

Occasional paper

43

Knowledge-based Hierarchies:
Using Organizations to
Understand the Economy

Luis Garicano
Esteban Rossi-Hansberg

October 2014

Abstract

We argue that incorporating the decision of how to organize the acquisition, use, and communication of knowledge into economic models is essential to understand a wide variety of economic phenomena. We survey the literature that has used knowledge-based hierarchies to study issues like the evolution of wage inequality, the growth and productivity of firms, economic development, the gains from international trade, as well as offshoring and the formation of international production teams, among many others. We also review the nascent empirical literature that has, so far, confirmed the importance of organizational decisions and many of its more salient implications.

Key words: Knowledge, hierarchies, wage inequality, international trade

This paper was produced as part of the Centre's Productivity and Innovation Programme. The Centre for Economic Performance is financed by the Economic and Social Research Council.

Acknowledgements

We thank Pol Antràs, Lorenzo Caliendo, Ferdinando Monte, Bob Lucas, and Steve Redding for useful comments. Ricardo Reyes-Heroles provided excellent research assistance.

Luis Garicano is an Associate of the Centre for Economic Performance. He is also Professor (Chair) of Economics and Strategy, Department of Management and Department of Economics at the London School of Economics and Political Science. Esteban Rossi-Hansberg is a Professor in the Economics Department at Princeton University, USA.

Published by
Centre for Economic Performance
London School of Economics and Political Science
Houghton Street
London WC2A 2AE

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means without the prior permission in writing of the publisher nor be issued to the public or circulated in any form other than that in which it is published.

Requests for permission to reproduce any article or part of the Working Paper should be sent to the editor at the above address.

1 Introduction

Most human endeavours, from the construction of the great infrastructure projects, to the development of key ideas and innovations, even the production of a simple pencil (as in Friedman's famous example), require the collaboration of teams of individuals. Understanding the organization of the skill, knowledge, and time of individuals into teams with common production goals is essential to comprehend virtually all human activities. However, most mainstream economic models still abstract from modelling the organizational problem that is necessarily embedded in any production process. Typically, these frameworks jump directly to the formulation of a production function that depends on total quantities of a pre-determined and inflexible set of inputs. In this survey, we argue that this simplification, although often practical, leaves important issues aside.

The importance of understanding organizations has been acknowledged at least since Knight's (1921) discussion of the role of managers in making decisions under uncertainty.¹ Some early attempts are Mayer (1960) and Ijiri and Simon (1964) who sought to explain the skewness in the distribution of pay and the firm-size distribution. Team theory, following on Radner (1993), studied the role of organizations as information processing devices: by processing information in parallel, hierarchy can reduce delay. Others, such as Geanakoplos and Milgrom (1991), studied decentralized resource allocation with limited information on unit costs. A line of research, following Calvo and Weillisz (1978), also studied hierarchies where managers at each level engaged in monitoring their subordinates. All of these papers study organization in isolation from the market and do not allow for the study of the interaction between organization and the economy, as they involve neither general equilibrium nor heterogeneity. Lucas (1978) and Rosen (1982) introduced models of equilibrium assignment and thus allowed for the joint study of the wage and firm size distribution.² Consequently, the literature of knowledge-based hierarchies that we and others have developed, and we review here, aims to incorporate organization into mainstream economics frameworks.

Knowledge is an indispensable input in production. So the ability to manage efficiently its acquisition, communication, and use is a determinant of the productive efficiency of an organization (e.g. a team, a firm, or even an economy as a whole). We conceptualize knowledge as the ability to solve the problems that naturally arise in any production process. The organizational problem arises because this knowledge is embedded in individuals who have limited time to work. Hence, although the knowledge they possess can be used repeatedly as an input in production, the individual that possesses it faces a time constraint that limits in practice how often this knowledge can be used. The theory of knowledge-based hierarchies recognizes that one way to relax this time constraint is to work in teams, where less knowledgeable workers deal with routine production tasks, thereby economizing on the time of experts and allowing them to specialize on giving directions on the harder tasks. The organizational problem then becomes the problem of determining who knows what, who do they communicate with, and how many

¹See Garicano and Van Zandt (2013) for a recent review of the literature on Hierarchies.

²Prescott and Visscher (1980) and Boyd and Prescott (1987) also incorporate organizations into equilibrium models.

workers of each type are required to minimize the cost of producing a certain output. That is, it becomes the problem of determining how to use knowledge efficiently.

Because the organizational problem optimizes jointly the knowledge and number of members of a production team, it links the marginal return of a given worker to the characteristics of the other workers in her team, and more broadly to the distribution of knowledge and wages in the population. In fact, labor market and firm organization are the two sides of the same coin. The link between the two is the “scale of operations effect”, whereby the marginal value of an agent’s ability is given by the amount of workers he manages. So one important application of the organizational problem described above is to understand the impact of economy-wide changes in technologies that affect the acquisition and communication of knowledge (ICT) on the distribution of wages in the economy.

Take for example the well known fact of “wage polarization”. The labor literature has argued that, during the last few decades, inequality at the top of the wage distribution has increased dramatically, while, concurrently, it has declined at the bottom. Simple theories that explain increases in wage inequality as a result of increases in the price of skill have a hard time explaining these facts. They are, however, easily rationalized by the theory of knowledge-based hierarchies as a result of large improvements in communication technology since the late 80’s. Better communication technology effectively relaxes the time constraint of experts by allowing them to deal with more tasks. This implies that less-knowledgeable workers, who earn less, become less differentiated since the purpose of their knowledge is to relax a constraint that is now less binding. The result is that experts, the superstars of the knowledge economy, earn a lot more, while less-knowledgeable workers become more equal since their knowledge becomes less useful. Moreover, communication technology allows superstars to leverage their expertise by hiring many workers that know little, thereby casting a shadow on the best workers who used to be the ones exclusively working with them. We call this the *shadow of superstars*.

In Section 2 we discuss this example in the context of the existing literature as way of motivation. The rest of this paper presents a more systematic survey of the work that we and others have done to incorporate organizations into equilibrium frameworks as well as the application of these models to a variety of issues including firm heterogeneity, economic development, and international trade and offshoring.

Thus in Section 3 we introduce the basic technology for knowledge utilization that relies

on the use of hierarchies. We capture the idea, best expressed by Harold Demsetz, that “those who are to produce on the basis of this knowledge, but not be possessed of it themselves, must have their activities directed by those who possess (more of) the knowledge. Direction substitutes for education (that is, for the transfer of the knowledge).” (Demsetz, 1988, p. 157). Instead of educating all agents, which is very costly, those in production learn the most routine tasks and ask for directions whenever they need help on less routine tasks. We start by setting up such a model, and by reviewing some of the basic results that will carry on throughout.

Organizations can be introduced in either an economy where agents are ex-ante identical or one where they are ex-ante heterogenous. We reserve the concept of skill (as opposed to knowledge) for an innate (or early acquired) exogenous individual characteristic. In frameworks where agents are ex-ante identical, skill is completely abstracted from. In models where agents are ex-ante heterogenous, skill can be either be synonymous with knowledge, when the acquisition of knowledge is not explicitly modeled, or it can determine the cost of acquiring knowledge, when knowledge acquisition is incorporated.

Knowledge-based hierarchies are efficient in either of these cases. When individuals are ex-ante homogeneous, they increase the utilization rate of knowledge. All agents obtains identical utility, although they acquire different knowledge depending on their role in the organization—wages just compensate workers for the cost of knowledge. But instead of having all production workers learn how to do something, most of them specialize in routine tasks and only a few agents specialize in exceptions and communicate this specialized knowledge to producers as needed. This ex-post heterogeneity is the outcome of the organizational problem. In Section 3 we review the basic frameworks that use homogenous agents, starting from Garicano (2000) which sets up the basic technology and the circumstances under which knowledge-based hierarchies are optimal. We then discuss Caliendo and Rossi-Hansberg (2012) which embeds this basic technology in a classic cost minimization problem, and then incorporates it into an equilibrium framework with heterogeneous firms. We finish Section 3 with a brief review of the dynamic problem in Garicano and Rossi-Hansberg (2012) of deciding when to exploit knowledge by using hierarchies to deepen knowledge in existing ideas, versus starting a new field, which requires discarding existing organizations.

In frameworks with ex-ante heterogeneous agents, hierarchical organization also becomes a device to allocate heterogenous agents into positions where they can optimally leverage their skill. Hence, these models involve an assortative matching problem between agents

with different skill levels. The wage function that results from this problem does not simply compensate agents for their cost of education, but it values the specific skills they possess relative to the rest of the individuals in the economy. As such, these frameworks can address the facts on the evolution of the labor market that we discussed above. Even though the resulting matching problem can be illuminating, it also comes at an analytical cost. So the benefits of using this variant of the framework certainly depend on the application one has in mind. In Section 4 we first discuss Garicano and Rossi-Hansberg (2004), who have ex-ante heterogeneous skilled agents but only two layers and no investment in knowledge. We then proceed to describe Garicano and Rossi-Hansberg (2006) who allow multiple layers, ex-ante heterogeneity and investment in knowledge. We conclude the section by studying the problems that appear when hierarchies operate in markets with double-sided asymmetric information, as in Fuchs et al. (2014).

In Section 5 we move on to applications of these frameworks that study international trade and the international organization of production. In particular, we discuss how the gains from a trade liberalization are affected by the changes in organization that they trigger. Because these changes in organization are mediated by the level of ICT, these technologies affect the gains from openness and the impact of international trade on wage distributions and firm productivity. We then turn to the implications of the theory of knowledge-based hierarchies on offshoring. Specifically, Antràs et al. (2006) show that the matching of individuals in international production teams can affect wages in countries with different skill distributions in ways that are consistent with the empirical evidence.

Section 6 describes the existing empirical work related to organization in knowledge base-hierarchies and its implications. One of the premises underlying the organizational problem is that the characteristics of production teams are not fixed or pre-determined by some exogenous force, but are active decisions by firms or market participants. As such, these decisions should change when, for example, firms decide to grow or shrink. Caliendo et al. (2014) show that, when French firms grow, wages at each layer react very differently depending on whether firms reorganize by adding new layers of management. Another central implication of the theories of knowledge-based hierarchies is that, within layers, more knowledgeable workers are matched with more knowledgeable managers, and so more knowledgeable managers also lead larger teams. Garicano and Hubbard (2013) show that this is, in fact, the case in U.S. law firms.

We finish the survey with a concluding section that points the reader to what we believe

are good directions for further research.

2 The New Organization of Work: Workers in the Shadow of Superstars

The distribution of wages and occupations have experienced dramatic changes over the last decades.³ In fact, the evolution of the wage structure provides us with a powerful tool to discriminate between alternative theories about the fundamental forces driving changes in the labor market, as well as the way heterogeneous workers interact in the labor market.⁴

The first central fact is that there has been a clear increase in the skill premium for workers between 1973 and 2009. Furthermore, the increase was concentrated on highly educated workers. For less educated workers the increase in the slope is small or perhaps reversed.


Beyond education, many other characteristics are important in determining the knowledge and skill levels of a worker. Perhaps the best measure of the marketable knowledge and skills of an agent is her wage. Hence, the labor literature has proposed to use wage percentiles, summarizing workers' rank in the overall distribution, as a measure of the rank of skill levels. The evidence shows that the winners from changes in the wage distribution between 1988 and 2008 were those in the bottom and top percentiles. The same was not true during the earlier 14 years where workers at the bottom of the distribution suffered important reductions in their wages. This phenomenon has been labeled "wage polarization". The finding, together with the convexification of the relationship between wages and schooling indicates that during the last three decades the mechanism causing these changes must have affected workers in the middle of the skill (or knowledge) distribution negatively, while it should have affected other workers positively. Figure 1 shows this unequivocally. After an initial increase, starting in 1986, inequality between the 50th and 10th percentile of the distribution has decreased slightly, while inequality in the top half of the distribution has increased markedly.⁵

³Acemoglu and Autor (2011) present a systematic assessment of the main trends in the data and the existing attempts to explain them with a unified framework.

⁴Katz and Murphy (1992), Acemoglu (1999), Autor et al. (2003), Lemieux (2006), Autor, Katz and Kearney (2006, 2008), Acemoglu and Autor (2011) present all the key facts.

⁵The figure relies on the Current Population Survey (CPS) data for the U.S. economy during the last few decades, as cleaned and provided by David Autor.

Figure 1: Wages Polarization



90/50 and 50/10 Wage Inequality

Simple theories of changes in the return to skill will clearly fail to explain these facts. First, whatever theory one postulates to explain the evolution of the wage distribution needs to predict that the return to skill of workers with less units of skill behaves differently than that of workers with units of skill above the median. That is, we need a theory where the price of a unit of skill depends on the number that the individual possesses. Second, we need a theory in which the baseline wage increases, but the return to skill decreases for the bottom percentiles while it increases for the top ones. Finally, whatever mechanism we propose has to be operative starting in the mid to late 80's, but not before.

The so-called “canonical model” explains the changes in the wage distribution as a race between technology and education.⁶ In this view, pioneered by Katz and Murphy (1992), as technology evolves it progressively becomes more skill intensive (that is, technological change is skill biased). This raises the return to skill as long as the increase in the supply of skilled workers does not compensate. Acemoglu, in a series of influential papers,⁷ gives precise conditions under which changes in technology will lead to an increase in the supply of workers without eliminating the increase in the returns to skill.

These explanations are successful in accounting for the increase in the skill premium in

⁶See the encyclopedic and informative book by Goldin and Katz, 2009, for an account of many of these facts using this approach.

⁷Acemoglu (2002) provides a nice survey of the skill-biased technological change literature.

the last few decades. However, since they speak only about the price of a unit of skill, rather than the price of a unit of skill as a function of the number of units an individual possesses, they cannot account for the distinct behavior of the wage distribution for low, medium, and high skill individuals.

Clearly, the salient fact in Figure 1 cannot safely be ignored so the literature has evolved to develop explanations that point to a different evolution of the price of skill across individuals. Two such explanations have been developed. First, there is the task-based approach initially proposed by Autor et al. (2003) and Autor et al. (2006) and forcefully pushed as the central explanation of these facts by Acemoglu and Autor (2011). Second, there are theories that emphasize the returns to managerial activities like Gabaix and Landier (2008) and initially Lucas (1978) and Rosen (1998). The latter set of papers include theories of organization and so we discuss them more fully below. Here, we only note that they fail to account for the behavior of the returns to skill for workers with low levels of skill. That is, these are theories that explain how changes in technology have increased the returns to managers, but cannot explain why they have decreased the returns to skill at the bottom of the distribution. As such, these theories are interesting as explanations of the emergence of superstar earners (Rosen, 1981), but not as explanations of the shadow they seem to have cast on middle-skilled workers.

The task-based explanation is the best known explanation of these patterns. These theories do not introduce organization explicitly. They postulate a structure of substitution between distinct tasks in the production process and exploit it to generate complementarity between tasks. A reduction in the cost of producing computers, for example, will reduce the need of some tasks in favor of others depending on the postulated substitution patterns assumed in the production function. A production function with constant elasticity of substitution between tasks, with a nesting structure that makes capital and routine tasks highly substitutable, is common.

One advantage of this approach is that the elasticities of substitution in the production function can be estimated. Harder to estimate and specify is the nesting structure between different tasks and factors in the production function. Not to mention the arbitrariness in defining the relevant tasks to include in the production function. A standard story in this literature is that the decline in the cost of capital equipment has decreased the price of routine tasks that are highly substitutable with this form of capital. In contrast, analytic and manual tasks are less substitutable, and so the improvement in technology leads to

higher demand and higher employment of workers performing those tasks.

A similar story can be told to explain the impact of offshoring rather than computer use. The tasks that are easily offshorable are the ones that can be specified in a set of well-defined rules. These are again the routine tasks done by workers in the middle of the wage distribution. The result is wage and employment polarization as illustrated in Figure 1.

The organization-based explanation of these facts that we propose is not in contradiction with this task-based approach. It is, to some extent, a distinct specification of the task-based approach in which tasks are hierarchical and the production function that links the different tasks is based on an explicit organizational problem. The flexible hierarchical nature of this theory has the potential to explain the set of distinct inputs used in production as well as their characteristics. Thus, it can explain for example why are the middle-skilled agents the ones that have lost in the last decades. Other task-based approaches also link these losses to agents specializing in routine tasks, but not why these are the agents with intermediate incomes. In that sense, in the theories reviewed in Autor and Acemoglu (2011), the fact that middle skilled agents are *in the shadow of superstars* is hard-wired in the theory, not endogenous to the evolution of the economy. Furthermore, because the theory of knowledge based hierarchies provides an explicit organizational theory, it can distinguish in a precise way between different types of information and communication technology (ICT) changes, and explain their distinct effects on worker's returns and firm performance.

2.1 Superstar Managers and Increases in the Scale of Operations

A well-known explanation for the increasing concentration of rewards at the top of the earnings distribution has to do with the increase in the effective market size of top talent: the scale of operations effect.⁸ The idea is that a good decision by a better manager increases the productivity of all her subordinates. As Rosen (1982) argues, a good soldier is useless in the wrong war. In this context, large returns to talent can be generated by the competition for better managers between organizations that can leverage the manager's knowledge by assigning them large teams. This argument can explain the phenomenon of "superstar" pay. As Manne (1965) first pointed out, a better concert pianist has higher earnings not only

⁸See Mayer (1960) and Simon (1964) for the first two accounts of this effect. Chandler (1977) offers a historical account of the growth of scale in the American corporation since the end of the XIX century and the appearance of hierarchy which, he argues, substituted for the old dominance of the market.

because consumers are willing to pay more for better performances, but also because he plays in larger halls.⁹

The first formal treatment which tries to capture these arguments is Lucas (1978). Lucas' starting point is Manne's (1965) suggestion that: "Observed size distribution is a solution to the problem: allocate productive factors over managers of different ability so as to maximize output". Lucas solves exactly that problem. He departs minimally from the neoclassical production model in that, essentially, a manager is required to run a firm formed by a set of undifferentiated workers and capital. A manager with skill α uses labor n^{10} to produce $\alpha f(n)$ units of output. Lucas' aim is not so much to have a theory of organization, but a theory of why organization matters.

This model generates a superstar effect for the best managers through increases in the scale of operations effects as a result of drops in coordination or communication costs. To see this, let $f(n) = \alpha n^\theta$, where θ governs the decreasing returns that result from larger spans of control. As illustrated in Figure 2, an increase in θ leads each entrepreneur to want a larger team so that $n(\alpha)$ shifts up. As a result, equilibrium requires that both the equilibrium worker's wage, w , and the cutoff skill between workers and managers, α^* , move up, to w' and $\alpha^{*'}$. The top managers can command potentially very large teams, so the marginal value of skill ($f(n(\alpha))$) at that high point in the distribution can be very large. Thus, competition for those top managers can lead to very high earnings –the managerial earnings function, R , becomes steeper as technology increases the marginal returns to managerial skill.


In contrast, as a result of the technological change, nothing happens to wage inequality among workers, since there is none. In this model, what is good for managers is also good for workers, who uniformly earn more (with the potential exception of managers who change occupation and become workers). That is, in this model superstars cast no shadow.

For wage inequality among workers to exist we need to model how their skill matters. Rosen (1982) develops a model of hierarchy where managers increase, like in Lucas, the productivity of the team, but they also need to spend time supervising each worker. In this model worker skill (and not just manager skill) matters and, under constant returns to scale, it enters in efficiency units. Thus, managers here do not care about the specific characteristics of each worker, but only about the aggregate amount of skills they hire. In

⁹Rosen (1981) 'superstar' model provides a formal model and a complete analysis of this scale of resources effects in the labor market. Even a top economist, who would be laboring in obscurity, can become an overnight star selling books worldwide thanks to the reach of information technology. This can happen even to economists who do not appear to believe in the importance of this mechanism, such as Thomas Piketty.

¹⁰In the original formulation of Lucas (1978) n refers to labor per unit of capital.

Figure 2: Lucas (1978): Hierarchy and inequality when worker skill is irrelevant



this respect Rosen shows that this technology can be seen as a generalization of Lucas' technology (Rosen also has a multilayer hierarchy). Namely, $y = \alpha f(Q)$ where now Q is the total number of “efficiency units” of labor supervised by a manager with skill α .


Assume again that $f(Q) = Q^\theta$. Figure 3 illustrates the impact of an increase in θ . As a result of this technological change, which increases the effect of scale, all managers supervise a larger team in terms of Q . Moreover, the price of skill increases (as the demand for it goes up) and thus inequality in differences (although not in ratios, unlike in the empirical evidence) among workers increases. All workers benefit from the new technology and so, as in Lucas (1979), superstars cast a light, not a shadow, on others.

Can these type of models explain the patterns of earnings at the top of the distribution? Gabaix and Landier (2008) develop and calibrate a version of these “superstar” models of Lucas and Rosen, relying on extreme value theory, where the production of the firm is the product of talent and firm size, as measured by the market value of the firm. The key implications of their model are that CEO pay is proportional to firm size, both cross-sectionally and longitudinally, and that for a given firm size CEO pay varies with median firm size across countries. Their empirical evidence is consistent with the model. They find that the firm's market value is correlated with CEO pay in the cross-section. Most strikingly,

in the time series, they show that managerial compensation is almost perfectly correlated with mean market capitalization.

Can these scale theories account of the evolution of pay in the rest of the distribution? The answer is clearly no. In Lucas all workers gain, but there is no inequality among workers. The assumption in Rosen's model that workers are perfect substitutes (only total skill matters) implies, counterfactually, that in equilibrium managers are indifferent between working with a few relatively skilled workers or many unskilled workers. Hence, assignment patterns between individual managers and workers are indeterminate. Crucially, this substitutability also means that the price of skill moves in a uniform direction for all agents, in the face of an increase in scale of operations effect. This implication is clearly counterfactual given the evidence in Figure 1 and motivates a model in which the units of knowledge possessed by workers with different levels of knowledge are not perfect substitutes.

Figure 3: Rosen (1982): Hierarchy and Inequality when only total efficiency units matter



2.2 Workers in the Shadow of Superstars

Work involves solving problems. Confronting a hard problem requires a sufficiently knowledgeable worker. Multiple less knowledgeable workers will not solve the problem if none of

them knows the answer. In this sense, the basic nature of knowledge implies its imperfect substitutability across heterogeneous individuals. If one worker on his own does not know how to program a robot, a team of ten similar workers will also fail.

This type of “problem solving” production function was proposed by Garicano (2000). He shows that, if matching problems to solutions is hard,¹¹ it leads to a knowledge-based hierarchy. Garicano and Rossi-Hansberg (2004 and 2006) embed the model of knowledge-based hierarchies in Garicano (2000) in a context where workers have heterogeneous skill, like in Lucas (1978). In this model, workers are involved in production and managers in problem solving. The key difference is that the production function displays complementarities: workers benefit from having good managers (they solve more problems and make their own time more productive); and managers benefit from having good workers (they need less help and allow managers to have a larger team). Thus, like in Lucas, manager skill determines the productivity of the team. Unlike in that work, here worker skill determines a manager’s team size. The result is an equilibrium that exhibits positive one-to-many sorting between managers and workers.

In our model, since worker skill affects the scale of operations, *the same mechanism that drives increasing inequality at the top, reduces inequality at the bottom*. Namely, *superstars cast a shadow on workers*, particularly on the ones that barely fail to become managers. To understand this note that increases in the scale of operations increase directly, through the mechanism proposed by Lucas, the return to skill for managers. The marginal value of managerial skill is team size, and thus any mechanism that increases team size necessarily increases also the marginal return to skill. The result, as we discussed above, is wage inequality among managers and the concentration of rewards at the very top. It also makes it more demanding to be a manager (raises the required cutoff), further contributing to the inequality between workers and managers.

The main novelty concerns inequality among workers. In a small scale world, a good worker is managed by a much better manager than a bad worker. The marginal product of the good worker is consequently much higher. In a world with large teams, workers are matched with only the best managers, which are similarly skilled, resulting in smaller difference in their marginal product. Hence, as the scale of operations increases, the marginal productivity of differently skilled workers becomes more homogeneous, and so the inequality

¹¹That is, workers “don’t know what they don’t know”. When an agent sees a problem that she cannot solve, she only knows that she does not know the solution.

between them declines.

An example may help. Think of a world with very small firms. The best managers, say Steve Jobs or Alfred Sloan, run small workshops. They match with the best worker in the economy and produce the best products. Now pick up a worker who is only marginally worse. He is assigned to a worse manager and together they make worse (or less) products. Thus, his marginal product is much lower. Now increase the scale of operations so Jobs or Sloan manage both workers. In this new world, the marginal product of both workers is only differentiated by their talent, not by the talent of their boss. The difference between them is smaller and so wage inequality is lower. In relative terms, the better worker loses from the enlarged “shadow of the superstars”.


Figure 4 illustrates the patterns of wage inequality that follows from an increase in the scale of operations (in this model a decrease in h , which parameterizes communication costs, as we explain in detail in the next section). Managers, like in Lucas (1978), experience a superstar effect. Workers are now matched with a smaller range of managers, and experience a reduction in inequality among them. Also, note the kink in the graph. It implies that the “middle classes”, the mid-skilled types, could actually be worse off in levels, not only relative terms, as a result of the technological change. Exactly what happened in the last few decades, as we showed above.¹²

Under what conditions will this mechanism be operational? Essentially, when better workers allow managers to lead larger teams. This requires a production function where managers’ and workers’ skills are complementary, so that positive sorting follows; one with asymmetric skill sensitivity like in Kremer and Maskin (1996), where workers of different skills are not perfect substitutes to one another; as well as scale effects like in Lucas (1978). As we shall discuss next, knowledge-based hierarchies have precisely these characteristics: production takes place in teams, as problems are solved either by workers or their managers (management by exception); there is positive sorting, since more skilled workers ask harder questions and thus require more knowledgeable managers; there are also scale effects since better workers use up less time of their managers and thus allow for larger teams; and naturally result in stratified sorting equilibria.

We move next to the formal part of our review, where we set up and solve these type of problem solving production, derive its properties, and show a variety of formulations and

¹²We develop this model formally below and generalize it along multiple dimensions, including introducing several layers.

Figure 4: Garicano-Rossi-Hansberg (2004): Hierarchy and Inequality



applications. We come back to the issues of inequality, with a more formal treatments of the ideas above, in Section 4.

3 Organization with Ex-ante Homogeneous Workers

3.1 The Basic Model of Knowledge-based Hierarchies

We start developing a basic model of an economy in which ex-ante identical agents acquire knowledge and select into different roles in the organization. Since all agents are identical before they acquire an education, they obtain the same utility independently of their final knowledge acquisition and level in the hierarchy. Thus, this model without ex-ante heterogeneity is inadequate to talk about heterogenous welfare effects, but will allow us to study the structure of the hierarchy and the knowledge acquisition decisions, and therefore wages, of the individuals that form them. Because individuals have no innate characteristics, there is no matching problem, so we simply need to determine the (hierarchical) position of each individual and the knowledge they acquire.

Production with problem-solving.

We introduce here the basic production technology proposed by Garicano (2000). Production requires time and knowledge. In order to produce, a worker must spend her unit of time and solve a problem. If she can solve it, she produces, and otherwise she does not. She can only solve the problems she knows, either because of her innate skills or because she has learnt them by incurring a cost. Thus if a worker draws a problem ζ and her knowledge is in interval $Z \in R^+$, he produces if $\zeta \in Z$.

Suppose that the cost of acquiring a knowledge set Z (learning all the problems in Z) is proportional to its size, i.e. the “number” of problems in it (formally, its Lebesgue measure). Let the cost of learning a unit length of problems be c . For example, it costs cz to learn $Z = [0, z]$.

Some problems are more common than others. The relative frequency of different problems is governed by a continuous distribution function F with density f . The expected output q of a worker is the probability that she confronts a problem she knows, so $q = \int_Z dF$. We assume without loss of generality that problems are indexed so that $f'(z) < 0$. Namely, the most common problems have a lower index. Then, a worker who can draw one problem per unit of time maximizes her expected net output, which is given by

$$q = \Pr(\zeta \leq z) - cz = \int_0^z f(\zeta)d\zeta - cz. \quad (1)$$

The optimal solution is implicitly given by $f(z^*) = c$, which equates the marginal value of acquiring knowledge to the marginal cost: she learns those problems which are “common enough” to justify investing in them.¹³

Communication and production in teams.

Organization allows different workers to acquire different intervals of knowledge and communicate it as required. This has the advantage of allowing workers to increase the utilization rate of knowledge, decreasing the per capita learning cost. The trade-off is that it is costly to match the problem with the worker who knows the solution and to communicate the answer. In this context, hierarchies appear endogenously when matching problems to those who know how to solve them (or “labeling” the problems) is very costly. Then, the optimal way to obtain help is to ask progressively more knowledgeable workers for help until someone

¹³Throughout this survey, we assume the regularity conditions for existence of a unique optimum are satisfied. If the density function $f(z)$ is nonincreasing, the second order condition is always satisfied and the solution is unique.

knows the solution, or until it is sufficiently unlikely that someone can solve it.¹⁴

The communication cost, or “helping cost” h , is the time spent away from production by workers communicating how to solve the problem. Garicano (2000) makes three assumptions about this cost. First, all the communication losses are incurred by the agent who is being asked and helps with the solution. Second, communications costs are incurred even when the agent asked does not know the answer, since she must figure out if she knows the answer and communicate it to the agent who asked.¹⁵ Third, $h < 1$, capturing the idea that agents communicating how to solve the problem are not directly involved, and therefore do not spend time, in producing on the basis of that knowledge.

The problem of the organization is then to determine the number of distinct classes of workers (or layers), as well as their measure, their knowledge, a sequence of classes to ask sequentially for help, and production and helping time so as to maximize output per capita, subject to the time constraint that workers use their whole unit of time either helping or producing. Garicano (2000) shows that in a given organization, any arbitrary original allocation of workers, knowledge, communication and time can be improved, and thus is not optimal, unless it has the following characteristics:

1. *Agents specialize in either production or solving problems. Furthermore, only one class is specialized in production.*
2. *Knowledge acquired by different classes does not overlap.*
3. *Production workers learn to solve the most common problems; problem solvers learn the exceptions. Moreover, the sequence of classes of agents to ask for help is ordered sequentially from the bottom to the top of the hierarchy.*
4. *The organization has a pyramidal structure, with each layer of a smaller size than the previous one.*

The key characteristic of this structure is management by exception. Production workers know solutions to common problems, and ask successively problem solvers who know increasingly exceptional problems. To understand this result, note that any arrangement

¹⁴The assumption that matching problems to knowledge is hard is realistic in situations in which knowledge is hard to codify, as for example when knowledge is tacit. Endogenous codes, where agents have to choose where to set up the boundaries between events, is studied by Cremer, Garicano and Prat (2007).

¹⁵Note that the same problem comes up again with probability zero. Thus, there is no learning involved in communication.

where agents in some layer know problems in an interval $[z, z + \varepsilon]$ which come up more often than the problems known by those asking them questions, $[z', z' + \varepsilon]$ (where $f(z) > f(z')$), can be improved upon. Simply swap these equally long intervals. Learning costs are unchanged, as all agents learn the same *mass* of problems as before. Production is unchanged, as the total amount of knowledge is unchanged. But communication costs are lower, as those asking questions are now less likely to confront a problem they do not know. Thus the crucial insight is that, by relying on “management by exception”, specialization in knowledge can be attained while minimizing communication costs.

Given these results, we can simply let the knowledge of workers in layer 0 be given by $[0, z_0]$, that of layer one managers by $[z_0, z_0 + z_1]$, and, generically, that of agents in layer ℓ by $[Z_{\ell-1}, Z_{\ell-1} + z_\ell]$ where $Z_\ell = \sum_{i=0}^{\ell} z_i$ is the total knowledge in the organization up to layer ℓ and z_ℓ is the measure of knowledge an agent in layer ℓ acquires.

For example, output in an organization with one layer of management and n_0 workers is simply given by $q = F(Z_1)n_0$. In words, production combines the (non overlapping) knowledge of the organization, $F(z_0 + z_1)$, with the production time of the n_0 workers with knowledge z_0 . In choosing the size of her team the manager faces a time constraint given by

$$1 \geq hn_0(1 - F(z_0)). \quad (2)$$

Namely, the unit of managerial time has to be sufficient to help the n_0 workers in her team who ask with probability $1 - F(z_0)$ at a cost of h units of her time. So better workers are more costly as they possess more knowledge, but they facilitate having a larger team, thereby leveraging the knowledge of the manager. Of course, the optimal organization might involve many layers of management, not just one as in this example. The next section characterizes the general problem using a cost minimization approach.

3.2 Optimal Scale, Knowledge and the Number of Layers

In the previous section we introduced the basic technology of knowledge-based hierarchies, some of its key properties, and illustrated the way the technology works with a simple two-layer hierarchy. Garicano (2000) goes further and studies the problem of a single organization that has potentially many layers of management. This problem was reformulated in a classic cost function approach in Caliendo and Rossi-Hansberg (2012), who then embed it in a standard heterogenous firm model based on Melitz (2003).

Consider the cost minimization of a firm that aims to produce q units of output with an organization like the one described above. Let the market price of a unit of work time be given by w , so that the wage of an agent with one unit of time and z units of knowledge is $cz + w$. The cost of producing q units of output with an organization with L layers, the wage bill of such an organization, is given by

$$C^L(q) = \min_{\{n_l^L, z_l^L\}_{l=0}^L} \sum_{l=0}^L n_l^L (cz_l^L + w), \quad (3)$$

subject to the organization producing at least q units of output,

$$F(Z_L^L) n_0^L \geq q, \quad (4)$$

and the time constraints of managers at each layer,

$$n_l^L = h(1 - F(Z_{l-1}^{l-1})) n_0^L \text{ for } l \leq L. \quad (5)$$

The last set of constraints say that the number of units of time of managers in layer l is given by the number of problems (or workers), n_0^L , times the fraction of them that has not been solved up to layer l , $(1 - F(Z_{l-1}^{l-1}))$, times the amount of time managers need to spend with each problem, h .

Given a number of layers, the solution to this problem determines the structure of the efficient organization to produce q units of output. Practically, it determines the knowledge levels and the number of agents in each of them that the organization needs to employ. Note that the problem above does not impose a constraint on the number of managers at the top of the organization, n_L^L . Absent such a constraint, the cost function is linear in q . This means that, given the number of layers, there is an optimal organization that is simply replicated to produced the desired level of output. A small corner shop will use a few minutes of the same organization used by a large corporation like Walmart. This is clearly an unappealing feature.

Larger organization use more complex organizations with more layers in their corporate hierarchy. As we review in Section 6, Caliendo et al. (2014) present systematic evidence that larger French firms, in fact, use more complex hierarchies. One natural way of incorporating this realistic feature is to add an integer constraint on the number of managers at the very top of the organization. Because of their tacit knowledge and skills, it is hard for CEOs

to replicate themselves. For example, we can assume, as in Caliendo and Rossi-Hansberg (2012), that each organization requires one agent at the top, so

$$n_L^L = 1. \tag{6}$$

With this constraint in place, one can show that the cost function is U-shaped in output, with a minimum efficient scale that increases with the number of layers and a total wage bill at the minimum efficient scale that decreases with the number of layers. So organizations that produce more use more layers and produce, typically, at a lower cost. Size gives an advantage by making a complex organization with many layers the efficient way to produce. In this sense, independently of any exogenous cost advantage firms might have, larger firms (say due to higher demand for a better product) will also have an endogenous cost advantage as a result of their more efficient organization. Clearly, such a large organization is only efficient if the firm produces a large quantity. For small production levels such an organization is extremely inefficient. More formally, the cost function is given by the envelope of $C^L(q)$ for all L , namely,

$$C(q) = \min_{L \geq 0} C^L(q). \tag{7}$$

In reduced form, one could think of this organizational problem as choosing a technology from a menu of pairs of fixed and marginal costs. Large firms or organizations choose options with larger fixed costs and lower marginal costs and small firms do the inverse. The organizational problem above goes beyond this reduced form choice by specifying the characteristic of the workforce within the organization that leads to each particular combination.¹⁶

Empirical Implications: How firms grow.

A full characterization of this problem for the case of an exponential distribution of problems, $F(z) = 1 - e^{-\lambda z}$, can be found in Caliendo and Rossi-Hansberg (2012). Suppose the managerial integer constraint at the top, $n_L^L = 1$, determines the boundary of the firm so that organizations form as firms. Then, one central implication of the organizational problem above is that when firms reorganize to grow due to, say, a positive demand shock –by adding a layer of management–, they reduce the knowledge and wages of all employees in preexisting layers. Namely, they employ workers that are less knowledgeable. In contrast, when they

¹⁶See Qian (1994) for an alternative hierarchical model with many layers using a theory of monitoring rather than knowledge-based hierarchies. See Chen (2014) for an equilibrium version of that model with heterogenous firms.

grow without adding layers, they upgrade the knowledge of their employees and pay them more. The intuition is simple. When the organization grows substantially it prefers to add one or a few experts (a new layer) rather than making all the existing workers learn more. The latter is too expensive since it involves training a potentially large number of workers. Hence, instead, the firm adds a layer of management and economizes on its wage bill by reducing the knowledge of everyone else. In contrast, when no layer is added, the firm has to increase the knowledge of everyone in order to increase its scale without violating the time constraints of its managers. Caliendo et al. (2014) present evidence that this implications are very much in line with the evidence for French firms, as we discuss in Section 6.

Empirical Implications: ICT and the organization of production.

A second set of implications concern the impact of information and communication technology on the hierarchy and the knowledge acquired. The theory distinguishes clearly between these two aspects of technology. The key trade-off is between acquiring knowledge and asking others for help. Technologies that make acquiring knowledge cheaper reduce communication and allow workers to do more without relying on help, while technologies that make communicating information cheaper increase the hierarchy.

As expected, both a decline in c and h shift down the cost function making production more efficient. Obviously, h only affects the cost function when the organization consists of more than just a layer of workers, since in that case no communication is required. So both of these changes are technological improvements. However, they have distinct effects on the equilibrium organizations used. A decline in c makes knowledge cheaper and therefore less useful to build a complex organization.¹⁷ After all, the whole point of the organization is to economize on the wage bill by using knowledge efficiently. Hence, the model predicts that such a decline would increase the scope of decision making by lower-level workers, increase the span of control of supervisors, increase the ratio of production workers to problem solvers, and reduce the number of layers of management.

In contrast, a decline in h implies that the mechanism to economize on wages by using a complex organization is more effective, and so more layers of managers are used. This makes some agents learn and earn less, as discussed above, but some others gain as they manage larger teams and can leverage their knowledge more. We discuss in Section 6 some evidence that information and communication technologies do indeed have different effects.

¹⁷In fact, for the exponential case the only relevant parameter is c/λ .

3.3 Growth: Exploiting and Exploring Knowledge

The models above study the hierarchical organization from a timeless perspective. Of course, in reality hierarchies are set up slowly over time. Garicano and Rossi-Hansberg (2012) extend the model above to analyze how this slow evolution affects technology use and innovation. We differentiate between two knowledge-generating activities: exploiting existing technology and innovating to develop new technologies. First, exploitation takes place as organizations undertake production over time and add new layers (new referral markets) of experts. That is, we model this process as the sequential emergence of a collection of markets for expert services where agents sell the problems they cannot solve to other agents. By allowing these new experts to leverage their knowledge about unusual problems, the new layers allow for more knowledge to be acquired and make production more efficient under the current technology. This process exhibits decreasing returns, since eventually most problems are well known and the knowledge acquired is less and less valuable.

Second, innovation is the result of agents' decisions about how much to invest to create radically new technologies. This investment process exhibits convex costs, so that the investment, if it happens, takes place smoothly over time. Of course, the economy's ability to exploit the new technology through organization determines the profitability of innovation investments. The rate of innovation, the extent of exploitation, and the amount of organization in the economy are jointly determined in the theory and depend on the cost of acquiring and communicating knowledge.

If it happens, progress in our model takes place in leaps and bounds. After a new technology is adopted, investment in innovation decreases and agents concentrate on exploitation as they first acquire the more productive pieces of knowledge about this technology and then the rarer ones. Radical innovation will not take place again until the current innovation has been exploited to a certain degree. Both the timing of the switch to a new technology and the size of the jump in the technology are endogenous, since agents must choose how much and when to invest in radical innovation. As long as the value of continuing with an existing technology is sufficiently high, the switch to the new technological generation does not take place. As in Arrow (1974), organizations are specific to a particular technology, and so adopting the new technology makes the knowledge acquired about the previous technology obsolete. It requires agents to start accumulating new knowledge and start building new organizations. Hence, inherent in new knowledge is a process of creative destruction (Schumpeter, 1942) whereby adopting a radical innovation makes the existing organization

obsolete.¹⁸

Progress may also come to a halt if agents decide not to invest in radical innovation. Specifically, the payoff of exploiting existing technologies may be such that agents optimally create very large organizations, composed of a large set of referral markets and a large number of different specialized occupations. Such organizations take a long time to build, and thus agents choose to postpone forever the moment in which a new technology would be exploited. The radical innovation process never gets started. When the current technology is already well exploited, agents do not have any development of the alternative radical technology to build on and prefer to exert their efforts on small advances of the existing technology. The result is stagnation.

Empirical Implications: ICT and economic growth.

ICTs determine the depth to which an innovation is exploited and, thus, the rate of growth. Consider first information technology. The main benefit of organization is that individuals can leverage their cost of acquiring knowledge over a larger set of problems, increasing the utilization rate of knowledge. Information technology advances (e.g., databases) decrease the need for organizational complexity, shorten the exploitation process, and unambiguously increase growth.

Better communication technology unambiguously increases welfare, but has an ambiguous effect on growth: for intermediate communication costs –and thus for intermediate cost of exploiting a particular technology– organizations may get stuck in an old technology forever. The intuition is quite clear. When communication technology is expensive, organizations do not grow very large and so the amount of “organizational learning” for a given technology is small; so organization does not give old technologies a considerable advantage over new technologies. Similarly, when communication technology is inexpensive, technologies are extensively exploited via organization and so technological improvements are very valuable. Thus, switching to a new technology is attractive and frequent. For intermediate values of communication costs, large organizations are useful for exploiting a technology, and the potential improvement in technology might not be valuable enough to induce switching. Hence, making organizations more efficient, by reducing communication costs from high to intermediate levels, shifts the balance of economic activity from investing in new innovations to exploiting more fully existing innovations, and this may reduce long-term economic growth,

¹⁸Previous models of creative destruction, following on the pioneering work of Aghion and Howitt (1992) and Grossman and Helpman (1991), do not take organizations into account– new products substitute for the old, but organizations play no role.

potentially all the way down to zero.

4 Organization with Ex-ante Heterogeneous Agents

4.1 Sorting and Inequality

We introduce here organizations that utilize the skills of heterogeneous workers and retake, more formally, the issue of wage inequality. We start with the simplest possible model (Garicano and Rossi-Hansberg, 2004), one where hierarchies only have two layers, and agents have exogenously given knowledge or, equivalently since here it is exogenously given, skill. This model generates the results described previously in Section 2.2.

A firm has n low skill agents (“workers”), with skill z_p and 1 high skill agent (a “manager”) with skill z_m . Agents are organized in a knowledge hierarchy as explained above. Specifically, the n workers attempt to solve one problem each. If they fail, they request help or directions from the manager and re-attempt their problem using those directions. Normalize z so that the distribution of problems $F(z)$ is uniform in $[0, 1]$.¹⁹ Let the distribution of skill in the population be given by $\phi(z)$.

The output produced by a team is then simply $z_m n$, where z_m is the knowledge or skill of its manager. This is subject to the time constraint of the managers which, as in Section 3 above, is given by: $hn(1 - z_p) = 1$. The span of the manager is limited by the knowledge of production workers; if production workers are more knowledgeable, they will need help less often, and managers will be able to supervise larger teams.²⁰ A manager with skill z_m maximizes her earnings $R = z_m n - w(z_p)n$ by his choice of the skill of the workers she employs, namely,

$$R = \max_{z_p} \frac{z_m - w(z_p)}{h(1 - z_p)}. \quad (8)$$

Solving for an equilibrium in this economy is a continuous assignment problem (see Sattinger, 1993) with two twists relative to standard assignment problems. First, who is assigned to whom is not a given, but an equilibrium outcome. In standard assignment

¹⁹This normalization is not in the original exposition in Garicano and Rossi-Hansberg (2004) but simplifies the exposition.

²⁰Note that there are no decreasing returns to scale in this production function directly, but only through the congestion in the time constraint of the manager. We thus have a trade-off between congestion and skill: a manager can increase team size if he so desires, but he needs to hire more skilled workers to do so. This implies that matching matters.

models this identity of the agents is assumed (e.g. by assuming one to one matching of pre-specified masses of agents with distinct types, or between agents and machines, industries, etc.). In contrast, here we are “marrying” a mass of workers with a mass of managers— *where those roles and masses are not given by assumption*. This property is the result of integrating an occupational choice, where agents can decide freely what job to take, and many-to-one matching, where the span of control of managers is an equilibrium outcome itself. Second, agents can decide not to be matched and produce on their own; their output in that case is just $F(z) = z$.

Optimality requires positive sorting, that is, workers with more skill must be assigned to managers with more skill. The reason is that the production function is supermodular in both skills. A better manager must leverage his higher knowledge over a larger number of workers, and that requires workers to be more skilled so that they can deal with more problems on their own. Technically this supermodularity is guaranteed by a positive cross-derivative of output $z_m/(h(1 - z_p))$ with respect to z_m and z_p .

To characterize the equilibrium in this economy we need to describe three objects: the allocation of agents to occupations (i.e. production workers, managers, and self-employed); the team composition (i.e., the matching between workers and managers and the optimal spans of control of managers); and the earnings function.²¹

Note that the slope of the wage equation must be such that the first order condition holds. So, for all workers ($z_p \leq z^*$, as we show below),

$$w'(z_p) = hn(z_p)(z_m - w(z_p)) = \frac{z_m - w(z_p)}{1 - z_p} \text{ for all } z_p \leq z^*. \quad (9)$$

Together with a boundary condition, equation (9) determines a unique wage function, $w(z)$, specifying the wage of each individual z_p for a given match (z_p, z_m) . The wage slope for those agents who do not join in a hierarchy (self-employed) is given by $w'_I(z) = 1$. Clearly, $w'(z_p) < w'_I(z_p) = 1$ for all workers (since for agents who are workers $w(z_p) > z_p$, and so $w'(z_p) < (z_m - z_p)/(1 - z_p) < 1$). Using the envelope theorem, we have that the earnings function of managers satisfies

$$R'(z_m) = n(z_p) > 1 = w'_I(z_m). \quad (10)$$

²¹Occupational choice works as in Lucas (1978) and Rosen (1982), but here there is a matching function that, together with the wage function, is pinned down in equilibrium.

This ordering of the slopes implies that an equilibrium with self-employed agents will be characterized by a pair of thresholds (z^*, z^{**}) , such that all agents with skill $z < z^*$ become production workers, all agents with skill $z > z^{**}$ become managers, and those in between are self-employed. Furthermore, optimal occupational choice requires that for a marginal worker, $w(z^*) = z^*$, for a marginal manager, $R(z^{**}) = z^{**}$, and $R(1) = 1/h$. Equations (9) and (10), together with these three boundary conditions, determine the full earnings function and one threshold, say, z^* .

Consider now the labor market clearing condition. Again, here the fact that one-to-many worker–manager matching takes place in equilibrium makes a difference. Instead of summing over all workers and over all managers, we need pointwise equalization of supply and demand. Let $m(z_p)$ denotes the skill of the manager assigned to a workers with skill z_p , which by positive sorting is an increasing function. The labor market clearing condition is then given by

$$\int_0^{z_p} \frac{1}{n(z)} \phi(z) dz = \int_{z^{**}}^{m(z_p)} \phi(z) dz \text{ for all } z_p \leq z^*. \quad (11)$$

Since (11) holds for all $z_p \leq z^*$, and $n(z) = 1/((1-z)h)$, we can differentiate with respect to z_p , to obtain $m'(z_p) = h(1-z_p)$, which, together with $m(0) = z^{**}$ and $m(z^*) = 1$, determines the equilibrium assignment function $m(z)$ and the remaining threshold z^{**} .²²

Empirical Implications: Superstars and their shadow.

The model provides some implications also found in Lucas (1978) and Rosen (1982). The main one is the correspondence between the firm size distribution and distribution of managerial earnings. More specifically, the implication that managerial wages grow faster than ability thanks to the scale effects generated by larger teams sizes (although here workers wages also grow more slowly with skill). The theory also provides some novel implications. It pins down, for the first time we believe, the occupational distribution in the economy: the less skilled agents are production workers, the medium skilled agents are independent or self employed, and the most skilled agents are managers or problem solvers. Moreover, the composition of each team is also pinned down. It is not just the case that more skilled managers have more subordinates, but those subordinates are also of better quality. This

²²An equilibrium with only two layers, as well as a positive mass of self-employed agents, exists only in part of the parameter space. For example, h needs to be sufficiently large. See Garicano and Rossi-Hansberg (2004) and Antràs, Garicano and Rossi-Hansberg, (2006) for details. Fuchs, Garicano and Rayo (2014) solve the (equivalent) planner’s problem, which has some additional measure theoretic complications.

matching in turn has implications for the wages of workers. As we shall see, all three properties (stratification, positive sorting, and skill-scale effects) are empirically tested in Garicano and Hubbard (2012).

As discussed extensively but informally in Section 2.2, a reduction in communication costs has the following consequences. First, it increases firm size for all managers, since team size satisfies the manager’s time constraint. Second, this leads to an increase in the managerial superstar effect, since the marginal return to managerial talent is, analogously to Lucas (1978), proportional to team size, $n(z_p)$. Third, and more unique to our framework, it reduces wage inequality among workers, as a result of the sorting effect. Intuitively, for any two workers z and z' , the skill of the managers they are matched with is closer to one another the larger are team sizes (formally, $w''(z) = h$).²³

Empirical Implications: Organizations and Power Laws

Geerolf (2014) develops a version of the theory described above but with exogenous skills and multiple layers. He shows that the distribution of spans of control for agents in firms with L layers is a Pareto distribution with coefficient $2^L/(2^L - 1)$. Hence, under minimal assumptions on the distribution of skills, as the number of layers increases the coefficient of the distribution of spans of control converges to 1. Using the empirical methodology in Caliendo et al. (2014) to identify layers (described in Section 6), he can account extremely well for the distribution of firm sizes with different numbers of layers in the French data. The key insight from this work is that the matching process with complementarities that we have discussed here can magnify even tiny skill differences and lead to Pareto distributions both in wages and firm sizes.

4.2 Endogenous Skill and Sorting

The previous setting abstracted from a variety of dimensions in order to gain tractability and deliver cleanly our results. Chief among these abstractions is that knowledge is exogenous, which forces us to abandon the critical distinction between information and communication technology. Furthermore, it simplifies the organization by determining exogenously that teams can have, at the most, a layer of managers and one of workers. Garicano and Rossi-Hansberg (2006) analyze the “full blown” model where agents are ex-ante heterogeneous, they

²³We discuss this further in Section 5 where we also present the exact mathematical expression for a case that can be solved in closed form.

can acquire knowledge endogenously, and they can form teams with an arbitrary number of layers.

Garicano and Rossi-Hansberg (2006) assume agents differ in their cognitive ability or skill so that higher ability agents incur lower learning costs. Specifically, the distribution of ability in the population can be described by a continuous density function, $\alpha \sim \phi(\alpha)$, with support in $[0, 1]$.²⁴ In particular, ability is defined so that the cost of learning to solve an interval of problems of length 1 is given by²⁵

$$c(\alpha; t) = t - \alpha. \tag{12}$$

The cost of knowledge acquisition is then $c(\alpha; t)z$. As required for comparative advantage, high ability types have a comparative advantage in knowledge acquisition. A decrease in t represents an improvement in information technology that decreases the cost of learning (e.g., a technology that decreases the cost of accessing knowledge, such as cheaper data-base storage and search).

The problem is somewhat tricky to solve, as the knowledge acquisition, the matching, and the number of layers (multiple hierarchies coexist), are endogenous in equilibrium. A given agent, with a given skill, could well be an entrepreneur, a middle manager, or a worker, and in each case would be matched with entirely different other agents. We do not go over the details of the construction of the equilibrium, the reader is referred to Garicano and Rossi-Hansberg (2006). Suffice is to say that the problem can be solved recursively. First, the first layer managers choose the knowledge of workers, then the second layer managers choose the knowledge of first layer managers, etc. Finally, entrepreneurs or CEOs (the agents for whom it is not profitable to pass problems to other experts) choose the knowledge of the managers one layer below, as well as their own knowledge.

Garicano and Rossi-Hansberg (2006) characterize the optimal organization as follows:

- **Equilibrium properties:** The equilibrium is unique, and it involves positive sorting, so that more skilled workers are matched with more skilled managers and more skilled entrepreneurs. The equilibrium wage function, $w(\alpha) : [0, 1] \rightarrow \mathbb{R}_+$, is increasing and convex. Furthermore, the knowledge function $z(\alpha) : [0, 1] \rightarrow [0, 1]$ is increasing. Oc-

²⁴Consistent with the notation in the previous sections, we reserve z for knowledge and introduce α for the exogenous skill, or ability, of an individual.

²⁵The linearity of the learning cost function $c(\cdot)$ in α and the restriction of the support to $[0, 1]$ is without loss of generality, since we can always scale α to fit these restrictions.

cupational stratification, as described above, still holds. The lowest skilled agents are workers, then the next ones are middle managers at progressively higher layers, and the top agents in the economy are entrepreneurs.

- **Knowledge substitution:** Organization allows for the substitution of the knowledge of less skilled workers for that of their supervisors. Thus, organization decrease the knowledge (and therefore the value) of less skilled agents. In contrast, organization increases the return from learning about difficult tasks (as the knowledge can be leveraged more). Thus, organization increases the knowledge acquired by more skilled agents. As a result organization decreases the marginal value of worker skill and increases the marginal value of the entrepreneur’s skill.
- **Wage inequality and organization:** It follows that organization makes low ability agents (and in particular workers) more equal, thereby reducing earnings-inequality *within* this group. In contrast, it makes high skilled agents more unequal, thereby increasing *within* earnings-inequality among them. Finally, it increases the gap between low and high skill agents, which results in increases in *between* earnings inequality.

Empirical Implications: ICT, wage inequality, and knowledge substitution.

The framework above allows for the separate analysis of each form of ICT on wage inequality. Communication technology (h) increases inequality at the top and reduces it at the bottom, as we have emphasized before. The introduction of endogenous knowledge accentuates this effect through the knowledge substitution effect, whereby the knowledge of subordinates is substituted by that of their supervisors. In contrast, information processing and knowledge acquisition technology (parameterized by t) increases within class inequality for both low and high skilled agents. Everyone acquires more knowledge which creates larger within class differences. The effects on between class wage inequality are thus ambiguous, but technological change unambiguously decreases firm size, increases the knowledge (“empowerment”) of workers, and reduces layers of management. Garicano and Rossi-Hansberg (2006) argue that the timing of innovation in these technologies is consistent with the evidence on the temporal evolution of firm sizes, wage inequality and CEO/worker earnings ratios.

4.3 Contracting with Double-sided Adverse Selection

In many of the previous discussions we have implicitly, and sometimes explicitly, equated hierarchies with firms. Of course, nothing in our setting requires the exchange of knowledge in a hierarchy to take place within a firm. Garicano and Rossi-Hansberg (2006) shows that, in a world with perfect information, consulting and referral markets can support the exact same equilibrium outcomes.

The above allocations put a high informational burden on agents. Absent some a priori knowledge about the persons conducting the exchange, both the difficulty of the questions posed to consultants and the knowledge of those consultants are hard to assess. A consultant could pretend to be more knowledgeable than he really is, and the agent asking for help could also misrepresent the difficulty of the question.

Such a setting with double sided adverse selection is studied by Fuchs et al. (2014). They show that in the presence of these frictions spot contracting is not efficient. However, they show that an ex-ante, firm-like contractual arrangement uniquely delivers the first best. This arrangement involves hierarchies in which consultants are full residual claimants of output and compensate producers via incentive contracts.

This characterization of the optimal ex-ante contract looks very much like some firms in the knowledge intensive sector (e.g. in professional services). There, the highest skilled agents hold all the equity and hire a number of less knowledgeable associates who are paid a combination of a fixed wage and a bonus for performance. Thus, the existence of these double-sided information problems shows why these knowledge-based hierarchies are often constituted as firms.

Relatedly, Acemoglu et al. (2014) consider the problem of crowdsourcing, where a principal who owns a set of unsolved problems contracts with agents (either his employees or external contractors) who attempt to solve these problems. They show that the first-best allocation of talent, which involves an endogenous hierarchy of problem solvers, is implementable regardless of informational asymmetries.²⁶

5 International Organizations: Trade and Offshoring

The theories above are useful to study a variety of issues concerning international trade and offshoring (the formation of international organizations and production processes). We have

²⁶See Fuchs, Garicano and Rayo (2014) for a comparison of both approaches.

studied these issues in a variety of papers including Caliendo and Rossi-Hansberg (2012), for the case of international trade, and Antràs et al. (2006 and 2008), for the case of offshoring.²⁷

5.1 Trade and Organization

Assessing the impact of the large increases in international trade that the world has experienced in the last several decades involves understanding the effect that reductions in trade costs have had on the characteristics of producers and their labor force. Bernard and Jensen (1997, 1999) initiated a large literature by pointing out the distinct characteristics of exporters relative to other firms. Melitz (2003) then provided the workhorse model to explain these observations using a mechanism that selected the best firms as exporters as the result of fixed exporting costs. In his model organization does not play a role, so the marginal cost of a firm remains constant as it starts to export. A firm is a technology that is immutable to the markets in which it sells.

This is not the case once we introduce organization into the problem. Parallel to Section 3.2, as an exporter grows due, for example, to a trade liberalization, it will do so by adding workers, sometimes layers, and will modify the knowledge and wages of its labour force accordingly. Furthermore, it will increase its productivity (but reduce its revenue-based productivity) as a result of that reorganization. The reverse is also true for non-exporters that tend to shrink or exit as a result of a bilateral trade liberalization, and will reduce their productivity.

The above reasoning implies that international trade, by changing the optimal production size of firms and therefore their organizational form, affects the efficiency of production of individual firms as well as aggregate productivity. Caliendo and Rossi-Hansberg (2012) show that the welfare gains from trade are declining in the communication costs, h , but increasing in the information costs, c . Trade allows firms to grow larger by selling to more markets. The larger scale makes using a large, more complex, organization economical. This channel leads to additional gains from trade. These additional gains are larger the more efficient is the organizational technology, namely, the lower are communication costs. In contrast, the additional gains are smaller if the cost of acquiring knowledge are smaller since, as we have

²⁷See also Gumpert (2014) for a discussion of multinational firms and how they determine the hierarchical organization of headquarters and affiliate plants. Marin and Verdier (2010, 2008), and Marin, Schymik, and Tarasov (2014) have used the monitoring-based theory of organization to study offshoring and the organization of international production. For a more general review on the large literature on organization and international production processes see Antràs and Rossi-Hansberg (2009).

underscored above, the purpose of organization is to economize on these costs. In fact, gains from trade asymptote to the ones in a standard Melitz model as c tends to zero since in that case the organizational problem disappears.

5.2 Offshoring and Organization

In Antràs et al. (2006) we address a different problem related to globalization: the formation of international production teams and its implications. Feenstra and Hanson (1997) and Goldberg and Pavcnik (2004) have shown that globalization has led to increases in wage inequality among workers in less developed countries. This has been viewed as puzzling in the international trade literature since trade, either in goods or in tasks, should lead to increases in the price of the abundant factor, as dictated by the standard Heckscher-Ohlin logic.²⁸ This is, however, not necessarily the case once we analyze the matching between workers in global production teams.

The model in Antràs et al. (2006) is simple enough to be solved analytically. Consider the model we used in Section 4.1 but with a uniform distribution of skills between 0 and u and assuming that parameters are such that there are no self-employed agents. In this case we can solve the system of two differential equations explicitly. The wage function for workers (agents with $z \in [0, z^*(h, u)]$) is given by

$$w(z) = z^*(h, u) - \sigma(h, u)(1 - z) + \frac{1}{2}hz^2. \quad (13)$$

The skill premium is just the derivative of equation (13) with respect to z , namely,

$$w'(z) = \sigma(h, u) + hz. \quad (14)$$

The skill premium has two parts. The first term is the baseline price of a unit of skill and is determined by the supply and demand of skills in this economy. This is what we call the “competition effect”. The second term is increasing in the number of units of skill an individual possesses and increases with communication costs: the “complementarity effect”. The source of this effect is that having more units of skill allows the worker to be matched to better managers and earn more per unit of skill. As we discussed above, this effect declines as teams get larger (which results from an improvement in communication costs). In the limit,

²⁸Grossman and Rossi-Hansberg (2008, 2012) analyze in detail the factor price implications of reductions in task trading costs in frameworks without hierarchical organization.

as h goes to zero, every agent works with the best manager and this source of inequality disappears.

Globalization allows this economy to form international production teams with other economies that might have different distributions of skills. Call the economy above, with a distribution of skills between 0 and u , the “South”. In order to economize on parameters, let the “North” be identical but with $u = 1$.²⁹ The result is a world economy with a mass of skills equal to 1 for $z > u$ and equal to $1 + 1/u$ for $z < u$.

The ability to form international production teams leads workers in the South to match with better managers in the North, which increases the price of a unit of skill in the South. Furthermore, for the worst workers, the “complementarity effect” increases as well due to the larger density of workers relative to managers. Finally, more Southern agents become workers so there is a compositional effect that increases worker’s wage inequality. Hence, consistent with empirical observations, wage inequality for Southern workers unambiguously rises. Firms in the South exit as a result of the formation of international teams.

In the North, the formation of international teams also leads to an increase in wage inequality among managers. Among workers, the baseline price of skill decreases (the relative number of unskilled workers in the world is larger than in the North) but they face the same enhanced “complementarity effect”. Thus, in the North, if communication technology is good enough so the first effect dominates, globalization can lead to a decrease in wage inequality at the bottom together with an increase in inequality at the top. This is exactly what happened in the U.S. in the last few decades, as we discussed in the context of Figure 1. *So superstars, and the shadow they cast on workers, can also be the result of reductions in offshoring costs.*

6 Empirics of Hierarchies

The importance of the hierarchical structure of organizations for worker, firm, and aggregate outcomes is not obvious. Even though the theory we discussed above can forcefully make the claim that worker, firm, and economy-wide performance, can be affected by the way agents organize in production teams, it could be the case that in reality firms just expand by replicating their operations. The returns to a individual’s skill and knowledge could be independent of the other agents she work with. The fact that the models above are consistent

²⁹Here we only study the implications from offshoring to a particular country. Antràs, Garicano, and Rossi-Hansberg (2008) study the role of local communication technology and the distribution of skills in making agents in target countries more attractive as production partners.

with the facts presented in Section 2 provides only indirect evidence in favor of these theories. We discuss below some more direct evidence that fails to falsify the mechanisms we have emphasized.

6.1 Firm Growth, Reorganization, and Wages

Caliendo et al. (2014) use French firm and worker level data to analyze the relationship between firm characteristics, firm growth, and a firm’s organization. To do so they use wages as their measure of worker skill and/or knowledge and divide the firm in layers using hierarchical occupational definitions. The hierarchical definitions identify a maximum of 4 layers per firm.³⁰

This data allows the authors to create a picture of the organization of French manufacturing firms between 2002 and 2007. Importantly, their study uses a large institutional dataset instead of more specialized datasets of a few firms, as in previous studies.³¹ The authors can follow firms, and the wages that they pay in each occupation, as they add and drop layers over time (namely, as they hire workers in a hierarchical occupation where they did not employ anyone before). Figure 5,³² presents the resulting representative hierarchies, divided by number of layers, for 553,125 firm-year observations. The results are quite stark. Firms with more layers are larger in terms of value added, total hours of work employed, and pay higher hourly wages on average. Furthermore, a vast majority of them are hierarchical in that higher layers employ less hours and pay higher hourly wages.

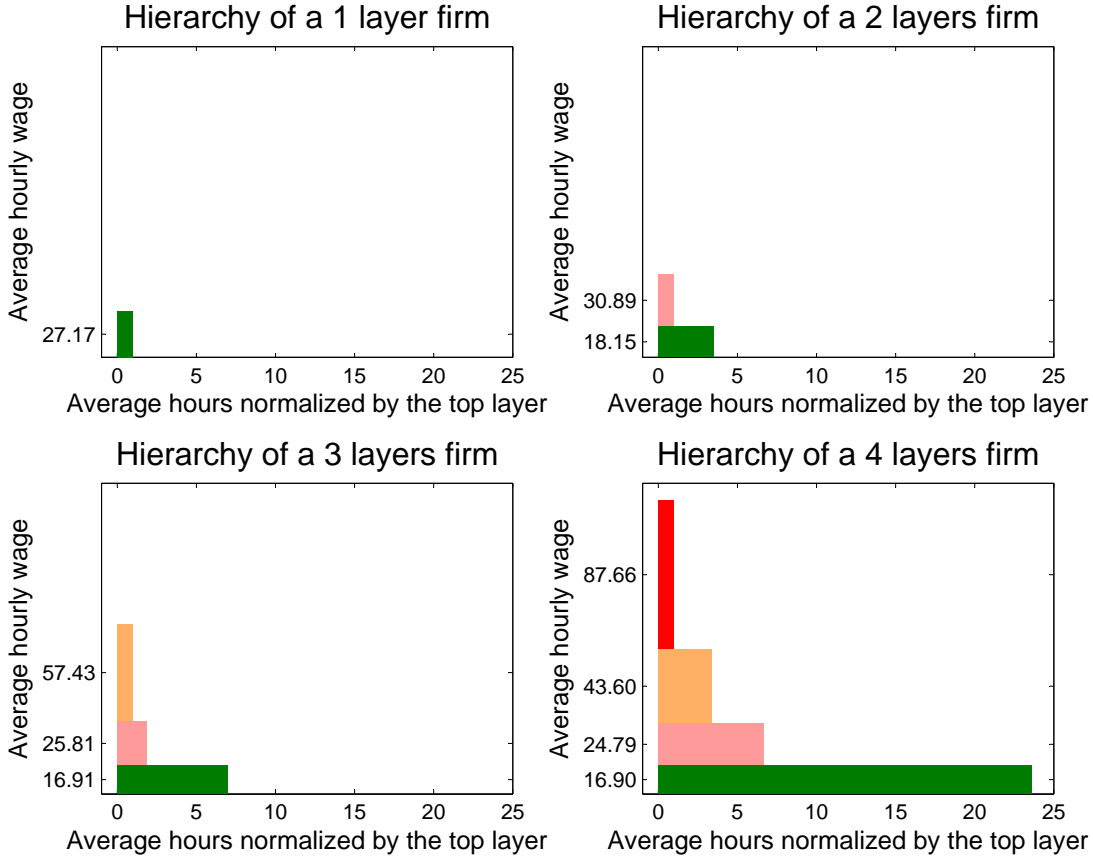
The results also show that transitions between layers are quite common, but mostly by adding one or dropping one layer. Furthermore, the probability of adding (dropping) a layer is increasing (decreasing) in a firm’s value added. Caliendo, et al. (2014) also show that when firms expand or contract substantially they tend to reorganize by adding or dropping layers. Dropping and adding layers is also associated with internal changes in the characteristics and wages of the employees of the firm. Figure 5, shows that firms with only one layer pay

³⁰The paper uses a match of the BRN dataset, which includes firm balance-sheet information, and the DADS dataset, which includes worker level occupation, hours of work, and earnings of salaried employees.

³¹The previous literature on organizational form has used very detailed information on a small set of firms. The best example is Baker, Gibbs, Holmstrom (1994) and Baker and Holmstrom (1995) who analyze one firm. Other studies like Rajan and Wulf (2006) and Guadalupe and Wulf (2010) study datasets with a few hundred firms. Caroli and van Reenen (2001) use larger surveys of English and French firms but do not focus on the hierarchical structure of firms.

³²Reproduced from Figure 5 in Caliendo, Monte, and Rossi-Hansberg (2014). The reader is referred to that paper for details. The figure shows the average wage in 2005 Euros in each layer (vertical axis) and the number of hours, normalized by the number of hours in the top layer (horizontal axis).

Figure 5: Representative Hierarchies in France



their labor force 27.17, 2005 Euros, per hour on average. In contrast, firms with 2 layers, pay the workers at the bottom only 18.15 Euros and firms with 3 layers only 16.91 Euros. That is, as firms add layers they hire workers in pre-existing layers that earn on average less. This is the case for all comparisons in Figure 5. Of course, these figures average over firms that are heterogenous in a variety of dimensions, and so these results might be generated by compositional effects. In their paper, Caliendo, et al. (2014) show that the reduction in average wages in all pre-existing layers is present for all transitions in firms with any number of layers when one uses time series evidence with firm fixed effects. Furthermore, as the theory predicts, this reduction (increase) in wages is only present when firms add (drop) layers but not when they expand or contract without reorganizing. In that case, as the theory above predicts, average wages in a layer grow when the firm expand and decline when it contracts.³³

³³Tag (2013) has replicated many of these qualitative results for Swedish firms, and ongoing projects have

Overall, these results show that layers of management vary systematically with firm characteristics. Firms seem to actively manage their production hierarchy. It is important then to condition on reorganization when analyzing changes in firm characteristics. Moreover, the impact of any economy-wide change that affects firm size will be mediated by these reorganizations. For example, the routinization of work (as firms grow and add more layers) reduces the knowledge and wages of workers in the middle of the wage distribution the most. Look at the top managers in 4-layer firms in Figure 5, who make on average €87.66 per hour. Their direct subordinates make only €43.6 rather than the €57.43 agents in the same layer make in 3-layer firms.

6.2 Sorting, Spans of Control, and the Distribution of Wages

Several papers have studied the effect of larger spans of control on wages. Smeets and Warzynski (2008) and Fox (2009) both find a positive effect of the number of subordinates, and the level of responsibility, on wages; the former for a large European firm and the latter for Swedish and U.S. firms.³⁴ Using 1992 Census data for U.S. Law firms,³⁵ Garicano and Hubbard (2007 and 2012) compare earnings among individuals who are at the same hierarchical rank, and show that those who work with more associates per partner (i.e. more workers per manager) earn more.

Garicano and Hubbard (2012) go further and present evidence that associate earnings are higher at offices where partner earnings are higher. Although this does not necessarily reflect that associates' and partners' ability is positively correlated (it could be driven by office-level demand shocks), this evidence is certainly consistent with equilibria with positive assortative matching. The paper also finds that lawyers are stratified as in the model in Section 4.1. Namely, independent lawyers earnings are in between those of associates and partners in large law firms.³⁶

also replicated them for Danish and Portuguese firms.

³⁴See also the classic study by Brown and Medoff (1989) that relates firm size and wages.

³⁵The data covers the entire legal industry in the US from the 1992 Census of Services, which had a questionnaire specific for law firms.

³⁶The exception are very large cities where associates of large law firms sometime can earn more than independent lawyers. Perhaps a sign that in large cities law firms have more than two layers with some associates operating in middle-rank positions.

6.3 Information Technology and the Organization of Work

Bloom et al. (2014) aim to separate the two distinct components of information and communication technology along the lines of our work here. Their test involves a key prediction of the theory: lower information costs shift decisions to lower layers, while lower communication costs shift them to higher ones. They test this using both decisions shifted from the bottom to the middle of the hierarchy and from the middle to the top. They rely on a proprietary firm survey on management practices which they merge with private sector data on use of ICT by 1000 firms. Although they do not make any causal claims for their mechanism, they find indeed that network technologies (a communication technology) result in autonomy shifting from plant managers up to the center and from workers up to plant managers. In contrast, the two IT technologies they study (ERP and CAD/CAM) lead to more decentralized decision making.

This is the first paper that, thanks to the rich data available, can evaluate empirically the direct implications of the knowledge-based hierarchy model concerning the different effects of information and communication technology.³⁷

7 The Road Ahead

We have argued in this survey that understanding equilibrium outcomes through the lens of theories that put at their core knowledge-based hierarchies can illuminate a variety of economic phenomena. The way in which individuals organize in teams to produce determines the productivity of these teams, as well as the rewards that individuals can command for their work. Once we view economics through this organizational lens, the evolution of technology, and its distinct forms, takes a different meaning. Communication technologies enhance the hierarchy by making it a better technology to economize on knowledge. Information technology makes knowledge cheaper and therefore organization less useful. We have argued that the evolution of wage inequality in the different parts of the income distribution can be well understood as a response to the important changes in ICT that we have observed in the last few decades. Communication technology is highlighting the advantages of superstars and is making the less skilled more equal, thereby hurting the middle classes. This is what

³⁷Studying the introduction of IT-based credit scoring models in a Colombian bank, Paravisini and Schoar (2012) also find evidence that improvements in information technology decentralize decisions by making managers dedicate more time to the hardest decisions.

we have dubbed, the *shadow of superstars*.

Of course, much work is still needed to develop the core theories, implications, and empirical evidence that ensue from this organization-based view of economics and, specifically, the view centered on knowledge-based hierarchies. Particularly promising, we believe, are three avenues. The first one involves dynamics and economic development. We outlined in Section 3.3 some of our work on the topic, but much more is needed. In particular, we need theories where agents decide on the timing of the development of heterogeneous hierarchies. Furthermore, we need to link hierarchies with the career paths of individuals and see whether progressive promotions within and across organizations can explain the facts on individual wage dynamics over the life-cycle.³⁸

Second, these frameworks should be used to evaluate a variety of tax policies, and other distortions, that affect the labor market, firm characteristics, and international transactions. The effect of firm size distortions, as studied in Garicano, Lelarge, and Van Reenen, (2013), and in the context of knowledge-based hierarchies by Torres-Coronado (2014), are good examples. There is room for much more work analyzing the effect of minimum wages, other forms of corporate taxation and subsidies, place based-policies that affect the spatial location of firms and therefore their organization, education and training subsidies, among many others policies.

Finally, a lot more empirical work is needed. The evidence in Section 6 establishes the basic empirical credibility of this approach. But much more empirical work is required to understand the effect that different economic shocks have on the organizations that we see in equilibrium and the implied distribution of wages. In particular, we need causal evidence of the effect of specific shocks, like trade liberalizations, policy reforms, or idiosyncratic firm-specific demand and productivity shocks, on organization. We also need better evidence of the effect that the large changes in technology, and specifically ICT technology, have had on the actual changes in organizations and labor market outcomes that we have observed. Can we establish an unequivocal causal link? Can we assess the magnitudes of these effects?

One could continue this list, of course. Can these organization-based theories explain the productivity gap between firms in developed and developing countries? Can they explain business cycle issues like slow recoveries or asymmetric cycles? We think that they can and that they will; but we are still far from getting there. In writing a survey like this, one is always left somewhat encouraged by the road travelled but also awed by the enormity of the

³⁸Some papers like Roys and Seshadri (2014) are starting to make progress on this front.

task ahead. Undoubtedly, we have focused unfairly on the part of this endeavour that we have contributed to. We apologize for this and encourage others to actively participate in this agenda.

8 References

1. Acemoglu D. 2002. Technical Change, Inequality, and the Labor Market. *J. of Econ. Lit.* 40:7-72
2. Acemoglu D. 1999. Changes in Unemployment and Wage Inequality: An Alternative Theory and Some Evidence. *Am. Econ. Rev.* 89:1259-78
3. Acemoglu D, Autor D. 2011. Skills, Tasks and Technologies: Implications for Employment and Earnings. *Handbook of Labor Econ.* Elsevier.
4. Acemoglu D, Mostagir M, Ozdaglar A. 2014. Managing Innovation in a Crowd. *NBER Work. Pap.* 19852
5. Aghion P, Howitt P. 1992. A Model of Growth Through Creative Destruction. *Econometrica*, 60:323–51
6. Antràs P, Garicano L, Rossi-Hansberg E. 2006. Offshoring in a Knowledge Economy. *Q. J. Econ.* 121:31-77
7. Antràs P, Garicano L, Rossi-Hansberg E. 2008. Organizing Offshoring: Middle Managers and Communication Costs. Helpman, E., D. Marin, and T. Verdier, *The Organization of Firms in a Global Economy*, Harvard University Press
8. Antràs P, Rossi-Hansberg E. 2009. Organizations and Trade. *Annual Rev. of Econ.*, 1:43-64
9. Arrow K. 1974. *The Limits of Organization*. Norton.
10. Autor D, Katz L, Kearney M. 2008. Trends in U.S. Wage Inequality: Revising the Revisionists. *Rev. of Econ. and Stat.*, 90:300-23
11. Autor D, Katz L, Kearney M. 2006. The Polarization of the U.S. Labor Market. *Am. Econ. Rev.* 96:189-94

12. Autor D, Levy F, Murnane R. 2003. The Skill Content of Recent Technological Change: An Empirical Exploration. *Q. J. Econ.* 118:1279-333
13. Baker G, Gibbs M, Holmstrom B. 1994. The Wage Policy of a Firm. *Q. J. Econ.* 109:921-55
14. Baker G, Holmstrom B. 1995. Internal Labor Markets: Too Many Theories, Too Little Evidence. *Am. Econ. Rev.* 85:255-9
15. Bebchuk L. 2005. The Growth of Executive Pay. *Oxford Rev. of Econ. Policy.* 21:283-303
16. Bernard A, Jensen JB. 1997. Exporters, Skill Upgrading and the Wage Gap. *J. of Inter. Econ.*, 42:3-31
17. Bernard A, Jensen JB. 1999. Exceptional Exporter Performance: Cause, Effect, or Both? *J. of Inter. Econ.*, 47:1-25
18. Bloom N, Sadun R, Garicano L, van Reenen J. 2014. The Distinct Effects of Information Technology and Communication Technology on Firm Organization. *Management Science*. In press
19. Prescott E, Boyd J. 1987. Dynamic Coalitions: Engines of Growth. *Am. Econ. Rev.* 77:63-67
20. Brown C, Medoff J. 1989. The Employer Size-Wage Effect. *J. of Polit. Econ.* 97:1027-59
21. Caliendo L, Rossi-Hansberg E. 2012. The Impact of Trade on Organization and Productivity. *Q. J. Econ.* 127:1393-467
22. Caliendo L, Monte F, Rossi-Hansberg E. 2014. The Anatomy of French Production Hierarchies. *J. of Polit. Econ.* In press
23. Calvo G, Wellisz S. 1978. Supervision, Loss of Control, and the Optimum Size of the Firm. *J. of Polit. Econ.* 86:943-52
24. Caroli E, van Reenen J. 2001. Skill-Biased Organizational Change? Evidence From A Panel Of British And French Establishments. *Q. J. Econ.* 116:1449-92

25. Chandler A. 1977. *The Visible Hand: the Managerial Revolution in Am. Business*, Belknap Press
26. Chen C. 2014. Management Quality, Firm Organization And International Trade. Unpublished manuscript, University of Hong Kong, Hong Kong
27. Crémer J, Garicano L, and Prat A. 2007. Language and the Theory of the Firm. *Q. J. Econ.* 122:373-407
28. Demsetz H. 1988. The Theory of the Firm Revisited. *J. of Law, Econ. and Org.* 4:141-61
29. Fox J. 2009. Firm-Size Wage Gaps, Job Responsibility, and Hierarchical Matching. *J. of Labor Econ.* 27:83-126
30. Fuchs W, Garicano L, Rayo L. 2014. Optimal Contracting and the Organization of Knowledge. *Rev. of Econ. Stud.* In press
31. Garicano L. 2000. Hierarchies and the Organization of Knowledge in Production. *J. of Polit. Econ.*, 108:874-904
32. Garicano L, Hubbard T. 2007. Managerial Leverage Is Limited by the Size of the Market: Theory and Evidence from the Legal Services Industry. *J. of Law and Econ.* 50:1-44
33. Garicano L, Hubbard T. 2012. Learning about the Nature of Production from Equilibrium Assignment Patterns. *J. of Econ. Behavior and Organization.* 84:136-53
34. Garicano L, Lelarge C, Van Reenen J. 2013. Firm Size Distortions and the Productivity Distribution: Evidence from France. Unpublished manuscript, LSE, London
35. Garicano L, Rossi-Hansberg E. 2012. Organizing Growth. *J. of Econ. Theory.* 147:623-56
36. Garicano L, Rossi-Hansberg E. 2006. Organization and Inequality in a Knowledge Economy. *Q. J. Econ.* 121:1383-435
37. Garicano L, Rossi-Hansberg E. 2004. Inequality and the Organization of Knowledge. *Am. Econ. Rev.* 94:197-202

38. Geerolf F. 2014. A Static and Microfounded Theory of Zipf's Law for firms and of the Top Labor Income Distribution. Unpublished manuscript, UCLA, Los Angeles
39. Grossman G, Helpman E. 1991. Quality Ladders in the Theory of Growth. *Rev. of Econ. Stud.* 58:43-61
40. Grossman G, Rossi-Hansberg, E. 2012. Task Trade between Similar Countries. *Econometrica.* 80:593-629
41. Grossman G, Rossi-Hansberg, E. 2008. Trading Tasks: A Simple Theory of Offshoring. *Am. Econ. Rev.* 98:1978-97
42. Guadalupe M, Wulf J. 2010. The Flattening Firm and Product Market Competition: The Effect of Trade Liberalization on Corporate Hierarchies. *Am. Econ. J.: Applied Econ.* 2:105-27
43. Gumpert A. 2014. The Organization of Knowledge in Multinational Firms. Unpublished manuscript, University of Munich, Munich.
44. Ijiri Y, Simon H. 1964. Business firm growth and size. *Am. Econ. Rev.* 54: 77-89
45. Kaplan S, and Rauh J. 2013. It's the Market: The Broad-Based Rise in the Return to Top Talent. *J. of Econ. Perspecti.* 27:35-56
46. Katz L, Murphy K. 1992. Changes in Relative Wages, 1963-1987: Supply and Demand Factors. *Q. J. Econ.* 107:35-78
47. Knight F. 1921. *Risk, uncertainty and profit.* Hart, Schaffner and Marx
48. Kremer M, Maskin E. 1996. Wage Inequality and Segregation by Skill. *NBER Work. Pap.* 5718
49. Lemieux T. 2006. Postsecondary Education and Increasing Wage Inequality. *Am. Econ. Rev.* 96:195-9
50. Lucas Jr. R. 1978. On the Size Distribution of Business Firms. *Bell J. of Econ.,* 9:508-23
51. Manne H. 1965. Mergers and the Market for Corporate Control. *J. of Polit. Econ.* 73:110-20

52. Marin D, Schymik J, Tarasov A. 2014. Trade in Tasks and the Organization of Firms. Unpublished manuscript, University of Munich, Munich
53. Marin D, Verdier T. 2010. Corporate Hierarchies and International Trade: Theory and Evidence. Unpublished manuscript, University of Munich, Munich
54. Marin D, Verdier T. 2008. Competing in Organizations: Firm Heterogeneity and International Tradein. *The Organization of Firms in a Global Economy*, E. Helpman, D. Marin and T. Verdier (eds), Harvard University Press
55. Mayer T. 1960. The distribution of ability and earnings. *The Rev. of Econ. and Statistics*. 42:189-95
56. Melitz M. 2003. The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity. *Econometrica*. 71:1695–725
57. Paravisini D, Schoar A. 2012. The Information and Agency Effects of Scores: Randomized Evidence from Credit Committees. Unpublished manuscript, LSE, London
58. Piketty T. 2014. *Capital in the Twenty-First Century*, Belknap Press
59. Prescott E, Visscher M. 1980. Organization Capital, *J. of Polit. Econ.* 88:446-61
60. Qian Y. 1994. Incentives and Loss of Control in an Optimal Hierarchy. *Rev. of Econ. Stud.* 61:527-44
61. Radner R. 1993. The organization of decentralized information processing. *Econometrica*. 61:1109-46
62. Rajan R, Wulf J. 2006. The Flattening Firm: Evidence on the Changing Nature of Firm Hierarchies from Panel Data. *Rev. of Econ. and Stat.* 88:759-73
63. Rosen S. 1982. Authority, Control, and the Distribution of Earnings. *The Bell J. of Econ.* 13:311-23
64. Rosen S. 1986. The Theory of Equalizing Differences. *Handbook of Labor Econ.* Elsevier
65. Roys N, Seshadri A. 2014. Econ. Development and the Organization of Production. Unpublished manuscript, University of Wisconsin, Wisconsin.

66. Sattinger M. 1993. Assignment Models of the Distribution of Earnings. *J. of Econ. Lit.* 31:831-80
67. Schumpeter J. 1942. *Capitalism, Socialism and Democracy*. Routledge
68. Smeets V, Warzynski, F. 2008. Too Many Theories, Too Few Facts? What the Data Tell us about the Link Between Span of Control, Compensation and Career Dynamics. *Labour Econ.* 15:687-703
69. Tåg J. 2013. Production Hierarchies in Sweden. *Econ. Letters.* 121:210-13
70. Torres-Coronado J. 2014. Size-dependent Firm Regulations and the Return to Skill in a Many-to-one Assignment Model. Unpublished manuscript, ITESM, Mexico City

CENTRE FOR ECONOMIC PERFORMANCE
Occasional Papers

- | | | |
|----|---|---|
| 42 | Chiara Criscuolo
Peter N. Gal
Carlo Menon | Dynemp: A Stata® Routine for Distributed Micro-data
Analysis of Business Dynamics |
| 41 | Nicholas Bloom
Renata Lemos
Raffaella Sadun
Daniela Scur
John Van Reenen | The New Empirical Economics of Management |
| 40 | Ghazala Azmat
Barbara Petrongolo | Gender and the Labor Market: What Have We Learned
from Field and Lab Experiments? |
| 39 | Christopher J. Boyce
Alex M. Wood
James Banks
Andrew E. Clark
Gordon D.A. Brown | Money, Well-Being, and Loss Aversion: Does an Income
Loss Have a Greater Effect on Well-Being than an
Equivalent Income Gain? |
| 38 | Nicholas Bloom | Fluctuations in Uncertainty |
| 37 | Nicholas Oulton | Medium and Long Run Prospects for UK Growth in the
Aftermath of the Financial Crisis |
| 36 | Phillipe Aghion
Nicholas Bloom
John Van Reenen | Incomplete Contracts and the Internal Organisation of
Firms |
| 35 | Brian Bell
John Van Reenen | Bankers and Their Bonuses |
| 34 | Brian Bell
John Van Reenen | Extreme Wage Inequality: Pay at the Very Top |
| 33 | Nicholas Oulton | Has the Growth of Real GDP in the UK been Overstated
Because of Mis-Measurement of Banking Output? |

32	Mariano Bosch Marco Manacorda	Social Policies and Labor Market Outcomes in Latin America and the Caribbean: A Review of the Existing Evidence
31	Alex Bryson John Forth Minghai Zhou	What Do We Know About China's CEO's? Evidence from Across the Whole Economy
30	Nicholas Oulton	Hurray for GDP
29	Stephen Machin	Houses and Schools: Valuation of School Quality through then Housing Market – EALE 2010 Presidential Address
28	John Van Reenen	Wage Inequality, Technology and Trade: 21 st Century Evidence
27	Barry Anderson Jörg Leib Ralf Martin Marty McGuigan Mirabelle Muûls Laure de Preux Ulrich J. Wagner	Climate Change Policy and Business in Europe Evidence from Interviewing Managers
26	Nicholas Bloom John Van Reenen	Why do Management Practices Differ Across Firms and Countries?
25	Paul Gregg Jonathan Wadsworth	The UK Labour Market and the 2008-2009 Recession
24	Nick Bloom Raffaella Sadun John Van Reenen	Do Private Equity Owned Firms Have Better Management Practices?
23	Richard Dickens Abigail McKnight	The Impact of Policy Change on Job Retention and Advancement
22	Richard Dickens Abigail McKnight	Assimilation of Migrants into the British Labour Market

For more information please contact the Publications Unit
Tel: +44 (0)20 7955 7673; Fax: +44 (0)20 7955 7595; Email: cep.info@lse.ac.uk
Website <http://cep.lse.ac.uk>