CHAPTER 1

Incentives and careers in organizations

Robert Gibbons

An outsider might be surprised to learn that modern labor economics has little to say about activities inside firms. After all, is not work (i.e., what workers do once they go through a firm's doors) one of the field's most natural areas of inquiry?

Let's take stock. Several research areas in labor economics end precisely when an employment relationship begins: unemployment duration and labor-force participation are examples, and even labor demand typically focuses on how many workers should be hired rather than on what the firm should then do with them. Other research areas in labor economics reduce the employment relationship to a wage, or at most a wage profile: on-the-job search, labor supply, and human-capital models of earnings, for example. Even research on the return to seniority more often focuses on econometric issues than on what actually happens during an employment relationship; similarly, research on training more often focuses on pre-employment government-sponsored programs than on skill development in firms. Simply put, modern labor economics contains little work on work.

The situation may be changing. In this chapter I describe theory and evidence on two aspects of some employment relationships: incentive pay and careers in organizations.[1] Most of the theory I describe is recent, emphasizing games and contracts rather than the workhorse theories of labor economics in the 1970s and 1980s, human capital and search. Much of the evidence is also new, at least in the sense of not having been part of the published discourse in labor economics over the last few decades. This same evidence is also old, however, both in the sense of sometimes referring to events long past (sharecropping in 1910 or a machine shop in Chicago around 1950, for example) and in the sense of sometimes being fairly well-known outside labor economics.

Because there is not much empirical work on employment relationships in labor economics, I draw on other fields – including accounting, human resource management, industrial relations, and organizational sociology – whenever my exposure allows. Unfortunately, data on employment relationships often must be collected virtually by hand. Doing the hard work of data collection and utilizing the microeconometric expertise that has become the hallmark of labor economics are crucial next steps for this emerging literature. In the meantime, I limit discussion of theory to classes of models that seem likely to deliver empirical implications (or, better still, have already done so).

1 INCENTIVE PAY

There are many senses in which pay may be linked to performance. Perhaps the simplest case is where workers' productivities differ and wages equal marginal products. More often, however, the phrase "pay for performance" connotes the provision of incentives. In this section I discuss the dominant model of incentive contracting, the principal-agent model.[2]

Several of the main issues can be illustrated quite simply in the context of sharecropping. Three standard sharecropping contracts are: wage labor, which imposes no risk on the agent; crop sharing, which shares risk between the principal and the agent; and fixed-payment land rental, which leaves the agent with all the crop risk. The classic agency model, which emphasizes the tradeoff between incentives and insurance, implies that where there is greater crop risk there should also be more risk sharing – more used of fixed wages and crop sharing rather than land rental. Higgs (1973) presents evidence consistent with this prediction: for both cotton and corn, and for two empirical measures of risk, a cross-sectional analysis of the southern states of the US for 1910 finds more risk sharing in states with greater crop risk. But Alston and Higgs (1982) document that Higgs's comforting finding obscures both (1) enormous variation within each of these main classes of contracts and (2) significant variation across the three classes of contracts even after controlling for risk.

Both the organization and the spirit of this section parallel this research on sharecropping. Parallel to Higgs's paper, I begin by summarizing the theory and evidence on the classic agency model. Parallel to Alston and Higgs's paper, I then explore five new issues, in the hope of accounting for some of the enormous richness in incentive contracting that the classic agency model simply chalks up as unexplained variation. I conclude that risk is a significant issue in incentive contracting, but that the principal-agent literature's initial obsession with its consequences distracted us from a host of equally important issues.

1.1 The much-studied tradeoff between incentives and insurance

The classic model in agency theory involves an agent who takes an action a to produce output of value y. The principal owns the output but contracts to share it with the agent by paying a wage contingent on output, $w(y)$. There is noise in the production function, so the agent's output is uncertain. Furthermore, the agent is risk averse. Paying a constant wage, independent of y, would provide the agent with full insurance but no incentive; selling the agent the firm for a fee of F (or, equivalently, paying the agent $w(y) = y - F$) would provide the agent with full incentives but no insurance.

An intuitive closed-form solution can be derived in the linear-normal-exponential case. The production function is linear, $y = a + \varepsilon$, where ε is a normally distributed noise term with zero mean and variance σ^2. The incentive contract is linear, $w(y) = s + by$, where the intercept s is the salary and the slope b is the bonus rate. The agent's utility function is $U(x) = -e^{-rx}$, where $r > 0$ is the agent's coefficient of absolute risk aversion and $x = w - c(a)$ is the agent's net payoff – the realized wage minus the convex disutility of action $c(a)$. The principal is risk neutral and so seeks to maximize the expected value of profit, $y - w$.

Given a contract $w(y) = s + by$, the agent's problem is to choose an action to maximize the expected utility

$$\int_\varepsilon e^{-r[s+b(a+\varepsilon)-c(a)]}\phi(\varepsilon)d\varepsilon = -e^{-r[s+ba-c(a)]}\int_\varepsilon e^{-rb\varepsilon}\phi(\varepsilon)d\varepsilon,$$

where $\phi(\varepsilon)$ denotes the normal density function. The agent's optimal action, denoted $a^*(b)$, solves $c'(a) = b$. The agent's maximized expected utility is therefore

$$\begin{aligned} &-e^{-r\{s+ba^*(b)-c[a^*(b)]\}}\int_\varepsilon e^{-rb\varepsilon}\phi(\varepsilon)d\varepsilon \\ &= -e^{-r\{s+ba^*(b)-c[a^*(b)]-(\frac{1}{2})rb^2\sigma^2\}}, \end{aligned}$$

so the agent's equivalent is

$$CE(s,b) = s + ba^*(b) - c[a^*(b)] - \tfrac{1}{2}rb^2\sigma^2.$$

That is, the agent's certainty equivalent from the contract $w(y) = s + by$ is the expected wage minus the cost of effort minus the cost of bearing risk. The principal's expected profit is

$$E^\pi(s,b) = (1 - b)a^*(b) - s,$$

so the total surplus (i.e., the sum of the principal's expected profit and the agent's certainty equivalent) depends on b but not on s

$$CE(s,b) + E\Pi(s,b) = a^*(b) - c[a^*(b)] - \tfrac{1}{2}rb^2\sigma^2 = TS(b).$$

We can now determine the efficient contract slope, denoted b^*: it is the slope that maximizes the total surplus $TS(b)$. If the parties agreed to a contract with some other slope then both parties could be made better off by switching to a contract with slope b^* and choosing an appropriate value of s to distribute the increased total surplus. The first-order condition for b^* is $a^{*\prime} - c'a^{*\prime} - rb\sigma^2 = 0$. Because $c'[a^*(b)] = b$, we have $a^{*\prime} = 1/c''$ and hence

$$b^* = \frac{1}{1 + r\sigma^2 c''}.$$

This result makes sense. Since r, σ^2, and c'' are positive, b^* is between zero (full insurance) and one (full incentives). Furthermore, b^* is smaller if the agent is more risk averse (r is higher) or there is more uncertainty in production (σ^2 is higher) or marginal disutility increases more quickly (c'' is higher).

This solution to the classic model is tidy but flawed: Mirrlees (1974) showed that the best linear contract, $w = s + b^*y$, is inferior to various non-linear contracts. In particular, a step-function contract (where the agent earns w_H if $y \geq y_0$ but $w_L < w_H$ if $y < y_0$) can perform very well, approaching the twin goals of full incentives and full insurance in the limit (as y_0 and w_L decrease in appropriate fashion, so that the agent almost surely receives w_H and yet has incentives from fear of w_L). Mirrlees's result prompted a decade of research on how the optimal contract depends on the details of the utility function and the conditional distribution of output given the agent's action. In brief, this work showed that the optimal contract in the classic agency model is extremely sensitive to these details. In particular, the optimal contract is linear only under very special assumptions about the utility function and the conditional distribution of output.

Holmstrom and Milgrom (1987) reinterpreted the classic agency model so as to rescue linear contracts. Rather than a single action (a) that influences a single outcome (y), Holmstrom and Milgrom envision a sequence of actions (say, one per day, over the course of a year) influencing a corresponding sequence of outcomes. There are no connections across days (i.e., the action a_t on day t affects that day's outcome, y_t, but has no influence on any other day's outcome) and all past outcomes are observed before the next day's action is chosen. The output y from the classic model is interpreted as the aggregate output for the year in the sequential-action model: $y = \Sigma y_t$.

Suppose that each day's outcome takes one of two values – say L or H. Then a *one-day* incentive contract is simply a pair of wages: w_H is paid if the outcome is H; w_L if L. Suppose that the agent labors under the same one-day contract for all the days of the year. If there are T days in the year and the agent produces H on N of these days then the aggregate output for the year is $y = TL + N(H - L)$ and the aggregate wage for the year is $w = Tw_L + N(w_H - w_L)$. Thus, $N = (y - TL)/(H - L)$ and

$$w = \frac{T(Hw_L - Lw_H)}{H - L} + \frac{w_H - w_L}{H - L} y = s + by.$$

That is, if the agent labors under the same one-day contract throughout the year then the aggregate wage is a linear function of the aggregate output. Given several other assumptions, Holmstrom and Milgrom show not only that it is optimal for the agent to labor under a constant one-day contract but also that the optimal slope in the aggregate representation of this contract (i.e., $w = s + by$) is b^*, just as in the classic agency model.

In my view, the main contribution of this Holmstrom–Milgrom model is not that it justifies linear contracts (by imposing quite strong assumptions), but rather that it alerts us to gaming as a natural consequence of non-linearity. For example, a step-function contract of the kind studied by Mirrlees (in the classic one-action model) induces no effort once the agent's aggregate output to date passes the hurdle y_0 (in the daily-action model). More generally, if the incentive contract for the year is a non-linear function of year-end aggregate output then the worker's incentives change from day to day, depending on the aggregate output to date. A growing body of evidence is consistent with this prediction: see Healy (1985) on bonus plans with ceilings and floors, Asch (1990) and Oyer (1995) on bonuses tied to quotas, Chevalier and Ellison (1995) on the effects of even modest convexities in smooth pay plans, and Ehrenberg and Bognanno (1990) on performance across rounds in professional golf tournaments.[3]

There is other evidence more closely related to the classic agency model. One basic question is "Does pay vary with performance?" For example, by the early 1980s, the received wisdom was that the compensation of chief executive officers (CEOs) in large US firms was closely related to the firm's size but unrelated (or even negatively related!) to its stock-market performance. Murphy (1985) noted, however, that if big firms pay higher salaries but small firms have superior stock-market performance (the "small-firm effect" from finance) then a cross-section regression of cash compensation (salary plus bonus) on stock-market performance will be biased downwards, unless there are adequate controls for firm heterogeneity. Murphy found that including fixed effects in a panel-data model produces a strong statistical relationship between CEO pay and stock-market performance.[4]

A second basic question is "Do incentives matter?" In brief, the answer is "Yes." For example, the evidence summarized above on the effects of non-linear incentive plans motivates this conclusion. Others have studied the proposition that steeper slopes create stronger incentives: $a^*(b)$ increases with b. Lazear (1996), for example, finds that the output of workers installing automobile windshields increased after a switch from hourly wages to piece rates.[5] Abowd (1990) and Kahn and Sherer (1990) estimate the sensitivity of managerial pay to current performance and then estimate the effect of this sensitivity on subsequent performance.[6] The results are generally consistent with the theory but are somewhat noisy, in keeping with having to estimate rather than observe the relation between pay and performance. Gaynor and Gertler (1995) use data on medical partnerships, where the sharing rule is included in the data but was chosen by the partners. Their instrumental-variable estimates again are consistent with the simple proposition that incentives matter. Finally, there is evidence that investors believe that incentives matter. Brickley, Bhagat, and Lease (1985) find that there is a significant increase in a firm's stock price (net of any movement in the market as a whole) when the firm announces a stock-based compensation plan; Tehranian and Waegelein (1985) present analogous evidence for announcements of accounting-based bonus contracts.

There is also evidence related to the main idea behind the classic agency model – the tradeoff between incentives and insurance. For example, there is evidence that the slope falls as risk or risk aversion increases. As noted earlier, Higgs (1973) presents evidence from sharecropping in 1910 that the slope falls as risk increases; Garen (1994) offers similar evidence for CEOs of large US firms. Gaynor and Gertler (1995) find that the slope of the sharing rule in medical partnership falls as the partners' risk aversion increases.

The tradeoff between insurance and incentives produces further predictions in a richer model with multiple performance measures, as follows. Suppose there is a second performance measure, $z = a + \mu$, where μ is a normally distributed noise term possibly correlated with ε. (Theoretical and empirical work in accounting often interprets y as the change in the firm's stock-market value and z as the firm's accounting earnings, but many other interpretations are possible.) Consider the contract $w = s + by + dz$. Holmstrom (1979) shows that the optimal contract uses both performance measures (i.e., $b^* \neq 0$ and $d^* \neq 0$) unless one is a sufficient statistic for the other. That is, $d^* = 0$ only if z contains no additional information about the agent's action beyond what is contained in y (i.e., $z = y + \phi$, or $\mu = \varepsilon + \phi$, where ϕ is independent of ε); likewise, $b^* = 0$ only if y contains no additional information beyond z (i.e., $y = z + \phi$, or $\varepsilon = \mu + \phi$, where ϕ is independent of μ). In Holmstrom's model, therefore, performance measures are simply signals about the agent's action, and a signal is not useful if it conveys no incremental information.

Some performance measures come from outside the firm, such as from other firms in the same industry. Consider n firms, each subject to a common shock (θ) and an idiosyncratic shock (ε_i). Suppose $y_i = a_i + \theta + \varepsilon_i$, where θ and $(\varepsilon_1, \ldots, \varepsilon_n)$ are independent normal noise terms. Let z_i denote the average of the $n - 1$ other firms' outputs (y_j). Then the pure own-performance contract $w_i = s + by_i$ subjects the agent to two noise terms, θ and ε_i, whereas the pure relative-performance contract $w_i = s + b(y_i - z_i)$ eliminates θ but subjects the agent to ε_i and to the average of the $n - 1$ other idiosyncratic error terms (ε_j). Holmstrom (1982a) shows that the efficient contract is $w_i = s + by_i - dz_i$, where $b^* > d^* > 0$. That is, the efficient contract reflects a tradeoff between eliminating the risk from θ (through the pure relative-performance contract) and avoiding the risk from the average of the $n - 1$ other error terms (through a pure own-performance contract). If the variance of θ is small then it is not worth introducing the risk from the $n - 1$ other error terms so d^* is close to zero; if the variance of θ is large then it is important that the contract filter out θ, even at the cost of introducing risk from the other error terms, so d^* is close to b^*.

Antle and Smith (1986) look for evidence of relative performance evaluation in CEO pay. Using data on 39 firms in three two-digit industries, they find weak support for the theory, even though they carefully compute the correlation in "output" for each pair of firms. Gibbons and Murphy (1990) use a less-sophisticated approach but a much larger dataset, including data on pay and performance from 1,000 firms, with performance comparisons computed from data on 11,000 firms. Gibbons and Murphy find stronger support for the theory: CEO pay depends on the firm's stock-market performance relative to the market as a whole and (additionally) on the firm's stock-market performance relative to its one-digit industry. Janakiraman, Lambert, and Larcker (1992) estimate separate regressions for each of 554 firms (as opposed to the pooled regression in Gibbons and Murphy). The mean of the firm-specific estimates in Janakiraman, Lambert, and Larcker is similar to the pooled coefficient in Gibbons and Murphy.[7]

In sum, there is a large body of theory and evidence related to the classic agency model. The theory has developed several insights, such as the role of linear contracts in deterring gaming and the interpretation of performance measures as signals of the agent's action. The evidence is broadly consistent with both the basic theory and its extension to multiple performance measures. But the literature does not explain (or even hint at) why paying for performance is so problematic for many firms.

1.2 Complications in real incentive contracts

The main idea behind the classic agency model is that there is a tradeoff between incentives and insurance, but the most striking single fact about real attempts to tie pay to performance is that it is a tricky business. The following examples are all too typical:

> At the H.J. Heinz Company, division managers received bonuses only if earnings increased from the prior year. The managers delivered consistent earnings growth by manipulating the timing of shipments to customers and by prepaying for services not yet received, both at some cost to the firm (Post and Goodpaster 1981). At Bausch & Lomb, the hurdle for a bonus was higher, often entailing double-digit earnings growth. Again, managers met their targets in ways that were not obviously in the best long-run interest of the firm (e.g., over half a million pairs of "sold" sunglasses were discovered in a warehouse in Hong Kong; Maremont 1995). At Dun & Bradstreet, salespeople earned no commission unless the customer bought a larger subscription to the firm's credit-report services than in the previous year. In 1989, the company faced millions of dollars in lawsuits following charges that its salespeople deceived customers into buying larger subscriptions by fraudulently overstating their historical usage (Roberts 1989). In 1992, Sears abolished the commission plan in its auto-repair shops, which paid mechanics based on the profits from repairs authorized by customers. Mechanics misled customers into authorizing unnecessary repairs, leading California officials to prepare to close Sears' auto-repair business statewide. (Patterson (1992))

In brief, "business history is littered with firms that got what they paid for" (Baker, Gibbons, and Murphy (1994, p. 1125)).[8]

I find it hard to relate the classic agency model to this evidence (and the larger body of evidence it represents). First, much of the evidence concerns non-linear contracts, whereas the classic model began with (and has recently returned to) linear contracts. Second, I see no necessary role for risk aversion in this evidence, whereas in the classic model the only reason to limit incentives is to provide insurance. Third, and most important, the performance measures used in these real incentive contracts differ from those envisioned in the classic model, as I describe below. In this subsection, therefore, I abandon the classic model, turning instead to five other issues in incentive contracting – performance measurement, implicit contracts, labor mobility, the ratchet effect, and career concerns. To emphasize that these five issues are departures from the classic model, I assume throughout this subsection that the agent is risk neutral. For lack of space, I give only brief attention to labor mobility, the ratchet effect, and career concerns; I focus on performance measurement and implicit contracts because I believe that together they offer an important complement to the classic agency model.

1 *Performance measurement* I lack the information to assess whether the incentive plans at Heinz, Bausch & Lomb, Dun & Bradstreet, and Sears were mistakes (as opposed to best responses to tough environments), but some of my colleagues in organizational behavior (OB) are less reticent. Kerr's (1975) classic title conveys his field's judgment: "On the folly of rewarding A, while hoping for B." Kerr's paper is so well known in OB that it has earned a place in the canonical MBA core course on organizations; in economics, in contrast, until recently there was no model that could even express Kerr's idea, not to mention evaluate or extend it.

Fortunately, Holmstrom and Milgrom (1991) and Baker (1992) now offer simple models of such distortionary performance measurement. Both emphasize the distinction between the agent's total contribution to firm value (henceforth denoted y) and the agent's measured performance (henceforth p). Even well-informed insiders may find it extremely difficult to assess an agent's total contribution to firm value, because total contribution includes aspects of performance such as the effects of the agent's actions on co-workers and the long-run effects of the agent's current actions. Furthermore, to enforce a contract contingent on the agent's total contribution, the parties would have to specify *ex ante* how y is to be measured *ex post* (so that a court would know what to measure if called in to enforce the contract).

These difficulties are assumed away in the classic agency model: the agent's total contribution is called "output," as though it could simply be counted at the end of the contract period, and contracts such as $w = s + by$ are assumed to be simple to write and enforce. The classic model may capture some employment relationships, where there are few interactions among co-workers and few long-run effects of current actions. Lazear's (1996) study of piece rates paid to workers installing auto windshields may be one example; more generally, Brown (1990) finds that piece rates are more likely to be used in jobs with a narrow set of routines than in jobs with a variety of duties.

In a vast array of jobs, however, the Holmstrom–Milgrom and Baker distinction between total contribution and measured performance seems crucially important. For example, Eccles and Crane (1988) describe how investment banks deliver a substantial fraction of a trader's compensation through a subjectively determined bonus, even though many objective aspects of the individual's performance are easily measured on a daily basis.

Similarly, Burtis and Gabarro (1995) offer a fictitious but persuasive account of the difficulties of performance evaluation in a law firm: nine objective measures paint a narrow and distorted picture (even when combined with four subjective assessments). Evaluating the performance of almost any manager or professional worker seems likely to involve similar issues – for example, see Greene and Schlesinger (1992) on incentive pay in a cable television firm. Finally, the recent enthusiasm for empowerment, participation, and self-managed teams suggests that difficulties in performance evaluation may become increasingly important for non-managerial workers as well.

Baker models the worker's contribution to firm value as $y = \theta a + \varepsilon$, whereas measured performance is $p = \mu a + v$. As in the classic model, ε and v are noise terms (independent of θ, μ, and each other), but θ and μ are features of the environment that are privately observed by the worker before choosing an action. As motivated above, Baker assumes that a contract contingent on y cannot be enforced, so the firm is reduced to contracting on p, through the linear contract $w = s + bp$. Because the agent's utility, $w - c(a)$, depends on p, the agent will be induced to take large actions when dp/da (i.e., μ) is large; because the firm's profit, $y - w$, depends importantly on y, the firm will value large actions when dy/da (i.e., θ) is large. Hence Baker's central insight: a good performance measure induces the agent to do the right thing at the right time (i.e., to work hard when doing so is valuable to the firm), so the quality of a performance measure depends on the correlation between dp/da and dy/da. Thus, whereas the classic model views a performance measure as a signal of the agent's action, Baker focuses on the value of the actions that a contract based on the performance measure will induce.

When measured performance omits important dimensions of total contribution, firms understand that they will "get what they pay for," and so may choose weak incentives in preference to strong but frequently dysfunctional incentives. In Kerr's terms, the Holmstrom–Milgrom and Baker models explore environments in which it might be necessary to reward A while hoping for B, but these models caution against over rewarding A.[9]

2 *Implicit contracts* A worker's total contribution to firm value may be impossible for a court to measure using a method specified *ex ante*, but well-informed insiders may nonetheless agree *ex post* on a particular worker's contribution (or at least on an estimate of this contribution). The great advantage of such *ex post* settling up is that the parties can take into account events that occurred during the contract period that were not foreseen (or were not articulated) *ex ante*. Thus, it might be possible for the