Lesson 5. Promotions and Tournaments

Aside from piece rates/commissions and delayed compensation, another way that firms use compensation policies to motivate workers is by awarding promotions (with associated wage increases) to the member of a work group who performs best. So one function of promotions is the incentives they provide to workers competing to get promoted. This is the function of promotions we will study this week. (Another function of promotions is to get the "right" people into supervisory or managerial tasks. This can be a little different and we'll consider it later in the course).

A. A Simple Mathematical Example

We begin with a simple mathematical example of a "tournament" in which two workers (agents) compete for a promotion. We'll work through, in turn, the:

1. The probability each worker (agent) wins the promotion as a function of both workers' effort levels
2. Optimal effort of each worker given the contest rules
3. Socially Efficient effort levels
4. Implementing Socially Efficient Effort Levels: Equivalence of tournaments and piece rates

In other words, we'll begin by showing that (in this simple example) *anything a firm can accomplish with a piece rate it can also accomplish with a tournament*. This is important because it is much easier for firms to run tournaments than to implement piece rates: Rather than measuring each worker's total output and rewarding him/her directly based on it, the firm only has to decide which worker's overall performance was best. It doesn't even have to be able to do this very accurately; in fact things can actually work *better* if there's some error in the firm's ability to measure performance.

Once we've accomplished the above task (demonstrating equivalence) we'll go beyond our simple mathematical example and argue that tournaments have other advantages that can make them strictly better than piece rates, plus some disadvantages that limit their usefulness in certain situations.

1. The probability of winning the promotion

Let's assume the outputs of the two workers in our work group are given by:

Output of worker 1: $q_1 = dE_1 + 0.5\varepsilon$

Output of worker 2: $q_2 = dE_2 - 0.5\varepsilon$

where ε is a random variable. In our example we'll assume ε~*uniform[-5,5]*; think of it as worker 1's "relative" luck. It can represent randomness both in the actual performance of the two workers or just in the firm's measurement of their performance. Graphically, the probability density function of ε looks like:

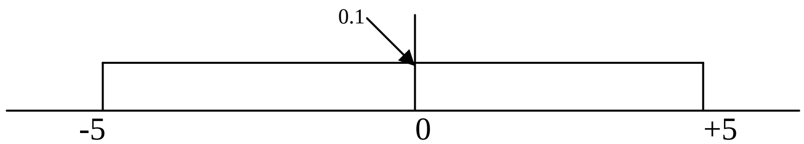

Essentially, "nature" picks a real number between –5 and +5. The chances this number lies between any two adjacent integers (e.g. between 2 and 3) are 10 percent.

By assumption, the worker with the highest level of measured performance wins the promotion. Thus, worker 1 wins if and only if: $q_1 > q_2 \Rightarrow dE_1 + 0.5\varepsilon > dE_2 - 0.5\varepsilon$

i.e. if: $\varepsilon > d(E_2 - E_1)$

(and if course worker 2 wins if and only if worker 1 loses).

What is the probability worker 1 wins, given the effort levels of *both* workers?
To understand this relationship let's graph it for the simple case where the productivity parameter, *d*, equals 1:

To get some reference points for the graph, note in turn that:

a) Suppose $E_2 = E_1$ (both workers exert the same effort). Then worker 1 wins if $\varepsilon > 0$. The probability of this is 0.5.

b) Suppose $E_2 = 0$ (worker 2 does nothing). Then
$\Pr ob(1\ wins) = prob(\varepsilon > -E_1)$, which depends on worker 1's own effort.

If $E_1 = 0 \rightarrow p_1 = 0.5$

If $E_1 = 1 \rightarrow p_1 = prob(\varepsilon > -1) = 0.6$

If $E_1 = 5 \rightarrow p_1 = prob(\varepsilon > -5) = 1$

c) Suppose $E_2 = 5$ (worker 2 works pretty hard). Then:

Prob (1 wins) $= prob(\varepsilon > 5 - E_1)$

If $E_1 = 0 \rightarrow p_1 = prob(\varepsilon > 5) = 0$

If $E_1 = 5 \rightarrow p_1 = prob(\varepsilon > 0) = 0.5$

If $E_1 = 10 \rightarrow p_1 = prob(\varepsilon > -5) = 1$

Putting it all together in a graph:

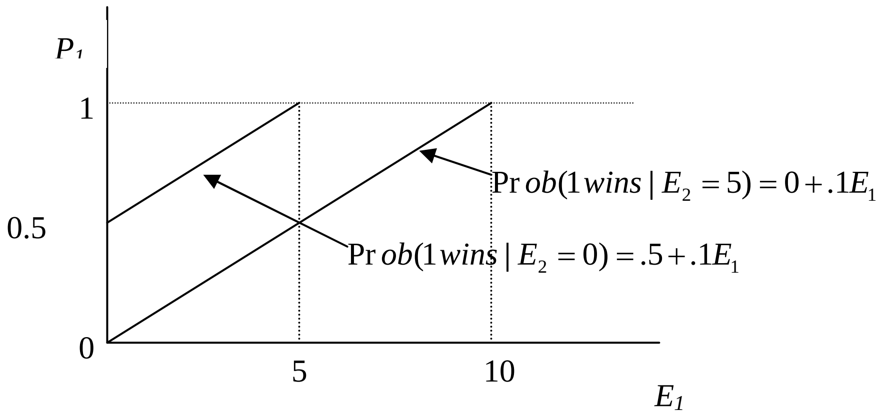

In general, for the case *d*=1: $p_1=0.5+0.1(E_1-E_2)$ and $p_2=0.5+0.1(E_2-E_1)$.

So the probability of winning the promotion:

- Increases with own effort, decreases with other's effort.
- Depends on relative effort only.
- Contest is fair: the probability of winning is the same if effort is the same.
- The effect of effort on the probability of winning decreases with the "noise" in measuring output. Suppose instead that ε~*uniform[-10,10]*. Then you get: $p_1=0.5+0.05(E_1-E_2)$ and $p_2=0.5+0.05(E_2-E_1)$.
- The effect of effort on the probability of winning also depends on workers' productivity (*d*). If instead of *d*=1, *d* was a higher number, the lines in the above graph would be steeper (holding the other worker's effort fixed, each extra unit of your own effort raises your chances of winning more).

In general, for a uniform distribution of "relative luck":

$$p_i(E_i, E_j) = 0.5 + \alpha d(E_i - E_j)$$

where $\alpha=1/R$ (*R* is the range of the distribution, which is also a measure of the amount of "noise" with which output is measured, or of the role of "luck" in the production process)

2. Optimal Individual Effort, given the contest rules

Let the "loser's" compensation be *a*.
Let the "winner's" compensation be $a+S$ (S is the prize "spread").
Let the disutility of effort be $E^2/2$.

Then in our example the expected utility of agent 1 is:

$$
\begin{aligned}
EU_1 &= p_1(E_1, E_2)[a+S] + [1 - p_1(E_1, E_2)]a - E_1^2/2 \\
&= a + p_1(E_1, E_2)S - E_1^2/2 \\
&= a + \{0.5 + \alpha\, d\, (E_1 - E_2)\}S - E_1^2/2
\end{aligned}
$$

And the agent's problem is $max_{E_1}\ EU_1$.
Taking the derivative wrt E_1, and treating worker 2's effort as given, the first-order condition for a maximum is $\alpha\, d\, S - E_1 = 0$. Re-arranging, the optimal effort is: $E_1 = \alpha\, d\, S$

Lessons:
- Effort increases with the prize spread, S, and with productivity, d.
- An increase of α (decrease of measurement noise) increases E_1
- (So you can always compensate for poorer measurement technology (low α) *by raising S.)*
- Effort is independent of a
- In this example (uniform distn of ε) each worker's optimal effort is independent of the other's effort. For other types of distribution (like a Normal) the other's effort *does* matter in one's own choice.
- Given a,d and S, both workers **provide the same effort**.

3. Socially Efficient Effort Levels

a) Finding the Socially Efficient Effort Level:

Remember that (expected) output is given by: $Q = dE_1 + dE_2 = d(E_1 + E_2)$
(why? according to the production function, worker 1's output is $q_1 = dE_1 + 0.5\varepsilon$. Since the expected value of ε is zero, the expected value of worker 1's output is just $q_1 = dE_1$.)

Disutility of effort is $E_i^2/2$ for i=1,2. so the efficient effort level solves:

$$\text{Max}_{E_1,E_2} \quad d(E_1 + E_2) - E_1^2/2 - E_2^2/2$$

FOC: $d - E_1 = 0$
$d - E_2 = 0$

In general, therefore, $E_i = d$ for $i = 1,2$

In what follows we'll do some numerical calculations for the **case where *d*=4**. Obviously in that case the size of the total "pie" to be divided between workers and firms is maximized when both workers choose $E_i = 4$.

b) Getting Workers to Pick the Socially Efficient Effort Level:

Agents' effort choices given the contest rules are: $E_1^* = E_2^* = \alpha dS = .1(4)S = 0.4S$
To induce the efficient level of effort $E_i^* = 4$ we therefore need $4 = 0.4S \Rightarrow S = 10$, i.e. a prize spread of 10.

Suppose (for the sake of argument—you'll see why below) we choose $a = 9$. Then we can work out:

Each worker's expected output = $dE = 16$

Each worker's expected income = $a + .5(S) = 9 + .5(10) = 14$

Expected Utility of each worker = $a + .5(S) - E^2/2 = 14 - 4^2/2 = 6$.

Profits per worker = output – worker's income 16-14 = 2

4. The Equivalence of Tournaments and Piece Rates.

Suppose that, instead of competing for a promotion, each of the two workers in the above example was compensated via a piece rate.

Under this kind of scheme, worker 1's income would be given by $Y_1=a+bdE_1$.

For the sake of argument, suppose that the firm chooses *b=1* (which we know leads to a Pareto-efficient outcome), and that *a=-2*.

Faced with this reward scheme and the production function used in this question, both workers will choose E^* = 4. (why? under the piece rate, worker 1's expected utility is therefore $a + bdE_1 - E_1^2 / 2$. Maximizing this yields the familiar result for piece rates that $E^* = bd = 1(4)$).

Now we can work out:

Each worker's expected output = dE = 16.

Each worker's total income = $a + dE$ = -2 + 16 = 14

Each worker's expected utility = $a + dE - E^2/2$ = 14 – 4^2 /2 = 6

Profits per worker = output – worker's income = 16 – 14 = 2.

Notice that this yields the same effort, the same worker utility, the same worker income, and the same profits as the tournament above. In general, any overall outcome (i.e. any combination of output, effort, worker utility and firm profits) you can generate via a tournament, can also be generatee via a carefully-chosen piece rate, and vice versa.

What does this mean?

1. Suppose you can't observe workers' output well, but you can observe *relative* output with some error (*you can rank workers*). As long as workers are risk neutral, you can do just as well with a tournament as with a piece rate. (in fact even if they are risk averse you may be able to do as well or even better—more on this later).
2. The efficient prize spread, *S,* increases with the degree of measurement error $1/\alpha$. So if you can't measure performance very well, you can still achieve efficiency by raising the prize spread.
3. In an efficient tournament, it is possible for workers with identical effort levels to receive very different rewards. In fact, in our simple example, both work equally hard. The one who actually receives the promotion is determined *purely by luck!* Despite this, the tournament is doing exactly what it was designed to do: inducing efficient levels of work effort among *both* the workers competing for the promotion.

B. Generalizations of the simple tournament model.

1) Risk aversion. If workers are risk averse and all shocks are idiosyncratic, it is not straightforward to rank piece rates versus tournaments.
2) Risk aversion and common shocks. In the previous example, the only kind of random shock was relative. But what about a common shock (like weather)? Under a piece rate we argued that it is optimal to give workers a break in bad times (lower production standards, higher *a*). A nice feature of tournaments is that they *automatically* do this, because they pay workers precisely based on relative performance, not absolute. Therefore tournaments are likely to be preferred to piece rates if workers' outputs are strongly affected by common shocks. A good example is the paper on chicken farmers by Knoeber.
3) More players. With only one prize, the odds of winning decrease as the number of players increase. Need to compensate with a bigger prize spread, or with prizes for second, third place, etc.
4) More levels. With more than one level, promotion to level 2 has two rewards: the prize of level 2 and the right to compete for prize 3. Some have argued that this helps explain why the wage increment for promotions is so much greater near the top of the corporate ladder than near the bottom.

C. Some Problems with tournaments.

1) Collusion: Agents might conspire to keep effort low. To avoid this the firm can reserve the right to hire an outside worker. Or it can reshuffle the "work groups" periodically, making it harder to collude. Again, the compensation of chicken farmers (Knoeber) is a good example. Collusion becomes less of a problem as the number of workers increases.

2) Sabotage: In a tournament, anything you can do to decrease your opponent's output raises your chances of winning. So tournaments can be vulnerable to sabotage. Short of actual sabotage, workers competing in a tournament may simply be less willing to help each other out and share information on the job. See the article by Drago and Garvey on this. Lesson: don't use tournaments (or limit the prize spread) when worker cooperation is an important aspect to production. (Or: you could even base promotion criteria on other workers' reports of "helpfulness").

3) Heterogeneous Ability: Tournaments with players of very different abilities will not be efficient, since *all* players can be discouraged. Nalbantian and Schotter demonstrate this in a laboratory experiment. Lesson: if possible, restrict tournaments to work groups with relatively homogenous abilities. If that's not possible, "handicap" the better players to give everyone an incentive.

D. Empirical Studies of Tournaments

1. Chicken Farmers: Knoeber

One can think of the market for broilers as a principal-agent problem, where the integrator (e.g. Purdue, Tyson, etc.) hires a group of growers to perform a service (raising chicks into chickens):

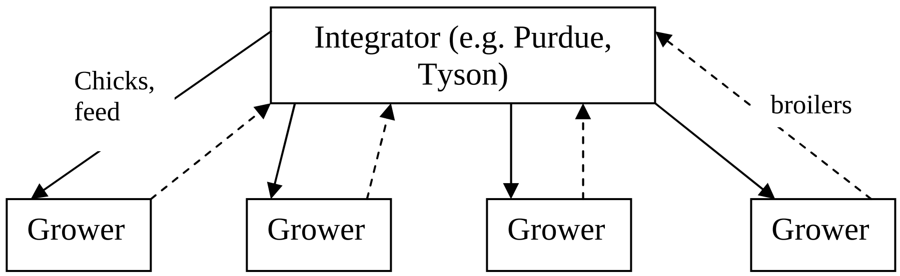

The integrator supplies chicks and feed. The growers provide "housing" and labor. Integrators pay the growers a certain amount per pound of broiler produced. Because the integrator supplies (and pays for) the feed and veterinary services, the integrator wants to encourage growers to use less feed and vet services per pound. It rewards low-cost producers by paying them a higher price per pound of broiler produced.

The interesting feature is that the price growers get depends only on their *relative* cost performance, compared to a dozen or so other growers in their area (about a 20-mile radius). This is a kind of tournament.

Why are tournaments used?

-it can't have anything to do with task assignment: growers are not promoted to be integrators. Therefore it's purely an incentive device.

-it can't be to economize on measurement costs: integrators already have and keep complete cost records.

-it's because growers are risk averse and there are important *common shocks*: mostly the breed of chicks, which is always changing. Some breeds are healthy and grow fast, others get sick and grow slowly. Weather and local disease outbreaks are other common shocks.

Knoeber calculates that payment by relative output eliminates about half the potential variance in growers' income, without compromising incentives in any significant way.

How are potential problems avoided?

-sabotage really isn’t possible
-collusion is avoided by randomly reshuffling the group of growers with whom your performance is compared. Aside from being in the same geographical area, members of the “groups” have no functional relationship to each other.
-because of the above, co-operation/info sharing isn’t really important to the production process. Also, the integrator supplies the latest production info.

Some other advantages of tournaments in this situation:

-the system adapts easily to technological progress affecting all growers the same way: no need to re-adjust the “piece rate” every time there is an improvement

-the system allows the integrator to experiment with genetic and other innovations without exposing the growers to additional risk

-the system (in fact *any* compensation scheme based on relative, not absolute performance) eliminates any temptation the *principal* might have to misrepresent the agents’ true performance.

2. Incentives for Helping on the Job: Drago and Garvey

While sabotage may not be a common problem, what about *lack of cooperation* among workers?

1988 survey of nonsupervisory eees at 23 workplaces in Australia

765 responding workers were allocated to work groups ranging in size from 2 to 80 (“work groups” do the same tasks in close physical proximity)

question asked about other work-group members:

> “do fellow employees refuse to let others use their equipment, tools, or machinery”: used to estimate each individual’s “helpfulness”

regression analysis “helpfulness” on

–Paid by a piece rate?
-paid a share of profits?
-promotion incentives

“Promotion incentives” i.e. the prize spread is measured as the variance in total pay levels within one’s work group that is not explained by differences in experience, seniority, etc. [wage “residual”].

Findings:

-workers paid by the piece are less helpful (but statistically insignificant)
-workers paid a share of profits help less (unexpected but also insignificant)
-higher promotion incentives reduce helpfulness

3. Asymmetric Tournaments: Schotter and Weigelt

So far, we have considered only:
-fair tournaments (the agent with the highest output wins)
-equally able contestants (same production function, or "*d*" for both workers)

Schotter and Weigelt study:
-two kinds of asymmetric tournaments
-two kinds policies to change asymmetric tournaments

Asymmetric tournament	**"Policy response"**
"Unfair" tournament (agent 1 wins if $q_1 > q_2 + k$: rules biased against agent 1 if $k > 0$)	*"Equal opportunity law"* (set $k = 0$: eliminate discrepancy in rules)
"*Uneven*" *tournament* (agent 1 has lower *d*, perhaps due to "pre-market" discrimination)	*"Affirmative Action law"* (introduces unfairness into tournament by biasing it against the more-able agent)

Implementing a "Tournament" in an Economics Experiment:

-agents don't actually work: they pick a "decision number", E, between 0 and 100
-agents are then charged an effort cost E^2/C_i
-agents then open an envelope containing their "luck": a random integer between **–a** and **+a**
-each agent then adds together his/her E and his/her "luck"
-we compare the totals of the two agents and award the payoffs (m, M) according to the contest rules.

Experimental Results:

Equal opportunity laws raise Effort, Utility and Profits

Effects of Affirmative action laws depend on the initial degree of disadvantage:

-for small initial asymmetries, AA laws reduce mean effort
-for large initial asymmetries, AA laws increase mean effort

Why? AA eliminates "dropout" and "discouragement". (agents "learn" in early rounds that it's pointless to try). [this occurs even though the theory in this case says they *shouldn't* give up—a property of the uniform distribution assumption]

Lessons:

Fair rules are always most efficient and profit-maximizing when contestants are equally able.

When possible, design separate tournaments to include homogeneous contestants.

When this is not possible, "handicaps" or "affirmative action" may increase the effort levels of *all* workers by preventing discouragement.