

Incomplete Contracts and the Internal Organization of Firms

Phillipe Aghion*
Harvard University

Nicholas Bloom**
Stanford University

John Van Reenen***
London School of Economics

We survey the theoretical and empirical literature on decentralization within firms. We first discuss how the concept of incomplete contracts shapes our views about the organization of decision-making. We then overview the empirical evidence on the determinants of decentralization and on the effects of decentralization on firm performance. A number of factors highlighted in the theory are shown to be important in accounting for delegation, such as heterogeneity and congruence of preferences as proxied by trust. Empirically, competition, human capital, and IT also appear to foster decentralization. There are substantial gaps between theoretical and empirical work and we suggest avenues for future research in bridging this gap (*JEL* O31, O32, O33, F23).

1. Introduction

Grossman and Hart (1986) developed the incomplete contracts approach to analyze the costs and benefits of vertical integration, which could explain why firms have boundaries, and why not all transactions take place within a single firm. The basic idea is that contracts cannot specify all states of nature or all actions in advance, or there are states of nature or

*Department of Economics, Harvard University, Cambridge, MA, USA, the CEPR and the NBER. Email: paghion@fas.harvard.edu.

**Department of Economics, Stanford University, Stanford, CA, USA, the CEP, CEPR and the NBER. Email: nbloom@stanford.edu

***Department of Economics, London School of Economics, London, UK, the CEP, CEPR and the NBER. Email: J.Vanreenen@lse.ac.uk

This paper was prepared for the “Grossman and Hart at 25” conference in Brussels. We are grateful for comments by our discussants, John Roberts and Bentley MacLeod, the editor and two anonymous referees, and research assistance from Megha Patnaik. We would like to thank the Economic and Social Research Council for their financial support through the Centre for Economic Performance. The empirical part builds heavily on joint work with Raffaella Sadun. A more detailed working paper version is available online as Aghion et al. (2013).

The Journal of Law, Economics, and Organization, Vol. 30, Supplement 1
doi:10.1093/jleo/ewt003

Advance Access published July 19, 2013

© The Author 2013. Published by Oxford University Press on behalf of Yale University.
All rights reserved. For Permissions, please email: journals.permissions@oup.com

actions that cannot be verified *ex post* by third parties, and which therefore are not *ex ante* contractible. They used this approach to develop theories of ownership and vertical integration. When contracts cannot specify all possible uses of an asset, the contract must *ex ante* leave some discretion over the use of the assets: in other words it must allocate ownership of the asset to one or the other party. The benefit of integration is that the owner avoids hold-up by the other party, which in turn will enhance her incentives to invest in the relationship. The cost is that the other party will tend to under-invest in the relationship.

In this article, we show how the incomplete contract approach can be used to think about the *internal organization* of firms. The first half of the article has a theoretical focus; we look at how formal and real authority are allocated between the firm's owner and its employees, between the top managers and subordinates, and between firms and financiers. We also examine how the allocation of authority affects communication within the firm.

In the second half of the article, we analyze some of the empirical literature, examining first the determinants of organization (focusing on delegation/decentralization) from the perspective of the Grossman-Hart approach and its extensions; and second the effects of decentralization on firm performance. Our emphasis is on looking at "stylized facts" from large-scale econometric studies of firms rather than at case studies. Case studies are helpful in suggesting theoretical approaches and mechanisms, but are poor for hypotheses testing as they are small in number and highly selective. As with the theory, we also look at decentralization *within* firms, rather than the more commonly studied issue of boundaries of the firm or vertical integration, which was the original motivation for Grossman and Hart (1986).¹ We find some support for aspects of the incomplete contracts approach in the importance of the congruence of preferences and firm heterogeneity for decentralization. However, we acknowledge that there are many other stylized facts from the empirical decentralization literature that may require alternative theoretical perspectives.

The structure of this article is as follows. Section 2 examines theory. We look at delegation and authority (Section 2.1), financial contracting (Section 2.2), and delegation and the informational content of decision making (Section 2.4).² Section 3 examines the empirical determinants of firm decentralization, focusing on some predictions of the theory (such as the importance of preference congruence as proxied by trust). Section 4 analyzes the effect of decentralization on firm performance. Section 5 concludes and suggests areas for future research.

1. See Lafontaine and Slade (2010) for a comprehensive review of the empirical literature on the issue of firm boundaries.

2. In Aghion et al. (2013) Appendices A and C we also discuss work on academia and multidivisional firms.

2. Theory

2.1 A Model of Delegation and Authority in Organizations

2.1.1 Basic Model. Real authority, that is, the ability to make decisions, requires information. But acquiring information in turn requires time and effort. Thus, for example, the CEO of a big holding company that consists of several horizontally integrated units, can only devote limited attention to each unit, which in turn implies that more real authority will lie with downstream agents in each unit. In fact, increasing the “span of control” is one way in which a top manager can commit to leave more real authority, and therefore more initiative, to her subordinates in various branches of activity. More generally, it is the design of the organization, together with the allocation of formal decision rights, that will determine how real authority is distributed within the firm.

The issue of real versus formal authority and of the implications of this distinction for the optimal design of firm organization is addressed by Aghion and Tirole (1997) using an incomplete contracts/property rights approach.

Their basic framework involves two parties: P (principal) and A (agent). It is assumed that *formal authority* can be allocated contractually (e.g. shareholders allocate authority to the board of directors). In turn, boards allocate authority to management—and management to different layers of management, and so on. In contrast, *real authority* is exerted by the party which has information; this may be the party with formal authority, but not necessarily so. Contractual incompleteness is again key to the whole analysis: contracts signed *ex ante* between P and A cannot specify particular project choices, as these are not verifiable by third parties.

After the contract is signed, both P and A can invest in information acquisition: by investing effort $\frac{1}{2}E^2$, P acquires information with probability E . Similarly, by investing effort $\frac{1}{2}e^2$, A acquires information with probability e . An important parameter in the analysis of the costs and benefits of delegating formal (or real) authority to A is the degree of congruence between P's and A's preferences. Let α denote the probability that P's preferred project is also A's preferred project (call this congruence between the two parties' preferences), and suppose that a party gets zero utility if the other party chooses her preferred project and preferences are not congruent. Finally, assume that an uninformed party will never pick a project at random as this might be too risky.

The timing of moves is as follows. First, the two parties sign a contract that allocates formal authority to one party, either P or A. Then, both parties invest in information acquisition, that is, P and A choose E and e , respectively. Then, if she acquires information, the party with no formal authority proposes a project to the party with formal authority. The party with formal authority then either picks her preferred project (if she herself has acquired information) or she picks the project proposed to her by the party without formal authority if she did not acquire information. It is in

this latter case that *real authority* differs from *formal authority*, since the project is actually chosen by the party *without* formal authority (the party with formal authority is uninformed and therefore can only rubber stamp the other party's project proposal).

P delegating formal authority to her agent A involves a cost and a benefit. The cost is that the agent may choose a project which the principal does not prefer. This is the *loss of control* effect. The benefit is that delegating formal authority to the agent encourages her to invest more effort (i.e., higher e) in information acquisition. This is the *initiative* effect. Which effect dominates depends upon the congruence parameter α : there exists a cut-off value α^* such that for $\alpha < \alpha^*$, the first effect dominates, and it is better for P to retain formal authority, whereas for $\alpha > \alpha^*$, the second effect dominates, and it is better for P to delegate formal authority to A.

Since preference congruence turns out to be so critical for decentralization decisions, we examine this in the empirical section where there does appear to be some compelling evidence for the importance of empirical proxies for congruence, such as *trust*. Below we show that this is a robust prediction of generalizations to the basic theoretical approach.

2.1.2 Extensions. Subsequent to Aghion and Tirole (1997), AT, several papers have analyzed the allocation of formal authority internally in organizations. We shall describe some of these attempts in the next subsections. At this stage, let us mention a first attempt by Hart and Moore (1999), HM. HM analyze the optimal allocation of authority in multi-layer hierarchies. Their model is one where by assumption, upstream agents are less likely to get new ideas (getting an idea in HM is like obtaining information in AT) due to their higher span of control. However, when they do get an idea, this idea has higher potential because of their greater span. HM then show that it is optimal to have "chains of commands," whereby whenever they have an idea, upstream agents (the "generalists") have priority rights over implementing the idea; only if they do not give an idea, can downstream agents (the "specialists") have their say on which action to implement. The intuition is that although upstream agents are less likely to get a new idea, having priority control rights makes sure that they are in control of all the assets downstream, which in turn allows them to fully realize the idea's potential. But if they fail to get a new idea, then the next downstream agent on each branch of the hierarchy should have her say if she has an idea, and so on, moving down in the hierarchy.

So far, we have concentrated on the allocation of formal authority within organizations. However, going back to AT, it could be that delegating formal authority to A is too costly to P, for example, because with arbitrarily small probability, A might take some very costly action. In that case P will always want to retain formal authority, but yet she may want to commit herself not to invest too much in information acquisition, so as to

preserve A's incentives to invest in e even though she keeps formal authority with herself. One way to achieve such commitment is through the choice of *span of control*. More specifically, by increasing the span of control, that is, the number of agents and activities under her supervision, the principal will commit to limiting how much effort she devotes to acquiring information on each particular activity. This in turn will encourage initiative by agents on each activity, as they anticipate that the principal will ignore their proposals less often (as she will not have acquired the relevant information). The choice of the optimal span of control by the principal at date zero is in turn subject to the same trade-off between the principal's loss of control and the agents' initiatives as above. This trade-off also underlies other features of organizational design, such as the role of intermediaries, the costs and benefits of having multiple principals on some activities, or the optimal combination of tasks within teams.

The idea that the design of organizations can serve as a commitment device to delegate (real) authority, is further explored by Acemoglu et al. (2007), who use it to test their theory of the determinants of decentralization on French and British firm-level panel data (see Section 3 for more empirical detail). The model, closely related to AT, is one in which the owner of a firm in a given sector can learn about the outcome of an investment decision through observing other firms in the same sector, or by relying on the superior information of downstream agents (or on downstream agents' effort to acquire information) within the firm. The more precise the public information acquired through observing other firms in the same sector, the less a firm needs to delegate control to its better informed agent. This simple observation delivers a rich set of predictions on the determinants of decentralization. In particular it suggests that older firms should delegate less, as these firms will have had time to learn from more predecessors in the same sector. It also suggests that greater the firm heterogeneity in the same sector, the more a firm in that sector should delegate control as what it observes from other firms is less likely to be relevant for its own choices (if firms are very different it is harder for them to learn from each other). Finally, it suggests that a firm closer to the technological frontier should delegate control more, as it is more likely to face problems that have not been solved before by other firms in the same sector. In the empirical section later in the text, we discuss the empirical tests of these three predictions and show that they all receive support in the data.

Even if decentralization was the efficient choice due to characteristics of the firm's environment, Baker et al. (1999) emphasize that delegation is often informal because the corporate head quarters (CHQ) must usually sign-off on decisions. The issue is whether the CHQ credibly commits to allowing the plant manager to effectively make important decisions and does not override the plant manager (in order to establish her reputation of not interfering). Thus, the extent of decentralization is the outcome of a repeated game between the CHQ and manager. Again, trust may facilitate

a cooperative outcome to this repeated game, which suggests that regions or countries that have higher levels of trust should enjoy greater firm decentralization.

Another reason why delegating formal authority within firms may be difficult is explained by Bolton and Dewatripont (2011): nothing prevents a principal (the owner of a firm) who delegates formal authority to an agent to revert her decision at any time. Bolton and Dewatripont point to the so-called Business Judgment Rules, which prevent courts from enforcing contracts between several parties within the same firm. On the other hand, there are instances where transfer of formal authority can be enforced nevertheless. First, if this transfer is accompanied by a transfer of information from the principal to the agent, then this information transfer guarantees irreversibility of the transfer of control (see Aghion et al. (2004) for a more detailed discussion of this point). Second, in a more dynamic context, principals may want to establish a reputation for not reverting control allocation decisions over time, precisely to keep the option of credible control transfers in the future. Third, as stressed by Hart and Holmstrom (2010), taking control back from an agent causes the agent to become aggrieved which in turn may induce the agent to “shade,” that is, to take unobservable actions which are damaging to the firm. Finally, as pointed out by Bolton and Dewatripont (2011), there are at least two examples of contracting situations where Business Judgment Rules do not apply: financial contracting, that is, contracts between a firm and its investors, which we discuss in the next subsection, and universities, where the faculty are protected by contracts and tenure commitments which grants them academic freedom (see Aghion et al. (2009) and Appendix A of Aghion et al. (2013)).

2.1.3 Summary and Empirical Implications. There are many other important theoretical aspects of firm delegation stemming from GH and AT. The next two sections will focus on finance and information, as space constraints prevent us from elaborating on other important areas such as academia and multidivisional firms (the Appendices of Aghion et al. (2013), summarize these contributions).

As a general point, although decentralization models in the GH-AT tradition generate a rich set of predictions, it is fair to say that relatively few of these have been subject to rigorous empirical examination. One main reason for this is that it is hard to develop empirical analogs of theoretical objects such as decentralization, information, and communication in one firm, let alone in large-scale databases suitable for econometric analysis. In Section 3, we focus on some areas where empirical progress has been made.

2.2 Financial Contracting and the Role of Contingent Control Allocations

Since the work of Modigliani and Miller (1958)—which suggested that the mix of debt and equity a firm has does not affect its value—economists

have wondered why most firms have some combination of debt and equity financing. Debt has certain tax advantages (interest is typically tax deductible for the firm but dividend payments are not), yet corporate debt was prevalent even before corporate income tax existed. Why would firms have debt, then?

Aghion and Bolton (1992) develop an incomplete contracts model that provides a rationale for holding debt. They argue that the mix of debt and equity financing divide up the states of the world where debt and equity holders have control of the firm's assets. When times are good, and cash flows are sufficient to meet interest payments, equity holders have control. When times are bad, debt holders get control.

Suppose that an entrepreneur needs to finance a project that costs $K = 10$. She does not have private wealth and thus needs funding from an outside investor. The investor cares only for monetary benefits, whereas the entrepreneur only draws private benefits from taking various actions. Actions are not verifiable by a third party and therefore cannot be contracted upon ex ante. Hence, all the initial contract can do is to allocate control rights between the two parties. The timing of the relationship between the entrepreneur and the investor can be described as follows. At the contracting stage, the two parties write a financial contract that allocates control rights. The contract must be "feasible," that is, it must satisfy the investor's ex ante participation constraint (she must get at least as much as her outside option in expectation). Then the state of nature is realized. Suppose there is a good state θ_g and a bad state θ_b , each of which can occur with probability $1/2$, and which state is realized is verifiable by a third party. Then an action must be chosen. Suppose that only two actions can be chosen: a_1 and a_2 . Action a_1 maximizes monetary revenue in all states, whereas action a_2 maximizes private benefits in all states. For example, let $\pi(a, \theta)$ denote the monetary profit from taking action a in state θ , and $B(a)$ denote the private benefit from taking action a in any state θ and suppose that

$$\pi(a_1, \theta_g) = 11, \pi(a_1, \theta_b) = 13$$

$$\pi(a_2, \theta_g) = 10, \pi(a_2, \theta_b) = 6$$

and

$$B(a_1) = 2 < B(a_2) = 4.$$

The first best involves a_1 being chosen in state θ_b and a_2 being chosen in state θ_g , since

$$B(a_1) + \pi(a_1, \theta_b) = 15 > 10 = B(a_2) + \pi(a_2, \theta_b)$$

whereas

$$B(a_1) + \pi(a_1, \theta_g) = 13 < 14 = B(a_2) + \pi(a_2, \theta_g).$$

We can now compare three governance structures, which correspond to three types of financial contracts. Entrepreneur control (e.g. as implemented through issuing nonvoting shares), would lead to action a_2 being chosen in all states, but this would violate the investor's participation constraint since

$$\frac{1}{2}(6) + \frac{1}{2}(10) = 8 < 10 = K.$$

The investor has cash which he might use to convince the entrepreneur to take action a_1 in state θ_b . However, if most of the bargaining power at the renegotiation stage lies with the entrepreneur, the prospect of ex post renegotiation will not help satisfy the investor's ex ante participation constraint. In this case, entrepreneur control is not ex ante feasible. How about investor control (e.g., as implemented through issuing voting equity)? In this case, the investor will choose action a_1 in all states, even though action a_2 is the first-best action in state θ_b , that is, the action that maximizes the sum of monetary and private benefits in that state. Now, can the entrepreneur renegotiate the action from a_1 to a_2 in state θ_g ? The answer is no, simply because the entrepreneur has no cash she can use to bribe the investor into changing his choice of action. Investor control satisfies the investor's participation constraint, since

$$\frac{1}{2}(11) + \frac{1}{2}(13) = 12 > 10 = K.$$

However it is not first-best, as total surplus would be maximized by having action a_1 taken in state θ_b and action a_2 taken in state θ_g .

Now, consider a contract that specifies a contingent allocation of control—to the entrepreneur in state θ_g and to the investor in state θ_b . This contract will lead to action a_1 being taken in state θ_b and action a_2 being taken in state θ_g . Contingent control can in turn be implemented through a debt contract that transfers control from the entrepreneur to the investor in state θ_b .

The idea that contingent control can help align incentives goes beyond financial contracting. For example Bolton and Dewatripont (2011) restate the Aghion-Bolton model as one of a headquarter who commits to replace a divisional manager only in some states of nature (e.g., in state θ_b for the above notation). Under this reinterpretation ($B(a_2) - B(a_1)$) can be seen as the private cost for the division manager of closing down his unit and thus having his employees laid off, rather than maintaining the unit in operation. This contingent arrangement reduces the expected monetary losses of the headquarters and simultaneously internalizes the private costs of closing down units. In a context where the information about θ could be manipulated by the division managers, for example, through account manipulation or through risky decision making (gambling for resurrection) so as to avoid being replaced, contingent control can help mitigate the problem by offering a guarantee to divisional managers. Another

guarantee is to give managers a stake in the firm's profits even when they are being replaced (see Bolton and Dewatripont (2011) and Garicano (2000)).

2.3 Delegation as a Way to Improve the Informational Content of Decision-Making

Dessein (2002) analyzes how the allocation of control can help incorporate the agent's information into decision-making in a situation where the agent has private information. In contrast to Aghion and Tirole (1997), there is no information acquisition effort by the agent or the principal, therefore in Dessein's model the allocation of authority is not so much a tool to motivate the agent (as in Aghion and Tirole) or to give a supplier incentives to make relationship specific investments (as in Grossman and Hart). The main insight of Dessein (2002) is that in a world with asymmetric information and contractual incompleteness, the delegation of authority from a principal to an agent is often the best way to elicit the agent's private information.

In Dessein's setting, the agent is assumed to be better informed, but with preferences over decisions that are not fully congruent with those of the principal. If the principal has authority (i.e., if she holds the decision rights), which Dessein refers to as "centralization," then the agent communicates his information "strategically" in order to tilt the principal's decision. Centralization thus results in information loss. However, although "delegating" control to the agent avoids this information loss, it also makes the agent's biased decision-making prevail. Delegation thus results in a loss of control. Dessein (2002) shows that for a broad range of parameters, the loss of control under delegation matters less than the loss of information under centralization. In particular, the smaller the agent's bias (i.e., the more congruent the principal's and agent's objectives are), or the larger the agent's informational advantage, or the more uncertainty there is, the more likely it is that delegation is optimal.

An important assumption in Dessein (2002) is that under "principal authority," the principal cannot commit to not taking the decision that she believes maximizes her expected utility. She can, however, commit to "delegate" control rights to the agent. This is consistent with the incomplete contracting assumption that actions or decisions are noncontractible, even when control allocation is contractible. Delegation can then be interpreted as a commitment device from the principal to use the agent's information in the way that best fits the agent's objectives. While delegation results in biased decision-making (this is the *loss of control* effect) it ensures that decision-making responds more to the agent's information. Compare this with centralized authority, which ensures an unbiased decision from the principal's point of view but encourages the agent to distort his information in order to influence the principal.

Information communication is modeled using Crawford and Sobel (1982)'s model of strategic communication. Since decisions are noncontractible and information is assumed to be soft, communication between

the agent and principal takes the form of “cheap talk”. However, strategic communication can be informative if there is sufficient preference alignment between the agent (the sender) and the principal (the receiver). Dessein’s contribution can be seen as one of introducing control-rights considerations in strategic communication games, with the idea that delegation induces communication by the agent to become less “strategic” although at the cost of noncongruent decision-making.

The result that delegation dominates when the agent’s bias is small is not fully obvious: communication between the agent and principal is also better under centralization in that case, as the agent gains less by distorting information when her preferences are more congruent with those of the principal. Yet Dessein (2002) shows that under general conditions, as long as the bias of the agent is sufficiently small, delegation is always strictly preferred over centralization. For some specific cases, such as the leading example in Crawford and Sobel (1982), Dessein (2002) obtains the striking result that communication never takes place at the optimum: as long as preferences are sufficiently congruent, the principal is strictly better off avoiding communication (and the resulting information distortions) altogether, by committing to delegate control to the agent.

Overall, this article provides an important rejoinder to the central insight of the property rights literature that the cost of integration (or centralization) are the reduced incentives to either make relationship-specific investments (Grossman and Hart) or to acquire information (Aghion and Tirole). In Dessein (2002), the cost of centralization (e.g., through integration), is the distortion and loss of information.

3. Empirical Evidence on the Determinants of Decentralization

Recent empirical interest in decentralization has been stimulated by the growth of a substantial body of evidence that documents persistence performance differences among firms, even in narrowly defined industries (see Syverson 2011; Gibbons and Roberts 2012 or Aghion et al. 2013, Appendix C). Could these differences be due to different organizational structures within firms? This section is structured as follows. Section 3.1 focuses on the measurement of firm decentralization, Section 3.2 on the impact of trust on decentralization, Section 3.3 on empirical implication of learning models for decentralization, and Section 3.4 documents more general “stylized facts” of decentralization. We formally discuss identification issues in more detail in Appendix D of Aghion et al. (2013).

3.1 Measuring Firm Decentralization

A key factor in any organization is who makes decisions. A centralized firm is one where decisions are all taken at the top of the hierarchy, and a

decentralized firm is where decision-making is more evenly dispersed throughout the corporate hierarchy.³

How can the concept of decentralization be implemented empirically? One way is to look at the organization charts of firms (the “organogram”) as graphical representations of the formal authority structure. One of the best studies in this area is Rajan and Wulf (2006) who use codified versions of charts of over 300 large US corporations during 1987–98 to examine the evolution of organizations (e.g., the number of people who directly report to the CEO as a measure of the span of control). Unfortunately, as Max Weber and Aghion and Tirole (1997) stressed, formal authority is not the same as real authority; the organogram may not reflect where real power lies.⁴

Observing whether a firm is decentralized into profit centers is useful, as this is a formal delegation of power—the head of such a business unit will be judged by the CEO on the basis of the unit’s profitability. Similarly if the firm is composed of cost (or revenue) centers this indicates less decentralization as only costs (or revenue) are likely to be under control of the manager. If the firm does not even delegate responsibility to a cost or revenue center this indicates a very centralized company. Acemoglu et al. (2007) use this distinction to classify firms in their empirical work on French and British firm panel data.

Whether a company is organized into profit centers is a rather crude indicator of decentralization. A better (but more costly) approach is to directly survey firms. Bloom et al. (2012b) measure decentralization between the plant manager and the corporate central headquarters (CHQ). They asked plant managers about their decisions over investment (maximum capital investment that could be made without explicit sign off from central headquarters), hiring, marketing, and product introduction (the latter three on a scale of 1–5). The data collection techniques follow the methodology of the Bloom and Van Reenen (2007) management survey.

Bloom et al. (2012b) constructed an empirical summary of decentralization combining these four measures into a single index by z-scoring each individual indicator and z-scoring the average (so that the index has zero mean and unit standard deviation). In this article we follow their approach but combine their data with additional waves of the World Management Survey conducted in 2007 and 2009 covering eight more countries.

3. An extreme case of decentralized “organization” is an idealized market economy where atomistic individuals make all the decisions and form spot contracts with each other. The origins of many of the debates on decentralization have their origins in the 1930s, over the relative merits of a market economy relative to a centrally planned one.

4. We focus on decentralization as distinct from managerial spans of control. These are separate concepts—the span and depth (number of levels) of a hierarchy is compatible with different power relationships between the levels. Nevertheless, there is some evidence that the move toward delayering over the past 20 years has been associated with decentralization (see Rajan and Wulf 2006).

Figure 1 shows that decentralization varies substantially across countries, with the United States, the United Kingdom, and Northern European countries being the most decentralized and Southern European and Asian countries the least. There is an even wider spread of decentralization across firms *within* every country, as shown in Figure 2.

Decentralization extends beyond just plant managers and the CHQ. For example, one can also consider the autonomy of workers from the plant manager following Bresnahan et al. (2002). Proxies for this dimension of decentralization include questions indicating greater worker control over the pace of work and over the allocation of tasks.

3.2 Trust (Congruence of Preferences)

The Aghion-Tirole approach offers a natural implementation of incomplete contracts to study decentralization within firms. One key parameter in fostering decentralization is the congruence of preferences between principal and agent. Other theoretical developments following Grossman–Hart such as Baker et al. (1999) also support this idea. Finding empirical proxies for the congruence parameter is challenging, but one possibility is to use measures of trust. In recent years, economists have started to take cultural factors more seriously in determining economic outcomes (Greif 1994; Guiso et al. 2006). This stems in part from the influence of Putnam's (1993) work on the importance of social capital. Empirically, generalized social trust as a proxy for social capital has been found to be associated with many positive economic outcomes (e.g., see Knack and Keefer 1997, on trust and growth or Guiso et al. 2009, on trust and foreign trade and investment).

Bloom et al. (2012b) examine the importance of trust, finding that higher levels of trust in the region where the headquarters of a plant is located is associated with a significantly greater degree of decentralization from the CEO to the plant manager (using the data underlying Figure 1). As in many other papers, trust is measured using the standard indicators from the World Values Survey which asks random samples of individuals about generalized trust. These external measures are matched to firm locations to obtain a measure of trust in the region. The authors also exploit the fact that their data contains many subsidiaries of multinational firms to construct measures of trust in the country of origin (the multinational's headquarters) and the country of the plant's location (where the affiliate and plant manager are located). Both of these seem to matter for decentralization, but the most powerful factor is the *bilateral* trust between country pairs, that is, the degree to which people from the multinational's parent country trust people in the country where the subsidiary plant is located. Multinationals locating in countries that are seen to be relatively highly trusted (after country location and origin dummies are removed), are significantly more likely to decentralize. For example, even though the United Kingdom is overall a relatively high trust country, multinationals headquartered in the United States (where Britain has a

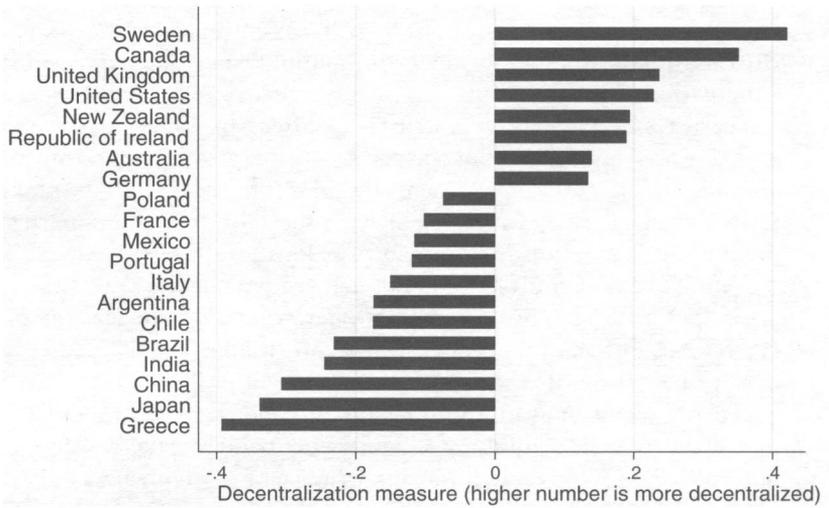


Figure 1. Decentralization is Higher in Scandanavia and Anglo-Saxon Countries, and Lower in Asia and Southern Europe. Source: Histogram of decentralization across firms within each country. The definition of decentralization follows Bloom et al. (2012b), with data updated to 2010 from www.worldmanagementsurvey.com.

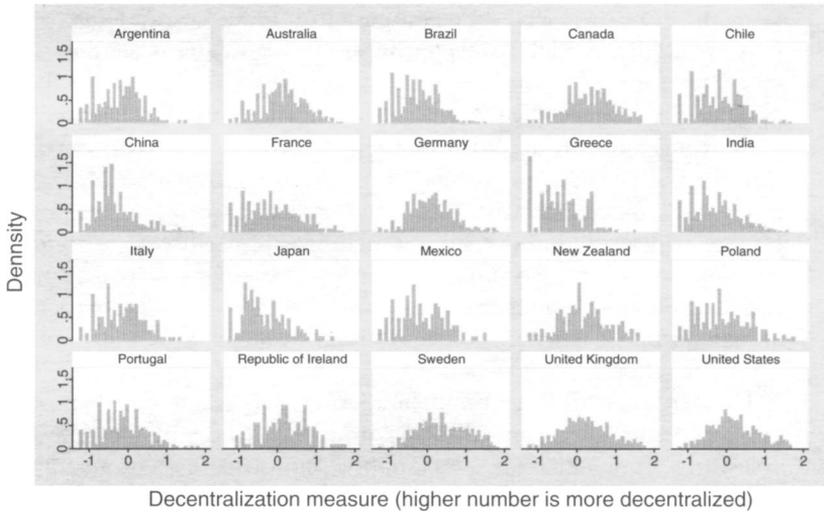


Figure 2. Decentralization Varies Heavily across Firms within Every Country. Source: Histogram of decentralization across firms within each country. The definition of decentralization follows Bloom et al. (2012b), with data updated to 2010 from www.worldmanagementsurvey.com.

relatively good trust reputation) tend to decentralize more towards their British affiliates than equivalent multinationals headquartered in France (where for historical reasons Britain has a relatively bad trust reputation).

These results suggest that trust can affect the internal structures of global firms and that some aspects of organization are transplanted abroad, as suggested by recent theories of international trade (e.g., Helpman et al. 2004). It fits well with the idea that the congruence of preferences is a major determinant of delegation.

Enforcement of contracts should further foster decentralization, and we do in fact observe more delegation where there is stronger rule of law.⁵ However, contracts are never perfectly enforceable, which leaves a role for trust to help generate more delegation. Recently, Bloom et al. (2013) ran field experiments on firms in India, and discovered that family size (in particular the number of adult male family members) was the key determinant of firm size, probably due to the importance of trust. Owners only trusted other family members to make major managerial decisions as they worried that outsiders would steal from the firm. Hence, the supply of (trusted) male family member time was typically the binding factor for firm growth.

3.3 Learning

In the theory section, we discussed an extension to the Aghion-Tirole approach when considering how firms could learn either from other firms or from themselves. Acemoglu et al. (2007) examine three predictions from their model (i) delegation should be greater when the industry is more heterogeneous (so it is harder to learn from others); (ii) the firm is close to the technological frontier (so that there are fewer other firms to learn from) and (iii) the firm is younger (so it has less experience to learn from its own mistakes). Acemoglu et al. (2007) measure decentralization using both formal measures of whether firms are organized into profit centers as discussed above (for French firms) and direct survey measures of the power managers have over hiring decisions (for British firms). In both samples they find decentralization is more likely in industries that are more heterogeneous and for firms that are younger or closer to the technological frontier.

These results are illustrated in Figures 3–5 where the y -axis has the average degree of decentralization in different bins of the relevant variables. In Figure 3, there is a reasonably clear positive relationship after the second decile between decentralization and heterogeneity (as measured by the dispersion of firm productivity growth). In Figure 4, decentralization appears to be higher among firms closer to the technological frontier (as measured by the distance of the firm's productivity from the leading firm in the four-digit industry). That firms closer to the frontier should delegate more may also explain why subsidizing higher education, in particular

5. See Bloom et al. (2012b). More generally on the importance of law and reputation on contract enforceability, see MacLeod (2007).

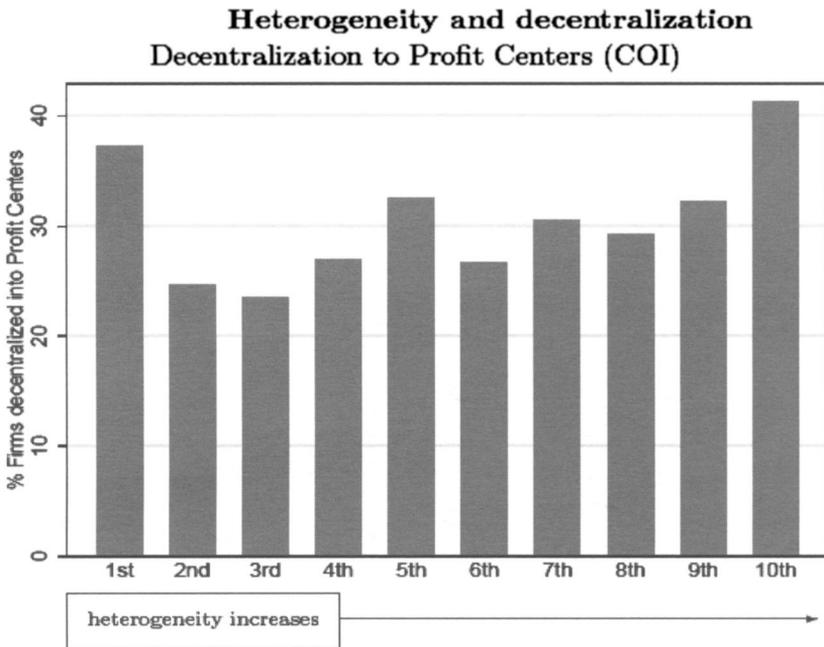


Figure 3. Decentralization is Higher When Heterogeneity is Greater. Source: Acemoglu et al. (2007).

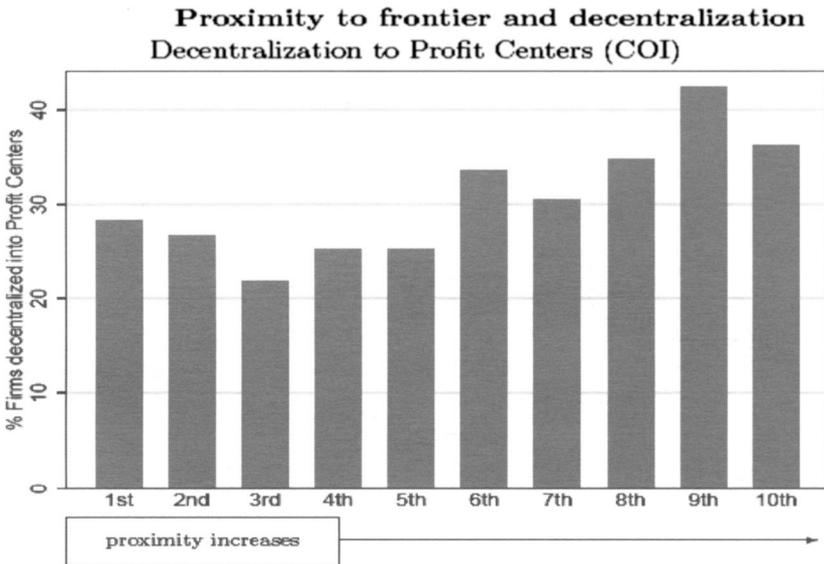


Figure 4. Decentralization is Higher When Plant is Closer to the TFP Frontier. Source: Acemoglu et al. (2007).

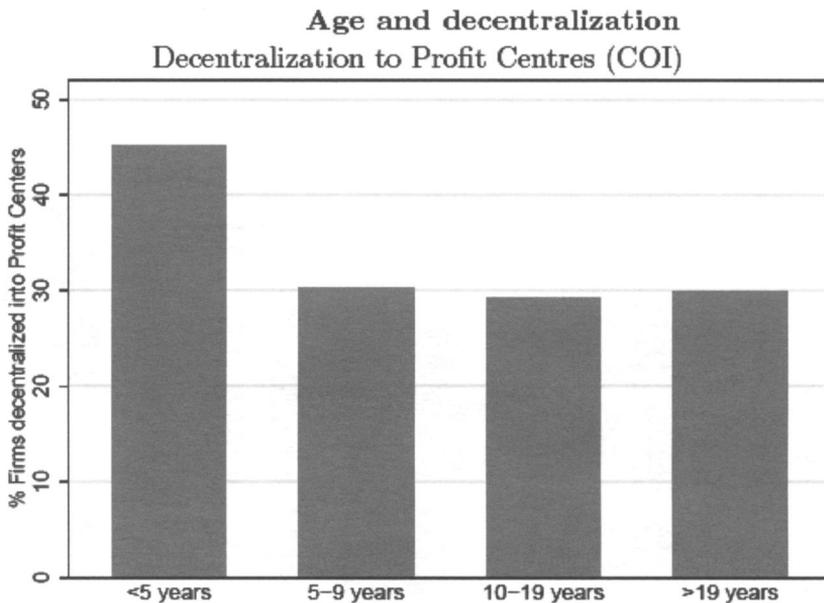


Figure 5. Decentralization is Higher in Younger Plants. Source: Acemoglu et al. (2007).

graduate education, is more likely to be growth-enhancing if universities are more decentralized (see e.g., Aghion et al. 2009). Figure 5 shows that older firms look more centralized than younger firms, possibly because they have learned better what to do, so there is less need of delegation to a local manager who is better informed but may not pursue the principal's interests.

3.4 Other Empirical Factors Influencing Decentralization

The development of the incomplete contracts approach as applied to firm decentralization appears to have some confirmation in the data. The congruence of preferences (as proxied by trust) and heterogeneity (making it harder to learn) both seem to foster decentralization. There are many other findings in the literature on the empirical determinants of decentralization. These are not so obviously implications of the Grossman-Hart approach, but it is worth considering them.

3.4.1 Firm Size and Scope. Some basic factors determine decentralization. All else equal, a larger firm will require more decentralization than a small firm. A sole entrepreneur does not need to delegate because she is her own boss, but as more workers are added, doing everything herself is no longer feasible. Penrose (1959) and Chandler (1962) stressed that decentralization was a necessary feature of larger firms, because CEOs do not have the time to take every decision in large firms (see also Geanakoplos and Milgrom

1991). Similarly as firms expand in their scope both geographically and in product space, local information will become more costly to transmit so this will also favor decentralization.

Most empirical findings support this. Bloom et al. (2012b) find that firm size and plant size are both associated with a significant increase in their decentralization index. Furthermore, plant managers in subsidiaries of foreign multinationals have more autonomy than similar plants of domestic non-multinationals. They interpret this as an indicator that managing at a distance is harder, inducing headquarters to give more autonomy to local managers.

3.4.2 Human Capital. Many models would predict that human capital should be associated with decentralization. For example, more skilled workers will have greater ability to take on more responsibility. When the environment changes due to new technologies and organizational change is required, skilled workers may be better at learning how to cope with the new organizational structure.

There is generally a robust and positive association of decentralization and skills. For example, Bloom et al. (2012b) measure skills by the proportion of people who hold a college degree and find this to be significantly correlated with decentralization. Caroli and Van Reenen (2001) examine the relationship between skills and organization in some detail, arguing in favor of “skill biased organizational change”: that is, increases in the supply of human capital will tend to increase delegation. To tackle the endogeneity problem, they use information on the differential price of skilled versus unskilled labor in the local market (as indicated by the wage differential between college educated workers and other individuals). They argue that this skill premium is partially driven by exogenous shifts in the supply of unskilled workers. For their sample of United Kingdom and French firms they find that regions where skill premia are higher have a lower probability of decentralization.

3.4.3 Information and Communication Technologies. Garicano (2000) formalizes the idea of the firm as a cognitive hierarchy. There are a number of problems to be solved and the task is how to solve them in the most efficient manner. The simplest tasks are performed by those at the lowest level of the hierarchy, and the “exceptional” problems are passed upwards to an expert. The cost of passing problems upwards is that communication is expensive. The benefit of passing the problem upwards is that it reduces the cognitive burden on lower level employees.

This framework was designed to address the impact of Information and Communication Technologies (ICT) on firm organization. Interestingly, information technologies have different implications for decentralization than communication technologies. Consider again the decentralization decision between the CHQ and plant manager. When communication costs fall (e.g., through the introduction of e-mail or company intranets),

it is cheaper for the plant manager to refer more decisions to the CHQ. So communication technologies should cause *centralization*. In contrast, technologies that make it easier for the plant manager to acquire information (e.g., Enterprise Resource Planning software, ERP like SAP) means that *decentralization* should increase. An example would be Lexis Nexis in law firms, which enables junior lawyers to quickly find relevant cases without consulting a more senior associate or partner.

Bloom et al. (2009a) test this theory and find considerable empirical support. Computer networks (reducing communication costs) significantly decrease decentralization to plant managers, whereas tools to help managers access more information (like ERP) significantly increase decentralization.⁶

3.4.4 Product Market Competition. Some authors such as Acemoglu et al. (2007) argue that a cause of the aggregate increase in more decentralized organizations is rapid technological change. An alternative explanation is that globalization and deregulation have increased the degree of product market competition which has in turn stimulated organizational change. The theory is ambiguous here. If competition has made swift decisions more important, then this will increase the salience of local knowledge, leading to greater decentralization under the framework discussed above (e.g., Aghion and Tirole 1997). Similarly if competition aligns the incentives of agents more with the principal, then the costs of decentralization may also fall. There are countervailing forces, however. For example, a larger number of firms in an industry aids yardstick competition, but it may also help learning in the Acemoglu et al. (2007) framework, which will reduce the need to decentralize.

The empirical evidence, however, seems more clear cut. Bloom et al. (2010) find a robust positive association between competition and decentralization using industry import competition, the inverse of the industry Lerner index, or simply the number of perceived competitors. A similar positive correlation was reported in Acemoglu et al. (2007) and Marin and Verdier (2008, 2009). All of these are cross-sectional studies, so the positive coefficient on competition could simply reflect unobserved variables correlated with competition. Guadalupe and Wulf (2010) try to tackle this endogeneity problem using the Rajan and Wulf (2006) panel dataset on the changing organizational structure of firms over time. They argue that the Canada-US Free Trade Agreement in 1989 constitutes an exogenous increase in competition for US firms in industries where tariffs were removed. Exploiting this policy experiment reveals that competition is associated with their proxy for decentralization.

6. The magnitude of the effect is substantial. An increase in ERP usage by 60% (the average difference in ICT between Europe and the United States) increases a plant manager's autonomy by an amount equivalent to the increase in US college graduates between 1990 and 2000.

4. Organizational Practices and Firm Productivity

How can researchers identify the effects of organizational structure (e.g., decentralization) on firm performance?⁷

4.1 Correlations of Performance and Organizational Practices: the Basic Identification Problem

Consider the basic production function as

$$q_{it} = \alpha_l l_{it} + \alpha_k k_{it} + a_{it} \tag{1}$$

where q is ln(value added), l is ln(labor), and k is ln(capital) of firm i at time t . Assume that we can write the TFP term a_{it} as

$$a_{it} = \alpha_0 + \beta m_{it} + u_{it} \tag{2}$$

where m_{it} is an organizational feature of the firm (such as decentralization), and u_{it} is an unobserved error. Together these equations imply

$$q_{it} = \alpha_0 + \alpha_l l_{it} + \alpha_k k_{it} + \beta m_{it} + u_{it} \tag{3}$$

This of course contains several assumptions. It assumes that the relevant organizational factor enters linearly, whereas organization could instead be affecting the coefficients on the other factor inputs, and many theories (e.g., of complementarity) would predict this. We discuss these below. We will assume that we can deal with the econometric problems in estimating the coefficients on the production function so that we have a consistent measure of total factor productivity (see Akerberg et al. (2006), for a discussion of recent contributions).

The traditional strategy is to assume that m_{it} is a firm fixed effect. So one approach is simply to recover average firm TFP under this assumption and project it on some cross-sectional measure of management m_i . This will indicate whether there is an association between the two measures, but not whether the relationship is causal. For example, Bloom and Van Reenen (2007) show that there is a robust relationship between TFP and their measure of management quality, but they interpret this as an “external validity” test of the quality of the management data rather than as any causal relationship.

An analogous strategy if there are time varying measures of organization is to treat all the correlated unobservables as fixed, that is, $u_{it} = \eta_i + \varepsilon_{it}$ with $E(m_{it}\eta_i) \neq 0$ but $E(m_{it}\varepsilon_{it}) = E(m_{it}\varepsilon_{it-1}) = 0$. Then the fixed-effect model estimated in (say) first differences would be

$$\Delta a_{it} = \beta \Delta m_{it} + \Delta \varepsilon_{it} \tag{4}$$

which can be consistently estimated by OLS.

7. Aghion et al. (2013) Appendix D discusses general identification issues in this literature in more detail.

There are a huge number of studies that have correlated various aspects of the firm's performance on various aspects of its organizational form (e.g., the survey in Lazear and Oyer 2012). The better studies use micro data and pay careful attention to measurement issues and need to control for many covariates. For example, Cappelli and Neumark (2001) and Black and Lynch (2001) examine various aspects of "high performance" workplaces, mostly relating to employee involvement, team work, and meetings. Both papers look across many industries and find no direct effect of these measures on performance (in contrast to many case studies). As we discuss later, Ichniowski et al. (1997), however, examined management practices and performance in 37 US steel mills over time and found a link between upgrading to bundles of modern organizational practices and improved performance, so the correlation evidence is mixed.

There remain several serious problems. First is the data constraint that measuring organization is hard and finding data with time series variation even harder. Second, the organizational proxies are measured with error, so this will cause attenuation towards zero if the measurement error is classical. This bias is exacerbated in first differences. Third, and most seriously, the factors that cause variation in the propensity to adopt organizational practices will also likely be correlated with those affecting TFP so the assumption is unlikely to hold in most cases. The bias could be upwards or downwards (e.g., if firms doing badly are more likely to innovate organizationally as argued by Nickell et al. 2001).

There is no simple solution to these endogeneity problems as we fundamentally need some exogenous identifying variation. Bloom et al. (2013) implemented a randomized control trial in Indian textile plants. The intervention was by high-quality management consultants to improve a range of management practices (as in Bloom and Van Reenen 2007) which appeared to dramatically improve productivity. Most of the quasi-experiments have been in labor economics. A good example is Lazear (2000) who looked at the introduction of a pay for performance system for windshield installers in the Safelite Glass Company. Lazear found that productivity increased by around 44%, with about half of this due to selection effects and half from the same individuals changing behavior. More recently, Bandiera et al. (2007, 2009) engineered a change in the incentive pay system for managers in farm. They have no contemporaneous control group, but can examine the behavior of individuals before and after the introduction of the incentive scheme. Productivity rose by 21% mainly with at least half due to improved selection (the managers allocated more fields to the ablest workers rather than to their colleagues).

4.2 Complementarities between Organizational Practices

One of the key reasons why firms may find it difficult to adjust their organizational form is that there are important complementarities between sets of organizational practices. Milgrom and Roberts (1990) build a theoretical structure where such complementarities (or more precisely,

super-additivities) mean that firms optimally choose clusters of practices that “fit together.” When the environment change so that an entrant firm would use this group of optimal practices, incumbent firms will find it harder—they will either switch a large number of practices together, or none at all.

This has important implications for productivity analysis. The effects of introducing a single practice will be heterogeneous between firms and depend on what practices they already use. This implies that linear regressions of the form of equation (3) may be misleading. To see this, consider two practices, m^1 and m^2 , whose relationship with productivity is such that TFP increases only when both are used together.

$$a_{it} = \alpha_0 + \beta_1 m_{it}^1 + \beta_2 m_{it}^2 + \beta_{12}(m_{it}^1 \times m_{it}^2) + u_{it} \quad (5)$$

One version of the complementary hypothesis is $\beta_1 > 0$, $\beta_2 < 0$, and $\beta_{12} > 0$, that is, the disruption caused by just using one practice (m_{it}^2) could actually reduce productivity. A regression that omits the interaction term may $m_{it}^1 \times m_{it}^2$ find only a zero coefficient on the linear terms.

The case study literature emphasizes the importance of complementarities. Testing for their existence poses some challenges, however, as pointed out most clearly by Athey and Stern (1998). A common approach is a regression of organizational practice 1 (m_{it}^1) on practice 2 (m_{it}^2) with a positive covariance (conditional on other factors) indicating complementarity. It is true that complements will tend to covary positively, but this is a very weak test. There could be many other unobservables causing the two practices to move together. We need an instrumental variable for one of the practices (e.g., Van Biesebroeck 2007), but this is hard to obtain as it is unclear what such an instrument would be, that is, could it be legitimately excluded from the second-stage equation? In classical factor demand analysis we would examine the cross price effects to gauge the existence of complements versus substitutes, for example does demand for practice 1 fall when the price of practice 2 rises (all else equal) indicating complementarity. There still remains the concern that the price shocks could be correlated with the productivity shocks, but such an assumption is weaker than assuming unobserved shocks to the firm’s choice of practices are uncorrelated. Unfortunately, such tests are particularly hard to implement because there are generally no market prices for the organizational factors typically considered.

An alternative strategy is to work straight from the production function (or performance equation more generally). Consider the productivity equation after substituting in multiple practices:

$$q_{it} = \alpha_0 + \alpha_l l_{it} + \alpha_k k_{it} + \beta_1 m_{it}^1 + \beta_2 m_{it}^2 + \beta_3(m_{it}^1 \times m_{it}^2) + u_{it} \quad (6)$$

Ichniowski et al. (1997) estimate a version of equation (6) using very disaggregate panel data on finishing lines in US steel mills, using 11 human resource practices (including incentive pay, recruitment, teamwork, job flexibility, and rotation). Their measure of productivity is based on

downtime—the less productive lines were idle for longer. They find that introducing one or two practices has no effect on productivity, but introducing a large number together significantly raises productivity. Although the endogeneity problem is not eliminated, the controls for fixed effects, looking within one firm and using performance data, helps reduce some of the more obvious sources of bias.

4.3 The Role of ICT Again

One of the key productivity puzzles of recent decades has been why the returns to the use of information and communication technologies appear to be so high and so heterogeneous between firms and between countries. For example, Brynjolfsson and Hitt (2003) find that the elasticity of output with respect to IT capital is far higher than its share in gross output (see also Stiroh 2002). One explanation for this is that effective use of IT also requires significant changes in firm organization. Changing the notation of (6) slightly, we could write

$$q_{it} = \alpha_0 + \alpha_l l_{it} + \alpha_k k_{it} + \beta_c c_{it} + \beta_m m_{it} + \beta_{cm}(c_{it} \times m_{it}) + u_{it} \quad (7)$$

where c_{it} is the log of the IT capital stock and m_{it} could be an organizational practice such as the degree of decentralization with the hypothesis that $\beta_{cm} > 0$. This is broadly the position of papers in the macro literature explaining the faster productivity growth of the United States than Europe after 1995 (e.g., Jorgenson et al. 2008).

Bresnahan et al. (2002) try to test this directly by surveying large US firms on decentralization and team work (for a cross section) and combining this with data on IT (from a private company Harte-Hanks), and productivity from Compustat. They find evidence that $\beta_{cm} > 0$, that is, that computer capital is more productive when firms have greater decentralization and team work. Bloom et al. (2012a) broaden the sample to cover both the United States and firms in seven European countries, and find evidence of complementarity of IT with people management. They also show that their results are robust to controlling for firm fixed effects. Careful econometric case studies (e.g., Baker and Hubbard 2004; Bartel et al. 2007) also identify differential productivity effects of ICT depending on organization form. Lemieux et al. (2009) show that one particular people management practice, performance pay, is becoming increasingly important in the United States and has a significant impact on widening inequality. They suggest that the spread of ICT innovations has facilitated the adoption of performance pay techniques.

4.4 The Role of Human Capital

One of the reasons for the renewed interest in organizational change by labor economists was the attempt to understand why technology seemed to increase the demand for human capital, thus contributing to the rise in wage inequality experienced by the United States, United Kingdom, and

other countries since the late 1970s. Many theories have been proposed (see Autor et al. 2003, for a review), but one hypothesis is that lower IT prices increased decentralization for the reasons outlined in Garicano (2000), and decentralization leads to an increase in inequality (Garicano and Rossi-Hansberg 2012). Further, decentralization is complementary with skills for at least three reasons. First, skilled workers are more able to analyze and synthesize new pieces of knowledge, so the benefits of local processing of information are enhanced. Additionally, skilled workers are better at communicating, which reduces the risk of duplication of information. Second, the cost of training them for multi-tasking is lower, and they are more autonomous and less likely to make mistakes. Finally, workers who are better educated may be more likely to enjoy job enrichment, partly because they expect more from their job in terms of satisfaction.

This has three main implications:

1. Decentralization leads to skill upgrading within firms. This is due to the fact that the return to new work practices is greater when the skill level of the workforce is higher.
2. A lower price of skilled labor will accelerate the introduction of organizational changes.
3. Skill intensive firms will experience greater productivity growth when decentralizing.

Caroli and Van Reenen (2001) find support for all three predictions. They estimate production functions (with the relevant interactions), skill share equations, and organizational design equations.

5. Conclusions

We began by surveying the theoretical literature on the organization of firms and the optimal decentralization of decision rights within firms. We discussed how the concept of incomplete contracts shapes this organization of decision making within firms. In particular, the inability to contract over all possible states of the world leads principals to delegate control to agents as a way to ex ante commit to letting agents expropriate some of the returns from costly activities, like collecting information on the best actions to take.

We then overviewed some of the empirical evidence on the organization and management of firms, focusing on decentralization. We looked at within firm organization, especially decentralization following the theoretical survey, and we also looked at econometric studies, focusing on large-scale firm databases rather than case studies. There has recently been a number of papers measuring management and organizational practices across firms and countries. Like productivity, decentralization varies a lot across firms and countries (e.g., Scandinavian and Anglo-Saxon firms are more decentralized than those from Asia and Southern Europe). A

number of factors highlighted in the incomplete contract theory are shown to be important for accounting for differences in firm organization—in particular heterogeneity and the congruence of preferences (as proxied by trust). Several other factors appear robustly positively correlated with decentralization, such as product market competition, human capital, IT and firm size.

In terms of future work, we see many areas of opportunity. There is a need to match up empirical work more closely with the theory. Until recently, comprehensive datasets on measures of organization across many firms and countries were unavailable. Now that this gap in the core data infrastructure is being covered, there is a great opportunity for testing some of the theories of organizational economics. This is challenging, first because many of the important aspects in the environment emphasized in the theoretical literature are hard to match into empirical counterparts, and second because identifying the causal relationship between organizational changes and firm outcomes like productivity and growth is difficult. Even with strong measurement and tight links to theory, it is essential to identify the direction of causality in some of stylized empirical results we have identified above—for example, do skilled managers enable more decentralized decision making, or are skilled managers attracted to more decentralized firms? We see the theory and empirics of the organization of the firm as one of the key growth areas in economics over the next 25 years.

References

- Acemoglu, S., P. Aghion, C. Lelarge, J. Van Reenen, and F. Zilibotti. 2007. “Technology, Information and the Decentralization of the Firm” 122 *Quarterly Journal of Economics* 1759–99.
- Akerberg, D., C.L. Benkard, S. Berry, and A. Pakes. 2007. “Econometric Tools for Analyzing Market Outcomes,” 122 *Quarterly Journal of Economics* 1759–99.
- Aghion, P., and J. Tirole. 1997. “Formal and Real Authority in Organizations,” 105 *Journal of Political Economy* 1–29.
- Aghion, P., M. Dewatripont, and J.C. Stein. 2009. “Academic Freedom, Private-Sector Focus, and the Process of Innovation,” 39 *The RAND Journal of Economics* 617–35.
- Aghion, P., M. Dewatripont, and P. Rey. 2004. “Transferable Control,” 2 *Journal of the European Economic Association* 115–38.
- Aghion, P., N. Bloom, and J. Van Reenen. 2013. “Incomplete Contracts and the Internal Organization of Firms,.” National Bureau of Economic Research. Working Paper 18842.
- Aghion, P., and P. Bolton. 1992. “An Incomplete Contracts Approach to Financial Contracting,” 59 *The Review of Economic Studies* 473.
- Athey, S., and S. Stern. 1998. “An Empirical Framework for Testing Theories about Complementarity in Organizational Design” National Bureau of Economic Research. Working paper 6600.
- Autor, D.H., F. Levy, and R.J. Murnane. 2003. “The Skill Content of Recent Technological Change: An Empirical Exploration,” 118 *Quarterly Journal of Economics* 1279–333.
- Baker, G., R. Gibbons, and K.J. Murphy. 1999. “Informal Authority in Organizations,” 15 *Journal of Law, Economics, and Organization* 56.

- Baker, G.P., and T.N. Hubbard. 2004. "Contractibility and Asset Ownership: On-Board Computers and Governance in US Trucking," 119 *Quarterly Journal of Economics* 1443–79.
- Bandiera, O., I. Barankay, and I. Rasul. 2007. "Incentives for Managers and Inequality Among Workers: Evidence from a Firm Level Experiment," 122 *Quarterly Journal of Economics* 729–73.
- . 2009. "Social Connections and Incentives in the Workplace: Evidence from Personnel Data," 77 *Econometrica* 1047–94.
- Bartel, A.C., C. Ichiniowski, and K. Shaw. 2007. "How does Information Technology Affect Productivity? Plant-Level Comparisons of Product Innovation, Process Improvement and Worker Skills," 122 *Quarterly Journal of Economics* 1721–58.
- Biesebroeck, J.V. 2007. "Complementarities in Automobile Production," 22 *Journal of Applied Econometrics* 1315–45.
- Black, S.E., and L.M. Lynch. 2001. "How to Compete: the Impact of Workplace Practices and Information Technology on Productivity," 83 *Review of Economics and Statistics* 434–45.
- Bloom, N., B. Eifert, A. Mahajan, D. McKenzie, and J. Roberts. 2013. "Does Management Matter? Evidence from India," 128 *Quarterly Journal of Economics* 1–51.
- Bloom, N., C. Genakos, R. Sadun, and J. Van Reenen. 2012b. "Management Practices across firms and countries," 26 *Academy of Management Perspectives* 12–33.
- Bloom, N., C. Propper, S. Seiler, and J. Van Reenen. 2010c. "The Impact of Competition on Management Quality: Evidence from Public Hospitals," Technical Report, National Bureau of Economic Research.
- Bloom, N., and J. Van Reenen. 2007. "Measuring and Explaining Management Practices across Firms and Countries," 122 *Quarterly Journal of Economics* 1351–408.
- Bloom, N., L. Garicano, R. Sadun, and J. Van Reenen. 2009a. "The Distinct Effect of Communication Technology and Information Technology on Firm Organization".
- . 2009b. "The Organization of Firms Across Countries", National Bureau of Economic Research. Working Paper.
- Bloom, N., R. Sadun, and J. Van Reenen. 2012a. "Americans do I.T. Better: US Multinationals and the Productivity Miracle," 102 *American Economic Review* 167–201.
- . 2012b. "The Organization of Firms Across Countries," 127 *Quarterly Journal of Economics* 1663–705.
- . 2010. "Does Product Market Competition Lead Firms to Decentralize?," 100 *American Economic Review* 434–38.
- Bolton, P., and M. Dewatripont. 2011. "Authority in Organizations: A Survey!" in E. Gibbons, and D.J. Robert, eds., *Handbook of Organizational Economics* Princeton: Princeton University Press.
- Bresnahan, T.F., E. Brynjolfsson, and L.M. Hitt. 2002. "Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence," 117 *Quarterly Journal of Economics* 339–76.
- Brynjolfsson, E., and L.M. Hitt. 2003. "Computing Productivity: Firm-Level Evidence," 85 *The Review of Economics and Statistics* 793–808.
- Cappelli, P., and D. Neumark. 2001. "Do High-Performance Work Practices Improve Establishment-Level Outcomes," 54 *Industrial and Labor Relations Review* 737–75.
- Caroli, E., and J. Van Reenen. 2001. "Skill-Biased Organizational Change? Evidence from a Panel of British and French Establishments," 116 *Quarterly Journal of Economics* 1449–92.
- Chandler, A.D. 1962. *Strategy and Structure: Chapters in the History of American Enterprise* Cambridge: MIT Press.
- Crawford, V.P., and J. Sobel. 1982. "Strategic Information Transmission," 50 *Econometrica: Journal of the Econometric Society* 1431–51.
- Dessein, W. 2002. "Authority and Communication in Organizations," 69 *The Review of Economic Studies* 811.

- DiNardo, J., and D.S. Lee. 2004. "Economic Impacts of New Unionization on Private Sector Employers: 1984-2001," 119 *Quarterly Journal of Economics* 1383-41.
- Foster, L., J. Haltiwanger, and C. Syverson. 2008. "Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?," 98 *The American Economic Review* 94-425.
- Garicano, L. 2000. "Hierarchies and the Organization of Knowledge in Production," 108 *Journal of Political Economy* 874-904.
- Garicano, L., and E. Rossi-Hansberg. 2012. "Organizing Growth," 147 *Journal of Economic Theory* 623-56.
- Garicano, L., and T. Hubbard. 2009. "Earnings Inequality and Coordination Costs: Evidence from US Law Firms" National Bureau of Economic Research. Working paper 14741.
- Geanakoplos, J., and P. Milgrom. 1991. "A Theory of Hierarchies Based on Limited Managerial Attention," 5 *Journal of the Japanese and International Economies* 205-25.
- Gibbons, R., and J. Roberts. 2012. *Handbook of Organizational Economics*. Princeton: Princeton University Press.
- Greif, A. 1994. "Cultural Beliefs and the Organization of Society: A Historical and Theoretical Reflection on Collectivist and Individualist Societies," 102 *Journal of Political Economy* 912-50.
- Grossman, S.J., and O.D. Hart. 1986. "The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration," *The Journal of Political Economy* 691-719.
- Guadalupe, M., and J. Wulf. 2010. "The Flattening Firm and Product Market Competition: The Effect of Trade Liberalization on Corporate Hierarchies," 2 *American Economic Journal: Applied Economics* 105-27.
- Guiso, L., P. Sapienza, and L. Zingales. 2003. "People's Opium? Religion and Economic Attitudes," 50 *Journal of Monetary Economics* 225-82.
- . 2006. "Does Culture Affect Economic Outcomes?," 20 *The Journal of Economic Perspectives* 23-48.
- . 2009. "Cultural Biases in Economic Exchange?," 124 *Quarterly Journal of Economics* 1095.
- Hart, O., and B. Holmstrom. 2010. "A Theory of Firm Scope," 125 *Quarterly Journal of Economics* 483-513.
- Hart, O.D., and J. Moore. 1999. "On the Design of Hierarchies: Coordination versus Specialization" National Bureau of Economic Research. Working paper 7388.
- Helpman, E., M.J. Melitz, and S.R. Yeaple. 2004. "Export versus FDI with Heterogeneous Firms," 94 *The American Economic Review* 300-16.
- Ichniowski, C., K. Shaw, and G. Prennushi. 1997. "The Effects of Human Resource Management Practices on Productivity: A Study of Steel Finishing Lines," 87 *American Economic Review* 291-313.
- Jorgenson, D.W., M.S. Ho, and K.J. Stiroh. 2008. "A Retrospective Look at the US Productivity Growth Resurgence," 22 *The Journal of Economic Perspectives* 3-24.
- Knack, S., and P. Keefer. 1997. "Does Social Capital Have An Economic Payoff? A Cross-Country Investigation," 112 *Quarterly Journal of Economics* 1251-88.
- Lafontaine, F., and M.E. Slade. 2010. "Inter-Firm Contracts: Evidence," in E. Gibbons, and D.J. Robert, eds., *Handbook of Organizational Economics*. 342-72. Princeton: Princeton University Press.
- Lazear, E. 2000. "Performance Pay and Productivity," 90 *American Economic Review* 1346-61.
- Lazear, E., and P. Oyer. 2012. "Personnel Economics," in E. Gibbons, and D.J. Robert, eds., *Handbook of Organizational Economics*. 520-74. Princeton: Princeton University Press.
- Lemieux, T., W.B. MacLeod, and D. Parent. 2009. "Performance Pay and Wage Inequality," 124 *Quarterly Journal of Economics* 1-49.
- MacLeod, W.B. 2007. "Reputations, Relationships and Contract Enforcement," 45 *Journal of Economic Literature* 597-630.

- Marin, D., and T. Verdier. 2008. "Power Inside the Firm and the Market: A General Equilibrium Approach," 6 *Journal of the European Economic Association* 752–88.
- . 2009. "Power in the Multinational Corporation in Industry Equilibrium," 38 *Economic Theory* 437–64.
- Milgrom, P., and J. Roberts. 1990. "The Economics of Modern Manufacturing: Technology, Strategy, and Organization," 80 *The American Economic Review* 511–28.
- Modigliani, F., and M.H. Miller. 1958. "The Cost of Capital, Corporation Finance and the Theory of Investment," 48 *The American Economic Review* 261–97.
- Nickell, S., D. Nicolitsas, and M. Patterson. 2001. "Does Doing Badly Encourage Management Innovation?," 63 *Oxford Bulletin of Economics and Statistics* 5–28.
- Penrose, E.T. 1959. *The Theory of the Growth of the Firm*. New York: Wiley.
- Putnam, R. 1993. "The Prosperous Community," 4 *The American Prospect* 35–42.
- Rajan, R.G., and J. Wulf. 2006. "The Flattening Firm: Evidence from Panel Data on the Changing Nature of Corporate Hierarchies," 88 *The Review of Economics and Statistics* 759–73.
- Stiroh, K. 2002. "Reassessing the Role of IT in the Production Function: A Meta-Analysis," Federal Reserve Bank of New York Working Paper.
- Syversen, C. 2011. "What Determines Productivity?," 49 *Journal of Economic Literature* 326–65.