THE CONTINGENT PRICING PROBLEM: SOME CONSIDERATIONS IN FORMULATING QUALITY INCENTIVES

Eugene P. Durbin

April 1965

COPY 2 OF 3 16p.
HARD COPY $. 1.00
MICROFICHE $. 0.50

DDC
RECEIVED
MAY 17 1965
DDC-IRA E

P-3124

ARCHIVE COPY

Approved for OTS release

THE CONTINGENT PRICING PROBLEM: SOME CONSIDERATIONS IN FORMULATING QUALITY INCENTIVES

Eugene P. Durbin*
The RAND Corporation, Santa Monica, California

INCENTIVE CONTRACTS

The participants in this conference are all aware of the increasing use of contractual incentives. In the last three years there has been greatly increased use of multiple incentive contracts, or contracts that make profit dependent on project cost and schedule outcomes as well as product performance. One extreme type of cost incentive contract is the firm-fixed-price (FFP) contract, in which after design specifications and target cost are negotiated, any cost deviation is the responsibility of the contractor. He alone benefits from any reduction in cost, and incurs the entire burden of a cost overrun. At the other extreme is the cost-plus-fixed-fee (CPFF) contract which was extensively used throughout the late 1950's. With this contract type the government assumes responsibility for any cost deviation from the agreed-on target. Between these polar contract types are the cost-plus-incentive fee (CPIF), and fixed-price-incentive (FPI) contracts which provide for the sharing of any cost deviations. Multiple incentive contracts base contractor profit on delivery schedule and product performance in addition to cost outcome. Incentive payments are usually awarded

---

*This paper was prepared prior to the time the author joined The RAND Corporation. Any views expressed in this paper are those of the author. They should not be interpreted as reflecting the views of The RAND Corporation or the official opinion or policy of any of its governmental or private research sponsors. Papers are reproduced by The RAND Corporation as a courtesy to members of its staff.

This paper is to be presented at the Conference on Quality Control and Industrial Statistics, Stanford University, May 7, 1965.

after development or after the production of a series of items, and payment is contingent on some estimate of product performance.

The expanded use of incentive contracts was in part due to the opinion of many government procurement agencies that CPFF contracts were sapping the efficiency of industrial contractors. CPFF contracts were felt to be instrumental in causing both cost overruns and the inclusion of excessive frills or "goldplating." The stated intent of the incentive environment is to induce the contractor to improve program management by sharing with him a portion of the cost saving resulting from his more critical management.

The primary problem faced by the government in formulating pure cost incentive contracts is that the government negotiator faces a better informed opponent across the bargaining table and cannot know the contractor's motivations. Studies have suggested that defense contractors are risk-averse, and when forced into accepting high sharing proportions, they bargain fiercely for loose target costs. In formulating multiple incentive contracts, additional difficulties arise. Contractors do not know their own cost-delivery-performance tradeoff surfaces precisely, nor do they exert perfect control over product performance or cost outcomes. Moreover, the typical contractor is motivated by many factors other than short-term accounting profit. Survival of the contractor's organization, long-term growth, long-term profit, and ability to keep key teams productively employed appear to influence program decisions more than current profit considerations. Unless the contractor perceives the potential current profit as very large, it can be a weak and uncertain motivator, and does not permit accurate prediction of contractor behavior.

THE CONTINGENT PRICING PROBLEM

The contingent pricing problem is an outgrowth of acceptance sampling studies and was first defined by Vernon M. Johns and Gerald J. Lieberman.[(1)] The connection with incentive contracting arises since true product performance cannot be determined directly in most cases, but must be inferred through measurements and tests, thus giving rise to measurement uncertainty analogous to sampling uncertainty.

To illustrate the ideas involved in both the contingent pricing problem and the quality incentive problem, consider that a government procurement agency (the consumer) purchases a new computer. Perhaps the mean number of failures per unit time is a performance index of interest. We assume that the manufacturer can precisely control the mean number of failures per unit time by properly designing the device, using high cost components, extensive checkouts, or other means. Thus a cost may be associated with the final true quality of the computer. Of course in any time interval the actual number of failures is a random variable. We consider that the consumer accepts the delivered computer, places the unit on test for a specified interval, notes the number of failures that occur, and then pays the manufacturer a specified fee based on the number of failures observed. This is an example of payment contingent upon the observed quality of a product.

Let us reformulate the problem in terms of the familiar acceptance sampling situation. We assume that at a known cost $h(p)$, the producer can control the probability, p, that each item produced is a defective item, and that subsequent to production, items are formed into batches or lots of N, and are delivered to the consumer. The consumer then samples $n(\leq N)$ from each batch at a cost of c per unit sampled, observes x

defectives, and pays $\varphi(x)$ based on the number of defectives observed. We may also assume that the consumer knows V(p), the expected value of the product when the manufacturer produces at an average quality level, p. The consumer may actually know V(p) for only a few values of p, or he may not know it at all. The pricing policies that we will derive are not dependent on the assumption that the consumer knows V(p).

BASIC MODEL

We have thus far postulated the existence of a production and procurement situation in which the manufacturer exerts a known degree of control over his product at a known cost. The consumer purchases batches of items at a unit price to be determined according to the results of inspecting a sample of the product. With the additional premise that the producer will select p to maximize expected profit, and that the consumer will choose $\varphi(x)$ and the sample size, n, to maximize his own total expected profit, we may define a motivating price strategy. Assume that the consumer desires production quality to be set at p*, and that he can select a sample size and construct a price schedule so that the producer's expected profit is maximized at p* and is at least as great as the agreed-on profit at p*. The producer will then be induced to produce at p* and the strategy selected is a _motivating_ strategy. Under these assumptions, Johns and Lieberman showed that a sample size of one is sufficient to allow the construction of a motivating strategy. Such a policy is shown on Fig. 1 and is obviously the least-cost policy for the consumer in the situation postulated. B. J. Flehinger and J. Miller of the IBM Corporation added the assumption that acceptable policies would require that payments, $\varphi(x)$, be bounded below.[2] They

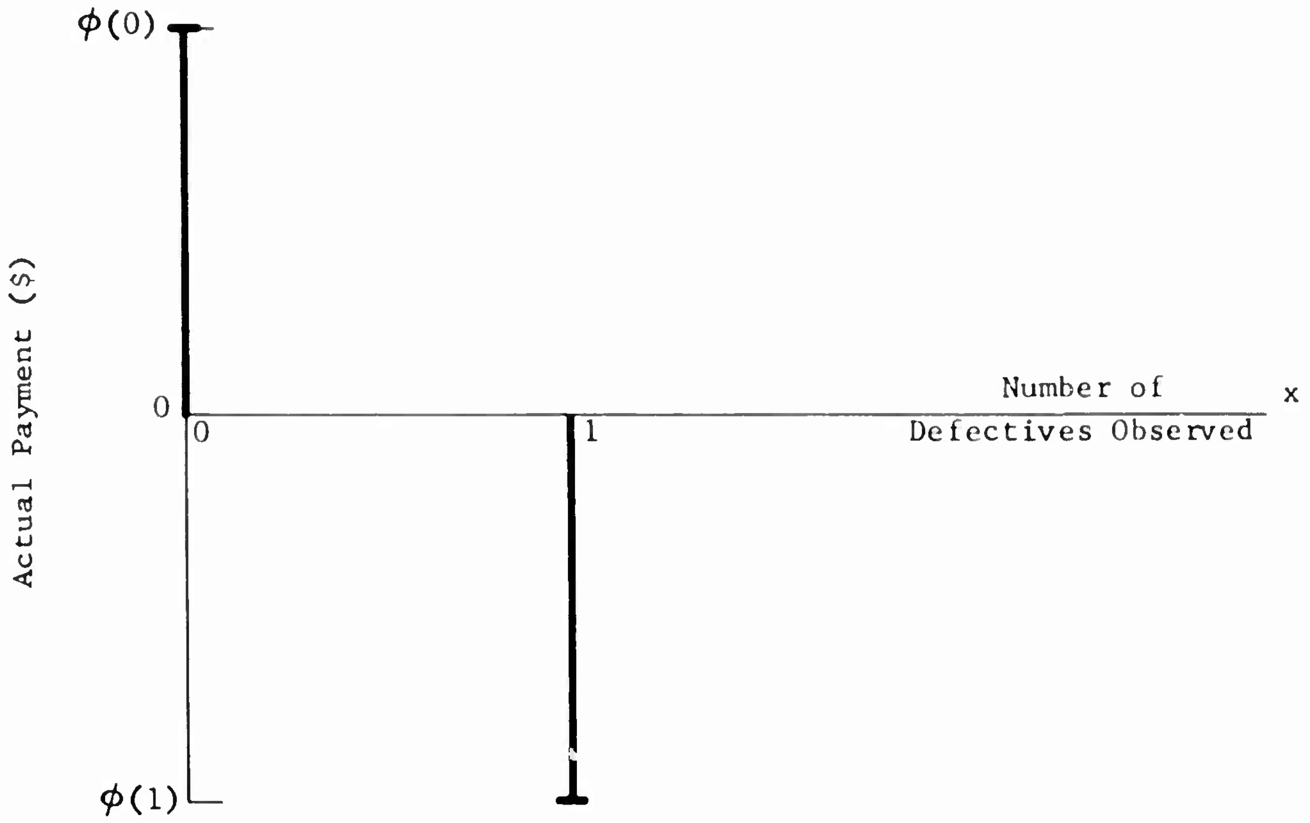

Fig. 1 -- Price Schedule Based on a Sample of One

Fig. 2 -- Price Schedule Resulting from Flehinger-Miller Formulation

chose zero as the lower bound. The consumer's optimal strategy in this case includes reasonably large sample sizes from an operational point of view. However, a sizeable payment is awarded the contractor if and only if zero defectives are observed in the sample, while if any defectives are observed in the sample, only the minimum payment is made. Figure 2 illustrates such a policy. There are two objections to this type of policy. First, even if the manufacturer produces at a high level of quality, sampling variation may cause only the minimum payment to be awarded an uncomfortably large number of times. And there is no way to control the risk of incorrect payment since both the payment $\varphi(x)$ and the sample size, n, are jointly determined by the requirements that producer expected profit be both maximized at p^* and also that the magnitude of the expected profit be not less than the value agreed on.

Of course, what is missing here is the realization that there are factors other than the level of current expected profit that concern the contractor, and one such consideration is the possibility of a large loss. To treat this factor completely, we must assume that both parties seek to maximize expected utility rather than expected profit. Acceptance sampling theory is faced with this problem, and the tactic often used to avoid constructing complete utility functions is to focus directly on relevant quality levels and probabilities of incorrect decisions at those levels. The common terminology that has arisen for these points and risks are the AQL (acceptable quality level) and producer's risk, and the LTPD (lot tolerance per cent defective) and consumer's risk. Analogous considerations in the contingent pricing situation lead us to assume that the producer is interested in being protected against receiving less than w when his quality is in fact acceptable, and that the consumer desires protection against paying more than v when quality is actually poor. Both parties must specify the levels of protection, α and β, they desire. The consumer specifies explicitly the level of quality he considers poor, and specifies implicitly the level of acceptable quality by his knowledge of the cost function, h(p), and his subsequent choice of $\varphi(x)$ and n.

We now write the consumer and producer expected profits:

(1) $$T_c\ (n,\varphi,p) = N\cdot V(p) - N\cdot g(p) - nc$$

(2) $$T_p\ (n,\varphi,p) = N\cdot g(p) - N\cdot h(p)$$

where g(p) is the expected value of $\varphi(x)$ given p for fixed n.

The assumptions concerning risk are included as:

(3) $$\mathrm{pr}\left[\varphi(x) \geq w \;\middle|\; p \leq p^*\right] \geq 1 - \alpha$$

(4) $$\mathrm{pr}\left[\varphi(x) \leq v \;\middle|\; p \geq p_b\right] \geq 1 - \beta$$

and we include

(5) $$M \geq \varphi(o) \geq \ \ldots \geq \varphi(n) \geq \ m$$

(6) $$N\ g(p^*) - N\ h(p^*) \geq 0$$

which are constraints on the form of the pricing policy, and set the agreed-on profit at zero.

Recall that the basic assumption here is that given an acceptable pricing policy, the producer will act to maximize expected profit by selecting p. Thus the consumer must select that $\varphi(x)$ and n which both induce the producer to select that p* desired by the consumer, and also achieve the minimum expected profit that the producer will accept at that p*. The "acceptability" of the policy is determined by the restrictions (3) to (6). Formally, this model is a two-person non-zero-sum game in which not all pricing strategies and sample sizes are permitted, but only those that assure the producer and consumer that their losses due to sampling variation will not be excessive.

SOLUTION PROCEDURE AND RESULTS

We now fix the quality level and sample size, say, at (n,p*), and ask if there is a price schedule, $\varphi(x)$ which satisfies the stated conditions. If so, we derive it and compute the consumer expected profit. By comparing consumer expected profit resulting from different quality levels and sample sizes we can select that (n,p) which leads to the maximum consumer expected profit. To derive the price schedule for each (n,p) we solve the following linear programming problem.

(7) $$\underset{\varphi_x}{\text{Minimize}} \sum_{x=0}^{n} \pi_x \varphi_x$$

(8) $$\sum_{x=0}^{n} \delta_x \varphi_x = h'(p^*)$$

(9) $$\varphi_{k_\alpha} \geq w$$

(10) $$\varphi_{k_\beta} \leq v$$

(11) $$M \geq \varphi_o \geq \ldots \geq \varphi_n \geq m$$

(12) $$\sum_{x=0}^{n} \pi_x \varphi_x \geq h(p^*)$$

where π_x is the probablity of x defectives and $\delta_x = d/dp\ (\pi_x)$.

It can be shown that a pricing strategy that satisfies (1) - (6) must be a solution to this linear programming problem. That this must be so is plausible since by specifying (n,p) we fix the variable part of the non-zero-sum game and arrive at a constant-sum game.

The transformation:

(13) $$y_x = \varpi_x - \varpi_{x+1}\,, \qquad x = 0, 1, \ldots, n-1$$

(14) $$y_n = \varpi_n - m$$

allows us to seek not the actual price schedule, $\varpi(x)$, but the jumps in price level, y(x). This transformation leads to a problem with only five constraints, thus allowing rapid computation of the price schedule. Furthermore, from the structure of this problem and the properties of solutions to linear programming problems, we can infer that an appropriate price schedule need have at most six distinct price levels. Such a policy is illustrated in Fig. 3. This policy maximizes producer expected profit at p*, and leads to the agreed-on expected profit (zero in this case) at p*. It assures the producer that he will receive no less than the negotiated price w at p* with probability $1-\alpha$, and assures the consumer that with probability $1-\beta$ he will pay no more than v when quality is as poor as p_b. The expected payment, g(p), resulting from this policy is shown as a function of the producer's choice of p in Fig. 4.

If there is a known function, V(p), we can systematically search for that value of p which maximizes consumer expected profit. We have experimented with a generalization of the Fibonacci search technique to seek an optimal (n*,p*) but there is very little that can be said in general about the optimality of a given policy. If one particular p* is specified we may find the minimum sample size, n*, for which a price

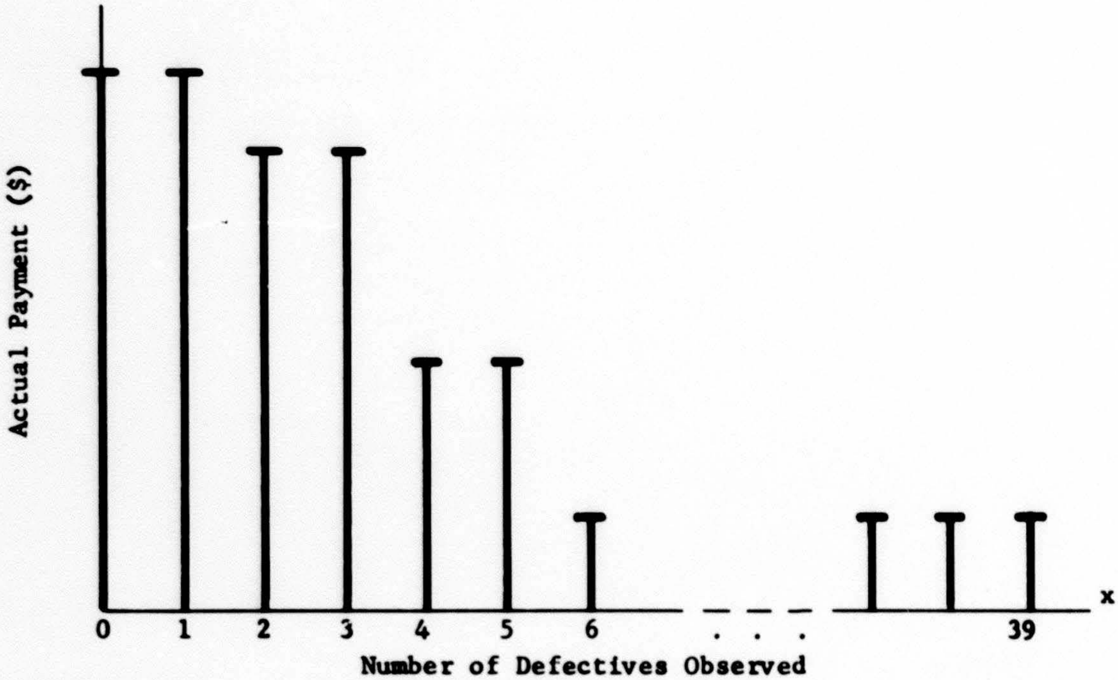

Fig. 3 -- Piecewise Constant Price Schedule

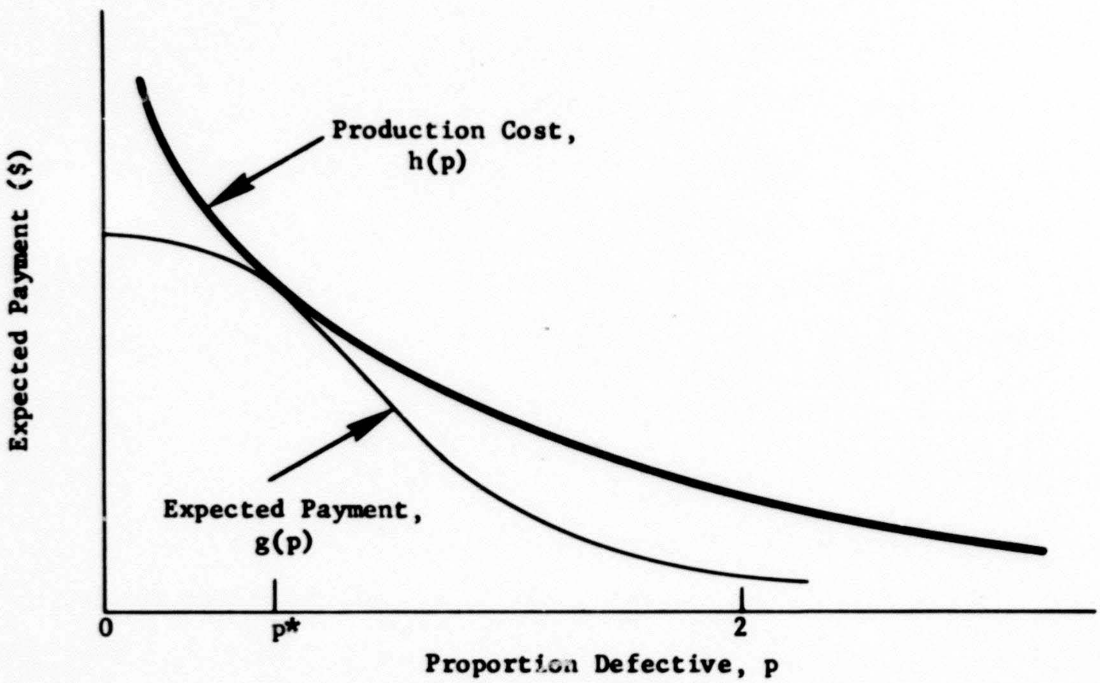

Fig. 4 -- Expected Payment Resulting from Piecewise Constant Price Schedule

schedule can exist at p^* satisfying the risk constraints (3) and (4). At n^* we can ensure that a price schedule $\varphi(x)$ exists by properly relaxing the upper and lower price bounds, M and m, and also relaxing the risk-averting price levels w and v. If prices are fixed, however, a sample size larger than n^* may be required to permit the construction of a satisfactory price schedule.

EXTENSIONS

The basic approach presented can be generalized to include a wider class of problems, and the basic assumptions can be weakened.

Policies of a particular functional form, such as the piecewise linear policies observed in practice, can be derived by suitably constraining the problem (1) - (6). We discuss the linear policy in a later section. In most operational settings the contractor must replace all defectives discovered during sampling inspection at no extra cost to the consumer. This situation can be treated by appropriately modifying the objective function and constraints of the basic model. Next we need not assume that the producer exerts precise control over the production process, but only that he can choose an average quality, μ, and that the true quality achieved on any run is a random variable with a known distribution dependent on this choice of μ. Similarly, the cost of process control may contain an additive random component with a known distribution. No essential difficulties are encountered in including any of these considerations.

Finally, we may consider situations in which the result of a test or inspection is a continuous rather than a discrete variable. This would occur in the case of the computer manufacturer if the quantity of interest were mean time between failures. Assume that the computer is put on test until a specified number of failures occur. The result of the test is total time, a continuous variable. The basic model can describe this situation, but with integrals replacing the summations. The linear programming problem becomes a constrained variational problem, linear in the price schedule, $\varphi(t)$. There are techniques available for solving such problems and the solution to this one is readily obtained by the Maximum Principle of Pontryagin. The structure of the resulting optimal price schedule is similar to that of the policies obtained through linear programming. In fact, linear programming is a practical method of approximating optimal policies in these continuous cases.

PIECEWISE LINEAR PRICE SCHEDULES

We define piecewise linear price schedules as those which pay maximum constant fee for x less than some point a, have the fee linearly related to the observed number of defectives from a up to some point b, and pay a constant minimum fee thereafter. Such policies are in common use since they are easily interpreted by the producer, and are intuitively more satisfying to both consumer and producer than price schedules with a series of discontinuities or jumps.

In the discrete case, we modify the constraint (5) to derive linear policies by requiring $\varphi(x)$ to be linear within a specified interval (a,b). Of course, we must then vary (a,b) to find a profit-maximizing policy, but this can be done in an efficient manner, making the numerical determination of such policies practicable. In the continuous case we can derive policies which are linear by using a quadratic criterion function and no risk constraints. This last result has the following interesting interpretation. Let the difference between the actual payment, $\varphi(t)$ and $h(\mu^*)$ be <u>excess profit</u> given that the producer selects μ^*. Observe that excess profit may be negative. We assume that the negotiated expected profit is zero, and that the consumer is equally unhappy when negative or positive excess profit results. If he seeks that policy which both minimizes the mean square excess profit at μ^* and also maximizes producer-expected profit at μ^*, thus motivating the producer to select μ^*, the resulting policy is, under very general conditions, a piecewise linear price function. Thus a linear policy is in a sense a "fair" policy.

To offset the advantages offered by the linear price schedules, they require a larger sample size to satisfy the same conditions as a basic schedule. How much larger the sample size need be depends on the risk levels and the prices associated with upper and lower bounds. As the conditions (3) - (6) become more restrictive, the ratio of the sample size required for a linear policy to that required for a basic policy increases. A fairly inclusive estimate is about 1.1 - 1.6 times the basic sample size. The increased sample size can be reduced by negotiating for lower prices w, accepting higher prices v, and attempting to keep the upper and lower price bounds, M and m, spread apart.

The linear payment schedule may also be misleading in regard to the profit-motivation it offers. The producer's cost function will generally be convex and monotonic, while the expected payment will deviate from linearity as it approaches the upper and lower bounds. Thus the expected profit will drop away more sharply on the superior side of the desired quality level p* than on the inferior side.

Recognizing this, it behooves the government to formulate linear payment schedules with at least two linear sections, constructed so that the policy $\varphi(x)$ and the expected payment are more concave in the region of p^*, lessening the temptation for the contractor to allow performance to slip slightly below p^*.

In this discussion we have indicated some of the factors involved in formulating incentive price schedules. The primary point is that government negotiators must allow for sampling variation and the fact that contractors will consider expected results as well as explicit payment schedules.

REFERENCES

1. V. M. Johns and G. J. Lieberman, Contingent Pricing Policies, Unpublished Research Proposal, Stanford University, 1961.

2. B. J. Flehinger and J. Miller, "Incentive Contracts and Price Differential Acceptance Tests," Journal of the American Statistical Association, Vol. 59, No. 305, March, 1964.