

EFFICIENT INCENTIVE CONTRACTS

MARTIN L. WEITZMAN

A so-called "incentive contract" is a linear payment schedule, where the buyer pays a fixed fee plus some proportion of audited project cost. That remaining proportion of project cost borne by the seller is called the "sharing ratio." A higher sharing ratio creates more incentive to reduce costs. But it also makes the agent bear more cost uncertainty, requiring as compensation a greater fixed fee. The tradeoff between incentives and risk in determining the sharing ratio of an efficient contract is the central theme of the present paper. A formula is derived that shows how the optimal sharing ratio depends on such features as uncertainty, risk aversion, and the contractor's ability to control costs. Some numerical examples are calculated from the area of defense contracting.

SUMMARY

This paper analyzes the widely used "incentive contract"—a linear payment schedule where the buyer pays a fixed fee plus some proportion of project cost. The remaining proportion of project cost borne by the seller is usually called the "sharing ratio." A higher sharing ratio creates more incentive to reduce costs. But it also makes the contractor bear more risk, requiring a greater fixed fee as compensation. The basic contribution of the present paper is a simple formula showing explicitly how the optimal sharing ratio depends on such features as uncertainty, risk aversion, and the contractor's ability to control costs. The formula is applied to some numerical examples drawn from the area of defense contracting.

Designing an efficient contract is an example of the so-called "principal-agent problem." While the basic theoretical issues are fairly well understood,¹ results are at a rather high level of abstraction, somewhat removed from the realm of practical application. At the other extreme is a rich body of descriptive material about an already institutionalized linear contract² (the "incentive contract"). In this paper I want to strike a middle position. Because a linear payment schedule has a simple structure, explicit formulae for an efficient contract can be derived that show clearly the tradeoff between risk-sharing and incentives.³ This has immediate applications, since the "incentive contract" is used in practice.

1. See, e.g., Holmstrom [1979], Ross [1973], Shavell [1979].

2. See, e.g., Fox [1974], Moore [1967], Peck and Scherer [1962], Scherer [1964b], and Williamson [1967].

3. The interested reader may want to contrast my approach with the earlier work of Berhold [1971], Cummins [1977], and Scherer [1964a].

CONTRACT TYPES

Two polar contract types⁴ have been in widespread use for a long time.

At one extreme is the "cost plus" contract (Cost Plus Fixed Fee = CPFF). This contract type pays actual costs plus a fixed dollar fee that is usually determined as some percentage of a cost estimate. Once set, the fee is fixed. The buyer additionally compensates (within limits) all legally allowable costs incurred by the contractor in fulfilling the project. The CPFF contract has the significant drawback of providing no incentive for cost reduction, which results in a well-known tendency to cost overrun.

The opposite extreme is the "fixed price" contract (Firm Fixed Price = FFP). Here the contractor agrees to fulfill the project for a fixed dollar price, which, once negotiated, will not be readjusted to include actual cost experience. With every dollar of cost saved ending up a dollar of extra profits, a strong incentive is created to reduce project cost. The disadvantage of the FFP arrangement is that the firm, bearing all the risk, must be compensated by a fee representing on average a high nominal profit rate.

The "incentive contract" falls between the polar extremes of CPFF and FFP. Sometimes called Cost Plus Incentive Fee = CPIF, sometimes Fixed Price Incentive = FPI (depending on which of CPFF or FFP was its conceptual antecedent), an incentive contract essentially pays a fixed fee plus some fraction of project costs.⁵ No matter how varied the way it is represented or the names of its different parameters, the principal's payment under an incentive contract can always be written in the reduced form:

$$(1) \quad K + (1 - \lambda)X,$$

where X is the accounting cost of the project, K is the fixed fee, and λ is the agent's share of project costs (the principal's share is $1 - \lambda$).

Note that $\lambda = 0$ is cost plus and $\lambda = 1$ is fixed price. Thus, CPFF and FFP are special polar cases of expression (1).

An incentive contract offers the possibility of striking a balance between the positive incentive effect of a high sharing ratio and the negative risk effect. The purpose of this paper is to formulate the

4. For more detailed information see Cummins [1977], Fox [1974], Moore [1967], and Scherer [1964a, 1964b].

5. In practice there may exist ceilings or floors, which are ignored because they are rarely penetrated. See Scherer [1964a], p. 258.

problem of selecting an optimal sharing ratio and show clearly how the solution depends on various factors.

A MODEL OF THE CONTRACTOR FIRM

Suppose that a contract is let for some well-defined project. Consider a firm that has agreed to undertake the project. Let X be the accounting cost of the project.⁶ It is supposed that the firm engages in some variety of economic activities. Total net profits attributed to the firm's *other* activities, excluding this project, are denoted ψ .

Suppose that, through bilateral negotiation between principal and agent, or by competitive bidding, or even by some other process, a mutually acceptable combination of fixed fee K and sharing ratio λ is obtained. The buyer agrees to compensate the seller by the terms of the incentive contract (1).

The contractor's total net profit is

$$(2) \quad \pi = K + (1 - \lambda)X + \psi - X,$$

which can be rewritten as

$$(3) \quad \pi = K - \lambda X + \psi.$$

Thus, the sharing ratio λ is like a "price" per dollar of cost overrun to the contractor.

Generally speaking, there is a tradeoff between X and ψ . Project costs can be made lower if net profits in the rest of the firm are also lowered.

There is a variety of discretionary actions the firm can take to influence jointly X and ψ . Unmeasured, hidden, or imperfectly imputed resources (such as managerial attention or some kinds of overhead) can be shifted between the contracted project and other activities, or an externality between X and ψ may be involved. Perhaps by doing more R&D or learning under the contract, the firm can trade higher costs on this project for greater long-run profits generated by strategic advances that spill over into future economic activities. Other stories can also be told.

An essential ingredient of the environment I am trying to model is uncertainty. The terms of the contract are established in an atmosphere where project costs are not precisely known.

6. In practice, it is difficult to measure precisely (or even to define) project-related costs, and some degree of arbitrariness is inevitable. Certain costs can be attributed directly to the project. But others, like overhead costs, cannot easily be assigned to one activity or another, and some project costs, like managerial attention, are hard to measure, observe, or monitor.

Once set, K and λ cannot be renegotiated. But after the uncertainty resolves itself, the firm will adjust discretionary action under its control to maximize total profits, given K and λ .

Let θ represent a state of the world bearing on this project. For example, materials might be more or less expensive, construction conditions may vary, research and development outcomes could differ, etc.

In state of the world θ , let

$$(4) \quad \psi = F_{\theta}(X)$$

be a function representing the maximum ψ attainable for a given X .

Under state of the world θ , in profit-maximizing equilibrium the firm will choose $X_{\theta}(\lambda)$ and $\psi_{\theta}(\lambda)$ to maximize

$$F_{\theta}(X) - \lambda X,$$

yielding the first-order condition,

$$(5) \quad F'_{\theta} = \lambda$$

or

$$(6) \quad -\lambda X'_{\theta}(\lambda) + \psi'_{\theta}(\lambda) = 0.$$

When $\lambda < 1$, condition (5) indicates inefficiency, because the marginal rate of transformation between project costs and revenue earned elsewhere is not one to one. Instead, a dollar saved on the project is only equivalent to the fraction λ of an extra dollar earned on other activities of the firm, due to partial reimbursement of project costs at $1 - \lambda$ to the dollar. The lower λ is, the less the incentive is to cut costs on the contracting activity at the expense of profits in the rest of the firm. On the other hand, when λ is higher, the agent must bear a greater share of cost risk, and K must be made sufficiently larger to induce the firm to accept the contract.

EFFICIENT INCENTIVE CONTRACTS

Parameters K and λ are fixed when the state of the world is uncertain. That level of expected utility that the firm obtains is a function of its bargaining strength and the kind of negotiating process that takes place. Whether determined by bilateral confrontation or competitive bidding, it is in the common interest of buyer and seller to negotiate parameter values that are Pareto efficient with respect to the given form of the contract. Even if forces making for Pareto op-

tinality are weak in practice, characterizing an efficient incentive contract is still an important normative issue. This is especially true because, as long as the optimal sharing ratio is relatively invariant to expected utility levels, the buyer can set it at the efficient value and limit negotiations to determining the fixed fee.

Assuming that the theoretical difficulties associated with representing collective choices are insignificant, let the agent's utility function be $U(\cdot)$ and the principal's utility function be $V(\cdot)$. For most applications V is more linear than U , because the government (or any large buyer) is likely to be less risk-averse than the contractor. In fact, the government is frequently assumed to be risk-neutral as a first approximation.

Let \bar{U} be a given level of the contractor's expected utility (it may be varied parametrically). The expected value operator over θ is denoted E .

An efficient incentive contract (λ^*, K^*) solves the problem,

$$(7) \quad \max_{\lambda, K} EV(-K - (1 - \lambda)X_\theta(\lambda)),$$

subject to

$$(8) \quad EU(K - \lambda X_\theta(\lambda) + \psi_\theta(\lambda)) = \bar{U}$$

$$(9) \quad F'_\theta = \lambda.$$

Formally speaking, (7)–(9) has a generic relationship to problems in the theory of optimal income taxation and competitive insurance contracts.⁷ However, besides being worthy of analysis on its own merits as an important economic issue, the efficient incentive contract has a special structure that allows a distinctive characterization.

THE OPTIMAL SHARING RATIO

Our main interest is in characterizing the sharing ratio λ^* , which is the solution of (7)–(9). Once λ^* is determined, the optimal fixed fee K^* can be calculated as a residual from equation (8).

With $X_\theta(\lambda)$ and $\psi_\theta(\lambda)$ defined as solutions of (9), condition (8) implicitly determines K as a function of λ . Writing $K(\lambda)$ as that value of K which satisfies (8) for a given λ , I can show that

$$(10) \quad EU(K(\lambda) - \lambda X_\theta(\lambda) + \psi_\theta(\lambda)) = \bar{U}.$$

7. See, for example, Mirrlees [1971], [1976], Sheshinski [1971], and Spence and Zeckhauser [1971].

The above equality must hold for all values of λ . Differentiating the left-hand side of (10) with respect to λ and setting the derivative equal to zero yields the condition,

$$(11) \quad E(K'(\lambda) - X_\theta(\lambda) - \lambda X'_\theta(\lambda) + \psi'_\theta(\lambda))U'_\theta = 0.$$

By (6), the last two terms of (11) drop out, allowing it to be rewritten in the form,

$$(12) \quad K'(\lambda) = EXU'/EU'.$$

Given that $K(\lambda)$ is determined by (10) with $X_\theta(\lambda)$ and $\psi_\theta(\lambda)$ solutions of (9), problem (7)–(9) becomes simply

$$(13) \quad \max_{\lambda} EV(-K(\lambda) - (1 - \lambda)X_\theta(\lambda)).$$

The corresponding first-order condition is

$$(14) \quad E(-K'(\lambda^*) + X_\theta(\lambda^*) - (1 - \lambda^*)X'_\theta(\lambda^*))V'_\theta = 0.$$

Substituting from (12), I can rewrite the above equation as

$$(15) \quad EV'EXU'/EU' + EXV' - (1 - \lambda^*)EX'V' = 0.$$

Now introduce the definitions,

$$(16) \quad \bar{X}_u = EXU'/EU'$$

$$(17) \quad \bar{X}_v = EXV'/EV'$$

$$(18) \quad \bar{X}'_v = EX'V'/EV'$$

$$(19) \quad \bar{\epsilon} = -\lambda\bar{X}'_v/\bar{X}_v.$$

Employing (16)–(19), I can coax condition (15) into the form,

$$(20) \quad \lambda^* = \bar{\epsilon} / \left(\bar{\epsilon} - 1 + \frac{\bar{X}_u}{\bar{X}_v} \right).$$

Equation (20) is the basic result of the present paper. Naturally, the first-order condition (15) can be written in a variety of forms, but expression (20) perhaps has the most intuitive economic interpretation.

From (16), \bar{X}_u is a weighted average project cost, where the weights are the agent's marginal utility of income in various states of the world times the probability of occurrence. As such, \bar{X}_u is the expected "real cost" of the project to the contractor.

Analogously, \bar{X}_v is the weighted average project cost, where the

weights are the buyer's marginal utility of income in various states of the world multiplied by the probability of occurrence. Given these same weights, \bar{X}'_v is the average derivative of project cost with respect to the sharing ratio.

The number $\bar{\epsilon}$ of (19) is an elasticity-like measure. It is akin to an average percentage cost reduction for a 1 percent increase in the sharing ratio, where the averaging is done in a special way. As such, $\bar{\epsilon}$ is some measure of the responsiveness of project costs to changes in the sharing ratio. If $\bar{\epsilon}$ is big, on average, X responds elastically to λ ; if $\bar{\epsilon}$ is small, response is inelastic. Note that when uncertainty is multiplicative, i.e.,

$$X_\theta(\lambda) = g(\theta)f(\lambda),$$

for some function $g(\cdot)$ and $f(\cdot)$, then $\bar{\epsilon}$ represents the cost elasticity with respect to λ in each state of the world.

Because higher costs are associated with lower income for the buyer and seller, X_θ is positively correlated with marginal utilities U'_θ and V'_θ . Thus, a ceteris paribus increase in the risk aversion or curvature property of the weights $U'_\theta(V'_\theta)$ results in larger $\bar{X}_u(\bar{X}_v)$ and a smaller (larger) value of λ^* by formula (20).

If we assume that the agent is far more risk-averse than the principal, that will tend to result in the condition,

$$(21) \quad \bar{X}_u/\bar{X}_v > 1.$$

Inequality (21) would certainly hold if the agent is risk-averse and the principal is risk-neutral. When V is linear or nearly linear, a ceteris paribus situation with costs less spread out will diminish \bar{X}_u/\bar{X}_v . In this sense, a less risky distribution of X_θ is associated with a larger value of λ^* .

For the same ratio of \bar{X}_u/\bar{X}_v satisfying (21), higher (lower) average cost elasticities $\bar{\epsilon}$ imply larger (smaller) values of λ^* . This conclusion accords with common sense. When the contractor firm has greater discretionary power to reduce project costs, it should be made to bear a greater share of those costs. Conversely, when there is little that a firm can or will do to cut costs, it may as well be freed from bearing the risk, in which case it can be paid a lower fixed fee.

Finally, it is useful to check out (20) for some extreme values. If $\bar{\epsilon} = \infty$, then $\lambda^* = 1$. If $\bar{\epsilon} = 0$, then $\lambda^* = 0$. If there is no risk aversion, $\bar{X}_u = \bar{X}_v = \bar{X}$ implies that $\lambda^* = 1$. If there is no uncertainty, $\bar{X}_u = \bar{X}_v = \bar{X}$ again implies that $\lambda^* = 1$. All of these results are intuitively plausible.

A RIGOROUS EXAMPLE

It is difficult to do rigorous comparative statics on (20) as it stands, because the right-hand side implicitly contains the sharing ratio. Strictly speaking, the analysis of the last section provides only some general insights into the nature of a solution and what it depends upon. To obtain a closed form expression for λ^* , we turn to a simplified example that preserves the basic features of the general case.

Suppose that there are but two states of the world. State 1 is good (low costs), and state 2 is bad (high costs). When project costs are low, the marginal utility of an extra dollar is one for both principal and agent. When project costs are high, the marginal utility of an extra dollar to the agent is β , whereas to the principal it is γ . The usual case is

$$(22) \quad \beta \geq \gamma \geq 1.$$

Assume that costs take the special multiplicative form

$$(23) \quad X_{\theta}(\lambda) = \begin{cases} \theta_1 \lambda^{-a} & \text{with probability } 1 - p \\ \theta_2 \lambda^{-a} & \text{with probability } p, \end{cases}$$

where

$$\theta_1 < \theta_2.$$

In each state of the world, the elasticity of cost reduction with respect to the sharing ratio is a .

Plugging the above specification into the definitions (16)–(19) yields

$$\begin{aligned} \bar{X}_u &= \frac{(1-p)\theta_1 \lambda^{-a} + \beta p \theta_2 \lambda^{-a}}{1-p+p\beta} \\ \bar{X}_v &= \frac{(1-p)\theta_1 \lambda^{-a} + \gamma p \theta_2 \lambda^{-a}}{1-p+p\gamma} \\ \bar{\epsilon} &= a. \end{aligned}$$

With the above values, (20) becomes

$$(24) \quad \lambda^* = \frac{a}{a-1+\mu},$$

where

$$(25) \quad \mu = \frac{[(1-p)\theta_1 + \beta p \theta_2] [1-p+p\gamma]}{[(1-p)\theta_1 + \gamma p \theta_2] [1-p+p\beta]}.$$

Let the mean of θ be

$$(26) \quad \bar{\theta} = \theta_1(1 - p) + \theta_2 p,$$

where

$$\theta_1 < \bar{\theta} < \theta_2.$$

The variance of θ is then

$$(27) \quad \sigma^2 = (1 - p)(\bar{\theta} - \theta_1)^2 + p(\theta_2 - \bar{\theta})^2.$$

Substituting (26) into (27), I can write the standard deviation of θ as

$$\sigma = \sqrt{p(1 - p)} (\theta_2 - \bar{\theta}),$$

or,

$$\sigma = \sqrt{(1 - p)/p} (\bar{\theta} - \theta_1).$$

Thus,

$$\theta_2 = \bar{\theta} + \sqrt{(1 - p)/p} \sigma$$

$$\theta_1 = \bar{\theta} - \sqrt{p/(1 - p)} \sigma.$$

Substituting for θ_1 and θ_2 from the above expression into (25) yields

$$(28) \quad \mu = 1 + \frac{\sqrt{p(1 - p)} (\beta - 1) \frac{\sigma}{\bar{\theta}}}{1 + p(\beta - 1)} \bigg/ 1 + \frac{\sqrt{p(1 - p)} (\gamma - 1) \frac{\sigma}{\bar{\theta}}}{1 + p(\gamma - 1)}.$$

From (24), with the definition (28) and the condition (22), the following comparative statics statements can be derived:

$$\frac{\partial \lambda^*}{\partial \beta} < 0 \text{ (increased loss aversion of the agent makes } \lambda^* \text{ lower)}$$

$$\frac{\partial \lambda^*}{\partial \gamma} > 0 \text{ (increased loss aversion of the principal makes } \lambda^* \text{ higher)}$$

$$\frac{\partial \lambda^*}{\partial a} > 0 \text{ (increased cost elasticity makes } \lambda^* \text{ higher)}$$

$$\frac{\partial \lambda^*}{\partial \sigma} < 0 \text{ (mean preserving increases in risk make } \lambda^* \text{ lower).}$$

FURTHER SIMPLIFICATIONS AND NUMERICAL ANALYSIS

To render the analysis more transparent, suppose that $\gamma = 1$ (the

principal is risk-neutral) and $p = 1/2$ (high and low project costs are equally likely). Then (28) becomes

$$(29) \quad \mu = 1 + (\beta - 1)/(\beta + 1) (\sigma/\bar{\theta}),$$

and (24) simplifies to

$$(30) \quad \lambda^* = \left(1 + \frac{\sigma}{\bar{\theta}a} \left(\frac{\beta - 1}{\beta + 1}\right)\right)^{-1}.$$

Let z be the ratio of the standard deviation of project cost (measured in terms of the mean) divided by the cost elasticity of response with respect to the sharing ratio:

$$(31) \quad z = \frac{(\sigma/\bar{\theta})}{a}.$$

In some sense, z is a heuristic measure of the ratio of the percentage of project costs not under control of the firm to percentage of costs under control.

Then (30) can be rewritten as

$$(32) \quad \lambda^* = \frac{1}{1 + z((\beta - 1)/(\beta + 1))}.$$

To get a feeling for likely values of λ^* , a few numerical experiments are tried.

It is, of course, difficult to pin down parameter values.

No one has done a formal study that would shed light on reasonable values of a . A number sometimes thrown around in casual observation of defense contracting⁸ is $a = 0.1$, which we take as a point estimate.

Cost variations have been formally studied, but it is difficult to reach definitive conclusions. For military hardware in the early or intermediate production stage, $\sigma/\bar{\theta} = 10$ percent is a figure that seems to emerge.⁹ Basic development work on big hardware systems at an early stage can make $\sigma/\bar{\theta}$ much higher, perhaps as much as the order of 50 percent or more. On the other hand, procurement items in a late stage of development are likely to result in $\sigma/\bar{\theta} = 5$ percent or less. As an "average" value, $\sigma/\bar{\theta} = 15$ percent might be selected.

The loss aversion coefficient β has to be an educated guess, depending greatly upon circumstances. I pick $\beta = 1.25$ as a point estimate and experiment with other values.

Table I summarizes the results of numerical experiments based

8. Moore [1967], p. 214, and Scherer [1964a], p. 264.

9. Moore [1967], p. 221, and Scherer [1964b], Ch. 8.

TABLE I
 λ^* AS A FUNCTION OF z AND β (IN PERCENTAGE POINTS)

β	z					
	$1/4$	$1/2$	1	2	4	8
1.1	98	96	92	85	73	58
1.25	97	95	90	82	69	53
1.5	95	91	83	71	56	38

on (31) and (32). Especially striking to me are the rather high values of λ^* that seem to emerge. My "average" specification $z = 1.5$, $\beta = 1.25$ yields $\lambda^* = 86$ percent. In most reasonable scenarios it seems difficult to escape the conclusion that λ^* ought to be above 50 percent, sometimes well above.¹⁰ This result is perhaps a bit surprising, especially in view of the low cost elasticities being assumed.

Of course, the model is a gross oversimplification of reality, and calculations based on it should be cautiously received. But at least this framework provides some basis for determining sharing ratios—which brings us to a concluding note. This type of model often does not admit a closed-form solution when modified to suit particular applications. Nevertheless, simulations are usually tractable and lead to significant insights into the nature of an efficient contract.

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

REFERENCES

- Berhold, M. A., "Linear Profit-Sharing Incentives," this *Journal*, LXXXV (Aug. 1971), 460–82.
- Cummins, J. M., "Incentive Contracting for National Defense: a Problem of Optimal Risk Sharing," *Bell Journal of Economics*, VIII (Spring 1977), 168–85.
- Fox, J. R., *Arming America: How the U. S. Buys Weapons* (Cambridge: Harvard University Press, 1974).
- Holmstrom, B., "Moral Hazard and Observability," *Bell Journal of Economics*, forthcoming.
- Mirrlees, J., "An Exploration in the Theory of Optimum Income Taxation," *Review of Economics Studies*, XXXVIII (April 1971), 175–208.
- , "Optimal Tax Theory: A Synthesis," *Journal of Public Economics*, VI (1976b), 327–58.
- Moore, F. T., "Incentive Contracts," in S. Enke, ed., *Defense Management* (Englewood Cliffs, NJ: Prentice-Hall, 1967), pp. 213–31.

10. This conclusion is somewhat softened if, instead of (23), the linear form $X_\theta(\lambda) = \theta - B\lambda$ is selected. In this case formula (32) becomes $\lambda^* = 1 - ((\beta - 1)/(\beta + 1))(\sigma/B)$. Associating $z = \sigma/B$, one can find that both formulas give similar results for z and β in the intermediate and low ranges. But for the higher values of z and β , the linear case yields decidedly lower sharing ratios. My own feeling is that the multiplicative specification is more reasonable than the additive, but this is not a closed issue.

- Peck, M. J., and F. M. Scherer, *The Weapons Acquisition Process: an Economic Analysis* (Cambridge: Harvard University Press, 1962).
- Ross, S., "The Economic Theory of Agency: The Principal's Problem," *American Economic Review*, LXIII (May 1973), 134-39.
- Scherer, F. M., "The Theory of Contractual Incentives for Cost Reduction," this *Journal*, LXXVIII (May 1964a), 257-80.
- , *The Weapons Acquisition Process: Economic Incentives* (Cambridge: Harvard University Press, 1964b).
- Shavell, S., "Risk Sharing and Incentive in the Principal and Agent Relationship," *Bell Journal of Economics*, forthcoming.
- Sheshinski, E., "On the Theory of Optimal Income Taxation," Harvard Institute of Economic Research Paper No. 172, February 1971.
- Spence, A. M., and R. Zeckhauser, "Income, Information, and Individual Action," *American Economic Review*, LXI (May 1971), 380-87.
- Williamson, O. E., "The Economics of Defense Contracting: Incentives and Performance," in R. N. McKean, ed., *Issues in Defense Economics* (New York: Columbia University Press, 1967), pp. 217-56.