

a) *Efficient effort level (Social Problem)*. The Pareto Optimal contract is the one that maximizes the sum of utilities of all the group's members, where the utility for each worker is defined as amount that worker i is paid (Y_i) minus the disutility of effort, i.e. $U_i = Y_i - C(E_i)$. Noting that the total amount the workers are paid must add up to the total amount of output that is produced $\sum_i Y_i = \sum_i Q_i$, the Social Problem can be expressed:

$$\text{Max}_{E_1, E_2, \dots, E_N} \sum_i U_i = E_1 + E_2 + \dots + E_N - \frac{E_1^2}{2} - \frac{E_2^2}{2} - \dots - \frac{E_n^2}{2}$$

With FOC: $1 - E_i = 0$, for every i

Therefore the efficient outcome is for each worker to supply one unit of effort.

b) *An individual worker's optimal effort level (Worker's Problem), assuming a simple proportional sharing rule*. Suppose now that each worker maximizes his/her individual utility, given a compensation schedule, or "output sharing rule". Further suppose this sharing rule has the form: $Y_i = Q/N \forall i$. 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 1's payoff: $Y_1 = \frac{E_1 + E_2 + \dots + E_N}{N}$

Individual 1's Disutility of effort: $C(E) = \frac{E_1^2}{2}$

So the individual's problem is:

$$\max_{E_i} \frac{(E_1 + E_2 + \dots + E_N)}{N} - \frac{E_1^2}{2}$$

with FOC: $\frac{1}{N} - E_1 = 0$

Individual 1's (privately) optimal effort is therefore given by $E_1^* = \frac{1}{N}$

The worker will work $\frac{1}{N}$ th as hard as in the efficient case!

Why? While the worker is still bearing the full cost of the additional effort, on the margin he is only getting $\frac{1}{N}$ of the extra output. This is called the "FREE RIDER PROBLEM"

or “ $\frac{1}{N}$ PROBLEM”. The total effort of *all* agents is $N \cdot \frac{1}{N} = 1$ instead of N , which is the Pareto-efficient level.

Note: 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 α_i of the total, as long as $\sum_i \alpha_i = 1$ (all the workers’ shares must add up to 100%).

Under the above rule in our example each worker’s will equal α_i , and the sum of all the workers’ efforts will equal $\sum_i \alpha_i = 1$, 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.*

Note: some other examples of free riding are:

- i) Group projects in class with a common grade.
- ii) Sharing a restaurant bill.
- iii) Non-point-source pollution problem

2. Remedies to the Team Production Problem

a) *An efficiency-inducing continuous reward schedule:* Recall from above that the individual reward was assumed to be $\frac{1}{N} \cdot Q$ or $\alpha \cdot Q$. Think of these as special cases of a linear reward schedule based on the group’s output: $Y_i = a + bQ$. We want to know the efficient b , and hence the problem becomes one of finding b to induce an efficient effort $E^* = 1$ by every team member. Hence individual 1’s problem becomes:

$$\max_{E_1} a + b \cdot (E_1 + E_2 + \dots + E_n) - \frac{E_1^2}{2}$$


with FOC: $b = E_1$. Therefore to get the worker to choose $E^* = 1$, we need to set $b=1$.

Summarizing our Result: *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.*

The way this scheme works is as follows:

- The “firm” or some agreed-on team member sets a target output M .
- If the group’s total output is M (which is the efficient output where $E^* = 1$) each worker gets paid \$1 and the firm breaks even.
- If total output is \$1 short of M , then the firm will dock *every* worker’s pay by \$1 and keep or destroy the difference.
- If total output is \$1 higher than M , then give *every* worker an extra dollar. In this case of course the firm loses $\$M-1$.

- b) *Discontinuous Reward Schedules*. This is the same basic idea, but much more draconian. Everyone is paid zero, or a very low amount if the target output is not reached. If the target, M is reached, all the output is distributed among the team members.



3. Practical Considerations

- a) When to use teams? As we saw above teams have a cost (“free riding”) and this should be considered in the decision to use them or not. But there are also benefits to teams such as complementarities among the team members. Obviously, if the benefits of the complementarities outweigh the costs from free riding then we should use teams.
- b) Examples of actual team incentives: Ways in which individual compensation is tied to group performance:
- Team bonuses
 - Profit Sharing
 - Stock Options

It’s unlikely the latter two of these will have much of an incentive effect because the “team” whose total output affects your compensation is your entire firm: the $1/N$ problem is likely to be very severe.

4. Evidence on Team Compensation

1. Experimental Evidence: Nalbantian and Schotter, “Productivity Under Group Incentives”

In a manner similar to Schotter and Weigelt’s experiment on tournaments (which we discussed last week), Nalbantian and Schotter conduct experiments on the relative performance of the various incentive schemes in groups of 6 subjects. In all cases, group output is the sum of the six members’ efforts plus a uniformly-distributed random error. Disutility of effort is a quadratic function of effort.

The schemes they implemented are:

1. “*Revenue Sharing*”: Each group member is paid $q/6$, where q is the group’s total output.
2. “*Forcing Contracts*” If group output falls below the target, R , every member receives a low, “penalty” wage, B . If it exceeds R , output is divided evenly, 6 ways.
3. “*Gainsharing*”: same as the forcing contract but with the target based on group performance in a previous period.
4. “*Competitive Teams*” there are two teams of 6, and the team generating the biggest total output gets a larger payment per worker than the other team (payments are equal within a team).
5. *Individual Monitoring* With some probability p , the firm checks individual effort levels, and dismisses workers supplying less than the efficient, target amount.

Main results:

-revenue sharing schemes induce free riding: there is more effort than theory predicts, but much less than the efficient level.

-forcing contracts don’t work very well. Why? There is a *coordination problem*: if you think even one member of your team will underperform or make a mistake, there’s no point in working hard yourself.

-“gainsharing” works better than forcing contracts.

-“competitive teams” work even better (note both the above schemes have “endogenously” defined targets)

-individual monitoring works just fine if the monitoring probability is high enough. (but note this could be expensive in practice).

Other result: “history matters”. Even in these totally anonymous situations, groups establish work “norms”. If they were in a low-effort-inducing scheme together to start, these low levels tend to persist.

2. A “Quantitative” Case Study: Knez and Simester, “Firm-Wide Incentives and Mutual Monitoring at Continental Airlines”

This is a (quantitative) case study of a firm where group incentives were introduced. We’ll discuss:

- a) Continental’s incentive scheme
- b) Why we would not expect it to work
- c) Evidence that it did work
- d) Why did it work? What’s special about Continental/Airlines? What’s missing from the simple economic model?

a) Continental’s incentive scheme

- before 1995, Continental was consistently one of the worst-performing airlines
 - since airline deregulation in 1978:
 - under bankruptcy protection twice
 - never made an annual profit when not under bankruptcy protection
 - ranked last among the 10 domestic airlines in
 - on time arrival
 - baggage handling
 - customer complaints
- in late 1994 a new senior management team introduced three policies:
 - changing airport managers
 - improving the flight schedule
 - a group incentive scheme

The group incentive scheme:

- \$65 to every hourly employee in every month Continental’s on-time performance was in the industry’s top 5 (starting in ’96: \$65 for top 3, \$100 for top 1)

b) Why we would not expect it to work

- Basically, free riding, aggravated by the fact that:
- Continental had about 35,000 hourly employees
 - These employees are very geographically dispersed

These factors aggravate both: -the “first-order” free-rider problem

-the “second-order” free-rider problem in mutual monitoring

c) *Evidence that it did work anyway*

- management and workers both say they think it made a big difference
- on-time performance improved dramatically (goal met 9 of 11 months)
- additional cash flow just from fewer missed connections: \$8 million/mo.; cost of bonus: \$3 million/mo.

- profits (millions):
 - 1992: \$-125
 - 1993: \$-199
 - 1994: \$-613
 - 1995: \$+224
 - 1996: \$+319
 - 1997: \$+385

But was this just coincidence, an improving nationwide economy, or (most importantly) other changes introduced by management such as the flight schedule and new airport managers?

Study design:

Measure *changes* in on-time departure performance in an airport/month before and after the bonus introduced. E.g. Newark August 1995 – Newark August 1994. Data for 32 airports from Jan1994-Nov 1996. (*diffs*)

Compare these changes between “treatments” (airports run by Continental employees) with those in a “control group” (airports where gate and ramp operations were contracted out: shouldn’t be affected by the bonus scheme). (*diff in diffs*)

Use multiple *regression* analysis to account for other, confounding factors that could affect performance *diffs*, including:

- level and change in Continental’s scale at that airport
- load factor (takes longer to load full planes)
- level of 1994 perf (could be harder to improve on hi initial perf.)
- employee tenure (resistance to change?)
- weather *diffs*
- on-time arrival rate
- change in airport manager (yes/no)

Results: performance improvements at outsourced airports were significantly smaller than at airports operated by Continental personnel.

d) *Why did it work?*

Employees *did* monitor each other and exert peer pressure when co-workers’ performance jeopardized reaching the goal.

Examples:

- employees started own performance review meetings
- employees contacted colleagues who called in sick, offering “assistance”
- employees chastised others for leaving their stations, loading bags slowly
- employees helped others improve, if this did not work reported others to supervisors

Why was this in individual workers’ interests?

- within airports*, teams are small, jobs standardized, and workers know each others’ jobs
- across airports*, teams are highly interdependent (because flight delays create a “domino effect”). This greatly magnifies the effects of underperformance by one particular airport: any individual airport’s performance could significantly affect the whole firm’s performance.
- when the scheme was first introduced, airport teams may have confused the effects of a better flight schedule with higher effort by other airport teams, and felt pressure to “measure up”. (helped solve the “coordination problem” in equilibrium selection).

Lesson: In the right circumstances, group incentives can be effective even in very large groups. When they work, they work because employees **monitor each other**. The circumstances that helped group incentives work in Continental’s case (small, highly interdependent teams performing standardized tasks) may however be quite special.

Also, teams may work better when the production function is Leontief: $q = \min(E_1, E_2, \dots)$.

Recent trend toward team production may reflect greater complexity of work—info sharing *among* workers now more important than previously?

3. An industry case study: Boning, Ichniowski and Shaw, “Opportunity Counts: Teams and the Effectiveness of Production Incentives”

-study of HRM and productivity in steel minimills: 34 production lines, observed monthly over 3 years. Very good output data.

-focus on two, potentially complementary practices:

- group incentive pay

- “problem-solving teams” [the sort that emerged spontaneously at Continental]

Results:

-group incentive pay raises productivity (by half a percentage point)

-problem-solving teams are only adopted when there is group incentive pay (what’s the incentive to contribute otherwise?)

-adding problem-solving teams to a group incentive scheme raises productivity even further (by up to 0.4 percentage points), but only when the production process is fairly complex.

-problem-solving teams are more likely to be introduced when production is complex.

Lessons:

Again, **group incentive pay can work**. Why? Mutual monitoring and (perhaps) humans’ natural proclivity to work in groups?

Human resource management (HRM) decisions can have productivity effects as great, or greater than many technical innovations.

Complementarities matter, both among HRM practices and the nature of the production process: careless, piecemeal HRM innovation can be ineffective or even counterproductive.

-