Joint Strike Fighter:
Optimal Design through Contract Incentives

Paul D. Collopy

DFM Consulting

ABSTRACT

The development and production contract for the Joint Strike Fighter can be structured to save the government an estimated $50 billion and improve the contractor's expected profits at the same time. The key is to incentivize the contractor to manage the weapon system design in a novel way that maximizes the collective benefit to the government and the contractor. This paper presents a method of distributed optimal design that is tied to a value-based contract. A method of quantifying weapon system value is presented: a **value model** used both for contracting and design. Given the value model, the contract is shown to be a superior prospect from the government perspective. The contract is also attractive from the contractor's point of view, and has no risk of loss.

INTRODUCTION

Acquisition Reform has removed many of the constraints on acquisition contracting, but at this early point in the process there are few if any precedents to guide contracting officers. The time is ideal for the full application of rational analysis to structuring defense acquisition contracts to optimize program results.

The largest near term opportunity for innovative contracting is the Joint Strike Fighter, a massive program to replace the DoD inventory of F-16 Falcons, A-6 Intruders, AV-8B Harriers, and other strike fighters. The contracts are expected to stretch out over thirty years and produce thousands of aircraft. The stakes are high: Total life cycle cost is estimated at $250 billion. Moreover, three services will depend on the JSF as the backbone of their aviation capability for the first half of the 21st century. An application of decision analysis tools to structure an incentive-based contract along the lines described in this paper could save the US several tens of billions of dollars.

DECISIONS AND DECISION MAKERS

To design the contract for development and production of the Joint Strike Fighter, a complex set of interrelated decisions and decision-makers must be considered. Figure 1 focuses on decision makers on three levels: the government Director of the Joint Strike Fighter Program Office, the contractor Project Manager, and the myriad design engineers who actually develop the weapon system.

Government

The Joint Strike Fighter Program Office is charged with formulating the requirements for the aircraft family, contracting for the development and production, and, working with the contractor, managing the quality and the price of the program. A key decision of the Program Director is construction of the development and production contract.

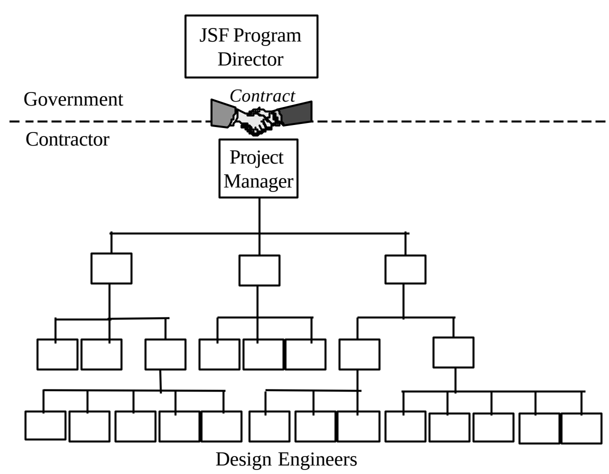

Figure 1: A complex set of decision makers

Contractor

The contractor will design and produce the Joint Strike Fighter aircraft, support equipment, and training systems to meet the specifications provided by the government. In return, the contractor will receive an agreed-upon price. Traditionally, the price has been based on reimbursing costs incurred by the contractor. Many contracts have added incentive payments based on contractor performance, and warranties have been proposed, which would tie contract price to performance of systems in the field.

The key decisions of the contractor's project manager considered in this paper are:

- whether to accept the contract
- whether to exercise contract options
- how to manage engineering design so as to maximize contract profit

Design Engineer

The vast majority of the decisions that determine the cost and performance (and therefore value) of the Joint Strike Fighter will be made by design engineers deep in the contractor's and subcontractors' engineering organizations. These decisions include material selections, subsystem architecture, and the detailed design of hundreds of thousands of parts in the aircraft.

SOLUTION

The key to structuring the JSF contract is the impact that contract incentives will have on the engineering design of the aircraft. Optimization of the aircraft design can profoundly change the value and cost of the production aircraft and benefit the development program as well. Thus, this section first presents a system for distributed global optimization in design. The system depends on a value model, which is described next. Finally, a contract is constructed, using the value model, which provides strong incentives for the use of the optimal design system, to the mutual benefit of the government and the contractor.

Global Optimization

Let v represent the value of the system (that is, the Joint Strike Fighter), let p represent the price paid to the contractor, and let c represent the cost incurred by the contractor. Let B represent the net benefit to the government, so that

$$B = v - p \tag{1}$$

and let Π represent the contractor's profit:

$$\Pi = p - c \tag{2}$$

The collective benefit to the government and the contractor is

$$B + \Pi = v - c \tag{3}$$

Thus, a design that maximizes v - c could, under a variety of contracting arrangements, increase both B and Π, a win-win situation. The contracting arrangements will be discussed later *(Contract Structure)*. First, though, it is necessary to understand how the designers can be directed to maximize v - c.

Distributed Optimization

Assuming for now that the contract will be one in which both B and Π increase in a monotonic fashion with v - c, then both the government and the contractor will want to manage design engineering in a way that maximizes v - c. A communication system will be described which allows every design engineer to optimize each component of the system in a way that globally optimizes v - c. The design is optimized in its **extensive variables**—other variables are specified in the traditional fashion and act as constraints.

Definition of Extensive Variables

An **extensive variable**, for the purpose of this paper, is one that measures a property of the whole system (such as weight) and also measures a property of many components of the system; and the system variable is an aggregation of the component variables. Specifically, if x is an extensive variable of the system and x_k, where k = 1, 2, ... n, are corresponding variables for n components, then it is required that

$$\frac{\partial v}{\partial x} = \sum_{k=1}^{n} \frac{\partial v}{\partial x_k} \tag{4}$$

Extensive variables of interest are those that measure properties of many components, where coordination and optimization are important and complex problems. Notice that the development and manufacturing cost, c, is similar to an extensive variable in that $\delta c \approx \sum \delta c_k$.

Local Application of the Value Model

Let $v(\vec{x})$ be the value of the system in terms of the vector $\vec{x}$ of all the extensive variables except cost, which is treated separately. (An example of the value vector v is shown in below under the heading *Value*

Modeling.) Let $\vec{x}_0$ describe the performance of the system projected at the time of the contract. Then, for $\vec{x}$ in the neighborhood of $\vec{x}_0$,

$$v(\vec{x}) \approx v(\vec{x}_0) + \nabla v(\vec{x}_0) \bullet (\vec{x} - \vec{x}_0) \qquad (5)$$

where ∇v is the gradient of v.

Then, by equation (4),

$$v(\vec{x}) \approx v(\vec{x}_0) + \nabla v(\vec{x}_0) \bullet \sum_{k=1}^{n} (\vec{x}_k - \vec{x}_{0k}) = v(\vec{x}_0) + \sum_{k=1}^{n} \left[\nabla v(\vec{x}_0) \bullet (\vec{x}_k - \vec{x}_{0k})\right] \qquad (6)$$

and since $c - c_0 \approx \sum_{k=1}^{n} (c_k - c_{0k})$,

$$v(\vec{x}) - c \approx v(\vec{x}_0) - c_0 + \sum_{k=1}^{n} \left[\nabla v(\vec{x}_0) \bullet (\vec{x}_k - \vec{x}_{0k}) - (c_k - c_{0k})\right] \qquad (7)$$

Define for convenience a single constant to collect all the constant terms:

$$A \approx v(\vec{x}_0) - c_0 + \sum_{k=1}^{n} c_{0k} - \nabla v(\vec{x}_0) \bullet \vec{x}_{0k} \qquad (8)$$

Then

$$v(\vec{x}) - c \approx A + \sum_{k=1}^{n} \left(\nabla v(\vec{x}_0) \bullet \vec{x}_k - c_k\right) \qquad (9)$$

The important consequence of (9) is that $v(\vec{x}) - c$ increases in a monotonic fashion with each term of the summation, so that every component design team that maximizes $\nabla v(\vec{x}_0) \bullet \vec{x}_k - c_k$ is maximizing $v(\vec{x}) - c$. Thus, with the right contract structure (one in which Π increases with v-c), it will be in the mutual interest of the government and the contractor to manage design teams so that they maximize $\nabla v(\vec{x}_0) \bullet \vec{x}_k - c_k$ (that is, value less cost).

Engine Inlet

	Status	**Gradient**	**Value**
Efficiency	90%	150,000	135,000
Weight	700	-130	-91,000
Reliability	1500	2.3	3,450
Maintainability	7.8	-340	-2,652
Maintenance Cost	500	-0.5	-250
Support Equipment	12	-15	-180
Radar Cross-Section	0.1	-1200	-120
InfraRed Signature	1.4	-50	-70
Manufacturing Cost	700	-1	-700
Design Value			**43,478**

Figure 2: Component Design Scorecard

Think Globally, Act Locally

Value misalignment is an endemic problem in the current spec-flowdown system, where every part has performance specifications and a cost target. An example:

Engineer A's part is over spec weight but under target cost. His best choice is to change to a lighter weight, more expensive material. Cost increases $5,000 and weight is reduced by 8 pounds. The design is now satisfactory. Engineer B's part is over cost but under on weight. A cheaper, heavier material increases weight 28 pounds, but reduces cost $2,000, meeting all goals. The net effect of both decisions on the system is +20 lbs. and +$3,000.

It is impossible for this problem to occur when both engineers are maximizing $\nabla v(\vec{x}_0) \bullet \vec{x}_k - c_k$, as they would using the scorecard illustrated in Figure 2.

The Status column on the scorecard lists candidate extensive variables (the elements of the vector $\vec{x}_k$ plus the value of c_k) for one system component, the engine inlet (the data shown is hypothetical). The center column is the gradient, $\nabla v(\vec{x}_0)$, augmented with -1, the coefficient of cost. This column is identical for all components. The elements of the third column, value, are the products of the same row elements in the first two columns, so that the Design Value, the sum of the third column, is equal to $\nabla v(\vec{x}_0) \bullet \vec{x}_k - c_k$.

Every design decision made by component design engineers can be assessed by its impact on the scorecard. The Status data can be determined for each design alternative, and the design value calculated. Generally, the alternative with the highest Design Value would be chosen. Thus, design engineers can evaluate every design decision according to what maximizes v - c for the overall system, that is, what provides the greatest collective benefit to the government and the contractor. The impact of coordinated distributed optimal design on the ultimate product should be substantial.

A side benefit is that project management can use data from scorecards to continually assess status of the overall system and quickly react to problems as they arise.

Value Modeling

The distributed optimization just described requires knowledge of the gradient $\nabla v(\vec{x}_0)$. A straightforward method of obtaining this gradient would be to construct a value model in the extensive variables, $v(\vec{x})$, and generate the gradient by sensitivity analysis around any known design $\vec{x}$. Such a model has actually been constructed by DFM Consulting for a general multi-role fighter in 1997. Following is a description of a value model for the Joint Strike Fighter.

Requirements

There are two requirements that such a value model must satisfy:

1. For any feasible design $\vec{x}$, the model must estimate the value $v(\vec{x})$ in dollars.
2. For any change in any of the extensive variables $(x_1, x_2, \cdots x_m) = \vec{x}$, the model must appropriately reflect the impact of the change on value, $v(\vec{x})$.

Structure

The essence of the value model is

$$\text{Value} = \text{Wartime Net Value} - \text{Peacetime Cost} \quad (10)$$

The second term in this equation, Peacetime Cost, is the Operations and Support portion of traditional Life Cycle Cost, which is well understood in the current state of the art and has already been modeled for the Joint Strike Fighter. Therefore, the discussion in this paper will focus on Wartime Net Value:

$$\begin{aligned}&\text{Wartime Net Value} = \\ &\text{Wartime Value} - \text{Wartime Cost}\end{aligned} \quad (11)$$

Wartime Value of a weapon system is the value of the military objectives that the weapon system can achieve in a conflict—objectives such as close air support, deep interdiction or air superiority. One paradox of military valuation is that a weapon provides value, in the form of deterrence, even when it is not used. The solution to the paradox is to evaluate the weapon's use in the prospective conflict upon which the deterrence argument is based, in the same way that a financial asset is valued according to the price it would command in the market, even if there is no intention to sell the asset. This is not an original notion:

> *The decision by arms is for all major and minor operations in war what cash payment is in commerce. Regardless how complex the relationship between the two parties, regardless how rarely settlements actually occur, they can never be entirely absent. [von Clausewitz, 1832, p. 97]*

The nature of these **intended campaigns**, will be explored, followed by a discussion of the terms of equation (11).

Intended Campaigns

A weapon system is developed with a particular set of campaigns in mind in which the weapon system plays a role or set of roles. Even when the weapon system is expected to deter the occurrence of such a campaign, the campaign must be thought through in detail prior to the system's design, so that the system presents the credible threat on which deterrence is founded.

For the intercontinental nuclear ballistic missiles developed in the 1950s and later, the intended campaign was a wholesale nuclear assault on the USSR, with attacks on cities or attacks on missile bases as variants. Although the architects of this system believed that its implementation would succeed in forestalling its use, the attacks were planned in great detail and tests were conducted to verify that the weapon system would accomplish the objectives of such an attack.

Similarly, the F-15 and F-16 fighters were developed primarily for a short, very intense, land campaign in central Europe. The campaign never occurred, in part because NATO was prepared with weapon systems designed and tested to be credible deterrents.

State of the art weapon system design includes extensive tactical and logistic simulation of the weapon system in the intended campaign. Data from such simulations form the basis for modeling Wartime Net Value.

Wartime Net Value

As stated in (11), Wartime Net Value is made of two components, Wartime Value and Wartime Cost. Both can be assessed on a per sortie basis for a strike fighter. For a particular type of sortie, such as deep interdiction, in a particular intended campaign, the mission geometry (range, altitude, and so on) can be defined by probabilistic distributions. From these, an expectation of cost can be calculated, including fuel and time from performance models, logistic costs from logistic models and aircraft losses from survivability models. The JSF program has already generated large amounts of this sort of data.

Wartime Value is a different phenomenon. Within the campaign there are many possible targets or objectives. Tacticians will arrange these roughly in order of value, as displayed on the Pareto chart in Figure 3.

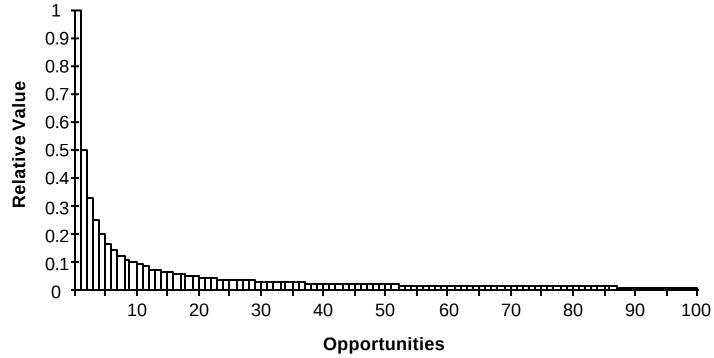

Figure 3: Pareto Chart of Value of Deep Interdiction Sorties

The combination of varying value of objectives and constant expectation of sortie cost is described in Figure 4. The important feature of Figure 4 is that there is a point of diminishing returns, where the value line crosses the cost line, such that only some of the sorties (those left of the crossover) are worth executing. The shaded area is the Wartime Net Value for this type of sortie in a particular intended campaign.

The challenge in this model is the determination of the Wartime Value appropriate to particular types of military objectives in particular intended campaigns. This problem can be solved by treating it as a pricing problem and imputing the value based on historical buying decisions.

Imputed Value

Engineers tend to find cost models straightforward, because they measure mostly physical attributes, like pounds of fuel, flight hours and pallets of support equipment, and multiply them by cost factors. Cost models are therefore similar to weight models or performance models in form. Pricing is completely different. Weapon system designers typically find the very concept of assigning a value to a military objective to be problematic. But they would also find the concept of assigning value to a loaf of bread problematic.

Former Programs

In fact, the approach to both pricing models is the same: look at what has been paid for similar commodities in the past. To price a loaf of bread, find a similar loaf at the local supermarket, and observe what it is selling for. To price a set of strike fighter objectives, find similar programs, such as the F-16 and F-18, set up the strike fighter value model with data for these programs, fix the fleet size actually purchased and the price paid for the fleet, and solve the value model for the value of the objectives. This, then, is the imputed value of the objectives based on the amount paid in historical programs. Like the price of the loaf of bread, this is in fact a lower bound—if the commodity could not have been obtained for this price, perhaps more would have been paid.

However, a cost analysis can yield an element of the upper bound, as follows: a certain number of aircraft were purchased in the historical program, which sets an upper limit on the number of sorties which could be flown in an intended campaign. Figure 4 illustrates that the value of the last sortie is ideally equal to the wartime cost. Thus, the historical data will imply an upper limit on the value of the last sortie as well as a lower limit on the shaded area, the net wartime value.

It is important to note, when imputing value with this approach, that the intended campaigns for the F-16 and F-18, for example, are different from the intended campaigns for the Joint Strike Fighter. This, and other insufficiencies in the imputed value method, can only be addressed by applying judgment to the results of the model.

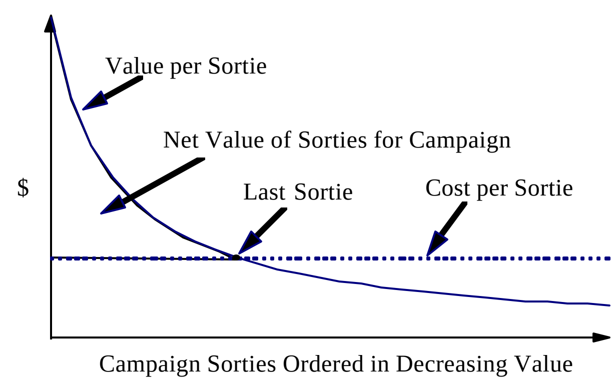

Figure 4: Wartime Value and Cost per Sortie

Judgment

Although the imputed value method provides rough guidance as to the value of military objectives and fixes some points on the value curves, it is unlikely to completely determine the curves for all the types of objectives in all the campaigns in which the Joint Strike Fighter is intended to be used. To finish defining the value curves (as in Figure 4), the JSF requirements team can exercise the model and judge whether reasonable numbers of sorties are being allocated in different situations. For example, if the model suggests flying only ten sorties against command, control and communication targets in a Persian Gulf campaign, the team might conclude that the objectives are undervalued, because in the actual Gulf War many more sorties were flown.

Comparison to Prior State of the Art

Wartime Value determined by the above methods will be imperfect and partly subjective. However, the alternative used today in effect sets Wartime Value to zero. Cost tradeoffs are done strictly on the basis of Peacetime Cost (life cycle cost), without quantitatively accounting for the benefits of, for example, low weight and fuel efficiency on range and payload capability in strike operations.

A good example of the improvement offered by the value model approach is in the quantitative assessment of stealth. The prior state of the art for cost based analyses was to count the reduction in combat losses due to improved survivability. The improvement was marginal, because the historical data do not show high loss rates. Squadrons avoid flying missions that put aircraft and crew in great jeopardy—the cost of such missions exceeds the value. The advantage of stealth, as demonstrated in the Gulf War and in Kosovo, is that aircraft can reasonably attack high value targets that would otherwise be too costly. The value model captures this quantitatively, allowing stealth to be traded rationally against weight, reliability, and (when v - c is maximized) manufacturing cost.

Contract Structure

The Joint Strike Fighter Program Office wants to construct a contract which provides incentives for the contractor to maximize v - p, the net value to the government, while providing a fair profit to the contractor. In particular, it is in the government's interest to focus the contractor's attention on v - c and the associated distributed optimization solution outlined above. Thus, the price formula which best produces the best incentives for the contractor to maximize v - p will directly incorporate the value model, $v(\vec{x})$.

Optimal Incentives

Given the model $v(\vec{x})$ described in equation (5) above, the Program Office wishes to construct a price function, $p(v(\vec{x}))$, so that when the contractor finds the most profitable design $\vec{x}^*$, which means

$$\text{for all } \vec{x},\ p\left(v\left(\vec{x}\right)\right) - c\left(\vec{x}\right) \le p\left(v\left(\vec{x}^*\right)\right) - c\left(\vec{x}^*\right) \quad (12)$$

then the contractor has also maximized the collective value to the government and the contractor:

$$\text{for all } x,\ v\left(\vec{x}\right) - c\left(\vec{x}\right) \le v\left(\vec{x}^*\right) - c\left(\vec{x}^*\right) \quad (13)$$

If p, v and c are smooth (continuous first order derivatives), then

$$\nabla p(\vec{x}^*) - \nabla c(\vec{x}^*) = 0 = \nabla v(\vec{x}^*) - \nabla c(\vec{x}^*) \quad (14)$$

Since c is unknown and could in itself locate $\vec{x}^*$ anywhere (imagine a deep pit), the only general solution is

$$\text{for all } \vec{x}, \nabla p(\vec{x}) = \nabla v(\vec{x}) \quad (15)$$

and this is only true if

$$\text{for all } \vec{x}, p(\vec{x}) = v(\vec{x}) - B \quad (16)$$

where B is a constant.

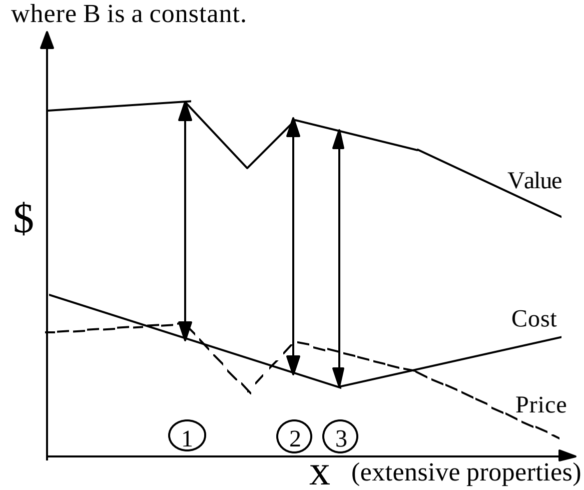

Figure 5: Optimal Contract Price

Figure 5 gives a visual illustration of the solution. Three possible optima are shown. Point 3 is the actual optimum. It is clearly due to the cusp in the cost curve at 3, but it is also due to the relative slopes (that is, divergence) of value and cost between 2 and 3. The price curve is an exact displacement of value, and the optimum of price minus cost is also at 3. Notice first that looking at the value curve alone, it would be impossible to identify 3 as the optimum, or even bound where the optimum is. Notice also that scaling price, or using any transformation of price to cost which changed the slope, say between 2 and 3, could make the optimum of price minus cost different from the optimum of value minus cost. Finally, the chart shows that p = v - B works even when v and c are not smooth.

The optimal pricing problem then reduces to the selection of B so that the contractor receives a fair profit and allowance for the contractor's attitude toward risk.

Risk

Decision maker's risk preferences are broadly described as risk neutral, risk averse or risk preferring. A typical defense contractor is somewhat more complex. With regard to profits, they can be modeled as risk neutral, but with regard to losses they are strongly averse. The defense community is accustomed to the protection of cost reimbursement, and contractors do not consider it wise to enter into arrangements in which they might lose money. Thus, protection from losses is assumed to be a minimum requirement for contractor participation. This is accomplished by giving the contractor a future option of receiving cost reimbursement whenever the value-based price is less than cost. That is, the contractor will choose, perhaps on a lot-by-lot basis, between accepting a price of v - B, or a price of c. The contractor will never lose money, and if the weapon system is designed so that v - B > c, the contractor stands to make money, perhaps a very large amount of money.

Contract Design

The proposed contract gives the contractor the option, on a lot-by-lot basis, of accepting a price of v - B, or a price of c. Ideally, B should be chosen so that the resulting price is a good deal for the contractor and the government. One such assignment is $B = V_1 - P_1$, where V_1 is the value, estimated at the time the contract is negotiated, and P_1 is the estimated price under a traditional cost plus incentive contract, estimated at the same time.

This is an enticing deal for the contractor, who will receive at the very least the currently estimated price, and is rewarded dollar-for-dollar for any improvement in v - c beyond the current projection. It is an excellent deal for the government, which has at least a

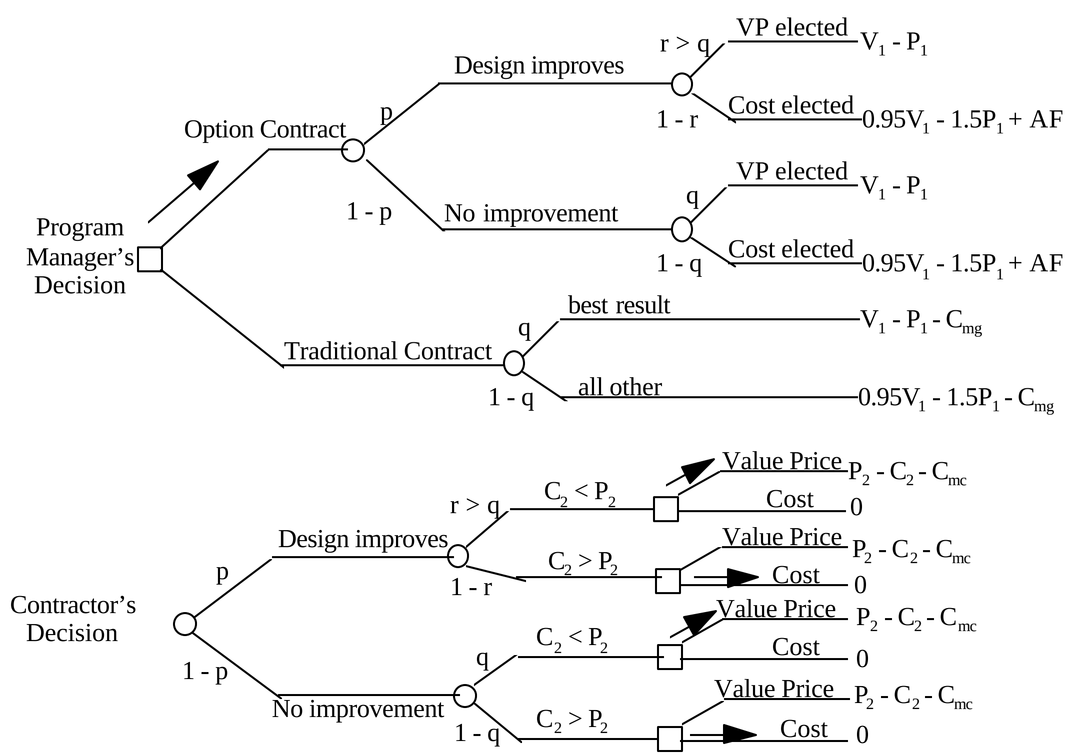

Figure 6: Decision Analysis of Contract Options

0.9 probability of receiving more benefit in any case where the contractor earns a profit, based on the distribution of outcomes plotted by Augustine [1983, p. 14].

Figure 6 shows the relative attractiveness of the value-based option to the government Program Manager and to the Contractor. Variables are defined as follows:

P_2 = actual value based price, V - (V_1 - P_1)

C_2 = actual cost to the contractor

AF = award fee incentive

C_{mc} and C_{mg} = costs of cost monitoring incurred by contractor and government on cost reimbursement contracts

p, q, r = probabilities (Augustine's chart [1983, p. 14] suggests q = 0.1)

Notice that even without assigning probabilities in Figure 6, the value priced contract is clearly a superior choice for the government, and the option of value pricing offers an attractive upside potential for the contractor with no downside potential.

Open Issues

This paper is a mere overview of the value-based approach to contracting. Among the issues which must be addressed to put this technique into practice are:

1. **Development cost**. Development might be funded by a guaranteed loan from the government to the contractor, which would be paid off through JSF sales. The value-price would include the $40 billion estimated Development Price in P_1. The loan would be forgiven if the government canceled the contract. The interest rate would be set at the government's discount rate, roughly the rate of the 30 year bond, so that the net value of the cashflow to the government is zero. Under this approach, the contractor would be strongly incentivized to minimize development costs.
2. **Learning curve**. The value-price should be higher on earlier units than on later units, in real dollars, because the earlier units will be more expensive to manufacture, and the optimal situation for all parties is when the contractor chooses to value-price every lot.

CONCLUSION—POTENTIAL SAVINGS

The potential savings are assessed in two steps. The first step calculates the current estimated program price. The second calculates the expectation of price increase and value decrease, and the savings achievable by locking in value minus price at current levels with the proposed contract structure.

JSF Development and Production Price

The estimated price of JSF development and production can be inferred from the following data [Department of Defense, 1995, p. E-13]:

Development Price = 15% of Life Cycle Cost

Production Price = 58% of Life Cycle Cost

Recurring Flyaway Price = 70% of Production Price = 41% of Life Cycle Cost

Average Unit Recurring Price = $35 million, in 1994 calendar year $

Lot Size = 2,807 aircraft

Using the Bureau of Labor Statistics estimate of change in the Producer Price Index from 1994 to 1998, the unit price estimate is $37 million in 1998 dollars. Multiplying by the lot size, the total Recurring Flyaway Price is $37 million x 2807 = $104 billion. Thus the total Life Cycle Cost = $104B / 41% = $250B. Therefore the Development and Production Prices, in 1998 dollars are 15% x $250B = $40B and 58% x $250B = $145B respectively. The total development and production contract price is estimated to be $185B in 1998 dollars.

Savings

Historical data on similar programs shows a mean cost growth in total development and production cost of 1.52 [Augustine, 1983, p. 19], or, on the JSF, a growth delta of 0.52 x $185B = $95B. Erosion in value is typically much smaller (Figure 7). Estimating total value conservatively at twice the original contract price (so that value exceeds actual life cycle cost, if barely), and estimating value erosion due to missed specifications during development at 5%, the loss is 5% x 2 x $185B = $18B. Thus, the total loss due to eroding cost and rising price forecast from traditional approaches is $110B.

If the contract described above is used, and if a strike fighter is designed so optimally that the contractor elects the value-priced contract option, the government can expect to save $110B. With no relevant knowledge as to which way it will go, the appropriate assessment is a 0.5 probability that the contractor will elect the value price option (the only option under which the contractor can earn a profit), and the expectation of savings is roughly $50B.

The conclusion is that value-pricing, and the optimal design to which it leads, provides as great an opportunity for government savings as any innovation in the Joint Strike Fighter program. In fact, the savings by themselves exceed the contract value of all but a few public programs in US history.

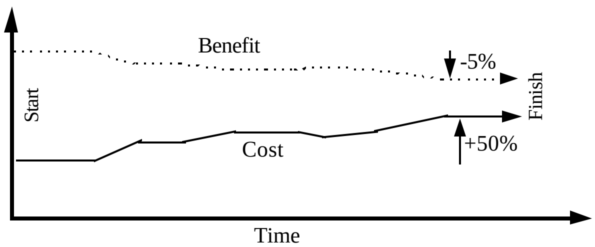

Figure 7: Typical Trends Of Cost Growth And Performance Erosion In Similar Programs

ACKNOWLEDGMENTS

This paper was initially developed in a seminar on decision analysis applications in national security, which was organized and directed by William J. Perry and John M. Shalikashvili at Stanford University in the spring of 1998. Two sessions of the seminar focused on Acquisition Reform, and these were led by Michael J. Lippitz. Dr. Lippitz encouraged the author to develop this paper for the seminar, where it was well received and strongly influenced by the seminar participants. In particular, the author would like to thank Bruce Floersheim, Norman R. Augustine, Michael J. Lippitz and William J. Perry for their constructive criticism and endorsement of this paper.

REFERENCES

Augustine, Norman R. Augustine's Laws. Second Edition, American Institute of Aeronautics and Astronautics, NY, 1983.

Department of Defense. JAST Operations Analysis Study Plan. Revision 1, Joint Strike Fighter Program Office, Arlington, VA, April 1995.

von Clausewitz, Carl. On War. Edited and Translated by Michael Howard and Peter Paret, Princeton University Press, Princeton, NJ, 1976. Originally published in German, 1832.