University of Heidelberg

Department of Economics

Discussion Paper Series | No. 468

Too Few Cooks Spoil the Broth: Division of Labour and Directed Production

Marisa Ratto and
Wendelin Schnedler

July 2008

Too Few Cooks Spoil the Broth: Division of Labour and Directed Production

Marisa Ratto
Université Paris-Dauphine (SDFi)

Wendelin Schnedler*
Department of Economics
University of Heidelberg

First Version: September 2003
This Version: July 2008

Abstract

How can a manager influence workers' activity while knowing little about it? This paper examines a situation where production requires several tasks, and the manager wants to direct production to achieve a preferred allocation of effort across tasks. However, the effort that is required for each task cannot be observed, and the production result is the only indicator of worker activity. This paper illustrates that in this situation, the manager cannot implement the preferred allocation with a single worker. On the other hand, the manager is able to implement the preferred allocation by inducing a game among several workers. Gains to workers from collusion may be eliminated by an ability-dependent, but potentially inefficient, task assignment. These findings provide a new explanation for the division of labor, and bureaucratic features such as "over"-specialization and "wrong" task allocation.

Keywords: specialization, job design, moral hazard, multitasking

JEL-Codes: D02, D86, M54, D23, L23, J23

---

*wendelin.schnedler@awi.uni-heidelberg.de
The authors wish to thank Guido Friebel, Paul Grout, Joop Hartog, Stefan Napel, Jörg Oechssler, Ulf Schiller, and Jan Zabojnik for their valuable comments. In addition, we thank participants of the Public Economic Theory conference in 2003 at Duke University, at the Royal Economic Society annual meeting 2004, at the Kolloquium für Personalökonomie in Bonn 2004, at the Fifth Symposium of the German Economic Association for Business Administration in Bonn 2004, and at seminars in Bristol and Erfurt, who contributed to this paper.

Such schemes as these are nothing without numbers.
One cannot have too large a party.

Jane Austen, Emma Vol. III, Chap. VI

1 Introduction

Division of labor has been a central theme in economic analysis since Adam Smith (1776) and a cornerstone of Frederick Winslow Taylor's "scientific management" (1911). Traditionally, the advantages of the division of labor have been attributed to the productive gains from learning by doing, economies of scale and comparative advantages. This paper argues that the division of labor has another advantage that has so far been neglected: it helps to influence the production method, when only the production result is verifiable.

In many production processes, little is known about how a particular result has been achieved. For example, software may run well due to careful engineering or due to extensive debugging; exam results can be improved by better teaching or by rehearsing test situations; a placement officer may be successful because she brushes up job seekers' interview skills or because she finds adequate vacancies. Although it is hard, or even impossible, to deduce how such outcomes are produced, it is often important to influence the decision of how to produce because some production methods are more desirable than others. Maintaining software is easier when it is well-designed; coping with test situations is only a secondary goal of the education system; job agencies prefer adequate and hence longer lasting placements.

This paper considers a multiple task principal-agent problem in which each task requires effort, the organizer of production (principal) prefers a specific allocation of effort, agents who work for her prefer a different allocation, and the only evidence of production is the final result. The principal cannot enforce effort, but she can prevent agents from exerting effort as she controls necessary assets. For example, she can withhold tools and material, or prohibit entry to certain production areas. However, the power of the principal is rather limited, as all tasks are required, and she has to allow at least one agent to use the task's respective asset. Can the principal ensure that production occurs in the desirable way despite such drastic limitations?

The paper shows that a single agent cannot be induced to produce in the desired way. The principal can pay the agent for a specific production result, and will even obtain this result, but the agent will focus effort according to his preferences.

The central result of this paper is that division of labor allows the principal to implement her desired production method, even though individual or task-wise effort cannot be verified. The intuition is simple. Each agent carries out only one task, and the principal pays on the basis of the overall result. Thus, an agent can only affect this result—and hence his pay—by changing his effort on his task. By adjusting the pay to a specific agent, the principal affects the effort of this agent and hence the effort exerted at this task. As this can be done separately for every agent, the desired allocation can be obtained. This result provides a novel explanation for the division of labor. Note that it is essential that the principal limits the agents' discretion over who carries out which task. Otherwise, agents will produce the result while performing their preferred tasks, which may not correspond to the desired production method.

A second finding is that the principal can implement her desired production method in a collusion-proof manner. Colluding agents have the same freedom to obtain a result as a single agent. Similar to a single agent, colluding agents can focus their effort on easy tasks, and reimburse the agents who are working on these tasks. The principal, however, can make such deviations unattractive by assigning easy tasks to unproductive agents. This finding explains inefficient task assignments.

This paper is closely related to the work of Dewatripont and Tirole (1999), but differs at a crucial point. In the setting of Dewatripont and Tirole, an agent can obtain the same result (for example, a verdict) by slacking on two conflicting tasks (for example, gathering less incriminating and exonerating evidence). In many production contexts, the same result can only be obtained if the reduction of effort on one task is compensated by an increase of effort on another task (for example, less careful programming is offset by more debugging). This paper's analysis covers such non-conflicting tasks. Specialization facilitates implementation, regardless of the nature of tasks. The implications of specialization are, however, very different. Where tasks are conflicting, incentive provision leads to competition. Competing agents impose a negative externality on each other. Where tasks are non-conflicting, agents have a joint interest in producing the result, and the externality is positive. Secondly, competition among agents invites sabotage activities, which are absent when tasks are non-conflicting. Finally, collusion is unavoidable with conflicting tasks, as agents gain by ceasing competition and slacking. In the situation analyzed here, collusion can sometimes be prevented by an appropriate assignment of agents to tasks.

The paper's findings also extend and complement the literature on multi-task principal-agent models that were pioneered by Holmström and Milgrom (1991). Recent contributions to this literature analyze which information about the activity of a single agent needs to be verifiable in order to obtain a desired allocation of effort across tasks (Baker 2002, Schnedler 2006 and forthcoming). While these models tackle the problem that effort allocation cannot be identified from the production result,[1] they do not consider the beneficial effect of labor division.

Implementing the desired production method requires that several agents be remunerated on the basis of a joint production result.[2] While the existence of teams has been attributed to interaction in the production function (Alchian and Demsetz, 1972) or due to task-wise additive separable and convex effort costs (Itoh, 1991), this paper contributes the alternative explanation that teamwork, rather than individual work, ensures a particular way of production. All team members are indeed engaged in productive work, and this distinguishes this paper's model from others in which incentives are improved because agents supervise their colleagues (see Strausz 1997 or Athey and Roberts 2001). The remainder of the paper is organized as follows. Section two sets up the model, and section three shows how the desired production method can only be implemented if labor is divided. Section four examines when and how the desired production can still be implemented by a specific task assignment when agents collude. Finally, section five concludes.

2 The Model

Consider a principal who wants to produce a good that involves two tasks. One task is easy (task E) and the other task is demanding (task D). Lacking the time or skills to do the tasks herself, the principal employs one or two

[1] Similarly, talent and effort cannot be identified in career concern models (Fama 1980, Holmström 1999, Gibbons and Murphy 1992) and Dewatripont, Jewitt, and Tirole (1999) show that rotating tasks can solve this identification problem.

[2] Of course, joint production has also costs. Notably, workers may free-ride on their colleagues' effort. In our analysis these costs do not feature because the principal acts as a budget breaker (Holmström, 1982).

agents to do the job. Each task involves an effort choice by agent i who carries out the task: e_i^D and e_i^E, where both are non-negative real numbers from a bounded interval. Agents are allowed to mix over pure effort choices, and the corresponding cumulative distribution function for agent i is denoted by F_i while $P_i(e_i = e)$ is the probability that i chooses e if e is a mass point.

In line with the examples from the introduction, the principal cares about how production is achieved. Denote the effort choices desired by the principal by e^{D^*} for the difficult task and e^{E^*} for the easy task. In the following, we examine the conditions under which these desired effort choices can be implemented.

The model assumes that effort is costly for the agent. In order to eliminate economies of scale and other incentives to specialize, the model assumes that costs are additively separable. The cost of agent i in task k, c_i^k , is a continuously differentiable function of effort that increases and is strictly convex. [3] The effort levels that an agent is willing to exert without incentives are denoted by e^{D^0} and e^{E^0}, and the respective costs are normalized to zero: $c_i^D(e^{D^0}) = c_i^E(e^{E^0}) = 0$. In order to make the implementation problem interesting, the model assumes that $e^{D^0} \neq e^{D^*}$ and $e^{D^0} \neq e^{D^*}$

We attach meaning to the labels "difficult" and "easy" of the two tasks by assuming that the marginal costs of both tasks are not identical given the desired allocation (e^{D^*}, e^{E^*}) and associating the difficult task with higher marginal costs:

$$\left.\frac{dc_i^D}{de}\right|_{e^D = e^{D^*}} > \left.\frac{dc_i^E}{de}\right|_{e^E = e^{E^*}}. \tag{1}$$

If the production method (the effort choices) could be verified by a court of law, the principal could hire a single agent and stipulate the desired effort levels in a contract. This contract would then work as a device to direct the agent's effort. In many settings, however, the only evidence is the result of production R, and it is not always clear how this result has been achieved. Accordingly, the model assumes that contracts about task-wise efforts cannot be written, and that anything that reveals how the result has been generated (for example, intermediate stages of production) cannot be verified.[4] In

[3] More generally, the agent may not exert effort but have a choice which affects himself as well as the principal. However, ordering these choices according to costs and speaking of effort greatly simplifies the illustration.

[4] It may well be possible that they can be verified later when the agent can no longer be held responsible—perhaps because he is working for another employer or retired.

short, the benefit of the principal is private information, and the action of the agent is hidden. The contractual environment is thus that of a "coping organization" according to Wilson (1989). Summarizing these considerations, the model assumes that the *only* verifiable quantity R is a function of both efforts. Notice that R is not stochastic. Still, R is imperfect as it confounds effort on both tasks and does not reveal exactly what the agent did.[5]

The principal may care about the production result R; for example, if she sells the product on a market. More importantly, she also cares about how the result is achieved. For example, neglecting the demanding task may lead to a higher probability of product breakdown, and damage the firm's reputation, or it may lead to more wear- and-tear of the production equipment, and necessitate costly repairs.[6]

In their seminal paper on advocates, Dewatripont and Tirole (1999) also suppose that the only verifiable variable, whether a culprit is convicted or not, is the result of two effort choices: searching for incriminating evidence, and searching for exculpatory evidence. These two efforts are conflicting. A larger R, say, conviction, can be obtained by increasing effort on one task, for example, searching harder for incriminating evidence, but also by decreasing effort on another task, for example, looking less hard for exculpatory evidence. In many production settings, it is not possible to improve a result by reducing effort. To the contrary, effort needs to be increased at least in one task to obtain a larger outcome. In these settings, tasks are non-conflicting, because the agent does not nullify the effect of effort in one task by increasing effort in another task. The paper examines such production settings, and assumes that the result R is an increasing (concave and twice continuously differentiable) function of effort exerted in the two tasks.

The result from the desired effort choice is denoted by $r^* := R(e^{D^*}, e^{E^*})$. An immediate consequence of the fact that the outcome increases and is continuous in both efforts is that the outcome r^* can also be achieved by some other effort choice. The crucial assumption of our model is that given the desired allocation (e^{D^*}, e^{E^*}), increasing effort on the easy task is at least as effective in generating a larger outcome than increasing effort on the de-

[5]The crucial mechanism that drives our results works also if the output measurement is noisy. However, eliminating noise simplifies the analysis.

[6]See the introduction for additional examples in which not only output but also the manner of production matters.

manding task:

$$\frac{\partial}{\partial e^E} R(e^D, e^E)\bigg|_{e^D = e^{D*}, e^E = e^{E*}} \geq \frac{\partial}{\partial e^D} R(e^D, e^E)\bigg|_{e^D = e^{D*}, e^E = e^{E*}}. \tag{2}$$

This assumption reflects the idea that often, there is an easier way than the desired one to produce an output. This paper will show that the presence of this opportunity may render the implementation of the desired production method impossible.

So far, the model has assumed that the production outcome is the only verifiable result of the agent's effort choice. By conditioning transfers τ on this outcome, the principal can influence the agent's choice. The model's final assumption is that the principal has another more basic, but very limited, method to control what the agent is doing. She can assign tasks to one or more agents. More specifically, she can prevent agents from carrying out a task. In practice, this could be done by withholding material or instruments needed for the task, restricting access to a location at which the task is performed, or refraining from training the worker to carry out the task. The most common form of prevention is prohibiting the agent from carrying out a task. In many circumstances, it is very simple to describe what the agent should not do, and much harder to describe what the agent should do. For example, there is, to our knowledge, no verifiable definition of good teaching practice, while it is relatively simple to spot whether a class is rehearsing a test. Whenever the principal controls an asset required for a task, she can prevent the agent from carrying out this task. If, however, the principal provides the necessary asset, this by no means implies that the agent will use the asset in the desired way.

3 Directing production

This section examines the organizational form within which the principal can implement the desired production method. The section first looks into the single-agent case, and then moves to the two-agent case.

3.1 Single-agent case

In the case where the principal only employs one agent, implementation is only possible if the agent is allowed to carry out all tasks. The implementation problem then consists in finding transfers τ at which the agent is willing to exert the desired effort allocation.

In order to implement the desired effort allocation, the agent has to be rewarded whenever the production result looks like it has been achieved by such allocation. In this case, however, he will produce this result in the cheapest, and not necessarily the most desired, way. Formally, the following result holds:

Proposition 1 (Window-dressing by a single agent)**.** *If there is a single agent i and the principal provides incentives to implement r^*, then the agent can profitably engage in window dressing, i.e. reduce effort on the demanding task and increase effort on the easy task. The desired production method cannot be implemented.*

Proof. To implement the desired production method, it is a necessary condition that the agent creates the result r^*. Whatever the incentive scheme used to implement (e^{D^*}, e^{E^*}) and hence r^*, the agent can obtain the same result and transfer by increasing effort on the easy task and reducing effort on the demanding task. The marginal reduction can be computed using the implicit function theorem: $\frac{d}{de^E} e^D = -\frac{\frac{\partial}{\partial e^E} R(e^D, e^E)}{\frac{\partial}{\partial e^D} R(e^D, e^E)}$. This derivative is larger than one for (e^{D^*}, e^{E^*}) by Equation (2), thus increasing e^E by one unit allows the agent to reduce e^D by at least one unit. The loss from increasing effort in the easy task is outweighed by the gain from reducing effort in the difficult task by Equation (1). So, deviating from (e^{D^*}, e^{E^*}) is profitable for the agent. □

The main message of this proposition is that any attempt by the principal to achieve the desired production method is doomed, because the agent always prefers a different effort allocation (refer to Figure 1). More specifically, the agent always slacks on the demanding task, and brushes up the appearance using the easy task. The crucial feature of the verifiable outcome R, which drives this result, is that it confounds the efforts of two (non-conflicting) tasks. This feature, together with the (local) preference for the easy task, suffices for the agent to engage in window-dressing rather than carrying out production in the desired way. This problem is not specific to deterministic environments. As long as the task on which the agent focuses effort matters, and as long as this effort cannot be identified by looking at the production result, window-dressing is possible.

The principal can circumvent this problem by assigning the easy task to another agent. This will be explored in the following sub-section.

Figure 1: Deviating from the desired production method by slacking on the demanding task ($-\Delta e^D$) and working harder on the easy task (Δe^E) leads to the same observable result (window dressing), while the agent gains (2)-(1).

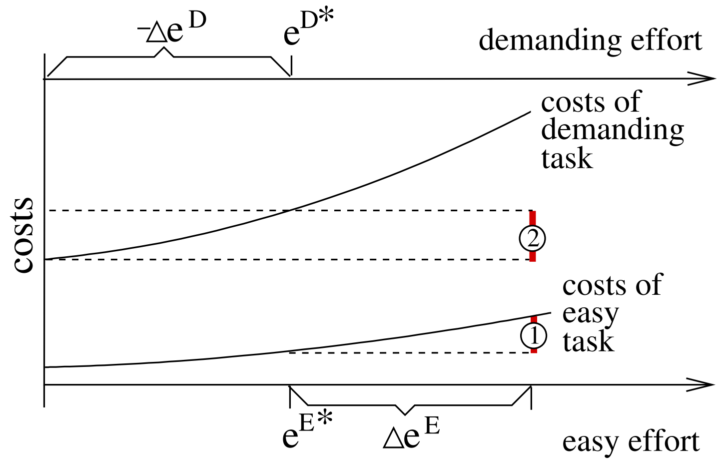

3.2 Two-agent case

In this section, we analyze the situation in which the principal hires two agents, and assigns them to the two tasks. This type of specialization enables the principal to implement the desired production method.

Proposition 2 (Implementation of the desired production method)**.** *If agents are assigned to different tasks (specialization), transfers can be chosen such that the desired production method (e^{D^*}, e^{E^*}) can be implemented as a (unique) Nash equilibrium.*

Proof. In order to show that (e^{D^*}, e^{E^*}) can be implemented as a (unique) Nash equilibrium, we consider two agents $i = 1, 2$, specify transfers, eliminate some strictly dominated strategies and compute the best-reply correspondences. For notational convenience, we focus on agent 1 and let him work on the demanding task (the analysis for agent 2 is completely analogous). Define the following transfer scheme:

$$\tau_1 := \begin{cases} c_1^D(e^{D^*}) & \text{if } R(e_1, e_2) = r^* \\ -K & \text{if } R(e_1, e_2) = R(e^{D^0}, e^{E^0}) =: r^0 \\ 0 & \text{else,} \end{cases}$$

where $K > 0$. For the transfers to agent 2 replace e^{D^*} by e^{E^*}.

Before determining best-reply correspondences given these transfers, we show that efforts that lead to negative payoffs are strictly dominated. Observe that agent 1 can always obtain a payoff of at least $-\delta$, where δ is an arbitrary small positive number, irrespective of the action chosen by agent 2. The respective strategy is to exert slightly more effort than e^{D^0}, say $e^{D^0} + \epsilon$. The same holds for agent 2. Since effort above the desired level e^{D^*} always leads to a negative payoff, these efforts are strictly dominated. In equilibrium the support of the strategy of agent 2 is hence a subset of $E_2 := \{e_2 \leq e^{E^*}\}$.

Given this observation about E_2, the payoff to agent 1 from a (possibly mixed) strategy F_1 given F_2 amounts to:

$$\int \int (\tau_1(R(e_1, e_2)) - c_1^D(e)) dF_1(e_1) dF_2(e_2) =$$

$$\int_{E_2 \setminus \{e^{E^*}, e^{E^0}\}} \int (\tau_1(R(e_1, e_2)) - c_1(e_1)) dF_1(e_1) dF_2(e_2) \tag{3}$$

$$+ P_2(e_2 = e^{E^*}) \cdot \int (\tau_1(R(e_1, e^{E^*}) - c_1(e_1)) dF_1(e_1) \tag{4}$$

$$+ P_2(e_2 = e^{E^0}) \cdot \int (\tau_1(R(e_1, e^{E^0}) - c_1(e_1)) dF_1(e_1). \tag{5}$$

Note that this payoff can at most be zero: the integrand of the integral in line (3) is zero only if $e_1 = e^{D^0}$ and otherwise negative; the integrand of the integral in line (4) is zero only if $e = e^{D^*}$ or $e = e^{D^0}$ and otherwise negative; the integrand in line (5) is always negative.

We use this payoff to compute the best-replies. We decompose the strategy space of agent 2 into three cases.

Case 1: $P_2(e_2 = e^{E^0}) > 0$. Agent 1's payoff is negative due to line (5) and he can profitably deviate to $e^{D^0} + \epsilon$ with ϵ chosen to be sufficiently small. Therefore, there is no best-reply for agent 1 in this case.

Case 2: $P_2(e_2 = e^{E^0}) = 0$ and $P_2(e_2 = e^{E^*}) < 1$. Line (3) implies that any strategy with $P_1(e_1 = e^{D^0}) < 1$ yields negative payoff and is strictly dominated by e^{D^0}. Thus, the best-reply in this case is e^{D^0}.

Case 3: $P_2(e_2 = e^{E^0}) = 0$ and $P_2(e_2 = e^{E^*}) = 1$. Only line (4) matters, and any strategy with $P_1(e_1 = e^{D^0}) + P_1(e_1 = e^{D^*}) < 1$ yields a negative payoff. Thus, the support of the best-reply is e^{D^0} and e^{D^*} in this case.

A completely analogous argument for agent 2 shows that (e^{D^*}, e^{E^*}) is the only fixed-point of the best-reply correspondences. □

The proof of this proposition is constructive and provides specific transfers that implement the desired production method. These transfers have to fulfill two conditions: (i) they need to work for the general class of result functions considered, and (ii) they need to implement the desired method as a unique Nash equilibrium. Both conditions are met by discontinuous transfers that reward the desired outcome r^*, and punish the outcome from low efforts r^0. The discontinuity implies that there is no best-reply if the other agent exerts the minimal effort, and transfers are reminiscent of a scheme used by Holmström (1982) in order to approximate the first-best solution in a single-task setting with noise (Theorem 3 in Holmström, 1982). Differing from Holmström's scheme, where the approximation requires ever larger punishments that occur with ever lower probabilities, the punishment here can be arbitrarily small. Moreover, the discontinuity is not essential for implementation as such (as in Holmström's argument), but only ensures uniqueness (see appendix). For specific result functions, it is possible to implement the desired allocation as a unique equilibrium even if transfers are continuous (for example, if the marginal effect of an agent on the result increases in the effort of the other agent—see appendix). The key ingredient in any mechanism to implement the desired allocation is that the principal gains an extra degree of freedom by having two agents: task separation allows her to set incentives separately for the two tasks.

The implementation of the desired production requires that the agent who is responsible for the demanding task has no access to the easy task. Otherwise, this agent could again profitably engage in window dressing.

The fact that the agent who carries out the demanding task has to be prevented from carrying out the easy task can entail direct costs, such as not being able to access certain areas or use certain tools. These costs may render it more difficult for him to carry out his tasks. Similarly, there may be a price for installing the necessary technology to ensure that the agent assigned to the demanding task is not working on the easy task. Moreover, hiring the additional agent can be costly, and finding an equally productive agent may not be possible. Proposition 2 provides a reason why specialization occurs despite all these costs, and in the absence of the advantages that are typically associated with specialization. Thus, it explains why a principal hires a less productive or otherwise costly agent, although the same observable result could be produced less expensively by a single agent.

The desired production method is achieved while the production technology remains the same. Moreover, the only contractual variable, the result of production,R, is also exactly the same. The crucial difference is that the principal employs two agents instead of one. Accordingly, the desired production may no longer be possible if the two agents act as if they were one agent. In other words, the two agents might collude, and thereby undermine the incentive scheme. This problem is explored in the following section.

4 Collusion

Collusion is only possible when agents have a larger contract space than the principal, so as to formally or informally contract on the easy task. For example, agents may engage in repeated interactions while the principal is being replaced.[7] We define the implementation of an effort choice e as *collusion-proof* when there is no contract among agents that stipulates transfers and effort choices different from e such that at least one agent is better off and none of them is worse off.[8]

If two agents collude, they will minimize their joint costs. If their marginal costs are identical, they are in the same position as a single agent. In this case, division of labor cannot help to implement the desired production method. However, the fact that all agents find one task demanding and the other one easy does not mean that all agents have to have identical marginal costs. There may, for example, be strong agents whose marginal costs for both tasks are lower than those of weak agents. This difference in marginal costs enables the implementation of the desired production method even when agents can collude. If a strong agent is assigned to the demanding task and a weak agent to the easy task, the gains from slacking on the demanding task may be more than offset by the loss from greater effort on the easy task. Figure 2 provides an example.

[7] Roy (1952) reports on workers' behavior in a production line where workers were able to enforce effort levels in order to restrict joint output, while the employer was unable to observe it.

[8] This definition of collusion-proofness is in line with Tirole (1986).

Figure 2: The gains to an agent with low marginal costs (2) from slacking are offset by the costs to the agent with high marginal costs (3) from higher effort, and collusion does not pay.

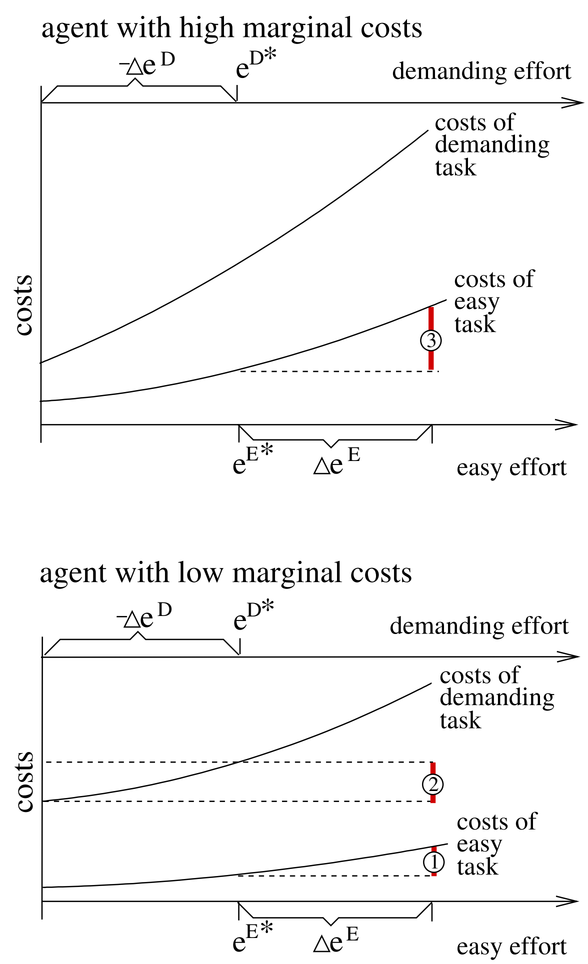

Proposition 3 (Collusion-proof implementation)**.** *A collusion-proof implementation of the desired production method is possible if and only if there is a (specialized) assignment of agents to tasks such that the desired way of production minimizes the costs of producing r^*, i.e.,*

$$(e^{D^*}, e^{E^*}) \in argmin_{e^D, e^E}\{c_i^D(e^D) + c_j^E(e^E) | R(e^D, e^E) = r^*\},$$

where i refers to the agent on the difficult task and j to that on the easy task $i \neq j$.

Proof. As a preliminary step, let us examine the problem of minimizing the costs of producing r^*:

$$\min_{e_i^D, e_j^E} c_i^D(e^D) + c_j^E(e^E) \text{ such that } R(e_i^D, e_j^E) = r^*. \tag{6}$$

Since $c_i^D + c_j^E$ is a continuous function, efforts are bounded and the side-constraint defines a closed set, the minimizer always exists. Let us denote the minimizer by $\underline{e}$.

The first step of the proof is to illustrate that it is impossible to implement $e^* := (e^{D^*}, e^{E^*})$ in a collusion-proof manner when this choice does not minimize the costs of producing r^*: $(e^{D^*}, e^{E^*}) \notin \text{argmin}_{e^D, e^E}\{c_i^D(e^D) + c_j^E(e^E) | R(e^D, e^E) = r^*\}$. Suppose a respective mechanism exists, then the agents can write a contract that stipulates $\underline{e}$ and leads to the same observable result r^*. The agent with a higher effort under this contract can be compensated for the additional effort by the other agent, and the latter is still strictly better off because the joint costs of production are lower for $\underline{e}$ than for e^*. Accordingly, the implementation is not collusion-proof.

The second step of the proof is to show that it is possible to implement e^* in a collusion-proof manner when it minimizes production costs $(e^{D^*}, e^{E^*}) \in \text{argmin}_{e^D, e^E}\{c_i^D(e^D) + c_j^E(e^E) | R(e^D, e^E) = r^*\}$. For this step of the proof, take the transfers from Proposition 2. The payoff from (e^{E^*}, e^{D^*}) is non-negative, whereas any result different from r^* does not yield a payoff that is larger than zero. So, agents cannot gain from a contract that stipulates efforts that lead to a result different from r^*. Within the effort allocations that lead to r^*, the desired production method has the lowest costs. Thus, there is no contract that leads to a weakly better payoff for both agents and to a strictly better payoff for at least one agent. The implementation is collusion-proof. □

Again, the ability of the principal to prevent agents from carrying out certain tasks is crucial here. If the strong agent has access to the easy task, agents would maximize their joint surplus by letting the strong agent do all the work. Then, window dressing by this agent would be a profitable deviation from the desired way of production.

It would be wrong to conclude from this proposition that the negative consequences of collusion can generally be avoided. This requires a specific relationship between the costs of agents and it may be difficult to find such agents. The key idea that is formalized by this proposition is that deviations from the desired allocation can be made more costly by assigning an agent with high marginal costs to the easy task, and an agent with low marginal costs to the demanding task. This idea allows the principal to obtain an allocation that is closer to the desired one, even if the desired allocation itself cannot be achieved.

Notice that the assignment of agents to tasks is entirely driven by the desire to implement a certain allocation. This assignment is based on the strong agent's marginal costs in the demanding task and the weak agent's marginal costs in the easy task. In particular, the assignment is independent of the relative marginal costs of agents with respect to tasks. Thus, agents may not be assigned to tasks where they have a comparative advantage. In other words, an inefficient task assignment may be the only way to ensure that the desired allocation (or an allocation close to it) can be implemented in a collusion-proof manner. This could explain why large organizations and bureaucracies that are designed to achieve a specific goal are often plagued by inefficient allocation of workers to tasks.

5 Conclusion

Consider a scenario in which the only evidence of an agent's performance is the final result of production, and there is no way of verifying how such result has been achieved. This paper has shown that in this scenario, it is not possible to direct a single agent's effort toward tasks in a desired way. However, this problem can be overcome by splitting the production process such that agents are assigned specific tasks, and are prevented from carrying out the tasks that other agents are assigned to.

Division of labor entails costs that should be traded off against its benefits. These costs include the costs of hiring workers, the costs of communication and coordination among workers (refer to Becker and Murphy 1992 or Bolton and Dewatripont 1994), the forgone benefits from task complementarities (Lindbeck and Snower, 2000), and the possibly higher compensation required by workers that like variation (i.e. workers with convex costs that are additively separable across tasks, see Itoh 1992). The costs of labor division increase further if such division has to be enforced.

The first main insight of this paper is that there are additional benefits to the division of labor that have so far been neglected, which explain why division of labor may occur where it is not expected ("over"-specialization). The existence of these benefits is relevant to the trade-off between a tayloristic and holistic organization of production. Lindbeck and Snower (2000) argue that new versatile technologies which make workplaces more flexible, such as computers, diminish the gains from specialization. Accordingly, production under new technologies should be less divided. However, this paper shows that if the versatility of the new technology renders it more difficult to determine how a production result has been achieved, then the diminishing gains from specialization may be offset by the incentive gains from the division of labor

The second main insight of this paper is that inducing a game between agents helps implement activities, even when agents are not played off against each other as in Dewatripont and Tirole's article on advocates (1999). The paper also explains how it is possible to prevent collusion among agents to undermine the implementation in this case. This requires, however, that agents' productivities differ in a specific way, and that the less productive agent be assigned to the easy task. For this assignment, the principal has to know the costs of the agents. If these costs are private information, the principal should implement a mechanism to elicit them. Our conjecture is that such a mechanism may be constructed by allowing agents to self-select tasks, and giving the agent assigned to the easy task the right to denounce an unproductive colleague assigned to the demanding task. Designing such a mechanism, however, is beyond the scope of this paper, and it is left to future research. Interestingly, the assignment of agents to tasks in a way that prevents collusion does not necessarily coincide with the comparative advantages of the respective agents. This may explain the apparent absurdity of situations where agents are prevented from carrying out a simple task and have to rely on a less able co-worker, although it would be cheaper if they were in charge of the whole production.

Both features, "over"-specialization and inefficient job assignment, are often associated with bureaucracies. Prendergast (2003) points out that bureaucracies may be regarded as an optimal solution to the problem of providing incentives, when important quantities are not tangible and cannot be contracted upon. This paper extends Prendergast's observation to hiring decisions and work assignments in bureaucracies: seemingly inefficient hirings and assignments can ensure the optimal provision of incentives in a sparse contractual environment.

References

Alchian, Armen A. and Harold Demsetz. 1972. Production, information costs, and economic organization. *The American Economic Review* 62, no. 5:777–795.

Athey, Susan and John Roberts. 2001. Organizational design: Decision rights and incentive contracts. *American Economic Review* 91, no. 2:200–205.

Baker, George. 2002. Distortion and risk in optimal incentive contracts. *Journal of Human Resources* 37, no. 4:728–751. Special Issue on Designing Incentives to Promote Human Capital.

Becker, Gary S. and Kevin M. Murphy. 1992. The division of labor, coordination costs, and knowledge. *Quarterly Journal of Economics* 107:1137–1160.

Bolton, Patrick and Mathias Dewatripont. 1994. The firm as a communication device. *Quarterly Journal of Economics* 109, no. 4:809–839.

Dewatripont, Mathias, Ian Jewitt, and Jean Tirole. 1999. The economics of career concerns, part II: Application to missions and accountability of government agencies. *Review of Economic Studies* 66:199–217.

Dewatripont, Mathias and Jean Tirole. 1999. Advocates. *Journal of Political Economy* 107, no. 1:1–39.

Fama, Eugene F. 1980. Agency problems and the theory of the firm. *Journal of Political Economy* 88, no. 2:288–307.

Gibbons, Robert and Kevin J. Murphy. 1992. Optimal incentive contracts in the presence of career concerns theory and evidence. *Journal of Political Economy* 100, no. 3:469–505.

Holmström, Bengt. 1982. Moral hazard in teams. *Bell Journal of Economics* 13, no. 2:324–340.

Holmström, Bengt. 1999. Manegerial incentive problems: A dynamic perspective. *Review of Economic Studies* 66:169–182.

Holmström, Bengt and Paul Milgrom. 1991. Multitask principal-agent-analysis: Incentive contracts, asset ownership, and job design. *Journal of Law, Economics, and Organization* 7:24–52. Special issue.

Itoh, Hideshi. 1991. Incentives to help in multi-agent situations. *Econometrica* 59, no. 3:611–636.

Itoh, Hideshi. 1992. Cooperation in hierarchical organizations: An incentive perspective. *Journal of Law, Economics and Organization* 8, no. 2:321–345.

Lindbeck, Assar and Dennis Snower. 2000. Multitask learning and the reorganization of work: From tayloristic to holistic organization. *Journal of Labor Economics* 18, no. 3:353–376.

Prendergast, Canice. 2003. The limits of bureaucratic efficiency. *Journal of Political Economy* 111, no. 5:929–958.

Roy, Donald. 1952. Goldbricking in a machine shop. *American Journal of Sociology* 57, no. 5:427–442.

Schnedler, Wendelin. 2006. Incentives, effort, and the efficient use of effort. http://www.schnedler.de/papers/ocs.pdf. 2008/07/03.

Schnedler, Wendelin. forthcoming. On the prudence of rewarding for a while benefiting from B. *Journal of Labor Economics* .

Smith, Adam. 1999. *An Inquiry Into the Nature and Causes of the Wealth of Nations*. London: Penguin. First published in 1776.

Strausz, Roland. 1997. Delegation of monitoring in a principal-agent relationship. *Review of Economic Studies* 64, no. 3:337–357.

Taylor, Frederick Winslow. 1911. *The Principles of Scientific Management.* New York: Harper.

Tirole, Jean. 1986. Hierarchies and bureaucracies: On the role of collusion in large organisations. *Journal of Law, Economics, and Organization* 2:181–214.

Wilson, James Q. 1989. *Bureaucracy: What Governmental Agency Do and Why They Do It.* New York: Basic Books.

Appendix

Proposition 4 (Implementation with continuous transfers)**.** *There is a continuous transfer scheme τ such that the desired production ($e_1^D = e^{D^*}, e_2^E = e^{E^*}$) can be implemented as a Nash equilibrium.*

Proof. Denote by $k(i)$ the task carried out by agent i. Consider the linear transfers: $\tau_i(r) = \frac{c_i'(e^{k(i)^*})}{\frac{\partial R(e^D,e^E)}{\partial e^{k(i)}}|_{(e^D,e^E)=(e^{D^*},e^{E^*})}} \cdot r$. Observe that the payoff of agent i is concave and has a unique maximizer for any e_{-i}. This best-reply is characterized by the first order condition: $\frac{c_i'(e^{k(i)^*})}{\frac{\partial R(e^D,e^E)}{\partial e^{k(i)}}|_{(e^D,e^E)=(e^{D^*},e^{E^*})}} \frac{\partial R}{\partial e^{k(i)}}|_{(e^D,e^E)} - c_i'(e^{k(i)}) = 0$. For $e_{-i} = e^{k(-i)^*}$, the maximizer is hence $e^{k(i)^*}$. □

Proposition 5 (Unique implementation with continuous transfers)**.** *If*

$$\frac{\partial\partial R(e^D, e^E)}{\partial e^D \partial e^E} \geq 0 \textit{ for all } e^D \textit{ and } e^E,$$

there is a continuous transfer scheme τ such that the desired production (e^{D^*}, e^{E^*}) can be implemented as a (unique) Nash equilibrium.

Proof. For notational simplicity, let agent i work on the demanding task (the definitions and analysis for the other agent are completely analogous). Define $\tilde{e}_{-i}(r)$ implicitly via $R(e^{D^*}, \tilde{e}_{-i}) = r$. Consider the following transfers τ_i such that $\tau_i(r)$ is constant for $r \geq r*$ and twice continuously differentiable for values $r < r^*$ with the derivative:

$$\tau_i'(r) = \frac{c_i'(e^{D^*})}{\frac{\partial R(e^D,e^E)}{\partial e^D}|_{(e^D,e^E)=(e^{D^*},\tilde{e}_{-i}(r))}},$$

where the equality also holds for the limes $r \to r^*$ with $r < r^*$. The second derivative of the transfer in the outcome is:

$$\tau_i''(r) = -\frac{c_i'(e^{D*})}{(\frac{\partial R(e^D,e^E)}{\partial e^D}|_{(e^D,e^E)=(e^{D*},\tilde{e}_{-i}(r))})^2} \cdot \frac{\partial\partial R(e^D,e^E)}{\partial e^D \partial e^E}\bigg|_{(e^D,e^E)=(e^{D*},\tilde{e}_{-i}(r))} \cdot \frac{\partial \tilde{e}_{-i}(r)}{\partial r}.$$

This second derivative is not positive because $\frac{\partial\partial R(e^D,e^E)}{\partial e^D \partial e^E} \geq 0$ and $\frac{\partial \tilde{e}_{-i}(r)}{\partial r} > 0$. Then, consider the payoff to agent i given a pure strategy e_{-i}: $\tau_i(R(e_i, e_j)) - c_i(e_i)$. Since τ_i is a concave function in R and R itself is concave in effort e_i, the first term is concave. Together with the observation that c_i is strictly convex, we obtain that the payoff given a pure strategy is strictly concave. Integrating this strictly concave function over the possible values for e_{-i} yields also a concave function and thus there is a unique best-reply. Using continuous differentiability of the integrand and the intermediate value theorem for integrals, the payoff given any mixed strategy of the other agent is equal to some pure strategy e_{-i}. Without loss of generality, focus is restricted to these strategies. First, consider the best-reply to $e_{-i} \leq e^{E^*}$. By construction, the derivative of the agent i's payoff is zero at e^{D^*}. Concavity implies that the first-order condition describes a maximum. The best-reply is hence e^{D^*}. Second, consider the best-reply to $e_{-i} > e^{E^*}$. Then, the maximal transfer, which occurs at r^*, can be achieved with an effort $e_i < e^{D^*}$. So, the best reply is below e^{D^*}.

If $k(i)$ denotes the task carried out by agent i, the best reply is $e^{k(i)^*}$ for $e_{-i} \leq e^{k(-i)^*}$ and below $e^{k(i)^*}$ for $e_{-i} > e^{k(-i)^*}$. So, the only fixed point is at (e^{D^*}, e^{E^*}). Since the best-reply to any strategy is a unique pure strategy, there are no mixed-strategy equilibria either. □

The condition, $\frac{\partial\partial R(e^D,e^E)}{\partial e^D \partial e^E} \geq 0$, establishes that the transfers suggested in the proof are concave in the production result. As a result, the agent's payoff is concave and the agent has a unique and pure best-reply. However, transfers do not have to be concave for the agent's payoff to be concave. Thus, violations of this condition are possible, and the desired efforts can still be implemented using the same continuous scheme.