3
Relative Compensation

Piece rates are incentive devices that operate independent of any social environment. Specifically, workers need not be working with anyone else to be motivated by a piece rate scheme. Nor is the presence of other workers even relevant in many contexts. Piece rate compensation is based on an individual's absolute performance rather than his performance relative to some standard or some other individual. But many, including myself, believe that most motivation is produced not by an absolute reward but by compensation that is based on relative comparisons. Specifically, managerial employees who move up the corporate ladder do so by being better than their peers, not necessarily by being good. The same is true of young professors in search of tenure.

Relative versus Absolute

Comparisons are key not only in tenure decisions but also in determining promotions in private enterprise. Since promotions carry with them higher salaries, higher status, and perhaps more interesting assignments, workers seek to get promotions. In the process of doing so, they exert effort in an attempt to outperform their neighbors. Thus relative comparisons can provide as effective incentives as a piece rate or output-based compensation scheme based on individual performance.

Further there are good reasons why firms may prefer to use relative compensation schemes. Two come to mind. The first is that it may be easier to observe relative position than it is to observe absolute position. For example, consider coal mining: Two workers are sent into a mine at the beginning of a day. At the end of the day they bring out their piles of coal. It may be easier to eyeball the two piles and note which one is larger than to determine how much each particular pile of coal weighs in an unbiased way. Second, relative comparisons difference out common

Noise that risk-averse workers may not like. For example, two salespeople may have a very poor day, not because they did not put forth sufficient effort but because the economy was bad, a condition over which they had no control. If relative compensation is used, the effect of the economy is the same on both individuals and so the individual who performs better will still end up receiving the higher level of compensation.

Relative compensation theory, or "tournament theory" as it has come to be called,[1] is the theory used to determine the size of a raise associated with a particular promotion. It has three essential features. The metaphor of a tennis tournament is useful. Suppose that a Swedish promoter organizes a tennis tournament between Stefan Edberg and Bjorn Borg.

First, note that prizes are fixed in advance and are independent of absolute performance. Suppose that Borg wins the match. He receives a fixed prize that does not depend on the amount by which he beats Edberg. Furthermore both players may do exceedingly well and play an extremely tight contest. The total prize money will not be affected, nor will the distribution of prizes between the individual players be altered.

In the context of the firm this means that there are wage slots that are fixed in advance. There may be one vice-president slot and four assistant vice-president slots. The person who becomes vice-president enjoys the higher wage associated with the vice-president slot. The tournament model, taken literally, implies that the wage that goes to the vice-president is independent of the amount by which he exceeds the performance of the assistant vice-presidents in winning the job.

Second, a player receives the winner's or loser's prize not by being good or bad but by being better than, or worse than, the other player. Again relative performance rather than absolute performance is key. Borg wins the top prize because he is better than Edberg, not because he is good. It is the comparison between the two that is essential. In the context of the corporation, a worker is promoted not because he is good, but because he is better than everybody else in his cohort.

The third feature of relative compensation theory is that the effort with which the worker pursues the promotion depends on the size of the salary increase that comes with the promotion. Suppose that the prize money consists of $500,000. The prizes could be split evenly so that the winner gets $250,000 and the loser gets $250,000. Under these circumstances neither player would be particularly keen on winning and not much effort would be devoted to the activity. If, on the other hand, the winner took $500,000 and the loser took zero, a great deal of effort would be put into the activity. This suggests that the larger the raise associated with the given promotion, the higher the level of effort exerted to win that promotion.

Why then not allow the size of the raise, or the "spread," to go to infinity? After all, it is possible to create any size spread desired simply by requiring workers to ante up some amount initially and then allowing the winner to take everyone's ante in addition to the prize money. For example, suppose that the prize money is $500,000 but that we would like to create a $1 million spread between winner's prize and loser's prize. To do so, simply require that each contestant put up $250,000 going into the contest. The winner takes his own ante of $250,000 back, plus the $500,000 prize money, plus the other player's ante. The loser simply forfeits his $250,000 ante. Thus the loser comes out with $250,000 and the winner comes out +$750,000, creating a spread of a million dollars. By changing the ante, any spread can be created.

But the optimal amount of effort is not infinite. Even if a principal could dictate the amount of effort that an agent put into a particular project, it would be a finite level because after some point the additional compensation necessary to pay for the additional effort would not be justified by the incremental output associated with that effort.

To make the point vivid, think back to ancient Rome, where the difference between the prize going to the winning gladiator and the prize going to the losing gladiator was about as large as one could imagine. But the problem was that gladiators had to be drafted into service because people did not voluntarily sign up for the job. While the spread was large, the average salary was not sufficient to induce people to risk their lives for the compensation. Given the amount of effort associated with the activity, and the size of the average prize (the prizes were life to the winner and death to the loser), gladiators rarely volunteered for the job. It would be quite possible of course, with the appropriate level of compensation, to induce some people to volunteer. For example, a sufficiently large payment to their heirs could induce some men to compete for their lives. After all, that is exactly what happens when individuals volunteer for military campaigns. While the level of effort is high, and the loser's prize is extremely negative, the gains from winning are sufficiently large to induce people to enlist. This is important because in the context of the firm, individuals must be induced to join the organization voluntarily. Not only must be spread be large to induce effort, but the average prize money must be sufficiently high to attract workers to come to the firm in the first place. Otherwise, workers will opt to enter some other activity.

The mathematics behind this story are quite simple and somewhat illuminating. Consider a firm that has only two workers and sets up two jobs: boss and operator. Workers compete against one another with the winner being designated boss and the loser being designated operator. The winner receives wage W_1, and the loser receives wage W_2. No wages are paid until after the contest is completed. The probability of winning the contest depends on the amount of effort that each individual exerts. Let the two individuals be denoted *j* and *k* and let *j*'s output be given by equation (3.1a) and *k*'s output by (3.1b).

$$q_j = \mu_j + \varepsilon_j, \tag{3.1a}$$

$$q_k = \mu_k + \varepsilon_k, \tag{3.1b}$$

where μ_j and μ_k are the effort levels of *j* and *k*, respectively, ε_j and ε_k are random luck components, and q_j and q_k are output. The problem can be split up into two parts. First, worker behavior is modeled. Second, the firm maximizes profits, taking worker behavior into account by setting up the optimal compensation scheme.

Equation (3.2) is worker *j*'s optimization problem.

$$\underset{\mu_j}{\text{Max}}\ W_1 P + W_2(1 - P) - C(\mu_j), \tag{3.2}$$

where W_1 is the boss's wage, W_2 is the worker's wage, and *P* is the probability of winning the contest, conditional on the level of effort chosen. Also $C(\mu_j)$ is the monetary value of the pay associated with any given level of effort μ_j. The first-order condition is

$$(W_1 - W_2)\frac{\partial P}{\partial \mu_j} - C'(\mu_j) = 0. \tag{3.3}$$

There is a corresponding problem for worker *k*. The probability that *j* defeats *k* is given by

$$\begin{aligned} P &= \text{Prob}(\mu_j + \varepsilon_j > \mu_k + \varepsilon_k) = \text{Prob}(\mu_j - \mu_k > \varepsilon_k - \varepsilon_j) \\ &= G(\mu_j - \mu_k), \end{aligned}$$

where *G* is the distribution function on the random variable $\varepsilon_k \quad \varepsilon_j$. Differentiating *P* with respect to μ_j yields $g(\mu_j \quad \mu_k)$. Since *j* and *k* are ex ante identical, there should exist a symmetric equilibrium where *j* and *k* choose the same level of effort. Thus at the optimum $\mu_j = \mu_k$, so (3.3) becomes

$$(W_1 - W_2)g(0) = C'(\mu_j). \tag{3.4}$$

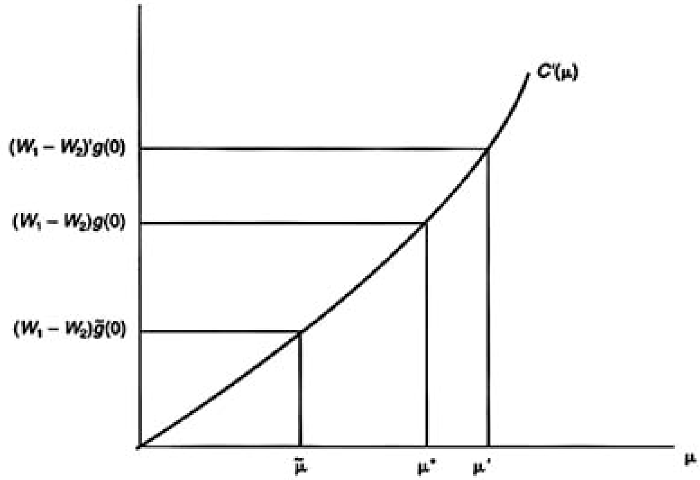

Figure 3.1

Equation (3.4) has two implications that are consistent with the tennis match story. First, the increase in W_1 W_2 implies a higher equilibrium level of effort, since $C'(\mu_j)$ is monotonically increasing in μ. A bigger raise induces workers to compete harder for the promotion.

In figure 3.1, the marginal cost of μ is plotted as $C'(\mu)$. The solution to the first-order condition is where $C'(\mu) = (W_1$ $W_2)g(0)$. This implies the solution $\mu = \mu*$. If the spread were raised to $(W_1$ $W_2)'$, the optimum would be at μ' rather than at $\mu*$. Note that $\mu' > \mu*$, since $C'(\mu)$ is necessarily increasing in μ.

A second implication is that the lower is $g(0)$, the lower is the level of effort exerted in equilibrium, since $g(0)$ is the measure of the importance of luck in this production environment. When luck is completely unimportant so that ε_k ε_j is degenerate, $g(0)$ goes to infinity. When luck is very important (when the distribution of ε_k ε_j has fat tails), $g(0)$ becomes very small. Thus, as the importance of luck increases, the amount of effort exerted for any given wage spread declines. Again, with reference to figure 3.1, if the spread were W_1 W_2, but the density function were $g(0)$ instead of $g(0)$ (and $g(0) < g(0)$), then the optimum level of μ would be only μ.

The logic behind this result is straightforward. If luck is the dominant factor in determining the outcome of the promotion decision, workers will not try very hard to win the promotion. In production environments where measurements of effort are noisy, large raises must be given in order to offset the tendency by workers to reduce effort. In fact, as will be shown momentarily, it is optimal to offset *any* reduction in effort induced by an increase in the importance of luck.

Let us now turn to the firm's problem. Given the workers' labor supply behavior, characterized by (3.4), the firm now wants to maximize expected profit, or equivalently, profit per worker,[2] since the number of workers hired is exogenous to this problem. The firm's problem then is

$$\max_{W_1, W_2} \mu - \frac{W_1 + W_2}{2} \tag{3.5}$$

subject to

$$\frac{W_1 + W_2}{2} = C(\mu). \tag{3.6}$$

Equation (3.6) is merely the condition that says that workers must be paid enough, on average, to induce them to apply for the job: $C(\mu)$ is the dollar value of the pain associated with the activity, whereas $(W_1 + W_2)/2$ is the expected wage that each risk-neutral contestant can expect to receive. Condition (3.6) merely says that the expected wage level must be high enough to induce workers to apply. Substituting (3.6) into (3.5) the maximization problem becomes

$$\max_{W_1, W_2} \mu - C(\mu), \tag{3.7}$$

with first-order conditions

$$\begin{aligned} \frac{\partial}{\partial W_1} &= (1 - C'(\mu))\frac{\partial \mu}{\partial W_1} = 0, \\ \frac{\partial}{\partial W_2} &= (1 - C'(\mu))\frac{\partial \mu}{\partial W_2} = 0. \end{aligned} \tag{3.8}$$

The solution to (3.8) implies that $C'(\mu) = 1$. This is the same efficiency condition that appeared in the last chapter of the piece rate problem (see equation (2.7) and recall that μ here is analogous to e in chapter 2). In other words, firms should set up a compensation scheme that induces workers to exert effort up to the point where its marginal cost is equal to its marginal benefit to the firm, namely \$1. Thus tournaments are efficient and induce the first-best level of effort. From this relation an optimal level of effort is determined, which after substituting into (3.6) gives the average wage necessary to attract workers to the firm. The wage spread is found by substituting the fact that $C'(\mu) = 1$ into equation (3.4) to obtain

$$W_1 - W_2 = \frac{1}{g(0)}. \quad (3.9)$$

Equations (3.6) and (3.9) are systems of two equations in two unknowns that solve for wage level and wage spread. As promised, the optimal wage spread varies inversely with $g(0)$, so the size of one's wage is increased to offset any increase in luck as reflected by a fall in $g(0)$. Of course the average wage does not change at all as a function of $g(0)$.

American CEOs have recently come under attack for their very high salaries, particularly in comparison to their European and Japanese counterparts. While their salaries may be too high, focusing on their salaries alone misses the entire point of the compensation structure. The CEO's salary is there not so much to motivate the CEO as it is to motivate everyone under him to attempt to attain that job. It is impossible to determine whether the CEO is overpaid simply by looking at the relation of CEO compensation to output. Earlier I argued that the structure of compensation was key; this is what I meant. It makes no sense to evaluate a job independent of the rest of the firm's hierarchical structure.

Tournament theory provides a way to think about the entire structure of compensation within the firm. In this sense it is very different from almost any other theory or set of theories of wages and compensation. Tournament theory implies a particular wage spread, or amount of variance in wages, within an organization. The size of the spread in this simple environment depends exclusively on the amount of noise associated with the production environment. If $g(0)$ is high, reflecting the fact that luck is unimportant in this particular environment, then the wage spread will be small. If $g(0)$ is low, so that luck is extremely important in this production environment, then the wage spread will be large. It is the amount of uncertainty that determines the wage spread and not the absolute output of the CEO.

Many have looked at the ratio of CEO salaries to the average salary of production workers. This is a measure of spread, although perhaps not a perfect one in this environment. But this ratio should depend only on the amount of noise in the environment, and not on the value of the CEO effort. It is misguided to attempt to prescribe some ratio of CEO compensation to production worker salary without understanding the fundamental role of each level's compensation in the entire structure.

The incentive role of salary is particularly important when one compares compensation structures across countries or across industries. Some industries may be riskier than others. The optimality conditions derived above imply that riskier industries should have larger wage spreads than less risky industries in order to induce workers to put forth the appropriate amount of effort. Thus CEOs in firms where demand or cost conditions vary significantly should be very well paid relative to production workers in those industries. In safer industries CEOs should be less well paid relative to production workers. Also note that this result does not depend on the existence of risk aversion. Throughout this chapter I have assumed that workers are risk neutral. It is the incentive aspect of a wage spread that induces high variance in noisy production processes and has nothing to do with distaste for risk.

If Japanese workers produced output in a less risky environment than American workers, then Japanese firms would optimally use a less skewed salary structure than American firms. European firms might find themselves somewhere in the middle. All of these would be optimal and would simply reflect the appropriate compensation structure.

It is important to think in terms of structures to understand compensation at all. For example, many American CEOs receive a large portion of their compensation in the form of stock and stock options. This is a natural response to the inability to use promotion as an incentive device for the top person in the firm. In the absence of relative comparisons, motivation of the top person must be provided through something akin to a piece rate. Stock price can be viewed as a proxy for CEO performance and is a natural instrument for CEO compensation.

Since tournament theory is new, it is difficult to provide specific evidence on the importance of the tournament story. But there are some empirical results. Main, O'Reilly, and Wade (1993) test tournament theory using data from large corporations in examining the relation of CEO compensation to that at the vice-president level. They attempt to find correlations, which they argue should be implications of the tournament model, between these ratios and a number of other factors. They find mild support for tournament theories, but they also argue that some of their results are inconsistent with the theory.

Knoeber (1989) examines compensation of chicken ranchers. He finds they are paid according to a tournament structure, which is a way to difference out common random components (the major one being weather) that affect chicken survival rates. Knoeber finds strong support for tournament theory in his paper and in a subsequent one (Knoeber and Thurman

1994), which also distinguishes between tournament theory and other relative compensation schemes that do not set up specific job slots. The distinction between tournament theory and other relative compensation theories will be discussed in depth in chapter 7.

Ehrenberg and Bognanno (1990) use data from the Professional Golf Association to determine whether prize money and its structure affect scores in golf tournaments. Surprisingly the predictions of tournament theory are borne out very clearly by their data, even in an activity where one would not think effort would be particularly sensitive to incentive pay. The Ehrenberg and Bognanno work is perhaps the best test of tournament theory, not because it is easily generalizable to the corporation but rather because the data are so well suited to testing the model.

Asch (1990) examines the response by navy recruiters to particular prize structures. She finds that they behave as tournament theory would predict and also points to some adverse consequences of prize structure pay. For example, she finds that recruiters substitute effort over time to win prizes and substitute quantity for quality when prizes are specified in terms of relative quantity goals as opposed, say, to more subjective performance evaluations.

Finally, Bull, Schotter, and Weigelt (1987) conduct experiments with college students. They find that the Nash equilibrium predicted by the tournament model is attained very quickly and that the predictions of the tournament model with respect to spread and variance are borne out in classroom experiments.

Industrial Politics

Tournament theory leads directly to a view of industrial politics (see Lazear 1989). Relative comparisons are key in tournament theory and this sets up an environment where competition between workers is likely. Indeed the way by which relative compensation motivates workers is by pitting one against the other, an action that provides incentives to put forth effort. But an adverse consequence of the competition is that workers will not want to cooperate with one another when they feel that they are competing for the same job.

In Lazear (1989) the implications of relative compensation for internal worker interaction are explored. Formally, the industrial politics model is quite similar to the one already presented. The difference is that equations (3.1a) and (3.1b) are changed to

$$q_j = \mu_j - \eta_k + \varepsilon_j,$$
$$q_k = \mu_j - \eta_j + \varepsilon_k. \quad (3.10)$$

where ŋ$_k$ is the harm *k* can inflict on *j* and ŋ$_j$ is the harm that *j* can inflict on *k*. In a relative environment, *j* does well not only by making himself look good but also by making *k* look bad. While it may be costly to engage in this type of sabotage, it pays for *j* and *k* to undertake some of the activity to the extent that it furthers their own relative positions. Workers do not want to cooperate with one another because their compensation depends upon "defeating" other workers within the firm.

Firms that recognize that workers are engaged in sabotage or other uncooperative behavior must adopt policies that mitigate the effects of these actions. Pay compression is a natural outgrowth of this situation. When the difference between the winner's salary and the loser's salary is reduced, two things happen. First, effort falls below the optimal level as workers respond to a smaller spread by decreasing their activities to obtain the promotion. This is bad. Second, workers reduce their anticooperative behavior because winning the contest is less valuable and engaging in sabotage or other adverse activity is costly to the individual who initiates the action. This is good. Thus pay compression, which works in the direction of equality, also enhances efficiency.

Tournament theory, then, implies salary compression, and the implications are very much in line with the standard stories told by personnel managers. Human resources managers often claim that salaries must be compressed to maintain internal harmony in a firm. If the difference between the winner's salary and the loser's salary is too great, morale suffers. One interpretation of the adverse impact on morale is that workers try too hard to disrupt one another's attempts to obtain the coveted position.

Individuals may differ in their ability to engage in sabotage or uncooperative behavior. Some people may be particularly good at attacking others, whereas some individuals may find it extremely costly to engage in this type of activity. Let us label the two types "hawks" and "doves," respectively. Hawks are individuals with an absolute advantage in attacking their neighbors.

Should personality should be a factor in the hiring decision? Intuition goes in both directions. Perhaps putting hawks in with doves makes the normally complacent doves hungrier, implying a positive effect on output. An alternative is that doves who are put in with hawks do not produce at their optimum level because hawks must be controlled by one means or another, and these control mechanisms have adverse effects on the effort level of doves.

The result is that the second argument wins out. If hawks and doves are put together, the optimal compensation strategy is to compress the wage spread to reduce the incentive of hawks to engage in negative activity. This reduces the effort not only of hawks but also of doves. However, if doves can be separated from hawks, the wage spread in a dove's firm can be set at its optimally higher level, implying that dove effort will not suffer. Furthermore the hawks in the firm will also have a compensation system that caters more closely to the direct interests of hawks. Total output is higher when different worker types are sorted into separate firms. Thus segregation rather than integration of worker types is best.

Many firms worry a great deal about their "culture" and whether a new employee fits nicely into it. Resources are devoted to screening an individual not only on the basis of ability level but also on the basis of personality attributes. Yet other firms argue that personality does not matter and that only ability is relevant. Indeed some firms (or academic departments) pride themselves on their willingness to ignore personality and to focus solely on quality. Which type of firm is doing the right thing?

This is simply a question of the existence of a separating equilibrium. Firms need not devote resources to screening workers if workers self-sort. But it can be shown that as a general matter, workers will not self-sort. Under most circumstances hawks will want to pass themselves off as doves for two reasons. First, the output in dovish firms is higher because no sabotage occurs. Second, hawks have an advantage in beating out doves for promotions. As a result firms cannot, as a general matter, afford the luxury of allowing workers to simply self-select.

Of course this does not mean that all firms must invest in screening. A firm that consists solely of hawks need not worry about doves attempting to contaminate the pool. But dovish firms must invest heavily in screening out hawks. I often joke that the University of Chicago, where I've spent most of career, does not screen on the basis of personality because it is already dominated by hawks. Any dove who is interested in coming there is welcome. He raises the output of the firm and is easy game. At less aggressive departments, hawks are discouraged from applying, and personality matters a great deal.

Tournament theory provides some implications for the compensation method appropriate for each level of the hierarchy. It is quite easy to show that individuals who have a comparative advantage in sabotage will be overrepresented in the higher levels of the firm. The reason is straightforward. Individuals can defeat their rivals in corporate challenges either by being able workers or by being good at making their rivals look bad. Unless the two characteristics are sufficiently negatively correlated, individuals who succeed are likely to have a higher than average level of ability and a higher than average level of aggressiveness. As one gets to the very top levels of firms, individuals are extremely able but also extremely aggressive. It is often remarked that CEOs tend to be among the least compassionate people in an organization. Tournament theory suggests that this result follows directly from having to fight the corporate war. Standard Bayesian updating shows that selecting on winning necessarily favors those whose underlying personal characteristics lean toward aggressiveness.

Since the high positions in organizations tend to be dominated by more aggressive individuals, it is necessary to reduce the incentives for those workers to compete with one another. Otherwise, all cooperation will be lost. Compensation at the top of the firm then should be based on absolute performance to a greater extent than compensation for middle-level managers. Since the upper ranks of the organization tend to be dominated by hawks, it pays to sacrifice some effort in order to prevent these extremely competitive individuals from killing each other off.

Another way to reduce the adverse affects of aggressive behavior is to set up the structure of the firm in a way that minimizes the consequences of such behavior. Two real world examples come to mind.

Before the breakup of AT&T, the president of the corporation was generally chosen from among the ranks of the presidents of the subsidiary operating companies. Thus the president of Illinois Bell or Pacific Telephone was more likely to become the president of AT&T than was the executive vice-president at the AT&T corporate office. One reason to select presidents from the field offices is that doing so prevents competition among vice-presidents at the AT&T head office. Competition among these individuals can be extremely destructive to the output of the firm, since their cooperation is highly desirable. While cooperation between the president of Illinois Bell and the president of Pacific Telephone might also be desirable, it is much less important to the output of the firm than is cooperation among individuals at the head office. Thus setting up competitions between individuals at the various subsidiaries who have little contact with one another is a better way to run an organization. It can be quite destructive to have individuals compete with one another when their cooperation is important.

This conclusion leads to a theory of organizational structure because it suggests that it may be better to have individuals compete across product lines than within function lines. More will be said about this in chapter 10.

Another example of how to structure a firm so as to minimize the consequences of aggressive behavior is provided by Dow Chemical Corporation. At Dow, individuals move from field offices to the corporate headquarters in Midland, Michigan. One executive told me that doing this creates a much more cooperative environment in the head office. Each individual knows that he will move into high-level jobs at the head office only after successful completion of fieldwork. This leads to competition among the field offices, but cooperation between field offices is not particularly important for the survival of the organization. On the other hand, cooperation at headquarters is much more important, so movements from the field to headquarters are more frequent than movements from headquarters back out to the field. If competition for the best field jobs occurred at the head office, then managers in Midland, Michigan, would be less cooperative with one another. Thus tournament theory not only provides a basis for thinking about industrial politics but also leads to a theory of hierarchical structure as well.[3]