ELSEVIER Journal of Accounting and Economics 19 (1995) 247–277

JOURNAL OF Accounting & Economics

Incentive compensation in a corporate hierarchy

Michael Gibbs

Graduate School of Business, University of Chicago, Chicago, IL 60637, USA

(Received May 1993; final version received June 1994)

Abstract

A theoretical and empirical analysis of within-job and promotion-based incentives for middle managers is presented, using personnel data from a firm. Within-job incentives are stronger than implied by previous studies. Evidence is provided that promotions sort employees by ability, and also generate incentives. Promotions are associated with large increases in lifetime earnings, as long as performance is sustained in the future. There is little evidence that the firm trades off within-job and promotion-based incentives as predicted. Instead, it appears to use a simple incentive scheme, resulting in declining incentives for those passed over for promotion.

Key words: Incentives; Compensation; Promotions

JEL classification: J33; J31; J30; J41; J44

1. Introduction

Large firms are often hierarchical, with relatively rigid organizational structures that do not change quickly. They are characterized as bureaucratic, slow moving, and filled with poorly motivated middle managers. Examples include the infamous Peter Principle, in which managers are promoted to their level of

This paper builds on and has benefited from joint work with George Baker and Bengt Holmström. I also appreciate comments from Ray Ball (the editor), Mike Bognanno, Charlie Brown, Joe Cooper, David Grossman, Paul Milgrom, Kevin J. Murphy, Canice Prendergast, the referee, and participants at talks at Baruch College, Northwestern University, SUNY-Buffalo, the University of Chicago, the University of Illinois, Wharton School, the 1993 meetings of the American Economic Association, and the 1993 KPMG/Olin Conference on Organizations, Incentives, and Innovation at the University of Rochester. Funding was provided by the Harvard Business School and the Graduate School of Business, University of Chicago.

SSDI 0165-4101(94)00384-H

incompetence (Peter and Hull, 1969). Such firms are very different from the entrepreneurial firm of economic theory, or the small firms from which much innovation seems to arise. How are managers motivated in large complex organizations? This is a key issue in understanding innovation in organizations. Innovation can arise or be stifled at all levels of the firm, from basic research to product positioning and marketing. Much of this work is carried out and overseen by middle management. Even though such employees may not be responsible for generating corporate strategy or fundamental new ideas, they implement corporate change.

A related question is how adaptable the incentive system is to changes in the firm's circumstances. If it is flexible and can change to reflect new organizational structures and environments, then the firm will be better able to cope with change. An incentive system that is difficult to adapt, on the other hand, may lead to organizational inertia.

In this paper I examine these issues by providing the most detailed look at incentive compensation for middle managers yet available, using personnel data from a typical large hierarchical firm. I document the use of within-job pay for performance in the form of raises and bonuses. I analyze the potential incentive effects of job assignments, including demotions, the threat of dismissal, and promotions. I also present a model of how a firm optimally combines these two general forms of incentive compensation. This is used as a framework for empirically examining how the firm manages the tradeoff between the two incentives.

The major empirical findings are as follows. First, pay for performance in the form of raises and bonuses is larger than suggested by most previous studies. Improving one's performance rating by one point (on a five-point scale) increases one's raise by about 4% for each year in which the performance improvement is sustained. In addition, the expected bonus increases by 1–2%, yielding a total annual increase of about 5–6%. These increases can lead to large differences in compensation between those with different performance after only a few years.

Bonuses provide incentives chiefly through the hope of earning one, rather than through variation in the size of the bonus. The threat of dismissal, and especially of demotion, does not seem to be an important source of incentives.

Promotions are a very important source of incentives. The rewards can be substantial, and promotion depends on performance. Promotions do entail an immediate one-time prize, though it is not enormous, but more importantly the opportunity to continue earning wage growth and further advancement through future sustained performance.

The evidence also suggests that this firm uses the promotion system to sort employees into different positions based on ability. Measures of ability affect both the likelihood of winning promotion and the prize earned on promotion.

Finally, I examine substitution between short-term incentive compensation and promotion incentives by focusing on samples of employees who are passed over for promotion. For these groups, performance falls as tenure in job level increases (as they are passed over for longer). This does not seem to be due to a change in their within-job incentive scheme, but rather from dwindling hopes of winning promotion. I find little evidence that this firm alters incentive schemes individualistically, as predicted by the theory, to balance incentives from these two sources. Instead, a reasonable conclusion is that this firm runs a relatively simple incentive compensation system that does not change as more is learned about employees in the same job. If such an interpretation is correct, then this incentive scheme does lead to declining performance for managers who are passed over, much like the Peter Principle, but based on incentives rather than incompetence. I conclude by discussing the implications of such an incentive scheme for agency theory and organizational innovation.

2. A model of incentives in hierarchies

In this section I provide a framework for analyzing the motivation of middle managers, before turning to the data. Middle managers generally work in hierarchies in which advancement appears to be an important part of the incentive system. In addition they are usually eligible for within-job merit raises and bonuses. An optimal incentive system should explicitly account for motivation from both. For example, it is sometimes asserted that those who are passed over, the 'corporate deadwood', are not well motivated by the potential for promotion, so that the firm should use stronger bonus or merit pay schemes for such managers. Similarly, it has been suggested that as a firm's growth slows it will have to shift toward greater use of short-term pay for performance because fewer promotion slots will be available (Baker, 1990). However, the combination of promotion-based incentives and short-term pay for performance has generally been ignored in agency theory. The model presented here is a first step in rectifying this.

First, some definition of terms. It is necessary to distinguish theoretically between rewards that are due to promotion, and those that are earned when there is no promotion. Promotion usually entails an immediate and largely permanent increase in salary. It also may make the employee eligible for additional rewards that come from the possibility of winning more promotions. This has an option value that is itself a reward (Rosen, 1986). In addition, in nonpromotion years real earnings and bonuses may vary from permanent levels depending on current performance. In the paper I will refer to 'within-job' rewards as merit raises and bonuses earned in the current job when there is no promotion, and 'promotion-based' rewards as those earned on promotion, including the present discounted value of immediate and permanent increases

in salary and bonus, and the option value of potential rewards from further advancement.

The setting of the model is the following. Employees combine ability α (which may include accumulations of human capital over time) and effort e to produce output $q = \mu \cdot \alpha \cdot e$, where μ is the marginal product associated with effort and ability. The firm measures performance with error $\varepsilon \sim N(0, \sigma_\varepsilon^2)$, so that the performance measure is $\mu \cdot \alpha \cdot e + \varepsilon$. Neither the firm nor employees know ability *ex ante*.[1] Ability is assumed to be distributed normally, $\alpha \sim N(\bar{\alpha}, \sigma_\alpha^2)$, with expected ability $\bar{\alpha}$ (the Appendix discusses relaxing this assumption).

The idea underlying this model is suggested by the findings in Baker, Gibbs, and Holmström (1994b). The firm sorts employees into different positions in the hierarchy based on learning over time about abilities. It is efficient to put higher-ability employees into positions of greater authority (larger μ) in the hierarchy. Empirically, most firms do this through promotions, which is the term used here. However, the assignment process might also include dismissals or demotions, which would entail negative rewards. All three will be examined in the empirical work. Sorting is modeled in a simplified way, by assuming that the firm sets an output standard, z, and promotes any and all employees whose performance measure exceeds this standard.[2] When this is so, the probability of winning promotion $p(e)$, given effort e, is

$$\begin{aligned} \text{prob}(\textit{promoted}) = p(e) &= \text{prob}(\mu \cdot \alpha \cdot e + \varepsilon > z) \\ &= 1 - F(z), \end{aligned} \tag{1}$$

where F is the cumulative distribution of measured output, which depends on the distributions of ability and measurement error (hereafter μ is normalized to 1, as differences in marginal product across jobs are not important in what follows).

Promotion signals to the external labor market that the employee is more able. This raises the employee's outside value, if ability has general rather than

[1] It might seem unreasonable that the employee does not know his or her ability. However, what matters is the match of the employee to the job's characteristics. In addition, the employee might not know his or her ability relative to peers, which is important if the firm runs a tournament.

[2] Another plausible rule is to promote the top performers regardless of their actual performance by running a tournament. The implications of a tournament version of the model would be identical to those derived below (Gibbs, 1994); the output standard rule is chosen solely for ease of exposition. This formulation also suppresses several interesting issues for the sake of simplicity. For example, when ability and effort are complements, the firm will want to elicit greater incentives in earlier periods to improve the signal-to-noise ratio in measuring ability, which helps sorting. Also, observing output over several periods and continually selecting out high performers means that the distribution of ability among stayers will change from a normal distribution, probably becoming skewed to the right. This will not change the basic results, as discussed in the Appendix.

just firm-specific value, which the firm must match in order to retain promotees. Even if all human capital is firm-specific, the firm must offer rewards to employees in order to elicit their investment in this human capital. The present value of this increase in lifetime earnings, plus the option value of potential future promotions, is the reward on advancement; it will be denoted by Δ.

The firm also uses within-job compensation to elicit optimal effort incentives, paying a base salary b_0 and a linear piece rate b based on the performance measure, so that income equals $b_0 + b(\alpha \cdot e + \varepsilon)$.[3] In principal, the firm would like to address the two objectives of sorting and incentives separately. However, noise in the performance measure makes it difficult to distinguish between the effects of ability and effort on performance. As long as promotions carry rewards, they will have incentive effects, because extra effort affects the chance of winning promotion in (1). Thus, the optimal incentive scheme must balance promotion-based and within-job incentives.

Promotions generate incentives here whether or not the firm desires to use them for this purpose. It is sometimes argued that firms explicitly design the structure of promotion rates and wages in the hierarchy in order to manage incentives, but there are questions raised by this view.[4] First, wages are influenced by the external market, so that the firm may be quite limited in the extent to which it can modify the promotion prize structure. The firm may have more control over the structure of the hierarchy, which is a key determinant of promotion rates and thus of promotion incentives. However, the hierarchical structure has important effects on productivity (Rosen, 1982). It impacts the distribution of ability across the organization, as well as the span of control of managers. Furthermore, some positions may be relatively fixed in number. For example, there is usually only one CEO, or one divisional head. The fixed nature of some job slots may mean that the firm has little control over promotion rates, at least in the short run. For these reasons the firm may find it quite costly to reorganize production in order to optimize promotion-based incentives, especially when there is the alternative of changing within-job incentives.

I will take a simpler approach, assuming that the hierarchical structure is set to optimize the distribution of ability, degree of centralization, spans of control,

[3] Linear pay for performance is not an essential assumption, though it provides tractability. On the other hand, Holmström and Milgrom (1987) have shown that if income effects are not too strong (as seems likely in the case of within-job incentive pay) linear schemes may be optimal. Moreover, linear or piece-wise linear schemes seem quite common empirically, perhaps because they are simpler to implement and enforce.

[4] It is also sometimes argued that the firm will use tournaments to manage incentives because ordinal evaluations are cheaper than cardinal evaluations. However, it does not follow that it must attach incentive compensation to job assignments to do so. It is possible to use relative performance evaluation within a job, without linking rewards to changes in job assignments. Malcomson (1984) gives a justification for attaching rewards to jobs based on enforceability of the incentive contract.

and other hierarchical production factors. This in turn determines an equilibrium wage distribution across jobs, driven by external labor market pressures and underlying learning about employee abilities. The firm then uses within-job pay for performance, which is relatively easy to adjust, to fine-tune individual incentives *conditional* on promotion-based incentives.

This presumption seems reasonable at least for the firm studied in the empirical sections below. The structure of the hierarchy, in terms of jobs, job levels, promotion rates, and relative wages, was remarkably stable over twenty years. Given that the firm tripled in employment and the data cover a recession and large variation in average wages over time, this stability is strong evidence that the hierarchy was costly to change. In any case, implications similar to those derived below would follow from a more general model in which the firm also alters the hierarchical structure to optimize incentives.

2.1. *The employee's optimization*

The employee maximizes expected income less the (convex) cost of effort $C(e)$. Expected utility is thus $EU(e) = b_0 + b\cdot\bar{\alpha}\cdot e + p(e)\Delta - C(e)$. The first-order condition is

$$C'(e^*) = b\cdot\bar{\alpha} + p'(e^*)\cdot\Delta. \tag{2}$$

The second-order sufficient condition, $-C'' + p''\cdot\Delta \leq 0$, is assumed to hold. This requires that the probability of promotion does not exhibit too-large increasing returns to effort. Since $p(e)$ is not concave, there is also a global constraint that expected utility is greater than the employee's best alternative. If these conditions are not satisfied, the employee will default on the promotion game and quit the firm (exert zero effort). Alternatively, this can be thought of as a situation in which the firm merely uses the piece rate so that p'' is set to 0, so that the second-order condition is sufficient and is satisfied. Such a case is a natural extension of the results derived below.

Eq. (2) illustrates the main idea of the model. Total incentives $C'^{-1}(b\cdot\bar{\alpha} + p'\cdot\Delta)$ are provided by the combination of within-job pay for performance, $b\cdot\bar{\alpha}$, and promotion-based pay for performance, $p'\cdot\Delta$. The first is the piece rate times the marginal product of effort (which includes an ability effect). The second is the prize from winning advancement times the marginal effect of effort on the chance of winning. The term p' is $\partial p'/\partial e = f(z)$, the height of a normal density curve. If the equilibrium promotion rate, given the standard z and optimal effort e^*, is $\rho = 1 - F(z)$, this height is taken at the $(1-\rho)$th percentile of the distribution.

2.2. The firm's optimization

Now consider the optimal within-job piece rate b^*. If competition drives the risk-neutral firm's expected profits to zero, then the value of output equals the employee's expected pay, so that $\bar{\alpha} \cdot e^* = b_0 + b \cdot \bar{\alpha} \cdot e^* + p \cdot \Delta$. Expected utility then reduces to $\bar{\alpha} \cdot e^* - C(e^*)$. The firm chooses a piece rate that maximizes this subject to (2), yielding the simple formula:

$$b^* = 1 - \frac{p' \cdot \Delta}{\bar{\alpha}}. \tag{3}$$

The intuition is simple: (3) implies from (2) that $C' = \bar{\alpha}$, which yields first-best incentives since all parties are risk-neutral. The piece rate adjusts for changes in promotion-based incentives so that total incentives are set properly. This is similar to the implications of Gibbons and Murphy (1992). The difference here is that performance affects the employee's reassignment within the firm.

Aside from noting that $b^* \leq 1$, the first observation about b^* is that it may be negative, if promotion incentives are very large. In such a case, the promotion system generates too much motivation given the effort cost, and the firm would use a negative piece rate to reduce total incentives. This is analogous to the possibility of negative incentive pay in Holmström's career concerns model (1982). In the firm studied here, bonuses are used in all jobs, so the possibility that $b^* < 0$ is ignored in the empirical work.

Note that in the presence of promotion-based incentives, b^* will be strictly smaller than it would otherwise be. If promotion rewards are large, then this may be an important part of the explanation for the common observation that within-job pay for performance seems 'low-powered'. This model predicts that those in the middle of the hierarchy where promotion opportunities are important, such as middle managers, will have smaller within-job pay for performance.

The main objective here is to find empirically testable predictions about pay for performance. Eq. (3) suggests that the firm will substitute between the two forms of incentive pay based on observable individual differences in expected ability. These expectations are updated as the firm sorts employees. The data used below allow several empirical measures of ability. Consider the effect of a change in expected ability $\bar{\alpha}$ on b^*:

$$\frac{\mathrm{d}b^*}{\mathrm{d}\bar{\alpha}} = \frac{p' \cdot \Delta}{\bar{\alpha}^2} + \frac{-p'}{\bar{\alpha}} \cdot \frac{\partial \Delta}{\partial \bar{\alpha}} + \frac{-\Delta}{\bar{\alpha}} \cdot \frac{\partial p'}{\partial \bar{\alpha}}. \tag{4}$$

Eq. (4) reveals that there are two general effects. First, there is a 'marginal productivity effect', the first term, which is positive. It is optimal to give more able employees greater incentives when effort and ability interact in production.

Expected ability also affects promotion incentives in two ways. Combined, these are an 'incentive substitution effect', the last two terms in (4). First, it affects

the prize, as higher ability employees will earn more over their careers. Thus the second term in (4) will be negative. Second, it alters the marginal effect of effort on the chance of winning. This appears in the third term in (4).

How does a change in expected ability alter the effect of extra effort on the chance of winning advancement? Since ability and effort are complements, we might expect that increasing expected ability raises the returns to additional effort in trying to win promotion. If so, then $\partial p'/\partial \bar{\alpha} \geq 0$, and the third term in (4) will be negative.

However, it is also true that if expected ability (and thus the promotion probability) is very high, extra effort makes little difference since the promotion is unlikely to be lost. Increasing ability could then lead to a smaller marginal effect of effort. For this reason, it is possible that for very high promotion rates the last term in (4) might be positive. For the employees in the firm examined in this paper, however, promotion rates are never very high, and thus the third term in (4) should be negative. These assertions are discussed in the Appendix.

We see that the combined incentive substitution effect should be negative. In other words, those with lower expected ability should generally be given larger within-job pay for performance as a result of lower promotion-based incentives. This is the intuition behind the pundits who urge greater short-term incentives for the deadwood who have been passed up by the promotion system. At the same time, though, the marginal productivity effect works in the opposite direction, reducing within-job pay for performance for those with lower ability.[5] This is the intuition of those who respond that incentive provision for the deadwood is not so important as they are less capable.

Which effect dominates is an empirical matter. In any case, within-job pay for performance should vary *individualistically* with differences in ability and promotion opportunities. This is the basic issue that will be examined. After a brief description of the data, I document the importance of both within-job and promotion-based pay for performance in the sample firm, and then turn to analysis of how the firm combines the two.

---

[5] A multitask model might predict that within-job and promotion-based incentives are complements. This would be the case if promotion incentives focus the employee too much on tasks that help win promotion rather than other tasks. Then larger promotion incentives would be accompanied by larger within-job incentive pay to balance out this effect. However, in managerial hierarchies job ladders are often organized around similar job tasks (especially in functionally organized firms). When that is so, the multitask effect seems less important unless one of the tasks that helps win promotion is influence activities, as in Milgrom and Roberts (1988). In any case, the empirical approach used below is to test the degree of substitution or complementarity between the two forms of incentives, so the findings shed light on both the single and multitask views.

3. Description of the data

The data used here are well documented elsewhere (Baker, Gibbs, and Holmström, 1994a,b), so I will describe them only briefly. They are the computerized personnel records for managerial (exempt) employees of a single large hierarchical firm over the period 1969–1988. Data include employee ID, age, sex, race, education, job title, country of employment, performance rating, salary, and bonus. Eight job levels from entry management (level 1) to Chairman–CEO (level 8) were clearly identified from analysis of the job titles. Seventeen titles out of over three hundred constitute a large share of employment. Levels 1–4 contain the bulk of managers, as the hierarchy narrows considerably from levels 4 to 5. All salary data are in local currencies, so this paper only looks at U.S. employees (92% of the sample). Most variables are available for the complete sample, but bonus data cover 1981–1988 only, so some of the analyses focus on that subperiod. Also, not all employees were given performance ratings in a given year.

An important feature of this firm is that it was remarkably stable in hierarchical structure over the entire period. Major job titles changed little in twenty years, and no new hierarchical levels were added. Levels contained similar proportions of employees in all years. Although the firm grew in employment by 8% per year, it did so very steadily. Moreover, job transition patterns between job titles and levels were very similar over the entire period. Finally, the firm experienced stable profits and revenue growth in all years except 1987, and no major restructurings. These characteristics are advantageous for the present study. The promotion system was in a rough steady state, despite firm growth. Thus, estimates of promotion rates and long-term rewards from promotion are easier to make and are more reliable. Moreover, the overall structure of the firm's compensation and career systems did not change, so the data can be pooled across years.

An important variable to measure is expected ability of employees. Several proxies will be examined. A logical one is the average performance rating earned by the employee while in his or her current position. However, the crude rating system and incompleteness of ratings used mean that this measure is not perfect and is not always available, so it is not used. Another proxy for expected ability is the average real raise earned while in the current job. This is based on the evidence in Baker, Gibbs, and Holmström (1994a,b). There we find that those with faster wage growth are more likely to be promoted, and are promoted more quickly. A second ability proxy that is used is tenure in the current job level.[6] This is based on the assumption that the firm sorts employees through the

[6] Years in job title may be a more precise measure because it controls for lateral transfers within a level. However, lateral transfers are very infrequent, so that the two measures are almost always the same. To make the tables clearer, years at level is used.

promotion system, so that expected ability should decline monotonically with how many times the employee has been passed over for promotion. A final ability proxy, suggested by the model, is the employee's probability of winning promotion the next year. This is measured by the fitted value in a logit predicting promotions.

The model glossed over the distinction between merit raises and bonuses. Theory gives us little guidance on this issue. We might think that raises represent a more permanent level of compensation based on the employee's market value and marginal product, while bonuses are more transitory and represent short-term incentives. However, this is an empirical question: the use of raises and the incidence and size of bonuses for incentives will all be examined.

4. Within-job pay for performance

This section considers within-job pay for performance, which has never been well documented quantitatively. This firm uses a very typical middle manager incentive system, comprised of a performance rating scheme combined with merit raises and bonuses. Table 1 presents summary statistics on pay, raises, and bonuses across major job titles in all levels. Mean salary is $52,506 in $1988 constant dollars – this is a relatively well paid sample. There is less variation in pay across titles in the same level than across levels, which corresponds to expectations. Pay rises with level, dramatically so at the top. The percentage differences in average salary across levels are quite similar to those found in Main, O'Reilly, and Wade's (1993) study of managers over many firms. These large differences are suggestive of the importance of promotion-based incentives, the topic of the next section. We also see in Table 1 that percentage raises are similar at all levels.

About 25% of managers receive a bonus, an average of about 10–40% of base salary. In the upper levels, bonuses tend to be very large. Similarly, the incidence of bonuses is higher in upper levels. This is broadly consistent with the theory in two ways. First, overall promotion rates are much smaller in these levels than below, independent of individual ability differences, so there may be an overall incentive substitution effect. Second, higher-level employees should be more able, and have larger within-job pay for performance by the marginal productivity effect.

These numbers do not control for performance. Table 2 provides summary statistics on performance ratings across major job titles and levels from 1981–1988. The firm uses a five-point rating scheme, where 1 is the best rating. This is like the performance rating systems seen in virtually every other documented case (Medoff and Abraham, 1980, 1981; Murphy, 1991). As usually observed, ratings are biased toward the best ratings, and highly concentrated, even more so in the upper levels. Only one employee, at level 1, ever earned

Table 1
Summary statistics on salary and bonus, by major title and level

To illustrate variation in pay across titles in the same level, this table shows only the 17 major titles, which represented 91% of U.S. employee-years. Minor titles that could be assigned to a level were included in the analyses in all other tables in the paper. Titles are disguised because of confidentiality of the data. '% raise' is the mean raise as percentage of last year's salary. '% receiving bonus' is the percentage of employees who received a bonus. '% bonus' is the mean bonus as percentage of last year's salary for those who received bonuses. All calculations are over 1981–1988. Standard deviations are equally-weighted means of annual standard deviations. All compensation data in all tables in the paper are in 1988 constant dollars.

						Positive bonus only	
Level	Title	*N*	Salary	% raise	% receiving bonus	% bonus	Std. dev.
1	Asst. J	6,021	$36,509	3%	12%	8%	$2,632
	Asst. K	1,898	$30,540	5	17	6	$912
	Asst. H	214	$32.726	4	7	5	$440
	All	8,133	$35,017	4	13	8	$2,187
2	M	3,448	$43,111	8%	17%	8%	$1,096
	L	2,991	$40,199	7	19	10	$4,160
	K	570	$36,637	6	22	10	$925
	I	559	$38,447	4	5	6	876
	J	423	$44,010	6	27	7	$1,288
	H	155	$40,258	5	19	7	$900
	All	8,146	$41,261	6	18	9	$2,227
3	G	7,469	$50,697	6%	31%	10%	$3,691
	F	361	$47,225	10	15	9	$2,516
	Sr. H	209	$53,156	4	23	7	$1,158
	All	8,039	$50,605	6	30	10	$3,579
4	E	7,366	$77,392	8%	40%	15%	$9,916
5	D	563	$126,322	7	38	30	$24,081
6	C	119	$173,385	6	39	36	$19,584
7	B	18	$374,740	11	50	38	$4,120
8	A	8	$561,663	9	50	38	$0
1–8	All	32,392	$52,506	6%	25%	12%	$4,741

Table 2
Within-job pay for performance

Summary statistics on the distribution of performance ratings and sensitivity of pay to ratings. Statistics are over 1981–1988. Levels 7–8 did not receive performance ratings. There was one rating of 5 (in level 1) during the period. Sample sizes are smaller than in Table 1 because ratings were given in 68% of employee-years (with similar fractions in most titles; somewhat fewer in levels 5–6). '% bonus' is calculated only over those who received a nonzero bonus. Cells with a • are empty, whereas those with a zero are not empty but round to zero.

		Performance rating				
Level	Statistic	1	2	3	4–5	1–5
1	% of ratings	25%	43%	30%	2%	100%
	% raise	9	6	3	−1	5
	% receiving bonus	18	14	7	5	13
	% bonus	7	8	9	25	8
	N	1,287	2,267	1,586	100	5,240
2	% of ratings	17%	55%	26%	1%	100%
	% raise	11	7	3	−3	7
	% receiving bonus	33	20	8	1	19
	% bonus	9	9	10	20	9
	N	997	3,130	1,512	72	5,711
3	% of ratings	27%	55%	17%	1%	100%
	% raise	10	7	3	−2	7
	% receiving bonus	53	28	9	3	31
	% bonus	10	8	11	26	9
	N	1,537	3,095	974	60	5,666
4	% of ratings	44%	46%	10%	1%	100%
	% raise	11	7	2	−3	8
	% receiving bonus	46	36	18	12	38
	% bonus	15	12	13	22	14
	N	2,249	2,361	496	33	5,139
5	% of ratings	45%	44%	10%	0%	100%
	% raise	13	8	3	6	10
	% receiving bonus	58	49	52	0	53
	% bonus	33	27	20	•	29
	N	98	96	21	1	216
6	% of ratings	52%	31%	17%	•%	100%
	% raise	11	9	−1	•	8
	% receiving bonus	56	56	44	•	54
	% bonus	40	39	17	•	36
	N	27	16	9	•	52

Table 2 (continued)

Level	Statistic	Performance rating 1	2	3	4–5	1–5
1–6	% of ratings	28%	50%	21%	1%	100%
	% raise	11	7	3	−2	7
	% receiving bonus	40	25	9	5	26
	% bonus	12	10	11	24	11
	N	6,195	10,965	4,598	266	22,024

a rating of 5 during 1981–1988. About 80% earned one of the two highest ratings. These ratings are not very precise or informative, and give employees little quantitative feedback about their performance. Though such rating systems seem to be far from perfect, most firms use them for their managerial employees. Subjective ratings are usually the *only* formal performance measures that are used, and raises and bonuses are often tied explicitly to them.

Table 2 also links pay changes to performance ratings. Improving one's rating by one point leads to about 4% additional real raise that year, with a little variation across levels. Bonuses rise about 2% in going from a rating of 2 to 1.[7] The incidence of bonuses increases by about 15% for each improvement in rating. This suggests that ratings have a greater effect on the likelihood of receiving a bonus than on the size of the bonus.

These numbers reveal a surprising degree of pay for performance within jobs in this firm. Overall, a one-point-higher rating leads to an average of 5–6% greater pay that year through salary and bonus. These are one-year increases, but much of the increase comes as the more permanent raise rather than as a one-time bonus. Sustained performance at a higher rating compounds this each year. For example, over five years, this would lead to about 35% higher compensation over someone who earned a one-point-lower rating each year.[8]

This level of pay for performance is larger than what is suggested by previous studies. For example, Medoff and Abraham (1980, 1981) present salary increase regressions for samples of managers from three firms. In the two with five-point

[7]Oddly, receiving a rating of 3–5 leads to slightly larger bonuses. However, it also reduces the likelihood of receiving any bonus at all, so that expected bonuses are smaller. The effect, visible across levels 1–4, is strongest for ratings of 4–5, but sizes here are small so it is difficult to say what drives it.

[8]These large numbers are not simply due to higher raises earned by those who received better ratings and were promoted. Removing those promoted from the calculations lowers the return to improved ratings only slightly.

rating scales, differences in raises between each rating ranged from 1–2%. In the third firm, which used an eight-point scale, the differences in raises between ratings were about 1.5%. In this latter case the range from high to low raises was about the same as in the firm studied here, with a 14% difference between the raises of the lowest and highest rated managers, so their third firm may be more comparable. Murphy's (1991) case study of Merck reveals differences in mean raises of about 1–2.5%, and 2.5–3% after Merck changed its merit pay system. In all cases the distribution of ratings is similar to those found here. Since we have so few studies of this kind, it is hard to say whether the firm studied here is very unusual or not. Along many dimensions the firm seems quite typical. The other firms are in manufacturing and pharmaceuticals, whereas the firm analyzed in this paper is in a service industry, but it is not clear that this would make much difference.

5. Promotion-based pay for performance

5.1. *Short- and long-term rewards from promotion*

This section documents the potential magnitude of incentives based on changes in job assignment. Before turning to promotions, consider the potential incentive effects of demotions and exits. Demotions do not seem to be a strong threat. For one thing, they almost never occur: there are less than two hundred demotions over the entire period. They could still be considered a strong source of motivation, if they were accompanied by very large penalties. However, this does not seem to be the case. The immediate effect of a demotion on real salary is close to zero, and many actually earn real raises when demoted. By the time demotees do leave the firm (the average tenure with the firm after demotion is 7.5 years), 36% have greater real salary than before the demotion. About 8% eventually are promoted to or above the level they were demoted from.

Now consider the threat of dismissal. We would expect that exit rates are lower for employees with worse performance ratings or lower wage growth (and for very high wage growth). There is some evidence for both, but only for the very rare ratings of 4–5, or for very poor wage growth. Moreover, the effects on exit rates are not large. In general, exit rates are surprisingly insensitive to tenure, performance, or other variables that are typically found to predict turnover. However, since we cannot observe what happens after employees leave, it is impossible to come to firm conclusions – they may experience large declines in lifetime earnings. All we can say for certain is that the firm does not seem to vary exits strongly with measures of performance, indicating that dismissals may not be as important an incentive instrument as compensation.

Now turn to promotion-based incentives. Table 3 provides measures of the potential magnitude of rewards earned from promotion. The first set of columns

Table 3
Short- and long-term rewards from promotion

Measures of the immediate and long-term rewards from earning promotion to each level. 'Lifetime earnings by final level' is the discounted (at 3%) present value of lifetime earnings (salary, but not bonus due to limited period of bonus data) of all employees who entered level 1 in 1970–1975, by the final level in which they were employed when they left the company, retired, or in 1988. No employees from this latter group advanced to levels 7–8. For those who left before retirement or remained with the firm in 1988, salary until retirement (assumed to be at age 65) is assumed to equal the salary in the final available year. '% difference from lower level' is the percentage change in lifetime earnings between adjacent levels, a rough measure of the long-term rewards from promotion.

	All promoted employees			1970–75 level 1 entrants		
Level	Raise	% raise	*N*	Lifetime earnings by final level	% difference from lower level	*N*
1	N/A	N/A	N/A	$1,125,677	N/A	346
2	$2,963	8%	3,989	$1,189,898	6%	309
3	$3,906	9	3,079	$1,307,796	10	315
4	$6,198	11	1,712	$1,606,329	23	320
5	$11,689	12	169	$2,493,127	55	13
6	$12,240	9	74	$2,319,793	−7	6
7	$70,338	26	8	•	•	•
8	$74,102	18	2	•	•	•
1–8	$4,212	9%	9,033	$682,520	N/A	1,309

N/A: not applicable.

show the short-term reward, in the form of the average raise earned by all promotees at each level. It is sometimes claimed that on promotion an employee earns a very large increase in pay. That is sometimes the case here, especially at the top levels (which have very small sample sizes), but is not generally so. The typical promotee earns a 9% raise (and bonuses similar in size to those reported in Table 1), not too different from those for all employees.[9]

Note that this raise is insufficient to explain mean differences in salary between levels. An important source of wage growth is that earned within jobs, not just in moves between jobs. Over their careers, 54% of real wage growth comes in the 80% of years in which employees do not earn promotion. Thus, a more appropriate measure of the rewards to promotion is to consider the long-term differences in earnings between those who earned a promotion and

[9]Table 1 shows an average raise of 6%, but the sample period is different from Table 3. Over similar periods the raise for those promoted is about 6% larger than for those not promoted.

otherwise similar employees who did not. The second set of columns in Table 3 provides a rough approach to this question. It presents the discounted present value of lifetime earnings of a quasi-cohort of employees who joined the firm at entry level (level 1) in nearby years (1970–1975), by the level at which the employee ended up at the firm. Thus, the percentage differences between these values in adjacent rows, also shown, provide some idea of the actual *ex post* long-term earnings differential that accrued to those who were promoted.

Such a measure is difficult to construct with precision. For the lifetime earnings calculations, I made several assumptions. The most important is that earnings until retirement were taken to equal the real salary the employee held in his or her last year with the firm. While this assumption may seem strong, there is some justification for it in that those who are passed over and never promoted again earn, on average, zero real wage growth in this firm (Baker, Gibbs, and Holmström, 1993). The assumption is problematic mostly for those employees who did not stay with the firm for long. For those who did, the imputed part of earnings is heavily discounted in the calculations. Though these numbers are crude, they are merely meant to get some handle on the potential magnitude of long-term rewards from promotion.

Long-term rewards as suggested by Table 3 are quite large, especially beyond level 2.[10] A promotion is associated with 6% to 55% larger total lifetime earnings. Even at level 2, this amounts to about two years of additional earnings. Thus, short-term rewards on promotion may seriously understate true rewards, especially if stated in dollars rather than percentages. This conclusion is consistent with those of Lazear (1992) and Main, O'Reilly, and Wade (1993), though it is difficult to make exact comparisons in the latter case because they focused on CEOs, while the numbers in Table 3 are for lower levels. Main, O'Reilly, and Wade find that the increase in long-term present value of salary on promotion to CEO (level 8 in this firm) is about $4.6 million, and that including stock and other forms of compensation increases it to about $6.2 million. Extrapolating prize growth from level 5 yields numbers in this order of magnitude. Thus, if promotion prospects are affected by effort, they should be an important driver to incentives.

5.2. *Promotions as sorting mechanisms*

I now present evidence that promotions are based on both current and past performance, and that the firm uses them to sort individuals by ability. Table 4 presents logits predicting who will be promoted. The logits include dummies

[10] Lifetime earnings are actually smaller for those who attained level 6 than for those who attained level 5. This is an artifact of small sample sizes.

Table 4
Ability proxies as predictors of promotion

Logits predicting whether or not the employee is promoted next year, over all years. A promotion is defined as a move to a job at a higher hierarchical level. 'Ave. % raise in level' is the geometric mean percentage raise earned by that employee in prior years in the current level.

Independent variables	Dependent variable: Promoted next year		
Intercept	1.76* (0.18)	0.49* (0.075)	0.71* (0.010)
Level = 2	− 0.69* (0.067)	− 0.52* (0.030)	− 0.44* (0.038)
Level = 3	− 1.43* (0.075)	− 1.18* (0.036)	− 1.03* (0.043)
Level = 4	− 4.43* (0.21)	− 3.47* (0.10)	− 3.55* (0.11)
Level = 5–7	− 3.97* (1.01)	− 2.29* (0.21)	− 2.22* (0.23)
Rating = 2	− 0.22* (0.068)		
Rating = 3	− 0.93* (0.11)		
Rating = 4–5	− 1.75* (0.52)		
Rating last year	− 0.27* (0.054)		
Rating 2 years ago	− 0.075 (0.046)		
Years at level = 2		0.73* (0.032)	
Years at level = 3		0.60* (0.041)	
Years at level = 4		0.44* (0.053)	
Years at level = 5 +		0.14* (0.048)	
Ave. % raise in level			0.036* (0.0028)
− 2 log-likelihood	9,957	37,794	23,840
N	15,267	48,039	29,544

Significance at 1% level indicated by asterisks (*). Standard errors are in parentheses. Controls for age, sex, and race were included.

for the employee's current level, to control for gross differences in promotion rates across levels.[11]

The first logit includes dummies for this year's performance rating, as well as continuous measures of ratings in the previous two years, to test whether sustained good ratings predict promotions. Current ratings are indeed significant predictors of promotion. The effect of receiving a 3 or 4 rating is even stronger than receiving a 2. Similarly, last year's rating is a significant predictor of promotion. However, the rating two years ago is insignificant; recent performance is weighted more than past performance.

The second and third logit focus on longer-term measures of expected ability rather than current performance. The second uses dummies for tenure in the

[11] Because of very large sample sizes, there is so much statistical power that most coefficients are statistically significant at better than 1%. In all tables only 1% significance is reported. Including year dummies in all regressions and logits in the paper does not change the conclusions.

current job, or how long the employee has been passed over for promotion so far. Promotion rates are lowest in the first year at the level, but peak in the second and decline thereafter. The tenure effects are significant predictors of promotion. This suggests that ability is important. If human capital accumulation were solely what was driving promotion, we would expect to see increasing promotion rates with tenure in level.

The third logit uses another proxy for ability, average percentage raise earned in previous years in the level. If wage growth reflects ability, then this should be positively related to the chance that one is promoted. This is indeed the case. However, the coefficient is small.

Taken together, these results suggest that promotions are based on performance. As long as performance measures, such as ratings, raises, and time to promotion, are affected by an employee's effort, then it follows that promotions will have incentive effects. If this is so, then given the magnitude of the rewards from promotion suggested in Table 3, it seems reasonable to conclude that promotions are a very important source of incentives in this firm.

But are promotions merely incentive schemes, or are they used to sort employees? Table 5 presents evidence on this question, in the form of regressions analyzing the determinants of the short-term reward earned on promotion. If promotions sort on ability, then measures of ability should predict not only promotions, but also the size of the rewards. This effect should be strongest for measures of ability that are more publicly observable, because such measures are more likely to increase the employee's external market value. On the other hand, if promotions are merely incentive schemes, then we would expect little individual variation in the rewards earned when promoted.

Table 5 indicates that performance and ability proxies are significant predictors of the size of the reward. A rating that is worse by one point reduces the percentage raise by 2–5%. A 1% larger average raise in the level before promotion increases the raise on promotion by 0.14%. This is not large, but it is more interesting considering that we might have expected this coefficient to be negative, due to measurement error, pay compression, or related effects that cause pay to regress to the mean. In fact, if a measure of relative pay before promotion is included in these regressions, it does have a negative coefficient, which is consistent with pay compression. What is happening is that those with faster pay growth are also promoted faster. Along these lines, the most interesting coefficients are those on tenure in the previous level, modeled as a quadratic. The effect of time to promotion on the raise at promotion is almost linear, and very large. The raise on promotion is reduced by 2% for each year it took until the promotion was earned, or 20% of the average raise for promotees. This is strong evidence that tenure in level is a good proxy for ability, especially as tenure in level is the most easily verifiable information available to the external labor market. Thus, it is also strong evidence that promotions are not merely incentive schemes, but also sorting schemes.

Table 5
Effects of ability proxies on rewards from promotion

Regressions predicting the immediate raise earned in promotion years, over all years. Rating is the rating earned in the promotion year, not the year prior to promotion as in Table 4.

Independent variables	Dependent variable: % raise in promotion year		
Intercept	19.65* (0.60)	20.74* (0.54)	13.99* (0.61)
New level = 3	0.35 (0.21)	0.96* (0.18)	0.82* (0.20)
New level = 4	2.83* (0.27)	3.94* (0.23)	3.09* (0.24)
New level = 5–8	2.09 (0.92)	5.33* (0.62)	3.63* (0.60)
Rating = 2	− 2.23* (0.20)		
Rating = 3	− 3.74* (0.45)		
Rating = 4–5	− 8.31* (2.78)		
Years at previous level		− 2.06* (0.14)	
(Years at previous level)^2		0.13* (0.015)	
Ave. % raise in previous level			0.14* (0.017)
F	68	121	100
Adj. R^2	0.23	0.28	0.31
Adj. R^2, regression w/out year dummies	0.12	0.14	0.15
N	5,533	7,831	5,076
Dependent mean	9.82	9.31	8.14

Significance at 1% level indicated by asterisks (*). Standard errors are in parentheses. Controls for age, sex, and race were included.

These results taken together are consistent with the structure of the model, that the firm uses within-job pay for performance, that it uses the promotion system at least in part to sort employees, and that promotions are an important source of incentives. How does the firm alter within-job pay for performance to reflect differences in promotion-based incentives due to changes in their expected abilities? That is the subject of the next section.

6. Interactions between within-job and promotion-based pay for performance

We have seen that the best empirical proxy for expected ability differences is tenure in level. This dovetails nicely with the emphasis by pundits on those who are passed over. Table 6 presents summary statistics related to within-job performance for samples of employees in levels 1–4 who have been passed over for promotion more and more times. Panel A simply selects from the entire

Table 6
Incentives for losers

Performance statistics by number of years in level (number of times passed over for promotion). Panel A is for all employees who stayed in U.S. positions over all given years at level, so that sample size declines as some move out of their level. Panel B is the subset of Panel A who stayed at the level at least 10 years, so that sample sizes are constant over years at each level, except for the last column, which includes individuals more than once if they stayed beyond 10 years in the level. In the last column the samples are identical in the two panels. For example, there are 536 individuals who stayed at least 10 years, and they had a combined 1,769 employee-years at tenure in level of 10 or more. Sample sizes in panel A do not match those implied by promotion and exit rates because of demotions and moves to jobs not assigned to a level. Because bonus data are available only for 1981–1988, bonus statistics are calculated only over that subperiod, and bonus data are not available in the first two years for those in panel B.

(A) All employees

		Years at level									
Level	Statistic	1	2	3	4	5	6	7	8	9	10 +
1–4	% raise	8.9%	5.1%	4.0%	2.9%	2.3%	2.0%	1.8%	1.7%	1.8%	2.1%
	% raise std. dev.	8.4	6.9	6.5	5.9	5.8	6.0	5.9	5.4	5.2	5.1
	% receiving bonus	21.8	26.3	26.0	25.8	23.7	22.7	23.5	25.7	24.4	18.9
	% positive bonus	9.6	10.4	11.2	11.2	11.3	12.5	13.1	14.0	13.8	13.0
	% promotion rate	15.4	23.9	19.1	14.9	12.3	9.6	10.8	7.8	6.2	4.6
	% exit rate	13.1	12.6	11.6	11.0	10.2	11.6	12.3	11.3	11.2	14.6
	N	18,023	11,965	6,847	4,314	2,919	2,068	1,468	984	700	1,769
	% rating = 1	45.9	29.9	25.5	23.2	22.1	21.0	19.0	21.0	18.1	14.6
	% rating = 2	47.9	53.5	52.8	50.5	48.3	48.7	48.0	48.4	50.1	44.9
	% rating = 3	6.1	16.2	21.1	25.0	28.4	29.1	30.8	28.0	29.7	37.0
	% rating = 4–5	0.2	0.4	0.7	1.4	1.2	1.3	2.3	2.7	2.2	3.6

(B) Repeat losers: Employees who stay in a level at least 10 years

Level	Statistic	Years at level 1	2	3	4	5	6	7	8	9	10 +
1	% raise	N/A	0.8	-- 0.7	− 1.1	− 2.0	− 2.4	− 1.4	0.0	− 0.4	0.0
(N = 65)	% raise std. dev.	N/A	4.8	4.0	4.0	4.6	4.6	4.0	4.2	4.5	4.2
	% receiving bonus	N/A	N/A	16.7	16.7	5.6	4.8	15.8	11.1	4.3	4.4
	% positive bonus	N/A	N/A	9.7	16.4	16.5	18.1	9.4	9.6	9.0	6.2
2	% raise	4.8	1.8	0.0	− 1.3	− 1.3	− 1.9	− 1.5	− 0.9	0.5	0.6
(N = 114)	% raise std. dev.	4.1	4.7	4.0	4.0	4.4	4.1	4.7	4.5	4.0	3.8
	% receiving bonus	N/A	N/A	16.7	16.7	13.6	13.3	3.5	8.8	12.4	9.3
	% positive bonus	N/A	N/A	7.9	7.5	11.6	10.5	8.2	8.9	8.9	8.4
3	% raise	5.4	1.5	0.1	− 0.9	− 1.6	− 0.9	0.0	0.2	1.4	2.0
(N = 144)	% raise std. dev.	5.8	4.2	4.6	5.1	4.4	4.6	4.4	4.3	4.7	4.9
	% receiving bonus	N/A	N/A	12.5	15.4	15.6	18.9	11.2	26.2	21.1	14.9
	% positive bonus	N/A	N/A	7.9	8.6	9.1	9.4	11.7	7.7	8.3	7.6
4	% raise	6.5	2.7	1.9	1.4	0.9	1.4	1.5	2.6	2.9	3.6
(N = 213)	% raise std. dev.	5.6	5.1	5.1	5.6	5.8	5.9	5.9	5.2	5.7	5.5
	% receiving bonus	N/A	N/A	81.0	65.2	71.2	68.9	60.7	54.0	45.7	29.9
	% positive bonus	N/A	N/A	16.7	16.2	16.2	18.4	18.1	18.0	16.8	15.7
1–4	% raise	5.9	1.9	0.7	- 0.1	− 0.6	− 0.4	0.1	0.9	1.6	2.1
(N = 536)	% raise std. dev.	5.4	4.8	4.7	5.1	5.2	5.3	5.2	4.9	5.1	5.1
	% receiving bonus	N/A	N/A	48.8	39.6	39.9	41.6	31.4	33.2	27.8	18.9
	% positive bonus	N/A	N/A	15.4	15.0	15.9	17.0	16.8	15.0	14.2	13.0
	% rating = 1	42.9	34.7	34.2	32.3	28.0	27.1	23.2	23.0	18.6	14.6
	% rating = 2	48.9	51.8	49.5	48.1	48.4	49.4	49.0	49.7	50.3	44.9
	% rating = 3	7.9	13.3	16.3	19.3	23.1	22.6	25.5	24.7	28.7	37.0
	% rating = 4–5	0.3%	0.3%	0.0%	0.3%	0.6%	1.0%	2.3%	2.6%	2.4%	3.6%

N/A: not applicable.

sample any employees who have been in their current level a given number of years. Thus, the sample declines in size with tenure, because employees are continually selected out through promotion, demotion, or exit. Because of these selection effects, ability of the group should decline with tenure, which makes it difficult to disentangle incentive and selection effects. To correct for this, panel B holds ability constant by selecting backwards, looking at the sample from the last column of panel A: those who remained in their level at least ten years. Thus the sample does not change in panel B.

Table 6 takes an eclectic approach to analyzing pay for performance. Both the raise and the standard deviation in raise are shown. In panel A, we see that both decline; the average raise declines substantially. The promotion rate declines dramatically (as claimed above, exit rates do not change much with tenure). Similarly, the distribution of performance ratings changes to many more 3's and many fewer 1's over time. All suggest that performance is declining for this sample, and that this is reflected in slower wage growth. But is the decline in performance due to selection of a lower ability group, or to lower incentives, or both? In panel A the incidence of bonuses declines only slightly, if at all, and if anything the size of the bonus increases a little. The incentive system may be substituting toward greater within-job pay for performance, especially in the form of bonuses, but it is difficult to tell from this table.

The model indicated that optimal total incentives should indeed fall for a group as their expected ability falls; this is embodied in the result that at the optimum $C' = \bar{x}$. To avoid this problem, panel B controls for differences in ability. Yet we still see similar patterns: raises fall, the incidence of bonuses declines, and performance ratings fall (again, the size of the bonus does not change significantly). None of these changes are as large as in panel A, because this is a lower ability group to start with. However, the fact that the patterns do arise indicates that performance is falling for the losers in panel B because their motivation has fallen. There are two potential explanations for this. Promotion incentives are surely smaller, as the probability of promotion declines steadily with tenure (and expected rewards should also decline, as indicated by Table 5). The firm may also be altering within-job pay for performance in either direction, as suggested by the model. If within-job incentives are increased, though, they must be outweighed by the fall in motivation from promotions.

The next two tables look at these issues more formally, analyzing how the sensitivity of pay for performance is adjusted by the firm for changes in expected ability. This is done by interacting performance ratings, measures of current performance, with ability proxies, in wage and bonus regressions. The objective is to see how the firm alters the reward it gives to each performance rating conditional on how the employee's perceived ability changes. Because this firm uses a typical performance rating system, and presumably ties rewards to ratings, changes in within-job pay for performance sensitivity should be reflected in changes in how sensitive raises or bonuses are to these interaction terms.

Table 7 follows this logic by considering the effects of rating–ability interactions on the size of the raise or bonus for employees who are not promoted. Three ability proxies are considered. The average raise in the level in prior years and years at level are as seen above. A new proxy is the employee's promotion probability, measured as fitted values from the second logit in Table 4. This measure of ability was suggested by the model.

Recall that the model was ambiguous in its prediction about how within-job pay for performance sensitivity changes as ability changes. Therefore, there is no presumption that the interaction coefficients will have either sign. However, if the marginal productivity effect dominates, then worse ratings should be penalized less as ability falls. This would mean that the interaction coefficients would be positive for proxies that are positively related to ability (promotion probability and average raise in level), and negative for years at level. If the incentive substitution effect dominates, the coefficient patterns should be the reverse.

The first three regressions in Table 7 are on the raise earned in nonpromotion years. All ability and rating variables are statistically significant. Interestingly, the coefficient on years at level is much smaller (in fact, negative) than the large one found in the middle regression in Table 5. This may be because the regression in Table 7 includes performance rating as well as years at level. However, this seems unlikely, because the rating coefficients do not change substantially from the first regression in Table 7, and because ratings distinguish only crudely between employees. An alternative explanation is that the regressions in Table 7 are for nonpromotion years, while those in Table 5 are for promotion years. This is consistent with the idea that the coefficients in Table 5 reflect external labor market effects that come about only when the employee earns a public signal by being promoted.

The interaction coefficients in the raise regressions in Table 7 are almost all economically trivial in size, and most are statistically insignificant. The largest ones, for years at level, show that the difference in raise from receiving a 2 rating instead of a 1 is 0.15% per year. Given the concentration in performance ratings, this effect is not large, but it does provide a little support for the view that the marginal productivity effect dominates.

We might expect to see trading-off between short- and long-term incentives reflected most strongly in bonuses. The second set of regressions in Table 7 predicts the size of the bonus (for those who receive one) when promoted. In these regressions, almost none of the coefficients are statistically significant, and almost all are economically insignificant. There is even less evidence that the firm adjusts incentives by changing the size of the bonus.

We have seen above that the size of the bonus seems to vary little across individuals. It may be that the firm is altering incentives through the likelihood of receiving a bonus rather than through the size of the bonus. To check this, Table 8 presents logits similar to the regression models from Table 7, predicting whether or not a bonus is earned in a non-promotion year. Again, the results

Table 7
Effects of ability proxies on sensitivity of salary and bonus to performance ratings

Regressions predicting the size of raises and bonuses for employees who are not promoted, including interactions between ability proxies and performance rating dummies. '% promotion probability' is the percentage probability that each employee will be promoted the next year, using fitted values from the second logit model in Table 4. Raise regressions are run over all years; bonus regressions are run over 1981–1988.

	Dependent variable					
Independent variables	% raise in nonpromotion year			log(positive bonus) in nonpromotion year		
Intercept	10.52* (0.46)	15.70* (0.25)	9.33* (0.14)	6.92* (0.11)	7.19* (0.062)	7.45 (0.043)
Level = 2	0.55* (0.12)	− 0.04 (0.094)	− 0.10 (0.092)	0.45* (0.033)	0.33* (0.030)	0.26* (0.033)
Level = 3	1.29* (0.17)	0.15 (0.1)	− 0.20 (0.093)	0.81* (0.042)	0.58* (0.029)	0.51* (0.031)
Level = 4	2.57* (0.25)	1.01* (0.12)	− 0.17 (0.10)	1.56* (0.062)	1.19* (0.030)	1.12* (0.031)
Level = 5–7	3.19* (0.43)	1.63* (0.39)	− 0.32 (0.25)	2.97* (0.093)	2.67* (0.080)	2.53* (0.057)
Rating = 2	− 2.57* (0.14)	− 2.93* (0.16)	− 2.73* (0.10)	− 0.14* (0.028)	− 0.060 (0.034)	− 0.12 (0.022)
Rating = 3	− 5.32* (0.18)	− 5.85* (0.20)	− 5.92* (0.11)	− 0.031 (0.055)	0.082 (0.063)	− 0.096 (0.038)
Rating = 4–5	− 11.01* (0.66)	− 9.61* (0.72)	− 10.62* (0.33)	− 0.28 (0.27)	1.12* (0.33)	0.74* (0.18)
% promotion probability	0.065* (0.0084)			0.0066* (0.0021)		
Years at level		− 0.84* (0.057)			0.13* (0.014)	
(Years at level)^2		0.027* (0.0042)			− 0.0069* (0.0012)	
Ave. % raise in level			0.074* (0.020)			0.04* (0.0045)

(Rating = 2)·(Prom. prob.)	−0.00038 (0.0062)			0.0031 (0.0015)		
(Rating = 3)·(Prom. prob.)	0.0035 (0.0082)			0.00064 (0.0028)		
(Rating = 4–5)·(Prom. prob.)	0.090* (0.035)			0.032* (0.0094)		
(Rating = 2)·(Years at level)		0.15* (0.041)			−0.019 (0.0085)	
(Rating = 3)·(Years at level)		0.25* (0.045)			−0.045 (0.013)	
(Rating = 4–5)·(Years at level)		0.18 (0.12)			−0.24 (0.10)	
(Rating = 2)·(Ave. raise)			−0.0026 (0.026)			−0.0094 (0.0063)
(Rating = 3)·(Ave. raise)			0.056 (0.037)			−0.0092* (0.013)
(Rating = 4–5)·(Ave. raise)			0.20 (0.12)			−0.029 (0.060)
F	379	392	439	238	238	312
Adj. R^2	0.33	0.35	0.32	0.55	0.56	0.58
Adj. R^2, regression w/out year dummies	0.11	0.12	0.10	0.53	0.54	0.52
N	22,644	22,644	25,725	4,121	4,121	4,094
Dependent mean	4.74	4.74	4.25	8.45	8.45	8.58

Significance at 1% level indicated by asterisks (*). Standard errors are in parentheses. Controls for age, sex, and race were included.

Table 8
Effects of ability proxies on likelihood of receiving a bonus in a nonpromotion year

Logits predicting whether or not an employee who was not promoted received a bonus, including interactions between ability proxies and performance rating dummies, over 1981–1988.

Independent variables	Dependent variable: Received bonus in nonpromotion year		
Intercept	− 2.66* (0.21)	− 1.57* (0.11)	− 0.81* (0.067)
Level = 2	0.66* (0.066)	0.42* (0.06)	0.39* (0.068)
Level = 3	1.39* (0.084)	0.88* (0.057)	0.85* (0.064)
Level = 4	1.93* (0.12)	1.10* (0.062)	1.10* (0.066)
Level = 5–7	1.98* (0.21)	1.21 (0.19)	1.62 (0.15)
Rating = 2	− 0.55* (0.062)	− 0.57* (0.073)	− 0.89* (0.054)
Rating = 3	− 1.65* (0.11)	− 1.76* (0.13)	− 2.10* (0.078)
Rating = 4–5	− 3.38* (0.63)	0.55 (0.83)	− 2.99* (0.35)
% promotion probability	0.030* (0.0044)		
Years at level		0.17* (0.030)	
(Years at level)^2		− 0.019* (0.0026)	
Ave. % raise in level			− 0.11* (0.010)
(Rating = 2)·(Prom. prob)	− 0.0039 (0.0032)		
(Rating = 3)·(Prom. prob.)	0.0013 (0.0054)		
(Rating = 4–5)·(Prom. prob.)	0.070* (0.025)		
(Rating = 2)·(Years at level)		− 0.0043 (0.018)	
(Rating = 3)·(Years at level)		0.034 (0.027)	
(Rating = 4–5)·(Years at level)		− 0.73* (0.26)	
(Rating = 2)·(Ave. raise)			0.078* (0.014)
(Rating = 3)·(Ave. raise)			0.13* (0.025)
(Rating = 4–5)·(Ave. raise)			− 0.068 (0.13)
− 2·log-likelihood	17,757	17,728	16,593
N	18,216	18,216	16,558

Significance at 1% level indicated by asterisks (*). Standard errors are in parentheses. Controls for age, sex, and race were included.

lend little support to the idea that this firm fine-tunes incentives based on individual circumstances. The coefficients are often insignificant, and all have little impact on the probability of receiving a bonus.

7. Discussion

The results reported above are certainly consistent with the pundits who complain about lack of motivation for middle managers. Even after controlling for ability differences, there are discernible declines in performance as employees

are passed over for promotion (which, incidentally, is concrete evidence that promotions provide incentives). The findings are also generally consistent with the framework of the model. There is evidence that the firm sorts employees by ability through promotions, that rewards depend on ability, and that promotions generate incentives.

However, there is little support for the main theoretical prediction of the model, that the firm will alter incentives individualistically. It is possible, of course, that the incentive substitution and marginal productivity effects are small in magnitude, so that incentives do not need to be changed much. However, this is belied by the observed decline in performance for losers. It is also possible that the two effects roughly cancel each other out, by coincidence, so that they are not picked up in the econometric analyses. To check for this possibility, regressions and logits like those in Tables 7 and 8 were rerun within individual levels and years (where sample size permitted). These regressions also failed to find economically significant changes in the sensitivity of within-job pay for performance due to differences in ability measures. It seems unlikely that such a coincidence would be so recurrent. Moreover, the model predicts that total incentives should not change even if substitution is small, but that is not what is observed.

A final possibility is that the firm does not fine-tune individual incentives, in either direction, as predicted by the theory. That is, the firm may run a relatively simple incentive scheme in which b^* does not change for individuals in the same job. This is similar to the conclusion in Baker, Gibbs, and Holmström (1994b), though from a somewhat different perspective. There we find that the firm appears to have a well-defined administrative wage policy that has discernible effects on compensation.

Why would a firm use a 'naive' incentive scheme as described above, despite what appear to be costs in terms of lost motivation as promotion opportunities fade? It may be that there are good contracting reasons to use simple non-individualistic policies. They are straightforward and economical to administer. They are simple to explain, and it is easy for employees to monitor the extent to which the firm reneges on any implicit part of the incentive contract. Equally important, such contracts are less susceptible to influence costs than are individualistic contracts (Milgrom and Roberts, 1988). For example, the lack of individualistic adjustments suggests a centralized merit pay system. It is harder for employees, and their supervisors, to manipulate such a system, since the only instrument left for this purpose is performance ratings. And, given the lack of fine distinctions in the rating system, even influence costs through such manipulation is limited.[12]

[12] This firm was in a relatively stable industry during this period. A simple incentive contract may have been adequate, given the costs of more individualistic schemes mentioned above, in such simple times. On the other hand, it seems easier to fine-tune incentives when the context is stable and simple.

8. Conclusions

This paper has provided a detailed look at the structure of incentives, both short- and long-term, for middle managers in a typical corporate hierarchy. Because such opportunities are rare, and the merit pay, bonus, and performance rating system seems exactly like what is used in most such settings, it is worth briefly summarizing the findings:

- Raises and bonuses are tied to ratings, and more strongly so than expected based on other studies. The effect on the raise plus bonus from improving one's performance rating by one point is to increase the raise by about 4% and the expected bonus by 1–2% each year that the rating is improved.
- Bonuses generate incentives from variation in the likelihood of receiving one, but much less so by variation in the size of the bonus if earned.
- Dismissals, and especially demotions, do not seem to be important instruments of the incentive system. Demotions almost never occur, and do not tend to bring large penalties. Exit rates do not vary strongly with performance.
- Patterns of promotions are consistent with the conclusions of Baker, Gibbs, and Holmström that the firm uses them to sort employees based on measured ability. The likelihood of receiving promotion, and the reward from promotion, both vary significantly with proxies for the employee's estimated ability based on previous performance. The best proxy for ability is tenure in job level, which indicates that those passed over have lower ability.
- Promotions appear to generate important incentives. The reward from promotion, especially over the long term, can be substantial. In addition, the likelihood of promotion depends on current performance ratings.
- Performance declines for employees who are passed over for promotion, even after controlling for ability differences. This implies that their incentives are falling. There is no evidence that the firm adjusts the terms of the incentive scheme based on changes in expected ability, as suggested by theory. Therefore, the fall in performance seems to be due to lower promotion-based incentives.
- The lack of evidence of trading-off between within-job and promotion-based incentives indicates that the firm runs a simple incentive scheme, probably centrally administered. The lack of individual variation in the size of the bonus also supports this conclusion. This is not consistent with the sophisticated approach taken in most agency models. However, it is consistent with the notion that certain kinds of contracting costs, such as observability and influence costs, lead to the use of contracts that are simpler in practice than in theory.

These results and the related results in Baker, Gibbs, and Holmström (1994a,b) allow us to get a fairly good view of the wage growth process over managers' careers. Stylized descriptions of managerial incentives usually say that wages have very little variation within jobs, and that most wage increases come through large jumps on promotion. In this firm, wage growth does not come primarily from large jumps at promotion. Instead, it comes from smaller jumps, plus substantial growth within levels. This wage growth is stronger for those who continue to move up in the firm. In part, this is because such employees have higher ability and earn larger raises and bonuses. It is also partly a result of the decline in performance of those who do not earn promotion, which seems to come from declining incentives due to reduced promotion prospects. Thus, stayers tend to slow down their wage growth, which eventually peters out to roughly zero growth, while movers continue to enjoy yearly increases plus gains from promotions. The wedge between the growth rates of these two groups drives much of the large differences in absolute pay that we observe between levels. Thus, studies that focus on average pay differences between levels, rather than individual gains at promotion, can give a misleading view of the wage growth process that overstates the short-term effect of promotions.

Does this firm run a tournament? It is difficult to answer this question. First, defining who competes against whom is not easy. Second, with only eight job levels and seventeen major titles, there is insufficient sample size, in terms of jobs and promotion slots, to formally analyze many of the predictions of tournament theory as has been done using samples from many firms. However, it is surely the case that the usual stylized version of tournament theory is not supported. There is much wage variation among those in the same job, and much wage growth within jobs. There are not very large one-time prizes earned on most promotions. Moreover, while there are large differences between lifetime earnings of those who are promoted and those who are not, these differences come about from *sustained* higher performance after promotion, not just from the promotion itself. This is evident in the larger weight given to more recent performance ratings in determining future promotions. Thus, promotions result in the opportunity to earn higher rewards at higher levels, but only if effort is maintained. Indeed, it makes sense that the firm must continue to motivate the winners as well as the losers. Thus, if there are tournament-like aspects to compensation in this firm, they are difficult to uncover, as the existing theory is inadequate for the rich wage dynamics that are observed.

The incentive scheme used by this firm seems to be simple by the standards of economic models. Yet, there are good practical reasons that firms might use such systems. It is interesting to consider the effects of such a simple system on firm innovation and adaptation. If this form of incentive system does not adjust easily to differences in individual circumstances, it is less likely to adapt quickly to changes in the firm's environment. Moreover, such a system does appear to lead to decreased motivation for managers who stay in the same job. As long as

the firm continues to grow, this may not be much of an issue. However, in more difficult times promotion slots may diminish. If individual incentives do not adjust easily and quickly, then the problems that led to poor firm growth could be exacerbated, until problems are sufficiently severe to justify the costs of large-scale change. This is consistent with the frequent observation that many large U.S. firms have reacted only slowly to poor economic performance and slowing growth in recent years. What was once optimal when the environment evolved slowly may be a barrier to organizational change under different circumstances.

Appendix

This appendix briefly clarifies the model; for more discussion see Gibbs (1994). Recall that $p = 1 - F(z)$, so that $p' = f(z)$, the density function of output. Given normal distributions, this is

$$p' = \bar{\alpha}\phi(\cdot)/\sigma_q, \tag{A.1}$$

where ϕ is the standard normal density and σ_q is the standard deviation of output. Since $p = 1 - F$, z is the $(1 - \rho)$th percentile of the output distribution, where ρ is the equilibrium promotion rate implied by the standard and optimal effort. This will be the same percentile when translated to the standard normal distribution as in (A.1). The normal density has a maximum at the 50th percentile, and is symmetric. Therefore, p' is largest for a promotion rate of 1/2. For lower promotion rates, p' increases in the promotion rate, but for larger promotion rates p' decreases symmetrically; in fact, these general properties will be true for any symmetric unimodal distribution. These results are due to the static nature of the model, but they do give some guidance for empirical implementation. In the data used here, the promotion rate is always well below 1/2. Even individual promotion probabilities, calculated as fitted values using the logit models in Table 5, yield maximum promotion rates for the highest ability employees of about 0.55, with the vast majority well below 1/2. Therefore, for the empirical work, it seems reasonable to consider only the case in which p' is increasing in promotion probabilities, and thus in expected ability.

Now consider an ability distribution that is not normal. In particular, sequentially sorting out the high end by promotion is likely to leave a distribution of survivors that is skewed to the right. As long as the distribution is unimodal, the general argument still applies. The maximum value of p' will be at the mode of the distribution; this is the point at which p' starts to decline with ability and promotion probabilities. If the distribution is skewed to the right, then the mode should be at a percentile below 50%. Thus, the mode would correspond to a standard z resulting in a promotion rate greater than 50% (at 1 minus the percentile of the mode). This makes it even more unlikely that we will

encounter cases empirically where higher promotion rates lead to lower p'. Thus, relaxing the model in this reasonable way will not change the implications that are tested.

References

Baker, George, 1990, Pay for performance for middle managers: Causes and consequences, Journal of Applied Corporate Finance 3, 193–204.

Baker, George, Michael Gibbs, and Bengt Holmström, 1993, Hierarchies and compensation: A case study, European Economic Review 37, 366–378.

Baker, George, Michael Gibbs, and Bengt Holmström, 1994a, The internal economics of the firm: Evidence from personnel data, Quarterly Journal of Economics, in press.

Baker, George, Michael Gibbs, and Bengt Holmström, 1994b, The wage policy of a firm, Quarterly Journal of Economics, in press.

Gibbons, Robert and Kevin J. Murphy, 1992, Optimal incentive contracting in the presence of career concerns: Theory and evidence, Journal of Political Economy 100, 468–505.

Gibbs, Michael, 1994, Promotions and incentives, Manuscript (University of Chicago, Chicago, IL).

Holmström, Bengt, 1982, Managerial incentive problems – A dynamic perspective, in: Essays in economics and management in honor of Lars Wahlbeck, Revised March 1983 (Swenska Handelshögkolan, Helsinki).

Holmström, Bengt and Paul Milgrom, 1987, Aggregation and linearity in the provision of intertemporal incentives, Econometrica 55, 303–328.

Lazear, Edward, 1992, The job as a concept, in: William Bruns, ed., Performance measurement, evaluation, and incentives (Harvard Business School Press, Boston, MA).

Lazear, Edward and Sherwin Rosen, 1981, Rank order tournaments as optimum labor contracts, Journal of Political Economy 89, 841–864.

Malcomson, James, 1984, Work incentives, hierarchy, and internal labor markets, Journal of Political Economy 92, 486–507.

Main, Brian, Charles O'Reilly, and James Wade, 1993, Top executive pay: Tournament or teamwork?, Journal of Labor Economics 11, 606–628.

Medoff, James and Katherine Abraham, 1980, Experience, performance, and earnings, Quarterly Journal of Economics 95, 703–736.

Medoff, James and Katherine Abraham, 1981, Are those paid more really more productive? The case of experience, Journal of Human Resources 16, 186–216.

Milgrom, Paul and John Roberts, 1988, An economic approach to influence activities in organizations, American Journal of Sociology 94, S154–S179.

Murphy, Kevin J., 1991, Merck & Co., Inc. (A), (B), & (C) (Harvard Business School Press, Boston, MA).

Peter, Laurence and Raymond Hull, 1969, The Peter principle (W. Morrow, New York, NY).

Rosen, Sherwin, 1982, Authority, control, and the distribution of earnings, Bell Journal of Economics 13, 311–323.

Rosen, Sherwin, 1986, Prizes and incentives in elimination tournaments, American Economic Review 76, 701–715.