

Contracting in Firms

Canice Prendergast¹

¹University of Chicago & NBER.

In many realms of economic life, the actions of individuals affect the welfare of others. Nowhere is this more relevant than in firms, where employees act on behalf of owners or shareholders to provide services for customers and clients. This separation of the interests of employees from those whose actions they benefit has generated a large literature on incentive contracting, where the overarching objective is the alignment of such interests. The early literature on agency theory, described in the first edition of this volume by Lazear (1987), conceptually mimics that on externalities - the other area of economics that deals with welfare consequences of actions on others - by showing a variety of ways in which the compensation of agents can be constructed to internalize the effects on one's actions on others. There are two ways of doing this. First, one could simply tell employees what to do and to penalize them if they fail to do so. In the literature, this is referred to as input monitoring. While this can sometimes help, it is often hard to monitor either what workers do, or the intensity with which they do so - a salesman on the road would be a good example. Similarly, while overseers can sometimes identify what it is that agents are doing, they may not know what they *should* be doing - a board of directors monitoring a CEO would be apposite here. Accordingly, the second solution to misaligned incentives is to design compensation plans such that the agent's pay depends on her contribution - "output" - so that the concerns of other parties are internalized.

A simple model can illustrate this point, and is useful to describe other complications that can arise. The agent is assumed to take some action ("effort") $e \geq 0$, which is unobserved by the principal. She is averse to taking effort. Consider a simple parameterization of the agent's utility function, where the agent cares about wages w and effort; assume that the agent has exponential utility $V = -\exp[-r(w - C(e))]$, where w is the worker's wage, $r \geq 0$ is the constant rate of absolute risk aversion, the worker's cost of supplying effort is $C(e) = \frac{ce^2}{2}$, and his reservation utility is U^* . To focus attention on the role of output contracting, assume that the principal cannot observe effort e (so monitoring of inputs is not possible), but instead only observes a signal on effort $y = e + \epsilon$,

where $\varepsilon \sim \mathcal{N}(0, \sigma^2)$, with σ^2 representing measurement error. Assume also that the principal chooses to reward the agent in a linear fashion on output - a piece rate: $w = \beta_0 + \beta_y y$. (There is a large literature on the optimal shape of compensation contracts - see Prendergast, 1999, and Gibbons, 1996, for an overview.) Then there is a simple solution to attaining efficient effort: choose the contract to internalize the benefit to others by setting $\beta_y = 1$. In words, efficiency arises when the agent is residual claimant on the benefit of others.

This solution, providing a simple prescription for how compensation contracts should be designed, is both simple and intuitive. And empirically false. There are, of course, some occupations where one can find evidence of such “high powered incentives”, where agents are essentially residual claimants on output. Indeed, the literature on agency theory is replete with references to such occupations - taxi cab drivers, franchisees, sharecropper farmers, and the self-employed. Yet these are exceptions; instead, “low powered incentives” in firms are more the norm. (See Prendergast, 1999, for details.) Consequently, one of the quandries of the literature has become why so few workers seem to have contracts where their pay is strongly linked to their performance, and much of the subsequent literature to that outlined in the last volume has been to identify relevant constraints on incentive contracting.

The earliest candidate to explain why high powered incentives are rare is that high powered contracts impose *risk* on workers (Holmstrom, 1979). Consider the contract that induces efficient effort above: $\beta_y = 1$. The objective of the firm is to maximize profits subject to the worker’s willingness to take the position. This implies that the fixed component, β_0 , is changed to guarantee that agents earn their reservation utility, so the principal’s objective becomes a surplus maximization exercise. When the worker is risk neutral, the fixed component is reduced sufficiently such that the total compensation cost is $U^* + \frac{c}{2}$. In words, the only cost that the employer incurs in addition to U^* is the effort cost. This is not true when the worker is risk averse. In the context

of the preferences V above, compensation costs increase when incentive contracts are used for two reasons - the cost of increased effort as above, but also a risk cost imposed on workers. Both costs are increasing in β_y . With exponential preferences and linear contracts, this tradeoff results in the optimal contract being $\beta_y^* = \frac{1}{1+rc\sigma_y^2}$. This approach to studying incentive contracting has become known as the “tradeoff of risk and incentives”, where firms tradeoff the benefits of great effort with higher compensation costs induced by a risk premium, such that the chosen level of effort becomes below the level that internalizes benefits to others. Only in the case where there is either no measurement error ($\sigma_y^2 = 0$) or risk neutrality ($r = 0$) does efficient effort arise.

At its most general, this costliness of exposing a worker to large degrees of risk (or its analog, liquidity constraints) surely explains some part of the absence of high powered incentives. In much the same way as financial assets with higher undiversifiable risk require higher expected returns, so also are risky jobs likely to demand higher compensation. Despite this, the empirical literature on how compensation contracts tradeoff such risk issues against higher effort has shown little evidence in its favor. There are two principal empirical implications of the theory. First, riskier environments should have lower incentives - β_y^* declines with σ_y^2 . There have been many studies on the relationship between risk and the strength of incentives in a variety of occupations. If anything, this literature suggests that the relationship between risk and the provision of incentives is positive rather than the negative relationship posited by this theory. See Prendergast, 2002, for details and an explanation as to why this may be. Second, the tradeoff of risk and incentives implies that compensation should not depend on measures that workers cannot control. Again, this has found little support in the data. For example, Bertrand and Mullainathan, 2001, have examined executive contracts in the US, and found little evidence that contracts reward executives any less for measures that they cannot control (say where an oil company’s profits change simply because the price of crude changed) than for those that they can (such as a merger). More evidence on this failure to filter out

uncontrollable factors concerns the infrequency of relative performance evaluation. Consider two salesforce workers (or executives) who carry out a similar job. If demand for the products that they sell varies for common reasons beyond their control, an efficient way of limiting risk exposure is to (at least partially) reward the workers on how well they do relative to each other. Yet empirically there is relatively little evidence of such benchmarking. For example, see Janakiraman et al., 1992.

A second limitation on incentive contracting arises when measures do not reflect the objectives of the principal. Workers often carry out a host of activities in their jobs, yet measures of performance may not reflect all these aspects. A good example of this would be measuring the performance of a teacher. While measures may be available on some component of what they do - such as test scores for a teacher - many important aspects may remain unmeasured. When contracts are designed on the subset of things that can be measured, there is a danger that they ignore the unmeasured aspects. For instance, there is evidence of teachers “teaching for the test” or cheating to achieve higher test scores (Jacob and Levitt, 2003). This phenomenon has become known as *multitasking* (Holmstrom and Milgrom, 1991), which becomes potentially important when there is no single measure which reflects the contribution of an agent. Accordingly, it is not surprising that a consistent empirical finding is that jobs which are described by firms as complex tend not to offer significant incentive pay (see Prendergast, 1999, for details).

Another limitation on the ability of firms to provide incentives to workers comes from team production. Measures of performance for most workers reflect not only what they do, but also the contributions of others. In itself, this does not change the calculus above in any conceptual sense, other than that the measurement error now includes the actions of others. As an example, assume that two agents (1 and 2) work on a team and that output measures the true contributions of both plus an error term - $y = e_1 + e_2 + \epsilon$. Efficient effort arises as before by setting $\beta_y = 1$ for each worker. However, there is now a potential problem of budget breaking, where marginal payments exceed

marginal output. In this example, when total output rises by \$1, compensation costs increase by \$2. In many firms - for instance partnerships - such budget breaking is not possible. If instead the principal can pay out no more than \$1 for every dollar extra on output, this naturally places an upper bound of $\beta_y = \frac{1}{2}$ on average for the agents. Hence, budget balancing places a natural limitation on firm incentives. This also leads to a free rider problem in teams, where maximum incentive compensation in an N member team mechanically declines as N increases. (This is known as the " $\frac{1}{N}$ problem".) There is also considerable empirical evidence (such as Gaynor and Pauly (1990)) on such free riding - mostly from legal and medical partnerships - illustrating how various measures of performance disimprove as the size of the team being rewarded gets larger.

Many measures of output are not denominated in dollar terms, but instead come in the form of evaluations by others. For instance, it would be difficult to measure the contribution of a social worker or a customer service representative without using feedback from supervisors or clients. Another limitation on contracts arises when such subjective measures can be corrupted by evaluators with vested interests. Two particular sources of such vested interests have been considered in the literature. First, information on performance often originates with clients as they are the only ones with first hand experience of the agent's efforts. For instance, compensation for many customer service representatives depends on client evaluations. When clients have relatively similar preferences to the principal - such as that the agent should be courteous and efficient - contracts based on evaluations can mimic the objective contracts above. Yet in other instances, the vested interest of clients can render incentives difficult to implement. A good example of this arises in occupations - such as police or immigration officers - whose objective is not necessarily to make their clients happy. Making pay depend on evaluations in these instances can be harmful as it gives agents incentives to keep clients happy when they should not, such as a police officer not arresting a suspect. In these cases, incentive contracting on evaluations typically needs to be

curtailed to avoid such incentives. See Prendergast, 2003, for details.

The second example of vested interests with subjective evaluations is where the principal has an incentive not to implement the (ex ante) efficient contract by renegeing on a promised payment to save costs. Thus, even though an agent exerts effort and performs well, the supervisor claims otherwise to keep costs down. This can arise either by outright lying or perhaps by manipulating whatever measures are available. A relevant example here is the movie industry, where actors are sometimes paid on the “net profits” of a film. As a result, there have been numerous court cases regarding firms using creative transfer pricing arrangements to reduce profits for very successful movies. See Carole Cheatham et al. (1996) for more details on this. Such incentives to renege are likely even worse when there are no objective measures of performance. Because the desire to renege is greater when discretionary incentive payments are higher, it follows that the only credible contracts often involve little incentives.²

It is incorrect to assume that the ability to manipulate measures of performance always mute incentives - sometimes it can result in incentive pay being inefficiently high. Consider again two occupations where agents are typically residual claimants - taxi drivers and sharecroppers. At first blush, it would seem odd that they have such extreme incentives. Aren't these as likely candidates for trading off risk against incentives as any? However, one characteristic of each of these occupations is that both have opportunities for hiding output, either by taking fares without using the meter (in the case of cab drivers) or selling crops privately (in the case of farmers). In both cases, the only outcome that makes this incentive irrelevant is to render them residual claimants, even if risk considerations would suggest otherwise.

Another issue that can constrain efficient incentives, yet has received almost no attention in

²Clive Bull (1987) considers a role for repeated interaction between the principal and agent as a means of reducing incentives to renege. While repeating the relationship can result in sufficient incentives for complete honesty by the principal, it remains the case that if the relationship's value is not sufficiently great, incentives must be muted to reduce incentives for cheating.

the empirical literature, is where agents hold *private information*. Take a specific instance - real estate agents. In Chicago, real estate contracts take a simple form - agents make 3% of the sales price of the house. Assume that my home is worth \$500. This linear contract not only offers only 3% on the relevant margin for improving the selling price of the house, but predominantly rewards the agent for selling the house for say \$450. Yet anyone could sell the house for \$450 and it seems highly inefficient to reward in this way. So why not renegotiate to something better? An example of such an improvement (modulo risk issues) would be to offer nothing on the first \$450, but to pay a piece rate of 30% on anything over \$450. In this way, the agent has more incentives on margin, yet breaks even relative to the original contract if the house sells for its original price.

One reason why such renegotiation does not arise is that the agent may privately know the true value of the home, while the owner believes it to be worth \$500 on average. Consider a homeowner who offers the new contract above to the agent. It is clear that the agent rejects the new contract if it is truly worth less than expected, and accepts it if worth more. But this implies that the agent earns *information rents* on average. As a result, on average the homeowner loses money from the renegotiation unless effort increases enough. This option available to the agent limits the ability of contracts to attain efficiency. Instead, in the usual monopoly fashion, the homeowner would offer a contract to tradeoff the efficiency gains of increased effort with inframarginal losses of the type described above, resulting in lower powered incentives.³

Much of the recent literature has been focused on how incentive contracts can cause adverse behavioral responses. Another possible mechanism for such responses is where *intrinsic motivation* can be crowded out by the use of incentive contracts. The premise of this literature has been that in many occupations, agents enjoy carrying out the activity or care about the outcomes of

³There is a large mechanism design literature on this topic that has largely been ignored by the empirical literature on incentives. See Laffont and Tirole, 1986, for example. Part of this is surely because of the empirical conundrum as to why mechanisms are so rare in reality.

their actions. As a result, they will exert effort beyond that which they can get away with even in the absence of incentive contracts. This, in itself, is not enough to limit incentive contracting. However, there is some psychological evidence that the extent to which agents enjoy their jobs becomes weaker when incentive contracting is used. In effect, they feel that they are only doing it “for the money” and hence lose interest. A commonly cited example is the willingness of people to donate blood, where the warm feeling from donating declines when payments are made. In some instances, this can imply that incentive contracting can *reduce* effort if these crowding out effects are strong enough. As a result, it can be optimal to provide no incentives even when effort is one dimensional. This area of research, whose empirical testing has largely been restricted to the laboratory, is still in its early stages and is likely to see much refining over the coming years. See Frey and Jegen, 2000, for a survey.

Another likely fruitful area of future research concerns non-monetary ways of motivating workers. This literature largely began as an exercise in how workers could be motivated to internalize the benefits of others, yet has almost exclusively become an exercise in how to motivate through monetary contracting. Yet it is clear that there are a myriad of means of motivating workers - sense of achievement, “doing good”, status, etc - that firms tap into. How such mechanisms operate, and the way in which they interact with monetary contracts, remain an unstudied topic of research, though see Besley and Ghatak, 2003, for some theoretical work on this issue.

It is worthwhile to note a caveat before concluding. The discussion above concerns the absence of *observed* incentive contracts. Yet workers often have unobserved carrots and sticks that can motivate them. For instance, many workers exert effort in the hope of attaining a promotion (Lazear and Rosen, 1981), or a better job offer (Holmstrom, 1982). Many of these mechanisms for inducing desired behavior are dynamic, where good performance today results in a greater likelihood of promotion, or better job offers in future. Such incentives are clearly important for

workers. However, it remains the case that explicit incentive payments remain limited even in those cases where the type of career concerns above are negligible. (For example, it is well known that promotion prospects become very limited for workers who remain in a job grade for a long period. Yet explicit incentives are no more common for those workers than for any other.) The interaction of unobserved (typically career) incentives with the more explicit set of piece rates and bonuses that have been considered above is surely of first order importance to firms, though it remains surprisingly unexplored in the literature. (See Baker et al., 1994, for an exception.)

To conclude, perhaps the most central foundation of modern economics is the idea that appropriate prices guide behavior in efficient ways. Despite this, one of the defining characteristics of the employment relationship in many firms is the absence of the kind of explicit prices where wages depend in a clear way on observed outcomes. The early incarnations of agency theory were concerned with returning prices in a way that could serve to fully internalize the effects of agents' actions on the welfare of their employers. Yet this initial optimism has now been tempered with a somewhat more nuanced view that shows tradeoffs that will ultimately help in defining more precisely the nature of labor market relationships.

REFERENCES

- Baker, George, Robert Gibbons and Kevin J. Murphy (1994)**, “Subjective Performance Measures in Optimal Incentive Contracts”, *Quarterly Journal of Economics*, 109, 1125-56.
- Bertrand, Marianne and Sendhil Mullainathan (2001)**, “Are CEOs Rewarded for Luck? The Ones without Principals Are”, *Quarterly Journal of Economics*, 116, 901-932.
- Besley, Tim, and Maitreesh Ghatak**, “Competition and Incentives with Motivated Agents,” 2003, mimeo, LSE.
- Bull, Clive (1987)**, “The Existence of Self-Enforcing Wage Contracts”, *Quarterly Journal of Economics*, 102, 147-59.
- Cheatham, Carole, Dorothy Davis and Leo Cheatham. (1996)**, “Hollywood Profits: Gone With the Wind?”, *The CPA Journal*, 32-34.
- Frey, Bruno and Reto Jegen (2000)**, “Motivation Crowding Theory: A Survey of Empirical Evidence”, in Zurich IEER Working Paper No. 26; CESifo Working Paper Series No. 245..
- Gaynor, Martin and M. Pauly (1990)**, “Compensation and Productive Efficiency in Partnerships” Evidence from Medical Group Practice”, *Journal of Political Economy*, 98, 544-74.
- Gibbons, Robert (1996)**, “Incentives and Careers in Organizations”, in *Advances in Economics and Econometrics: Theory and Applications*. David Kreps and Ken Wallis (eds.). Cambridge University Press, Cambridge, UK.
- Holmstrom, Bengt (1979)**, “Moral Hazard and Observability”, *Bell Journal of Economics*, 10, 74-91.
- Holmstrom, Bengt (1982)**, “Managerial Incentive Problems: A Dynamic Perspective”, mimeo, Kellogg School of Management.

- Holmstrom, Bengt and Paul Milgrom (1992)**, “Multi-Task Principal Agent Analyses: Linear Contracts, Asset Ownership and Job Design”, *Journal of Law, Economics, and Organization*, 7, 24-52.
- Janakiraman, Surya N., Richard A. Lambert, and David F. Larker, 1992**, “An empirical investigation of the relative performance evaluation hypothesis”, *Journal of Accounting Research*, 30, 53-69
- Laffont, J.-J. and J. Tirole (1986)** , ”Using Cost Observation to Regulate Firms,” *Journal of Political Economy*, 94:614-641
- Lazear, Edward (1987)**, “Incentive Contracting”, in *The New Palgrave Dictionary of Economics*, First Edition, Macmillan.
- Lazear, Edward and Sherwin Rosen (1981)**, “Rank Order Tournaments as Optimal Labor Contracts”, *Journal of Political Economy*, 89, 841-64.
- Brian A. Jacob and Steven D. Levitt (2003)**, “Rotten Apples: An Investigation Of The Prevalence And Predictors Of Teacher Cheating,” *The Quarterly Journal of Economics*, 118(3), 843-877.
- Prendergast, Canice (1999)**, “The Provision of Incentives in Firms,” *Journal of Economic Literature*, 1999, 37, 7-63.
- Prendergast, Canice, (2002)** “The Tenuous Trade-Off Between Risk and Incentives?”, *Journal of Political Economy*, 110, 1071-1102.
- Prendergast, Canice, (2003)** “The The Limits of Bureaucratic Efficiency”, *Journal of Political Economy*, 111, 929-59. .