, consider production functions that are supermodular in individuals' abilities, but in which production is asymmetrically sensitive to individuals' abilities, such as the function proposed by Kremer and Maskin (1996):

$$y = z_1^{\theta} z_2^{1-\theta}, \tag{1}$$

with $\theta > 1/2$, so that production is more sensitive to the ability of the individual assigned task 1 than task 2. A straightforward implication is that if differently talented individuals work with one another, the more talented one will be assigned to the more valuable task 1, and the less talented one to task 2. And as in the "symmetric complementarity" case discussed above, equilibrium assignment of individuals to each other in this case involves positive sorting: looking across productive units, the more talented the agent in task 1, the more talented the agent in task 2. However, in this "asymmetric complementarity" case, there is now a drawback associated with self-matching, and assignments that fully exploit complementarities might not be an equilibrium outcome because they would assign skilled individuals to tasks that are not skill-sensitive. Kremer and Maskin show that equilibrium assignment may involve either self-matching or "cross-matching," depending on the support of the distribution of skills. In the cross-matching outcome, "stratification by occupation" results, in which all agents in the economy above a given ability threshold work in task 1, and all agents below it work in task 2. Equilibrium assignment in the cross-matching outcome sharply differs from that in the self-matching outcome; for example, in the cross-matching outcome, the most able individual in the economy works with not with the next most able individual, but rather an individual in the middle of the ability distribution.

The following proposition summarizes the main known results of this literature.

Proposition 1 *Matching and Stratification. The equilibrium assignment of individuals to groups and positions depends on the nature of the production function.*

1. *Production functions that are symmetrically sensitive to indivdiuals' skills and in which skills are complementary imply self-matching: strong segregation of individuals into teams by skill (Becker, 1981; Kremer 1993).*

2. *Production functions in which individuals' skills are substitutes imply negative assortative matching. In the symmetric case, equilibrium assignment implies that the most and least able individuals in the economy work in the same team. (Grossman and Maggi, 2000)*

3. *Production functions that are asymmetrically sensitive to individuals' skills and in which skills are complementary may either imply self-matching or cross-matching, depending on the distribution of individuals' skills in the economy. Under cross matching, stratification by occupation may result in which there is an economy-wide correspondence between individuals' skill and their position. (Kremer and Maskin, 1996)*

Kremer and Maskin's (1996) analysis is motivated by empirical shortcomings of the equilibrium assignment patterns implied by symmetrically-skill-sensitive production functions: self-matching implies little or no within-firm heterogeneity in individuals' ability or earnings, and, if one labels the two positions as "managerial" and "production," a very weak economy-wide correspondence between individuals' ability and their organizational position. Their analysis leads them to argue that a production function involving asymmetric skill sensitivity is necessary in order to account for within-firm heterogeneity in individuals' abilities and wages and to generate an outcome in which individuals' positions correspond closely to their ability.

Although Kremer and Maskin's (1996) discussion of their production function refers to managers and workers, their production function actually involves only two agents. Below we propose that the very reason for asymmetric skill sensitivity is that one agent's talent, the manager's, affects the productivity of all of the workers he manages, and derive some testable implications of a class of hierarchical production functions with this feature.

2.2 Hierarchies and Scale of Operations Effects

A long-standing literature, starting with Simon (1957), and including papers by Mayer (1960), Lucas (1978), Calvo and Weillisz (1978, 1979), Rosen (1982) and Waldman (1984), has proposed that the reason that the distribution of income is more skewed than the underlying distribution of skills lies in how resources are allocated to individuals. In these models, higher-ability managers raise the productivity of the resources they are assigned more than lower-ability managers. As a result, in equilibrium, more able managers are allocated more resources, and this leads the marginal value of their ability to increase faster than if they were working on their own.

Generally, papers in this literature propose production functions where managerial human capital interacts with productive resources. Production functions in this literature have the generic structure:

$$y(z_h) = g(z_h)n \tag{2}$$

where z_h is managerial human capital and n is the span of control of the manager, which, depending on the model, may be the number of workers (Lucas, 1978), efficiency units of labor, i.e., total units of skill managed (Rosen, 1982), or physical capital. In these models, managerial human capital z_h shifts up the marginal product of the workers or capital they are assigned, but managers' span of control is generally limited implicitly or explicitly by managers' time. Equilibrium assignment patterns involve scale of operations effects, which follow from the complementarity between managerial human capital and productive resources.

The main equilibrium result from this class of models follows in the next proposition.

Proposition 2 *(Lucas, 1978, Rosen 1982) Production functions of the form (2) involve scale-of operations effects: more skilled managers are assigned more resources to manage in equilibrium. As a result, the distribution of earnings is more skewed than the distribution of skills.*

Models in this literature have generally assumed perfect substitutability among the resources managed by the manager, so that only the quantity of resources, and not the quality of which they are composed, matters. For example, in models where productive resources are human capital, either only the total number of workers matters (as in Lucas (1978)) or workers of different skill are perfect substitutes (as in Rosen (1982)).[6] But absent an element of imperfect substitutability between workers of different skill, these models do not allow for a full analysis of either the equilibrium assignment of

[6] In Waldman's (1984) more general model, no restrictions on the interaction between managers and workers skills are imposed, but that allows only to characterize the correlation between ability levels and hierarchical position and the fact that the wage distribution is more skewed to the right than the ability distribution. The specific model he analyzes does not allow for complementarities between worker and manager skill and as a result has equilibria with workers more skilled than managers.

individuals to each other or of earnings distributions; if skilled and unskilled workers are perfect substitutes, in equilibrium managers should be indifferent between working with a few relatively skilled workers or many unskilled workers. Assignment patterns between individual managers and workers would then be indeterminant. A more complete analysis of assignment patterns requires combining production with imperfect substitutability of the form in (1), with scale effects of the form in (2). We turn to this next.

2.3 Hierarchical Production, Scale Effects, and Imperfect Substitutability

We develop a simplified version of the hierarchical production function analyzed by Garicano (2000). Suppose that individuals are endowed with unidimensional skill z and one unit of time, and that production involves applying skill to time:

$$y(z,t) = g(z)t \tag{3}$$

This production function has a similar form to (2), in that greater human capital shifts up the marginal product of the time to which it is applied. For simplicity of exposition, let $g(z) = z$. Then the output of individuals who work on their own, applying their human capital to their own time endowment, is $y(z,1) = z$.

Suppose in addition that individuals instead can work as managers, and by doing so, apply their human capital to others' time endowment. Let z_h be the skill of the individual assigned the managerial position and z_l be the skill of the individuals assigned the worker position(s). The production function becomes:

$$y(z,t) = z_h t = z_h n(z_l) \tag{4}$$

Following Garicano (2000), we assume in addition that $n\prime(z_l) > 0$: more skilled workers take up less of managers' time, thereby allowing managers to apply their skill across more workers.

This production function has several key elements. First, individuals of different skills are not perfect substitutes to one another; unlike in Rosen (1982), this extends to individuals acting as workers as well as managers. Second, managers' and workers' skills are complementary. Third, output is asymmetrically sensitive to managerial and worker skill; the asymmetric sensitivity follows naturally from the fact that teams are formed by workers and a manager, whose skill increases the productivity of all of the workers to which this skill is applied. Last, unlike in Kremer and Maskin (1997) there is a mechanism through which managers can exploit scale effects associated with their human capital.

Equilibrium assignment under this "hierarchical production function" has three important characteristics, as discussed in Garicano and Rossi-Hansberg (2006). First, it involves positive sorting, which follows directly from the complementarity between managerial and worker skill. Intuitively, a more highly-skilled manager has a comparative advantage in working with more highly-skilled workers, since such workers allow the managers to apply their human capital to a greater amount of worker time. Second, since $n' > 0$, positive sorting implies that there exist scale of operations effects: more highly-skilled managers manage larger teams. Third, equilibrium assignment never involves self-matching. To see this, note that an agent with skill z_1 who works on his own earns at least z_1. A team of $n+1$ such agents working together in a hierarchy with one acting as manager and n acting as workers earns $z_1 n(z_1)$, which is less than they would earn if each worked on its own, $(n+1)z_1$. When workers are identical, the team produces less than all the workers would produce on their own, and thus it is not formed. Equilibrium assignment therefore must involve some cross-matching, though it need not involve strong stratification. Intuitively, the fact that n must be larger than 0 for teams to be formed makes the production function strongly asymmetrically sensitive to the skill of managers relative to workers; as a result, highly-skilled individuals should always be managers rather than workers.

Garicano and Rossi-Hansberg derive a stronger result in the context of a model where production involves problem-solving. In their formulation, problems are differentiated by their difficulty, and z is the fraction of problems an individual with skill z can solve; thus, in hierarchical production, managers and workers can solve a fraction $z_h$and z_l of problems, respectively. Managers expend their time endowment helping workers with problems workers cannot solve; letting h be the time cost per problem the manager incurs when helping workers, managers' time constraint is:

$$n(1 - z_l)h = 1 \tag{5}$$

Garicano and Rossi-Hansberg show that under these assumptions, equilibrium assignment necessarily involves strong stratification: that is, in equilibrium there must exist some skill level such that all agents of skill below a given level are workers, and all of those above that level are managers.

Proposition 3 *Equilibrium assignment with hierarchical production functions has the following properties.*

1. *Positive sorting. More highly skilled managers work with more highly skilled workers.*
2. *Scale of operations effects. More highly skilled managers manage larger teams.*
3. *Self-matching never obtains; there is always some cross-matching. In addition, strict stratification necessarily obtains, regardless of ability distribution in the population, when production involves problem-solving and takes the form in Garicano and Rossi-Hansberg (2006).*

The difference between our hierarchical production function, which takes the general form $g(z_h)f(n(z_l))$, and one like that analyzed by Kremer and Maskin (1997), which takes the general form $g(z_h)f(z_l)$, is that the skill of the workers enters only through the size of the production team a manager can manage. This yields a specific form of asymmetry that enourages cross-matching and, in certain problem-solving contexts, necessarily leads to strict stratification by skill.

In closing this section, we note that production functions of this form can capture a wide array of interactions between workers and managers, as long as (1) managerial skill raises the productivity of each worker and (2) better workers require less managerial intervention. These interactions need not be related to problem solving. For example, a model of monitoring (such as the models by Calvo and Weillisz (1978, 1979) or Qian (1994)) could naturally be reformulated in this way: the skill of better managers affects the productivity of each worker through better monitoring; better workers have a lower cost of effort, and thus require less monitoring. If so, as above, better managers would oversee larger team of higher-skilled workers. In a model of information processing, such as those proposed by Radner and Van Zandt, or more broadly of coordination, a better manager can process larger amounts of information, and thus increase the output of the entire team by making better decisions; better workers can allow the manager to economize on his own information processing ability by processing more of the information themselves. Finally, suppose managers are training subordinates. In this case (up to now not actually considered by the literature on hierarchies), again a better manager can increase the output of all workers by improving their training; smarter, more qualified workers learn faster and require less training, again economizing on managerial time and, thus, allowing for a larger span of control to the manager.

This production function is most applicable in human capital intensive industries, where the most important inputs are individuals' skills and where organizational structures are designed to exploit these skills. It is thus natural to think that optimizing the utilization of human capital is an important concern in the production of legal services; we turn to an analysis of earnings patterns in this industry.

3 Earnings Patterns in Legal Services

3.1 Data

The data are from the 1992 Census of Services, which is part of the 1992 U.S. Economic Census. Like the rest of the Economic Census, these data contain establishment-level (i.e., law-office-level) data on revenues, employment, payroll, and geographic location. In addition, the Census asks a large sample of law offices questions that are specific to the industry. Answers to these questions provide data on the number of partners, associates, and non-lawyers that work in the office, total pay to associates, and total pay to non-lawyers. They also provide data on the number of lawyers that specialize in particular fields (e.g., corporate law, insurance law) and the share of revenues that are derived from clients that are individuals, businesses, and governments.[7]

These data have several aspects that lend themselves to an analysis of equilibrium assignment. They cover an entire, well-defined human-capital-intensive industry in which organizational positions have a consistent ordering across firms, and allow us to construct estimates of individuals' earnings at the organizational position*office level at a large number of firms. This allows us to explore how individuals' earnings are related to others with whom they work, their organizational position, and characteristics of the firm and

[7] We describe these data in more detail in Garicano and Hubbard (2007, 2009) and show the survey form for the law-office-specific data in Garicano and Hubbard (2009).

market in which they work. Data that allows one to connect individuals' earnings with firm characteristics is not common, and it is even less common to be able to connect earnings with individuals' organizational position. These data have shortcomings, however: whether they contain information about organizational positions depends on firms' legal form of organization, they do not directly report partners' earnings, and at best they provide information on earnings at the organizational position*office level rather than the individual level. We next discuss these shortcomings and how we address them.

Responses and Firms' Legal Form of Organization Responses to some of the Census' questions have different meanings, depending on the office's legal form of organization. The reason for this is that all lawyers are legally considered associate lawyers at offices that are legally organized as "professional service organizations" (PSOs) such as limited liability corporations. This is true even though lawyers at these offices distinguish among themselves in the same way they do at offices legally organized as partnerships: some are partners and others are associates. The variables the Census collects thus differ between PSOs and partnerships.

Table 1 summarizes these differences. The data report the number of lawyers (and non-lawyers) regardless, but distinguish between partners and associates only at partnerships. The data report payroll of all lawyers at PSOs (since all lawyers are legally associates), but only the payroll of associate lawyers at partnerships. The data do not directly report the earnings of partners at partnerships, since these individuals are legally owners rather than employees; their earnings are not considered payroll. The data contain revenues, as reported from tax forms, for all offices, but not non-payroll operating expenses. Other surveys conducted by the Census indicate that much of the latter is accounted for by rent and fringe benefits, costs that are positively correlated with the size of the office.[8]

The data on partnerships are advantageous because they are disaggregated within establishments; they distinguish between partners and associates. This disaggregation is important for our analysis, both because it allows us to examine the implications of the models described above and more generally because it brings the analysis closer to the individual level. However, they do not report partners' earnings, which therefore must be

[8]Bureau of the Census (1996).

estimated. We next describe how we do so.

Estimating Partners' Earnings Partners are the residual claimants on a law firm's proceeds: their revenues less expenses. Thus, partner earnings can be depicted by the identity:

$$\begin{aligned} \textit{partner earnings} &= \textit{revenues} - \textit{associate earnings} \\ &\quad - \textit{nonlawyer earnings} - \textit{operating expenses} \end{aligned}$$

This can be rewritten as:

$$\begin{gathered} \textit{partner earnings} + \textit{operating expenses} = \\ \textit{revenues} - \textit{associate earnings} - \textit{nonlawyer earnings} \end{gathered}$$

The data on partnerships contain the variables on the right side of this expression. Thus, we observe the sum of partners' earnings and operating expenses; the task is to distinguish between these.

The above identity also implies:

$$\textit{operating expenses} = \textit{revenues} - \textit{lawyer earnings} - \textit{nonlawyer earnings}$$

The observations of PSOs contain all of the variables on the right hand side, and thus allow us to impute operating expenses for each of these offices.

Our approach for estimating partners' earnings is to use the operating expense data from the PSOs to develop estimates of operating expenses for each of the partnerships in the data. By the equation above, estimates of operating expenses for each partnership imply estimates of partners' earnings at each partnership.

Operating Expenses at PSOs As one would expect, there is a strong relationship between revenues and operating expenses at the PSOs in our sample. In a simple regression of operating expenses on revenues, the coefficient on revenues is 0.43; a dollar increase in revenues is associated with a 43 cent increase in operating expenses. The R-squared of this regression is 0.80, indicating a raw correlation of 0.89.

Table 2 reports operating expenses as a fraction of revenues, averaged across all PSOs and within various office size categories. Averaged across all offices, this ratio is 0.47. It is greater for smaller than larger offices, declining from 0.49 in single lawyer offices to 0.37 in large offices. This decline suggests that at least some operating expenses are fixed rather than variable.

Combined, this evidence indicates that on average operating expenses increase with revenues by a factor of 0.43, but this varies between 0.37-0.49 depending on the size of the office. We use this as a basis to develop estimates of operating expenses, and thus partners' earnings, at partnerships. We evaluate various estimates by comparing the earnings distributions they imply to earnings distributions constructed directly from data.

Comparing Estimated and Actual Earnings Distributions A first step is to compare distributions from actual data and the estimates using only the PSOs: if we apply the procedure to the PSOs, do we obtain a distribution close to what we started from?

Table 3 reports the results from this exercise. We compute the distribution of lawyers' earnings across offices, weighting each office by the number of lawyers, among PSOs. The median is \$93,000; the 25^{th} and 75^{th} percentiles are \$57,000 and \$132,000, respectively. The other columns report these percentiles when using estimates of operating expenses and imputed lawyers' earnings:

$$\begin{aligned}&\textit{imputed lawyer earnings} =\\&\textit{revenues} - \textit{nonlawyer earnings} - \textit{estimated operating expenses}\end{aligned}$$

The second column assumes that offices' operating expenses equal 43% of revenues, following the regression result above. The median is close to that in the first column, but it compresses the distribution at both ends. The third allows this fraction to differ with the number of lawyers at the office, declining with office size as suggested by Table 2. The resulting distribution is similar to that in the previous column. The final column introduces a fixed element of operating expenses, assuming that operating expenses equal \$80,000 plus 34% of revenues. This does a better job of matching the distribution than the assumptions in the other columns.

A second step is to compare estimates of the earnings distribution among lawyers in partnerships and proprietorships with those generated from other Census data that contain individual earnings data: the Census' Public Use Microdata Sample (PUMS).

The PUMS data contain individual-level observations from the 1990 Census of Population. We use the 5% State Sample. Among other things, the Census asks individuals their occupation, the industry in which they work, their usual hours of work, the number of weeks they worked in the previous year, and their business and salary income in the previous year. We extract observations of full-time lawyers working out of law offices.[9] We convert all dollar amounts to 1992 dollars to make them comparable to those reported in the Census of Services data.

A drawback to the PUMS data is that the earnings data are top-coded. Individuals' business income is top-coded if it exceeds $90,000; their salary income is top-coded if it exceeds $140,000. Thus, earnings distributions derived from PUMS reflect actual responses only below $90,000, which is approximately $102,000 in 1992 dollars. About two-thirds of lawyers in the PUMS have earnings less than this level.

The first column of the bottom panel of Table 4 reports quantiles of lawyers' earnings distribution generated from the PUMS data. We report these for the 10^{th}-60^{th} percentiles because the earnings data are top-coded above these levels. The median lawyer in our PUMS subsample earned $73,515.

The second column reports estimates derived from the partnerships and proprietorships in our Census data. We assume that operating expenses are 43% of revenues, and impute partners' earnings. The distribution generated by this method tracks that generated by the PUMS data quite closely; all of the quantiles reported here are within $3,000 of each other and several are within $1,000. The third column allows operating expenses to vary with the number of lawyers in the office; this matches the middle percentiles well, but corresponds less closely at the other quantiles. The third column assumes that operating expenses equal $80,000 plus 34% of revenues. The distribution generated from this assumption performs much worse than the others at the low end.

Comparing the upper and lower panels of the table, while estimates of operating expenses that contain a fixed element do very well in reproducing the earnings distribution of lawyers across offices organized as PSOs, those

[9] We extract observations of lawyers who worked out of law offices (rather than as judges or as in-house counsel), and eliminate those reporting that they were not in the labor force, whose usual hours were less than 40 hours per week, and who worked fewer than 46 weeks during the previous year. We also eliminate individuals younger than 25 or older than 70 years.

where expenses are proportionate to revenues do better in tracking the distributions generated from PUMS. We therefore have conducted our analysis using all three assumptions depicted in Table 3. Our results are similar, and thus do not depend crucially on which one is preferred. The results we present and discuss below apply the simplest of the three assumptions, that operating expenses equal 43% of revenues.

Aggregation of Individuals' Earnings Our data do not allow us to distinguish among associates or among partners who work at the same office. This aspect of our data limits our analysis of equilibrium assignment patterns: we cannot examine the matching among partners and among associates. In other work, we find evidence suggestive of positive assortative matching across firms within these organizational positions. In Garicano and Hubbard (2005), we use data from the "blue page" listings of law offices throughout Texas from the Martindale-Hubbell directory of lawyers, and show that partners work disproportionately with other partners who obtained their degree at a similarly-ranked law school, and with other partners with similar experience levels. Similar patterns hold for associates.

Our discussions of earnings patterns and what they imply about the nature of human-capital-intensive production will downplay assortative matching among partners and among associates, simply because we cannot investigate it empirically here. We suspect that there is positive assortative matching within organizational positions, and that it might take the form of self-matching, but further research with individual-level earnings data is necessary to determine whether this is the case. Such research would lend further insights on equilibrium assignment and the nature of human-capital-intensive production in this context.

Our analysis will also tend to understate earnings heterogeneity across lawyers, because at best we can examine earnings at the organizational position*office level rather than at the individual level. This issue will arise most prominently when we investigate earnings distributions across markets, the part of our analysis that revolves least around comparisons of conditional means. We will discuss its likely impact on our results at that point in the paper.

3.2 Earnings Patterns of Lawyers

Associates and Partners' Earnings Are Positively Correlated Our first evidence comes from simple regressions of average associate earnings within an office on average partner earnings within an office, using offices with at least one associate. Results are in Table 5. Panel A reports the coefficient on ln(partner earnings) in four regressions. In the first, there are no controls. The coefficient is positive and significant. The point estimate of 0.349 indicates that, on average, associate earnings are 35% higher where average partner earnings are 100% higher. The second column includes a vector of field controls; this vector includes the share of lawyers in the office that specialize in each of 13 fields (e.g., corporate law, probate law). The third and fourth control for geographic market differences. In the third, we include a vector of five dummies that correspond to the employment size of the county in which the office is located;[10] in the fourth, we instead include county fixed effects. The fifth column controls for the office's scale in terms of partners by including $partners$, $partners^2$, and $partners^3$. The coefficient on *ln(partner earnings)* decreases when including the field and market controls, indicating that part of the raw correlation captures cross-field and cross-market differences in average earnings. The result in the fourth column, which includes county fixed effects, indicates that associates' and partners' earnings are positively correlated within as well as between markets. The coefficient decreases only slightly when we control for the office's scale in terms of partners, indicating that the correlation between partner and associate earnings does not reflect that both partners and associates tend to earn more in offices with more partners. Throughout, the coefficient on *ln(partner earnings)* remains positive and significant; the coefficient in the last column indicates that on average, associate pay is 17% higher at offices where partner pay is twice as high.

The other panels report results when we split the sample between "business client offices," offices where over 50% of revenues come from business or government clients, and "individual client offices." These results indicate that regardless of whether offices serve businesses or individuals, associate and partner earnings are positively correlated. The magnitude of this relationship is somewhat greater for offices that serve individuals than businesses.

Table 6 reports results from even narrower subsamples, which confine the

[10]These correspond to the following employment size categories: 20,000-100,000, 100,000-200,000, 200,000-400,000, 400,000-1,000,000, and greater than 1,000,000.

analysis to lawyers in particular fields or geographic markets. The first row reports results from specifications that use only field-specialized offices: for example, offices where all lawyers specialize in insurance law.[11] The positive, significant relationship holds among offices where all lawyers specialize in insurance, negligence-defense, or negligence-plaintiff, though not patent law. The bottom row presents results for four large urban counties: New York County (Manhattan), Los Angeles County, Cook County (Chicago), and Harris County (Houston). The coefficient on *ln(partner pay)* is positive and significant in each case, indicating that the correlation appears within these large urban counties.

These results show strong evidence that associate earnings are higher at offices where partner earnings are higher. While they need not necessarily reflect that associates' and partners' ability is positively correlated – a positive correlation in earnings could be driven by office-level demand shocks (everyone receives a bonus in good years) – they are consistent with production functions that generate positive assortative matching in equilibrium. Such functions include those in which individuals' skills are complements, but not those in which they are substitutes.

Associates' and Partners' Earnings Are Positively Correlated with Associate/Partner Ratios We next investigate whether associates' and partners' earnings are higher at offices where the associate/partner ratio is higher. The first panel in Table 7 reports results from specifications where we regress ln(partner earnings) on ln(associates/partner). In each of the specifications, the coefficient on ln(associates/partner) is positive and significant. In the second row, the dependent variable is instead ln(associate earnings). Once again, the point estimates are positive and significant. Using the results from the last column and comparing offices where one has an associate/partner ratio that is twice as high as the other, average partner pay is 28% higher and average associate pay is 10% higher at the office with the higher associate/partner ratio. The elasticity between partner earnings and the associate/partner ratio is about three times that between associate earnings and the associate/partner ratio.

We have run analogous specifications using the business and individual

[11] We use these four fields because they contain the greatest number of field-specialized offices. See Garicano and Hubbard (2009) for an in-depth analysis of law firms' field boundaries.

client subsamples, looking at field-specialized offices and within New York, Los Angeles, and other large counties. We find very similar results.

Our results provide evidence consistent with a key implication of hierarchical production functions: that comparing earnings among individuals who are at the same hierarchical rank, those who work in groups with more lower-level individuals per upper-level individual earn more. This result also shows that the correlation between associate and partner earnings reported in the previous subsection do not just reflect transitory earnings shocks, unless these shocks also lead associate/partner ratios to change.

Cross-Matching and Stratification An important implication of hierarchical production functions is that equilibrium assignment patterns should involve cross-matching, and under some assumptions, necessarily lead to strong stratification. In this context, the latter would imply that all associates should be less able than any partner. The evidence above suggests that more able associates work in offices with higher associate/partner ratios, as do more able partners. Thus, in this context strong stratification requires in particular that partners in offices with low associate/partner ratios have higher ability than associates in offices with high associate/partner ratios.

We investigate this using a simple model. We assume that $w_i = z_i + \varepsilon_i$, where w_i is lawyer i's earnings, z_i is lawyer i's skill, and ε_i is an i.i.d. shock. We classify lawyers according to whether they are partners or associates, and the associate/partner ratio of their office. Regarding the latter, we create four categories: less than 0.5, between 0.5 and 1.0, between 1.0 and 2.0, and greater than 2.0. This divides lawyers into eight categories. We refer to the associate categories as A1-A4, and the partner categories as P1-P4. We then examine the ordering of lawyers' earnings across these categories. An ordering corresponding to occupational stratfication would be: A1, A2, A3, A4, P1, P2, P3, P4. An ordering corresponding to self-matching would be: A1, P1, A2, P2, A3, P3, A4, P4.

Our specifications take the form of ordered logits, where:

$$
\begin{aligned}
P_1 =& 1 - \Lambda(\beta w_i - \alpha_1) \\
P_j =& \Lambda(\beta w_i - \alpha_{j-1}) - \Lambda(\beta w_i - \alpha_j),\ j = 2, ..., N-1 \\
P_N =& \Lambda(\beta w_i - \alpha_N)
\end{aligned}
$$

P_j is the probability that lawyer i is in position j in the specified ordering.

For the occupational stratification ordering, position 1 is A1, position 2 is A2, and so on. These probabilites are a function of lawyer i's earnings w_i, and thresholds α_j. We estimate this model using different orderings, and compare orderings' explanatory power using Vuong's (1989) non-nested hypothesis test.

Our earnings data are at the level of individuals who work at the same office and organizational position; at offices with both partners and associates, there are two observations. We weight each observation by the number of lawyers the observation represents. We let the thresholds α_j vary across fields and counties, allowing them to be linear functions of the share of lawyers in the office who are in each of the 13 fields in our data, and a vector of county fixed effects. This allows relationships between earnings and organizational position to vary across fields and across markets. We impose the constraint $\alpha_j > \alpha_{j-1}$ so the model is well-defined.

Table 8 reports Vuong test statistics when comparing the occupational stratification specification with other specifications. Under the null hypothesis that specifications fit the data equally well, the Vuong test statistic is distributed N(0,1). Like the previous two subsections, here we use only data from offices with at least one associate. From the log-likelihood values, the occupational stratification specification fits the data better than that in the second row, in which associates at offices with high associate/partner ratios "outrank" partners at offices with low associate/partner ratios. The Vuong test statistic of 7.25 is easily greater than the critical value of 1.96 for a size 0.05 test, indicating that one can reject the null that the specifications fit the data equally well in favor of the alternative that the occupational stratification specification fits better. This test reflects that, controlling for market size and lawyers' fields, associates at offices where the associate/partner ratio is high earn less than partners at offices where this ratio is low. Associates not only tend to earn less than partners in their office, but also than partners more generally.

Table 8 also reports test statistics when comparing the occupational stratification specification with a specification that uses the "self-matching" ordering; this is in the third row. The results indicate that the former fits the data significantly better than the latter, which is not a surprise given the results reported in the previous row.

Table 9 expands the analysis by including partners at offices with no associates – "unleveraged partners" – in the analysis for the first time. Occupational stratification implies that such individuals should rank above all associates but beneath all other partners. We examine this by comparing this to other orderings, particularly those in which unleveraged partners rank lower than associates.[12]

The first set of results use our entire sample of lawyers in partnerships and proprietorships. The occupational stratification specification does not fit best; specifications in which unleveraged partners are outranked by associates in offices with high associate/partner ratios fit significantly better. The rest of the table explores this result further. The next two sets of results split the sample according to the employment size of the county. We use 400,000 as a threshold; counties with more than 400,000 employees include only the most populous counties in the U.S.[13] These results indicate that the occupational stratification specification fits significantly better than the other specifications for counties less than 400,000 employees, but significantly worse than the other specifications for counties with more than 400,000 employees. The final two sets of results explore the large counties further, splitting the sample between business and individual client offices. These results indicate that the occupational stratification specification fits worse than the other specifications for the business client offices, but not significantly so. In contrast, its fit is significantly worse for the individual client offices.

To sum up, our results provide clear evidence of cross-matching throughout our sample, and of strong stratification in most of it; the only exception to the latter concerns unleveraged partners in very large cities. The occupational stratification specification outperforms other orderings when examining offices with at least one associate. It also does so when including unleveraged partners and confining the analysis to all but the nation's largest counties. Outside of very large cities, unleveraged partners tend to earn more than all categories of associates, but less than leveraged partners. In contrast, the occupational stratification specification does not outperform orderings in which unleveraged partners are outranked by at least some associates when including the nation's largest counties, and it performs significantly worse when looking at individual client offices. In the nation's largest counties, associates tend to earn more than unleveraged partners.

[12] In results not shown here, we always strongly reject specifications in which unleveraged partners outrank leveraged partners.

[13] As noted in the table, only about 40 counties were above this threshold as of 1992. Counties that are near this level include Hillsborough County, FL (Tampa) and Orange County, FL (Orlando).

Earnings Distributions and Local Market Size In this section, we examine the distribution of lawyers' earnings across differently-sized local markets. This inquiry is motivated by our interest in how the assignment of individuals to markets reflects the equilibrium assignment of individuals to each other. One possible pattern in the assignment of individuals to markets is a simple one implied by Rosen (1981): in situations where there is limited substitution between the quality and quantity of human capital, "superstar effects" could lead individuals' skill and the size of the market in which they work to be positively associated throughout their respective domains.[14] However, if the equilibrium assignment of individuals to each other involves cross-matching, one would not expect such a pattern. Individuals who tend to work in the largest markets would include not only those with the greatest skill, but also individuals in the middle of the skill distribution whose comparative advantage is working under experts. Under cross-matching, skill and market size would not be positively associated throughout their respective domains, even in the presence of "superstar effects." At some point in the skill distribution, as an individual's skill increases, their comparative advantage would change from being a worker supporting a highly-skilled, highly-leveraged manager (who works in a large market) to being a low-leverage manager who works in a small market. When the equilibrium assignment of individuals to each other involves cross-matching, this could lead the relationship between individuals' skill and the size of the market in which they work to be non-monotonic.

Figure 1 depicts how the earnings distribution across lawyers varies with market size. We construct the Figure in the following way. We first compute earnings deciles across our entire sample, and classify lawyers according to the decile in which they fall. We then classify lawyers according to the employment size of the county in which they work. We then construct histograms that characterize the distribution of lawyers across earning deciles, within each of the six market size categories. We show these distributions across earnings deciles rather than earnings because it provides a useful benchmark: if the earnings distribution is the same across markets, then the histograms would depict a uniform distribution within each market size category. Departures from uniform indicate earnings ranges in which lawyers are over-

[14] Rosen writes that an important implication of his analysis is that "it is monetarily advantageous to work in a larger overall market; and it is increasingly advantageous the more talented one is...the best doctors, lawyers, and professional athletes should be found in the largest cities." (1981:855)

and under-represented within these market size categories. Actual earnings distributions are highly positively skewed; the fact that earnings ranges are much wider in the upper than lower deciles is a manifestation of this skewness.

This Figure shows an interesting pattern. Although higher-earning lawyers tend to work in larger markets, earnings and market size do not appear to be positively associated throughout their domains. Within market size categories, the earnings distributions tend to be bimodal, with each of the modes increasing as market size increases; there is a lower hump that moves from the 1st to the 8th decile, and an upper hump that moves from the 7th to the 10th decile as one moves from the upper to the lower panels.

Table 10 depicts a regression version of this Figure and tests whether the relationships depicted in the Figure are statistically significant. The Table reports results from six regressions. These regressions take the form:

$$y_i = \alpha + \delta_2 D_{i2} + ... + \delta_{10} D_{i10} + \gamma Z_i + \varepsilon_i$$

In the first column, y_i is a dummy variable that equals one if lawyer i works in a county with fewer than 20,000 employees and zero otherwise, D_{ij} is a dummy variable that equals one if lawyer i's earnings are at least decile j, and Z_i is a vector including the share of lawyers in the office who are in each of the 13 fields in our data.[15] The other columns contain analogous specifications using the dummy variables that equal one if lawyer i works in each of the five other market size categories we construct. The sum of the coefficients in the rows equals zero by construction, since the estimates in any one of the rows are implied by the other five. Like in our analysis of stratification, our observations are at the office*organizational position level, and all specifications weight observations using the product of the number of lawyers the observation represents and the Census sampling weight associated with the office. The variables of interest in these specifications are the δ_i's, which indicate whether the share of lawyers in decile i is greater or less than that in decile $i-1$.

These regressions indicate that the patterns depicted in Figure 1 are statistically significant for the most part, and are robust to controlling for systematic differences in lawyers' earnings across fields. The main exception to this is in the smallest local markets, where the upper mode suggested by

[15] Including the latter controls for cross-field differences in lawyers' earnings, but the patterns in the coefficients change little when excluding this vector.

the top panel of Figure 1 is not statistically significant. Looking at the coefficients in the first column of Table 10, the probability a lawyer works in a small market declines steadily through the first five earnings deciles, but there are no significant differences thereafter. The coefficients in the rest of the columns indicate that the main patterns in the rest of Figure 1, which depict densities that alternately increase and decrease, are statistically significant. For example, the coefficients in the second column indicate a significant decrease, then a significant increase as one moves down the table. Similar statistically significant changes in sign appear in the other columns as well.

We think these patterns are interesting, though they are admittedly not dispositive. It would be far better to conduct this analysis with individual-level earnings data. Some of the clustering of earnings may be due to the fact that our office-level data forces us to ignore heterogeneity in earnings among associates and among partners who work in the same office. This would be a particular problem in situations where much of the earnings heterogeneity across associates and across partners within local markets is within rather than across firms. We have investigated this by conducting a similar analysis using lawyers data from the PUMS database described above. The problem with using the PUMS data for this exercise is that it is top-coded above $102,000 1992 dollars, and therefore allows us to construct earnings distributions only for roughly the bottom two-thirds of the distribution. We analyzed these data, and found a similar pattern to that in our data: the within-market-size earning distributions exhibit a mode that increases with market size, similar to the lower mode in Figure 1 though less pronounced. The fact that it is less pronounced might reflect the difference between using individual- and office*organizational position data. This evidence leads us to believe that the patterns we depict are not just an artifact of aggregation, though aggregation might exaggerate these patterns. Top-coding prevents us from investigating whether, like in our data, there is an increasing upper mode when using the PUMS data.

In closing this section, we note that the cross-matching implied by hierarchical production functions implies other interesting non-monotonicities. For example, it implies that as individuals' talent increases, neither the worker/manager ratio nor the talent level (as proxied by the skill of the most able member) of the group in which they work increase monotonically, even though such production functions imply that worker skill, managerial skill, and the number of workers per manager are strictly complementary. This is because individuals' assignment to positions changes at some point as their skill increases – they change from being a worker to being a manager – and when this happens, their equilibrium assignment to each other changes dramatically. Individuals go from working with people at the top of the skill distribution to working with people at the bottom of the skill distribution.[16] We think these non-monotonicities could lend interesting insights into individuals' career progressions in some contexts, and one could test for them using longitudinal data that follows individuals over time. Finally, we note that these sorts of non-monotonicities follow straightforwardly from the principle of comparative advantage, but are only evident when one combines analysis of the production function with an equilibrium model. The analytic power of Lucas' (1978) and Rosen's (1982) idea lies precisely in its exploitation of this principle, which is underutilized in an organizational economics literature that generally does not exploit equilibrium conditions.

4 Conclusiones

This paper shows that earnings patterns contain a wealth of information about assignment patterns and about production functions. Production functions determine the way agents sort with each other in teams, and these two forces together, sorting and production, determine, together with supply conditions, earnings patterns. Thus one can ask, looking at empirical earnings patterns, what sorting patterns, and thus what production functions may be generating them.

The contribution of this paper towards unravelling this relatinship is twofold. First, it reviews and integrates the theoretical literature on the link between production functions and allocation of talent. Specifically, it integrates theories of one-to-one matching, where sorting patterns are determined by the existence of complementarities or substitutabilities in production, with theories of hierarchical production, where there may exist many-to-one sorting. On hierarchical matching, we show that under some relatively general conditions having to do with the existence of managerial time constraints, hierarchical matching functions lead to cross matching, with positive sorting, where the worst agents belong to a category (e.g. workers) the best to another (e.g. managers) and there is positive sorting between the two categories (e.g. better managers are matched with better workers). Moreover,

[16] From these individuals' perspective, this might involve the end of their apprenticeship; our analysis depicts the conditions in which apprenticeships are an equilibrium outcome.

we suggest that under these conditions, better teams not only have better, but also more, workers per manager.

The second contribution is to show empirically that these patterns account reasonably well for the earnings patterns in one particular industry, the law. Specifically, using confidential census data on law firms, we show that these earnings are consistent with the three key theoretical implications of hierarchical matching: scale of operations, positive sorting, and stratification.

Understanding the production functions that may have generated the observable earnings patterns may be useful in understanding technological change and its determinants. In Garicano and Hubbard (2012), we show how one can structurally estimate the underlying parameters of a hierarchical production functions using earnings data. We hope, in future work, to use such structural estimates to illuminate how changes in earnings patterns and in earnings inequality are related to changes in such parameters due, for example, to improvements in communication technology.

5 References

Becker, Gary. *Treatise on the Family.* Cambridge: Harvard University Press, 1981.

Bureau of the Census, *1992 Census of Service Industries, Subject Series: Capital Expenditures, Depreciable Assets, and Operating Expenses.* Government Printing Office: Washington, 1996.

Calvo, Guillermo A. and Stanislaw Weillisz. "Supervision, Loss of Control, and the Optimum Size of the Firm." *Journal of Political Economy*, Vol. 86, No. 5 (October 1978): 943-952.

Calvo, Guillermo A. and Stanislaw Weillisz "Hierarchy, Ability, and Income Distribution." *Journal of Political Economy*, Vol. 87, No. 5 (October 1979): 991-1010.

Garicano, Luis. "Hierarchies and the Organization of Knowledge in Production." *Journal of Political Economy*, Vol. 108, No. 5 (October 2000): 874-904.

Garicano, Luis and Thomas N. Hubbard. "Hierarchical Sorting and Learning Costs: Theory and Evidence from the Law," *Journal of Economic Behavior and Organization*, Vol. 58, No. 2 (October 2005): 349-369.

Garicano, Luis and Thomas N. Hubbard. "Managerial Leverage Is Limited by the Extent of the Market: Theory and Evidence from the Legal Services Industry." *Journal of Law and Economics*, Vol. 50, No. 1 (Febuary 2007): 1-44.

Garicano, Luis and Thomas N. Hubbard. "Specialization, Firms, and Markets: The Division of Labor Within and Between Law Firms." *Journal of Law, Economics, and Organization*, Vol. 25, No. 2 (October 2009): 339-371.

Garicano, Luis and Thomas N. Hubbard. "The Returns to Hierarchy," Mimeo, Northwestern University (2010).

Garicano, Luis and Esteban Rossi-Hansberg. "Organization and Inequality in a Knowledge Economy." *Quarterly Journal of Economics*, Vol. 121, No. 4 (November 2006): 1383-1436.

Grossman, Gene M. and Giovanni Maggi. "Diversity and Trade." *American Economic Review*, Vol. 90, No. 5 (December 2000): 1255-1275.

Kremer, Michael. "The O-Ring Theory of Economic Development." *Quarterly Journal of Economics*, Vol. 108, No. 3 (August 1993): 551-575.

Kremer, Michael and Eric Maskin. "Wage Inequality and Segregation by Skill." NBER Working Paper No. 5718, August 1996.

Legros, Patrick and Andrew Newman. "Monotone Matching in Perfect and Imperfect Worlds." *Review of Economic Studies,* Vol 69, No. 4, (July 2002): 925-942.

Lucas, Robert E., Jr. "On the Size Distribution of Business Firms." *Bell Journal of Economics*, Vol. 9, No. 2 (Autumn 1978): 508-523.

Mayer, Thomas A. "The Distribution of Ability and Earnings." *The Review of Economics and Statistics,* Vol.42, No. 2, (May 1960): 189-195.

Qian, Yingyi. "Incentives and Loss of Control In an Optimal Hierarchy." *Review of Economic Studies*, Vol. 61, No. 3 (July 1994): 527-544.

Radner, Roy and Timothy Van Zandt. "Information Processing in Firms and Returns to Scale." *Annales d'Economie et de Statistique*, Vol. 25-6 (January-June 1992): 265-298.

Rosen, Sherwin. "The Economics of Superstars." *American Economic Review*, Vol. 71, No. 5 (December 1981): 845-858.

Rosen, Sherwin. "Authority and the Distribution of Earnings." *Bell Journal of Economics*, Vol. 13, No. 2 (Autumn 1982): 311-323.

Sah, Raaj and Joseph Stiglitz, "The Architecture of Economic Systems: Hierarchies and Polyarchies," *American Economic Review*, Vol. 76, No. 4 (September 1986): 716-727.

Sattinger, Michael. "Assignment Models of the Distribution of Earnings." *Journal of Economic Literature*, Vol. 31, No. 2 (June 1993):831-880.

Simon, Herbert A. "The Compensation of Executives." *Sociometry*, Vol. 20, No. 1 (March 1957), 32-35.

Vuong, Quang H. "Likelihood Ratio Tests for Model Selection and Non-Nested Hypotheses." *Econometrica*, Vol. 57, No. 2 (March 1989): 307-333.

Waldman, Michael. "Worker Allocation, Hierarchies, and the Wage Distribution." *Review of Economic Studies*, Vol. 51, No. 1 (January 1984): 95-109.

Table 1
Data Reported for Legal Services Establishments in the 1992 Census of Services
by Legal Form of Organization

	PSOs	Partnerships and Proprietorships
Lawyers	Yes	Yes
Partners	No	Yes
Associates	No	Yes
Lawyers' Earnings	Yes	No
Partners' Earnings	No	No
Associates' Earnings	No	Yes
Non-Lawyers	Yes	Yes
Non-Lawyers' Earnings	Yes	Yes
Revenues	Yes	Yes
Operating Expenses (other than payroll)	No	No

Table 2
Operating Expenses As a Fraction of Revenues
Professional Service Organizations Only

	Operating Expenses/ Revenues	N
All Offices	0.47	12844
1 Lawyer	0.49	2576
2-3 Lawyers	0.48	3386
4-7 Lawyers	0.45	3656
8-20 Lawyers	0.41	2471
21-67 Lawyers	0.38	679
68 or More Lawyers	0.37	76

Source: 1992 Census of Services.

This table includes offices that are legally organized as Professional Service Organizations such as Professional Corporations and Limited Liability Corporations only.

Operating expenses equal revenues less lawyers' earnings less non-lawyers' earnings.

Table 3
Comparison of Lawyers' Earnings Distributions Using Actual Data and Estimates
Offices Legally Organized as Professional Service Organizations

Source of Lawyers' Earnings	Actual	Estimated Op. Expenses = 43%*Revenues	Estimated Op. Expenses = 49%*Revenues if 1 Lawyer Office Op. Expenses = 44%*Revenues if 2-7 Lawyer Office Op. Expenses = 38%*Revenues if 8+ Lawyer Office	Estimated Op. Expenses = $80,000+34%*Revenues
Percentile				
5th	0	38	35	-3
10th	24	46	43	26
25th	57	66	65	62
50th	93	91	92	95
75th	132	122	125	132
90th	181	169	171	183
95th	241	214	218	241

Source: 1992 Census of Services.

All earnings are reported in thousands of 1991 dollars.

Table 4
Comparison of Lawyers' Earnings Distributions Using Actual Data and Estimates
Compares Estimates From PUMS Data with those Using Offices Legally Organized as Partnerships or Proprietorships

Source of Lawyers' Earnings	Actual Lawyers in 1990 PUMS 5% State Sample	Estimated Partnerships and Proprietorships 1992 Census of Services Op. Expenses = 43%*Revenues	Estimated Partnerships and Proprietorships 1992 Census of Services Op. Expenses = 49%*Revenues if 1 Lawyer Office Op. Expenses = 44%*Revenues if 2-7 Lawyer Office Op. Expenses = 38%*Revenues if 8+ Lawyer Office	Estimated Partnerships and Proprietorships 1992 Census of Services Op. Expenses = $80,000+34%*Revenues
10th	29	27	24	-20
20th	40	40	37	21
30th	51	52	50	42
40th	61	63	61	57
50th	74	73	72	70
60th	88	85	84	85

Source: 1992 Census of Services.

All earnings are reported in thousands of 1992 dollars.

Table 5
Regressions of Associate Pay on Partner Pay
Partnerships and Proprietorships with at Least One Associate

Dependent Variable: ln(associate pay)

Panel A: Business and Individual Client Offices (N=5365)

ln(partner pay)	**0.349** **(0.008)**	**0.256** **(0.009)**	**0.190** **(0.008)**	**0.176** **(0.008)**	**0.173** **(0.008)**
R-squared	0.24	0.37	0.48	0.65	0.67

Panel B: Business Client Offices (N=3480)

ln(partner pay)	**0.308** **(0.009)**	**0.235** **(0.009)**	**0.162** **(0.009)**	**0.141** **(0.009)**	**0.138** **(0.009)**
R-squared	0.25	0.33	0.48	0.64	0.66

Panel C: Individual Client Offices (N=1885)

ln(partner pay)	**0.331** **(0.015)**	**0.252** **(0.016)**	**0.214** **(0.016)**	**0.189** **(0.017)**	**0.196** **(0.016)**
R-squared	0.21	0.31	0.37	0.70	0.72
Controls	None	Specialty Shares	Specialty Shares, Market Size Dummies	Specialty Shares, County Dummies	Specialty Shares, County Dummies, Partners, Partners**2, Partners**3

Bold indicates rejection of the hypothesis b=0 using a one-tailed t-test of size 0.05.

Table 6

Regressions of Associate Pay on Partner Pay

Field-Specialized Offices, Large Urban Counties

Partnerships and Proprietorships with at Least One Associate

Dependent Variable: ln(associate pay)

Panel A: Field-Specialized Offices

Field:	Insurance	Negligence-Defense	Negligence-Plaintiff	Patent
ln(partner pay)	**0.097** **(0.031)**	**0.152** **(0.035)**	**0.270** **(0.033)**	0.041 (0.041)
R-squared	0.27	0.28	0.43	0.30
N	214	155	239	130

Panel B: Large Urban Counties

County:	Manhattan (New York)	Los Angeles	Cook (Chicago)	Harris (Houston)
ln(partner pay)	**0.138** **(0.019)**	**0.104** **(0.027)**	**0.227** **(0.038)**	**0.285** **(0.059)**
R-squared	0.53	0.31	0.71	0.38
N	372	349	156	92

Specifications in Panel A include market size dummies, partners, partners**2, partners**3 as controls.
Specifications in Panel B include specialty shares, partners, partners**2, partners**3 as controls.

Bold indicates rejection of the hypothesis b=0 using a one-tailed t-test of size 0.05.

Table 7
Regressions of Partner Pay and Associate Pay on Associates/Partner
Partnerships and Proprietorships with at Least One Associate

Panel A: Dependent Variable: ln(partner pay), N=5365

ln(associates/partner)	**0.375** **(0.012)**	**0.312** **(0.011)**	**0.264** **(0.012)**	**0.285** **(0.014)**	**0.283** **(0.014)**
R-squared	0.15	0.30	0.33	0.50	0.51

Panel B: Dependent Variable: ln(associate pay), N=5475

ln(associates/partner)	**0.190** **(0.009)**	**0.156** **(0.008)**	**0.087** **(0.008)**	**0.055** **(0.008)**	**0.103** **(0.008)**
R-squared	0.07	0.30	0.43	0.59	0.62
Controls	None	Specialty Shares	Specialty Shares, Market Size Dummies	Specialty Shares, County Dummies	Specialty Shares, County Dummies, Partners, Partners**2, Partners**3

Associate pay is associate payroll within the office divided by the number of associates.
Partner pay is (revenues - payroll - overhead) divided by the number of partners, where overhead equals 0.43*revenues.
The number of observations differs between the two panels because ln(partner pay) is undefined when partner pay is negative.

Bold indicates rejection of the hypothesis b=0 using a one-tailed t-test of size 0.05.

Table 8
Vuong Tests of Occupational Stratification

Lawyers in Partnerships and Proprietorships With at Least One Associate

Ordering								-LogL	Vuong Test Statistic
A1	A2	A3	A4	P1	P2	P3	P4	17936	
A1	A2	A3	P1	A4	P2	P3	P4	18142	7.25
A1	P1	A2	P2	A3	P3	A4	P4	19823	15.03

This table reports Vuong test statistics when comparing the occupational stratification specification in the first row to alternative specifications. The null is that the specifications fit the data equally well. Under the null, this statistic is distributed N(0,1). See Vuong (1989) for details.

The specifications are ordered logits, where lawyers are classified according to their occupational position and the associate/partner ratio in their office. The categories A1-A4 correspond to associates in offices where this ratio is less than 0.5, at least 0.5 but less than 1.0, at least 1.0 but less than 2.0, and greater than 2.0, respectively. The categories P1-P4 correspond to partners classified analogously.

The ordered logits predict lawyers' classification as a function of their earnings. All specifications allow threshold "alphas" to vary with specialty shares and county employment size dummies (see text for how these are defined).

The unit of observation is at the occupation*office level (partners or associates at a given office). N=10,950, which reflects that there are two observations for each of the 5475 partnerships and proprietorships with at least one associate in our data.

Table 9
Vuong Tests of Occupational Stratification

Lawyers in Partnerships or Proprietorships

Ordering	-LogL	Vuong Test Statistic	-LogL	Vuong Test Statistic	-LogL	Vuong Test Statistic	-LogL	Vuong Test Statistic	-LogL	Vuong Test Statistic
A1 A2 A3 A4 P0 P1 P2 P3 P4	24580		15110		8842		6104		2057	
A1 A2 A3 P0 A4 P1 P2 P3 P4	24501	-2.08	15154	1.73	8769	-3.51	6079	-1.81	1979	-4.56
A1 A2 P0 A3 A4 P1 P2 P3 P4	24256	-3.01	15206	2.59	8620	-3.21	6017	-1.71	1953	-4.23
A1 P0 A2 A3 A4 P1 P2 P3 P4	24263	-2.09	15305	3.63	8584	-2.76	6016	-1.39	1966	-2.78
P0 A1 A2 A3 A4 P1 P2 P3 P4	24414	-1.04	15492	5.98	8565	-2.81	6002	-1.52	1973	-2.72
<u>Sample:</u> Offices	All		All		All		Business Client		Individual Client	
Counties	All		< 400K Employment		> 400K Employment		> 400K Employment		> 400K Employment	

This table reports Vuong test statistics when comparing the occupational stratification specification in the first row to alternative specifications. The null is that the specifications fit the data equally well. Under the null, this statistic is distributed N(0,1). See Vuong (1989) for details.

The specifications are ordered logits, where lawyers are classified according to their occupational position and the associate/partner ratio in their office. The categories A1-A4 correspond to associates in offices where this ratio is less than 0.5, at least 0.5 but less than 1.0, at least 1.0 but less than 2.0, and greater than 2.0, respectively. The categories P1-P4 correspond to partners classified analogously. P0 is partners at offices without associates.

The ordered logits predict lawyers' classification as a function of their earnings. All specifications allow threshold "alphas" to vary with specialty shares and county employment size dummies (see text for how these are defined).

The unit of observation is at the occupation*office level (partners or associates at a given office). N=14,918, which reflects that there are two observations for each of the 5475 partnerships and proprietorships with at least one associate in our data plus 3,968 offices with partners but not associates.

Individual client offices are offices where at least 50% of revenues come from individuals. All other offices are business offices.
Approximately 40 counties had >400K employment as of 1992; counties with approximately 400K employees include Hillsborough County, FL (Tampa) and Orange County, FL (Orlando).

Table 10
Regressions of Market Size Dummies on Lawyers' Earnings Decile

	Dependent Variable Dummy Variable That Equals One If Office Is Located In a County Where Number of Employees Is:					
	0-20,000	*20K-100K*	*100K-200K*	*200K-400K*	*400K-1M*	*Over 1M*
Earnings Above 1st Decile	**-0.027**	0.009	**0.060**	**0.048**	**-0.165**	**0.075**
	(0.009)	(0.012)	**(0.011)**	**(0.015)**	**(0.016)**	**(0.014)**
Earnings Above 2nd Decile	-0.013	-0.022	-0.016	0.028	**0.099**	**-0.076**
	(0.009)	(0.012)	(0.011)	(0.015)	**(0.016)**	**(0.014)**
Earnings Above 3rd Decile	**-0.024**	-0.018	0.015	0.067	**-0.036**	-0.005
	(0.009)	(0.012)	(0.011)	(0.015)	**(0.016)**	(0.014)
Earnings Above 4th Decile	**-0.029**	**-0.035**	**-0.046**	0.013	**0.072**	0.024
	(0.009)	**(0.012)**	**(0.011)**	(0.015)	**(0.016)**	(0.013)
Earnings Above 5th Decile	0.014	0.024	-0.008	**-0.141**	**0.053**	**0.057**
	(0.009)	(0.012)	(0.011)	**(0.015)**	**(0.016)**	**(0.014)**
Earnings Above 6th Decile	-0.012	**0.025**	-0.016	**0.057**	**-0.093**	**0.038**
	(0.009)	**(0.012)**	(0.011)	**(0.015)**	**(0.016)**	**(0.013)**
Earnings Above 7th Decile	0.006	-0.017	**0.039**	**-0.043**	**-0.066**	**0.081**
	(0.009)	(0.012)	**(0.011)**	**(0.015)**	**(0.016)**	**(0.014)**
Earnings Above 8th Decile	-0.010	-0.016	0.001	0.003	**0.077**	**-0.054**
	(0.009)	(0.012)	(0.011)	(0.015)	**(0.016)**	**(0.013)**
Earnings Above 9th Decile	0.004	-0.018	**-0.043**	**-0.062**	**0.035**	**0.083**
	(0.009)	(0.012)	**(0.011)**	**(0.015)**	**(0.016)**	**(0.013)**

This table reports results from six regressions; these are the regression versions of Figure 1. In the first column, the dependent variable is a dummy that equals one if the lawyers the observation represents work in a county with fewer than 20,000 employees. The independent variables are a series of dummies that indicate where in the overall distribution these lawyer lie. These dummies are defined so the coefficients represent changes relative to the previous category. The -0.024 point estimate in the third row, first column indicates that, moving from the 3rd earnings decile to the 4th lowers the probability a lawyer works in a very small market by 2.4 percentage points.

The row sums of the point estimates are zero by construction; the point estimates any one of the rows are implied by the other five.

All regressions include a vector of controls that include the shares of lawyers in the office who are in each of the 13 Census-defined fields. The estimates presented here differ little when excluding these controls.

Figure 1
The Distribution of Lawyers Across Earnings Deciles by Market Size Category

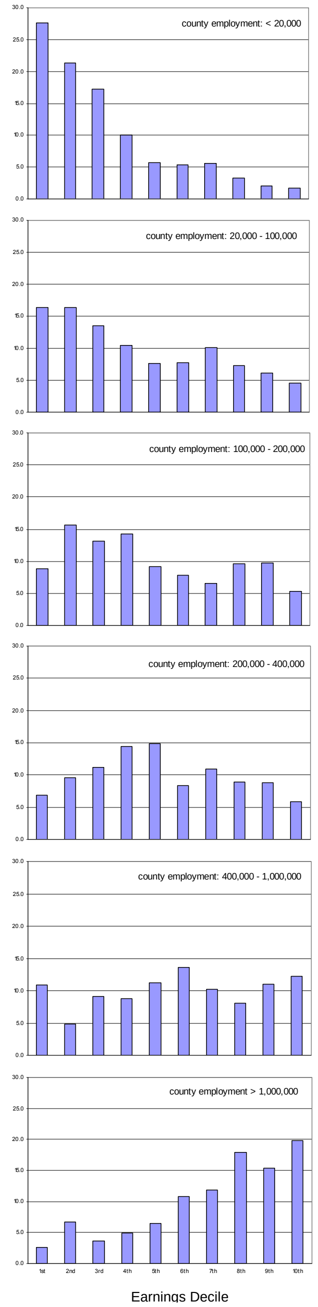

This Figure depicts the how the distribution of lawyers across earnings deciles varies across local markets of different sizes.

We developed this Figure in the following way. First, we computed earnings deciles across all markets, and assigned associates and partners within each office to earnings deciles accordingly. Then, we computed and plotted frequency distributions of lawyers across these deciles within market size categories. The first bar of the top panel indicates that 27.6% of lawyers in counties with less than 20,000 employees have earnings that put them in the 1st decile, when earnings deciles are calculated across all markets. If earnings distributions are identical across differently-sized local markets, these frequency plots would depict uniform distributions.

These plots indicate that, although higher-earning lawyers tend to work in larger markets, earnings and market size do not appear to be positively associated throughout their domains. Instead, these frequency distributions tend to be bimodal, with both modes increasing as one moves from smaller to larger local markets.

The mean earnings within earnings deciles are:

1st:	14,263	6th:	78,741
2nd:	33,057	7th:	91,908
3rd:	46,069	8th:	111,594
4th:	57,553	9th:	144,278
5th:	68,142	10th:	293,814