Garibaldi, 218 – 222

13 Teams and Group Incentives

13.1 Introduction

Team-organized production is a key feature of many organizations. Most job advertisements explicitly look for workers that possess team-specific skills, and are motivated to work in teams. Indeed, team production and group-based incentives pose a great challenge to personnel economics. As it turns out, a simple approach to the problem would argue that organized team production is likely to lead to a reduction in productivity. This is the so-called $1/N^{th}$ problem, whereby effort put forward by workers under team production is likely to be less than what would be exerted under individual-based incentives. If this is indeed the case, all we can do under team production is to find ways to preserve the right incentives, even though it appears that overall productivity will hardly be larger than what it would have been if individual-based incentives were in place.

Yet, to fully understand team production, and all the emphasis that real life organizations do place on teams, we have to move forward and explicitly recognize that there exist benefits and costs linked to team production. After all, if so many firms are interested in workers able to work in teams, it has to be the case that team-specific skills (such as communication skills) and the possibility of mutual learning among team members may increase the total surplus generated by the workforce. The personnel economics literature is still growing in this dimension. Many good questions exist, but answers are often limited. In this context, we prefer to offer a case study approach.

Does the adoption of a team increase or decrease productivity? How does team composition affect productivity? Are teams more productive if members are homogeneous or heterogeneous? As we argued above, personnel economics has no full answer to this problem. What we can do is to bring empirical research to the topic, and analyse in close detail a case study on the adoption of team production in the garment industry. Surprisingly, and contrary to what the $1/N^{th}$ problem would predict, we will see that in the specific study analysed, individual productivity under group incentives has increased, rather than decreased. This suggests that there is still a lot to be learnt in the area of team production.

13.2 The Team Production Problem

We begin by presenting the simple team production problem, or the $1/N^{th}$ problem, as it is often defined in the literature. Consider a group of individuals of size N that work together on a specific project. The price of the output is normalized to 1 and the production is such that the total revenue produced by the entire team depends on the effort of every team member via the simple, additive production function:

$$Q = e_1 + e_2 + e_3 + \cdots + e_N$$

where e_i is the effort of individual i.

We can think of a team of software developers engaged on a project with a time deadline. There is an agreed payment and a target delivery date for the final product, but each day of delay reduces the fee received by the entire team by $X. Total output is the simple sum of the effort of all individuals.

Note that, according to the above example, (i) all team members are equally productive, and that by assumption (ii) the production function is linear, so that there are no complementarities between the team members' efforts. While complementarities make team production more advantageous, neither of these assumptions affects the main results we will come up with below.

Finally, assume that each agent's disutility of effort is given by:

$$C(e) = \frac{e_i^2}{2} \qquad \forall i$$

We now ask two questions in turn. First, what is the efficient effort level for each member of this team? Second, given a simple revenue-sharing compensation scheme, how hard will each team member actually work (the 'worker's problem')?

Efficient effort level. To calculate the efficient effort of each and every team member we have to consider the benefits and the costs associated with increasing effort of each agent. Total output is given by $\sum_i e_i$ while the total cost of this effort is $\sum_i C(e_i)$. This suggests that the efficient effort level is the solution to

$$Max_{e_1,e_2,\ldots,e_n} = \sum_i e_i - \sum_i C(e_i) = e_1 + \cdots + e_N - \frac{e_1^2}{2} - \cdots - \frac{e_n^2}{2}$$

where the maximization is to be taken for each and every member

$$1 - e_i = 0 \qquad \forall i$$

Therefore the efficient outcome is for each worker to supply at the point in which the marginal benefit of effort to aggregate team production (which is 1) is equal to the marginal cost of effort, which in this case is equal to e_i.

Worker's problem. Beside the efficient level, we know that each individual solves their own maximization problem by maximizing their utility. The utility of the individual depends, as always, on the compensation scheme and on the cost of effort, so that

$$U_i = compensation_i - \frac{e_i^2}{2}$$

The most obvious compensation scheme is one in which the total output from the team is shared proportionally among its team members. In other words, each worker maximizes his or her individual utility, given a compensation schedule, or 'output sharing rule'. Further suppose this sharing rule implies the following compensation, $compensation_i = \frac{\sum e_i}{N}\ \forall i$. In other words the team members agree to share the total revenues of the team equally. In this situation, how hard will people actually work?

Individual i's pay-off: $Y_i = \frac{e_1+e_i+\ldots e_N}{N}$; Individual i's disutility of effort: $C(e_i) = \frac{e_1^2}{2}$. So the individual's problem is:

$$\max_{e_1} \frac{e_1 + e_i + \ldots e_N}{N} - \frac{e_i^2}{2}$$

with the following first-order condition:

$$\frac{1}{N} - e_i = 0$$

An Individual's (privately) optimal effort is therefore given by $e_1^* = 1/N$.

How can the previous conclusion be explained, and why is the individual optimal effort so different from the efficient level? The intuition is as follows. While the worker is still bearing the full cost of his or her additional effort, the compensation scheme implies that the marginal benefit of one additional unit of effort is only $1/N$. This is called the '**free rider problem**' or the '**1/N problem**'. The total effort of all agents is $N \times \frac{1}{N} = 1$ instead of N, which is the efficient level.

Note that this result has nothing to do with the assumption that the group's output is shared equally among its members. It is true for any fixed sharing rule, where worker i receives a share γ_i of the total, as long as $\sum \gamma_i = 1$ (all the workers' shares must add up to 100 per cent).

Under the above rule in our example each worker's effort will equal γ_i, and the sum of all the workers' efforts will equal $\sum \gamma_i = 1$, rather than N which is needed for efficiency. Summarizing our result:

> Any group compensation rule that shares group output according to a fixed rate induces sub-optimal effort due to free riding. The problem increases in severity as the group size, N, increases.

The free rider problem associated with team production is quite common in reality, even outside strict personnel economics examples. Among students, when group projects are associated with a common grade, it is very difficult to induce efforts by all members, and there is always a tendency to rely on other people's effort and ability. A somewhat similar problem takes place when a group of friends decide to share restaurant bills. Each individual has a tendency to consume expensive meals, so as to rely on other friends' contributions.

13.3 Team Norms as Remedies to the Team Production Problem

In the above example the individual reward was assumed to be $1/N\ Q$ or γQ where γ can be any share as long as $\sum \gamma_i = 1$. Such a revenue-splitting game can be thought of as a special case of our linear compensation scheme, which is based on a fixed amount α plus a bonus component β. Think of these as special cases of a linear compensation schedule based on the group's output which can be written as

$$compensation_i = \alpha + \beta Q$$

We can ask what should be the value of β that induces the efficient level of effort, which we know should be $e^* = 1$. In other words the problem becomes one of finding β to induce an efficient effort $E^* = 1$ by every team member. The individual problem is

$$Max\ \alpha + \beta[e_1 + \cdots + e_n] - \frac{e_1^2}{2}$$

Hence individual i's problem yields

$$e_i = \beta$$

Therefore to get the worker to choose $e^* = 1$, we need to set $\beta = 1$. This is not surprising, since we know very well that effort is maximum when the individual is the full residual claimant, as the following conclusion highlights.

> To induce efficient effort by every member of a team, each team member's individual compensation must increase by 1 for every 1 increase in the group's output. Each individual must be the full residual claimant.

This makes clear how difficult it is to find such a scheme. The difficulty with the previous comment is that the output of all individuals must increase by $1 for every $1 increase in the group's output. It sounds quite hard.

One way to obtain this scheme is that the 'firm' or some agreed-on team member sets a target output M, which we can call **team norm** M.

- Everyone is paid zero if the output M is not reached, and the firm keeps the difference.
- The output is split evenly among team members if the output norm is reached, while the pay is zero otherwise.
- If the team norm M is the output when all individuals produce at the efficient level $e^* = 1$, then each worker gets paid \$1 and the firm breaks even.

As you can see, establishing a team norm is the way to keep the incentive right for workers, since that marginal increase of producing at the efficient level is exactly \$1. But the very difficult task is to establish the proper norm, and to enforce it.

13.4 Teams in Reality

Many firms use teams. Using teams involves both costs and benefits. The existing literature argues that teams are desirable for three main reasons: (i) to make possible gains in complementarities in production among workers; (ii) to facilitate gains from specialization by allowing workers to accumulate task-specific human capital; (iii) to encourage gains from knowledge transfer. Yet teams also have costs. The most important cost is the free rider problem, which was analysed above. So it is difficult to say *ex ante* what is the impact of teams on output. Let's see some key questions. Does the adoption of a team increase or decrease productivity? How does team composition affect productivity? Are teams more productive if members are homogeneous or heterogeneous? In a case study of the garment industry at Koret it is possible to answer some of these questions.

Worker heterogeneity Theoretically, worker heterogeneity in teams has two advantages. It facilitates mutual learning and can influence the group production norm. (i) *Mutual learning* suggests that more able workers are able to teach less able workers to be more productive. If the mutual learning effect is significant, then teams that are initially more heterogeneous in ability will perform better. (ii) *Team norm.* A relationship between worker heterogeneity and team performance could also be the result of forming a team norm, a concept that we formally defined above.

Worker decision to join a team Let us imagine a context in which workers are given the possibility of joining a team, and let us also imagine a setting in which non-team jobs are available in the workplace. Then each worker will solve a cost–benefit analysis over the decision to enter team production. Obviously, the highest-ability worker will join a team only if he or she obtains an additional source of surplus from team production. In terms of output performance, high-ability workers are bound to suffer, since they are going to share production with less productive co-workers. This suggests that for a high-ability worker to join the team, there must be additional reasons, which can be conceptualized in two ways: (i) productivity gains may derive from multi-skill abilities; (ii) socialization within the team may compensate high-ability workers in terms of income. As we see in the case study below, these phenomena do exist in reality.

13.5 A Case Study: Production at Koret

The case study is based on weekly productivity reports from a Koret garment manufacturing facility in Napa, California. The establishment produces 'women's lowers' including pants, skirts, shorts, etc. Prior to 1995, production was organized with individual piece rate. Over 1995–7, production organization shifted to team production. The organization change was introduced in response to demand by retailers that companies make just-in-time deliveries. Such demands required a more flexible production system, and many firms in the industry responded to such demands with team-based production.

Garment production is done in three stages. First, cloth is cut into pieces that conform to garment patterns. Second, garments are assembled by sewing pieces together. Third, garments are finished by pressing. The case study we are considering focuses on the sewing operation.

The change that we analyse is a passage from progressive bundling system production (PBS) to module production (MP). In PBS sewing operations are broken down into distinct operations. PBS is a piece rate scheme while MP is a team-based remuneration. Sewers are paid on the basis of *individual piece rates* according to a standard set for the operation they undertake. The standards are usually set by the management in accordance with the unions. Quality is evaluated by randomly selecting six out of forty garments and the sewer's name is recorded. In this setting piece rate is appropriate and possible.

In 1994 the plant manager began experimenting with flexible teams, which in the garment industry are called module production (MP). The manager asked for volunteers. After joining a team sewers could return to PBS if they preferred. The data are described in Table 13.1. In module production, each team includes six or seven team members who work in a U-shaped work space of 4 × 8m. The close proximity of workers and machines reportedly facilitated