

A VALUE-BASED APPROACH FOR COMMERCIAL AIRCRAFT CONCEPTUAL DESIGN

Jacob Markish[†], Karen Willcox[‡]
Massachusetts Institute of Technology

Keywords: *aircraft design, value, dynamic programming, flexibility, uncertainty*

Abstract

An analytical tool is presented to perform program-level valuation of commercial aircraft designs. The algorithm used expands upon traditional Net Present Value methods through the explicit consideration of market uncertainty and the ability of the firm to react to such uncertainty through real-time decision-making throughout the course of the aircraft program.

The algorithm links three separate analytical models—performance, cost, and revenue—into a system-level analysis by viewing the firm as a decision-making agent facing continuous choices between several different “operating modes.”

An optimization problem is set up and solved using a dynamic programming approach to find a set of operating mode decisions that maximizes the firm’s expected value from the aircraft project. The result is a quantification of value that can be used to make program-level design trades and to gain insight into the effects of uncertainty on a particular aircraft design.

1 Introduction

Numerous methods have been developed to contribute to the aircraft design process. Many have focused on one or more technical disciplines, such as structures or aerodynamics, while others have addressed program-related parameters, such as cost, revenue, or schedule [1,2]. There have also been significant advances in the state-of-the-art for multidisciplinary technical analysis and optimization [3,4,5]. Few, however, have combined the technical and program-related analyses into a single design framework. The purpose of such a framework would be to yield a design process that uses program value as its objective function. Program value is

defined here as the amount a buyer would be willing to pay for the opportunity to invest in and build a particular aircraft design. Because a commercial aircraft is ultimately designed to generate value for its manufacturer, it is reasonable, albeit perhaps idealistic, to design the aircraft and its program to maximize its value as opposed to minimizing its weight or cost, or maximizing its revenue.

The purpose of this paper is to develop a framework that enables design for value. The framework consists of a quantitative tool used to quantify the value of the entire program associated with a given aircraft design. The tool relies on and links several analytical models that describe the performance, cost, and market-related characteristics of the aircraft in question. While the models are not described in depth here, this paper focuses on the algorithm used to link them and find a measure of program value.

The remainder of this paper is organized as follows. Section 2 formulates the aircraft program design problem. Section 3 introduces a stochastic dynamic programming (DP) approach, which is first described in general, theoretical terms, as applicable to a wide variety of problems, and then described in the context of a variation of DP, referred to as the “operating mode” framework. Next, Section 4 applies the DP approach specifically to the solution of the problem described in this introduction. Several variations and extensions are made to the basic DP approach, and a complete valuation algorithm is synthesized. It should be noted that the algorithm depends critically upon a set of models constructed to characterize the behavior of the endogenous and exogenous variables described above—roughly speaking, performance, cost, and revenue. These models are detailed in [6]. Finally, Section 5 uses a brief example to demonstrate the application of the algorithm, and Section 6 concludes.

[†] jmarkish@mit.edu

[‡] kwillcox@mit.edu

2 Problem Formulation

The basic problem to be solved may be broken up into three parts: endogenous variables—those that are internal to the aircraft development process and which may be controlled; exogenous variables—those that are external to the aircraft development process and may not be controlled; and a statement of the problem objective.

2.1 Endogenous Variables

The producing firm has a set (or “portfolio”) of aircraft designs, any of which it may choose to develop and bring to market. To bring a concept to market, the firm must go through several phases: detail design, tooling and capital investment, testing, certification, and finally production. Each phase entails some required expenditure of time and money, and the firm may decide, within reason, when to execute each phase. As for the aircraft designs, each is defined by a set of component parts (e.g., inner wing, outer wing, fuselage bay, etc.). Some parts may be common across several aircraft. This commonality results in potential cost impacts in both development and production.

2.1 Exogenous Variables

Given that an aircraft design is in production, the evolutions of sale price and quantity demanded per unit time are unaffected by any decisions made by the firm. Sale price evolves according to a steady growth rate, while quantity demanded evolves as a stochastic process, characterized by parameters such as drift rate and volatility. Each period that an aircraft design is in production, as many units are built and sold as are demanded, up to the maximum production capacity of the plant.

2.3 Problem Objective

Given the above endogenous and exogenous variables, which describe the aircraft development process and the market for the aircraft, the objective is to find a set of optimal decision rules governing (1) which aircraft to design, (2) which aircraft to produce, and (3) when; as a function of the demand level and the aircraft built to date at any given time. Achieving this goal will necessarily yield the overall program value, because program value is the objective function used to find the optimal decision rules.

3 Dynamic Programming (DP)

3.1 General Theory

A stochastic dynamic programming problem may, in general, be framed in five parts¹:

1. *State variables* continuously evolve and completely define the problem at any point in time.
2. *Control variables* are set at any point in time by the decision-maker, and generally impact the evolution of the state variables.
3. *Randomness*. One or more of the state variables is subject to random movements, and as such, involves a stochastic process.
4. *Profit function*. The goal of the dynamic programming method is to maximize some objective function, in this case the program value. The value is, in general, a function of certain “profits” incurred every period. These profits are functions of the state variables.
5. *Dynamics* represent the set of rules that govern the evolution of the state variables, including the effects of randomness, the effects of control variables, and any other relationships.

The problem is further defined by a time horizon (which may generally be finite or infinite), and a sequence of time periods of length Δt , which together comprise the time horizon. The objective, then, is to find the optimal vector of control variables as a function of time and state, such that the total value at the initial time (beginning of the time horizon) is maximized. Equivalently, the objective may be stated recursively, as an expression for the value at any time, t , as:

$$F_t(s_t) = \max_{u_t} \left\{ \pi_t(s_t, u_t) + \frac{1}{1+r} E_t[F_{t+1}(s_{t+1})] \right\} \quad (1)$$

where $F_t(s_t)$ is the value (objective function) at time t and state vector s_t ; π_t is the profit in time period t as a function of the state vector s_t and the control vector u_t ; r is some appropriate discount rate (addressed in Section 3.2 below); and E_t is the expectation operator, providing in this case the expected value of F at time $t+1$, given the state s_t and control u_t at time t . Note that the expectation operation for next period is affected by the control decision and the state in this period.

¹ This view of dynamic programming was suggested by D. Bertsimas, MIT Sloan School of Management.

The above is known as the Bellman equation, and is based on Bellman’s Principle of Optimality: “An optimal policy has the property that, whatever the initial action, the remaining choices constitute an optimal policy with respect to the subproblem starting at the state that results from the initial actions” [7]. In other words, given that the optimal value problem is solved for time $t+1$ and onward, the action (choice of u) maximizing the sum of this period’s profit flows and the expected future value is also the optimal action maximizing value for the entire problem for time t and onward.

The Bellman equation can therefore be solved recursively or, for a finite time horizon, iteratively. For a time horizon of T , this is done by first considering the end of the horizon, at time t_T . At this point, there are no future states, and no future expected value of F . Therefore, equation (1) reduces to

$$F_T(s_T) = \max_{u_T} \{\pi_T(s_T, u_T)\} \quad (2)$$

The optimal control decisions, u_T , given the final state, s_T , are readily found. This process is repeated for all possible final values of the state vector, s_T . Next, it is possible to take one step backwards in time, to $t = T - 1$. Now, equation (1) may be applied to find the optimal control decisions, u_{T-1} , because the expectation term, $E[\dots]$, is easily calculated as the probability-weighted average of the possible future values of F_T . Again, the optimal control values, u_{T-1} , are found for each possible value of s_{T-1} . At this point, the procedure is repeated by taking another backward timestep to $T - 2$, and continuing to iterate until the initial time, $t = 0$, is reached. At this point, the value F_0 is known for all possible initial values of the state, s_0 , and it is the optimal solution value.

3.2 Specific Application: Operating Modes

It is possible to extend the general DP framework presented above to a specific application useful for the valuation of projects. The application is centered around the concept of “operating modes,” and has been demonstrated by several authors to be useful in modeling flexible manufacturing systems [8,9]. Much of this description is based upon their work.

Consider a hypothetical factory, which at the beginning of any time period may choose to produce output A or output B. Let the prices for which it can sell each of the outputs be different functions of a single random variable, x , so it may be more profitable in some situations to produce one output

than the other. However, each time the factory switches production from A to B, or vice versa, a switching cost is incurred. Thus, it may not always be optimal to simply produce whichever output yields the higher profit flow in the current period. If there is a high probability of a switch back to the other output in the future, it may be preferable to choose the output with the lower profit this period.

This example lends itself well to the dynamic programming formulation. In this case, the control variable u_t is the choice of output, or “operating mode,” for the period beginning at time t : A or B. The state vector, s_t , consists of two elements: the random variable, x , and the operating mode from last period, m_t . The operating mode m will have one of two possible values, say 0 or 1, representing output A and B. Depending on the value of m_t , the control variable choice u_t may result in payment of a switching cost. Specifically, the Bellman equation may be re-written for this example as:

$$F_t(x_t, m_t) = \max_{u_t} \left\{ \pi_t(x_t, u_t) - I(m_t, u_t) + \frac{1}{1+r} E_t[F_{t+1}(x_{t+1}, u_t)] \right\} \quad (3)$$

Note that the state vector s has been separated into its two components—the random variable x and last period’s operating mode m . Here, the profit function term from equation (1) has been replaced by the difference between a profit flow and $I(m_t, u_t)$ —the switching cost from mode m to mode u . This will equal zero if $m_t = u_t$ and there is in fact no switch made, and will be nonzero otherwise. Note also that the future value, F_{t+1} (for which the expectation is found), is a function of the future random variable, x_{t+1} , and of the current control decision, u_t , because u_t will become the “operating mode from last period,” m_{t+1} , at time $t+1$. In other words, $m_{t+1} = u_t$, because as soon as the control decision (u_t) is made, the mode in which next period will be entered (m_{t+1}) is set.

As before, this equation can be solved iteratively by starting at the final time period and working backwards. Thus, the value of the factory project is found at time $t = 0$ as a function of the initial value of x and the initial operating mode m_0 . The value is arrived at by finding the optimal set of decisions u_t for all times t (starting at $t = T$) as a function of random variable x_t and “operating mode from last period,” m_t .

One additional point regarding equation (3) bears discussion: the selection of an appropriate discount rate, r . This is a nontrivial task; in fact, the selection of a discount rate is traditionally one of the

most difficult and sensitive steps in capital budgeting. For an in-depth discussion of discounting as it applies to this valuation technique, refer to [10].

While the factory example is very simplistic, the operating mode framework can be extended to any number of random variables and any number of modes. Operating modes may be chosen to represent not only production modes, but other decisions, such as waiting (doing nothing), abandoning the project for a salvage value, or investing in capital equipment to have the option of going into production at a later period. Each of these possible modes would have its own profit function associated with it, and its own set of switching costs to and from all other possible modes.

The operating mode framework forms the foundation for the approach outlined below, which applies dynamic programming to the aircraft program valuation problem.

4 Applying DP to the Aircraft Design Problem

4.1 Overview

The dynamic programming approach described above is adapted here to solve the problem of optimal decision-making in managing an aircraft program. As this problem is solved, the net value of the program is found—just as with any optimization problem, finding the value-maximizing independent variable(s) necessarily involves finding the associated maximum value.

The approach is presented in several steps: first, a connection is made to the general dynamic programming theory introduced in Section 3.1; second, a further connection is made to the specific application of dynamic programming to an “operating mode framework”; and finally, the entire algorithm is summarized.

4.2 Connection to General Theory

The aircraft program valuation algorithm is briefly overviewed here in the context of the five parts that frame a dynamic programming problem.

1. *State variables.* For each new aircraft design being simultaneously considered, two state variables exist: quantity demanded, which evolves stochastically, and the “operating mode from last period” (as introduced above) for that aircraft.

2. *Control variables.* For each aircraft design, one control variable exists: the choice of operating mode for the current period.
3. *Randomness.* For each aircraft design, one state variable exists with random characteristics: the quantity demanded. It evolves from a given initial value as a stochastic process.
4. *Profit function.* The profit function during each period is the sum of profits associated with the operating modes for each aircraft, less any switching costs incurred during that period. For production operating modes, the profits are simply revenues less recurring costs; however, other modes exist for which the profit functions represent non-recurring costs.
5. *Dynamics.* There are two types of state variables in this formulation: quantity demanded, which evolves as a stochastic process; and operating mode, which evolves as dictated by the control variables (operating mode decisions).

The time horizon, as defined in this application, is 30 years, which is a typical valuation timeframe for an aircraft program. For purposes of simplicity and computation time constraints, the time period length, Δt , was selected to be 1 year. For the same purposes, the maximum number of aircraft designs to be simultaneously considered by the algorithm was set to two.

The objective of the problem, then, is to find the vector of optimal control variables (operating modes), as a function of time and state, that maximizes the value of the program at time $t = 0$. This value is consistent with the definition proposed in the introduction: it is the price a potential buyer would be willing to pay for the opportunity to invest in the project defined by the aircraft design(s) and associated existing capital equipment.

4.3 Connection to Operating Modes

Whereas the operating modes introduced in Section 3.2 were simply modes of production, this formulation extends the operating mode framework to represent each phase of the lifecycle of an aircraft program. The purpose of this extension is to model the significant time and investment required to develop an aircraft, before any sales are made. Therefore, the non-recurring development process, which may last as many as six years, is represented as a chain of “operating modes.” Clearly, none of these modes entails a positive profit flow. Rather, each has some negative “profit” associated with the

non-recurring investment for that particular phase of the aircraft development cycle. The only incentive for the firm, and the optimizer, to enter one of these modes is the opportunity it creates to switch to the following development mode in the following period, and so on until the production mode is reached. A graphical representation of the operating modes for a single aircraft design is shown in Figure 1. The diagram is similar in concept to a Markov chain, where the arrows represent possible transitions between modes. In fact, the arrows connecting the modes represent switching costs that are finite—if two modes are not connected by an arrow, the associated switching cost is infinite.

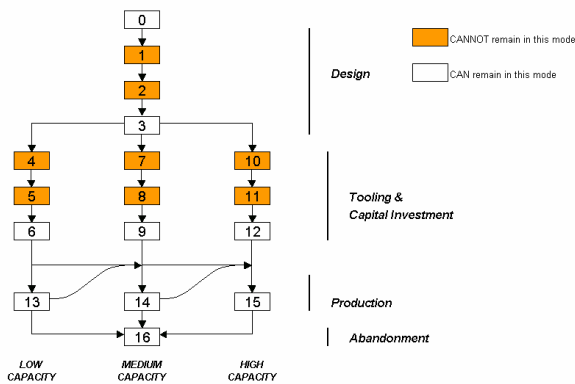


Figure 1. Operating mode framework for a single aircraft

Mode 0 represents the initial conditions: the firm is waiting to invest. Modes 1 through 3 represent roughly the first half of the development effort, primarily detailed design. Note that several operating modes are shaded. The shading indicates an infinite cost *not* to switch to a different mode. In other words, it is impossible to remain in a shaded mode for more than one period. Thus, once the firm commits to a detail design effort, it is assumed impractical to stop halfway through. However, it is possible to stop before the second half of development—here, mostly tooling and capital investment—begins. Once this development stage is initiated, a capacity choice must be made: a low, medium or high capacity production line. This determines the maximum demand level that may be satisfied with sales every period. Once the capacity choice is made, the firm must continue to switch modes annually until it reaches mode 6, 9, or 12, at which point it is ready to enter production. Recall that each time period has a duration of 1 year—therefore, if the firm does not wait midway through the development process, an aircraft design takes 6

years to bring to market². The production modes are 13, 14, and 15, corresponding to a low-, medium-, and high-capacity line. Note that each mode will produce exactly as many units as demanded each period, up to a maximum that depends upon the mode. The actual values for maximum capacity are parameters and easily changed. While in production (or waiting to enter production), it is possible to invest in additional tooling and expand the capacity of the production line. However, it is assumed impossible to reduce capacity—that is, the scrapping of tools has little to no salvage value, due to the high specificity of the tools to their product. Finally, an abandonment mode exists to model any salvage value associated with permanently shutting down the program. If the salvage value is positive, the switching costs to enter mode 16 will be negative.

The above overview of the operating mode framework has made no mention of the process by which the actual switching costs are to be found. The determination of switching costs is based upon the adaptation of a development cost model and a manufacturing cost model (as described in [10]) to the operating mode framework.

Development cost. The development cost model generates a time profile of the non-recurring expenses associated with the development of a given new aircraft design. The profile depends upon the aircraft’s characteristics, and also upon the existence of any of the aircraft’s component parts in other aircraft which have already been designed. Because of the inclusion of two aircraft designs in the valuation algorithm, the switching cost calculation proceeds as follows for each of the two aircraft designs. The sequential switching costs from mode 0 through mode 9 (medium production rate) are calculated by discretizing the non-recurring cost time profile into 1-year segments. This discretization is done as a step function of the operating mode of the other (remaining) aircraft design: if the other aircraft has not yet been fully developed (i.e., the other aircraft mode is less than 13—production), the baseline non-recurring cost profile is used. However, if the other aircraft has already been fully developed (i.e., the other aircraft mode is at least 13—production), the non-recurring cost profile is calculated with any commonality effects included. Specifically, if the aircraft share any common components, the development cost and

² In fact, the 6-year baseline duration may be altered, as discussed below, depending on previous design experience.

time are both reduced to reflect the savings resulting from a pre-existing design. See [6] for details. If the commonality effects are significant enough to result in a cost profile shorter than 6 years, one or more development modes are skipped (the candidate modes for skipping are 1, 2, 7, and 8). Finally, once medium capacity development process switching costs have been defined (as a step function of the other aircraft's mode), the switching costs corresponding to the tooling/capital investment part of the development process are scaled by a "low capacity" and a "high capacity" scaling factor to find the switching costs corresponding to the low and high capacity decisions (modes 4,5,6 and 10,11,12, respectively).

Manufacturing cost. As with development cost, to account for two aircraft designs present in the valuation, the switching costs associated with manufacturing are found for each aircraft as functions of the operating mode of the other aircraft. Switching costs associated with manufacturing have two components. The first component is switching from a "ready to produce" mode (6, 9, or 12) into the corresponding production mode (13, 14, or 15, respectively). The second component is switching from one production line capacity to another. Both of these components are sometimes involved in a single switch (e.g., 6 to 14, or 6 to 15); however, they are calculated separately and simply added together as necessary.

The costs of switching production line capacity are calculated as the product of a scaling factor³ and the cumulative difference in cash outflows between the two development processes associated with the production capacities in question. For example, the cost to switch from low to medium capacity (involved in either a "6 to 14" switch or a "13 to 14" switch) equals the difference in total development cost between low capacity development (3,4,5,6) and medium capacity development (3,7,8,9).

The other component of manufacturing-related switching costs is the initial switch into one of the three production modes (13, 14, or 15). For any given aircraft, the unit cost will generally fall as production starts and continues due to the learning curve effect. Eventually, it is reasonable to assume that unit cost approaches an asymptote and stabilizes at what is referred to here as "long-run marginal cost" (see Figure 2). To exactly model the effect of

the learning curve with dynamic programming would be impossible, because knowledge of unit cost requires a knowledge of how many units have been built to date. This information is not part of the state vector⁴. Therefore, once in a production mode, all aircraft are produced at their long-run marginal cost. Uncorrected, this assumption would introduce a significant error into the total cost incurred by the firm, of a magnitude approximated by area *A* in Figure 2.




Figure 2. Learning curve effect and long-run marginal cost (LRMC).

To account for this discrepancy, the switching cost to enter production is set equal to exactly the area *A*—that is, the total extra cost expected to be incurred during the production run of the aircraft over and above the long-run marginal cost. Because this extra cost will be incurred gradually and with certainty over time, the risk-free rate is used to find the expected present value of these cash flows, assuming a production rate equal to baseline demand. The switching cost is thus set to equal the present value of the cash flows represented by area *A*.

The entire above process, for both development cost and manufacturing cost, is conducted for both aircraft designs, resulting in a set of switching costs for each that is a function of the operating mode of the other. To use the symbology introduced in Section 3.2, the process finds the switching costs $I(m_i, u_i | m_j)$ for each aircraft i , where the "other" aircraft is j , for each prior operating mode m_i and control variable decision u_i . Then, the switching cost from any operating mode vector $[m_1, m_2]$ to $[u_1,$

³ The scaling factor, set to a value greater than 1, is meant to represent the additional costs incurred due to disruption of a pre-existing production line.

⁴ One possibility would be to include units built to date as an additional state variable, but computation time would suffer drastically as a result.

$u_2]$ is simply equal to the sum of $I(m_1, u_1 | m_2)$ and $I(m_2, u_2 | m_1)$. This set of data is stored in the switching cost matrix.

4.4 Stochastic Process Dynamics

A crucial component of the algorithm is the model describing the nature of the unpredictable behavior of the stochastic process representing the development of the market for commercial aircraft. This model is explained in detail in [10]. For the purposes of this paper, it is sufficient to state that the annual quantity demanded is represented as a stochastic variable that evolves according to a random walk, modeled as a binomial tree process. Each time period, the variable may either increase or decrease by a specified amount with a certain associated probability. Thus, if there are two such variables, representing the evolution of the market for two distinct aircraft, there are four possible outcomes each time period. A transition probability matrix is constructed linking possible initial states to final states in this framework. These transition probabilities are used to find the expected future value of the project as a probability-weighted average, expressed as the expectation term $E[\dots]$ in equation (3).

4.5 Algorithm Summary

The dynamic programming algorithm presented here uses a backwards-iterative solution scheme to find the set of optimal decisions maximizing the net discounted (present) value of the cash flows from the program. The decision-making facet of the approach is this algorithm's way of modeling managerial flexibility: the ability of the firm to control the program as it (and the market) evolves. Some specific decisions that were modeled include the decision to wait, to design, to invest in tooling at one of several capacity levels, to produce, and to abandon. In addition, the inclusion in the framework of multiple aircraft (in this work, exactly two) models product flexibility as well—the decision to produce one design over another, if not both; and the associated timing.

5 Example

While a complete demonstration of the program valuation tool described above is beyond the scope of this paper, a brief example is presented here to illustrate the mechanics of the algorithm and to highlight its distinguishing features. Although the algorithm has the capability to analyze two aircraft

designs simultaneously, the example considers a single design for simplicity and conciseness.

The notional vehicle used is based on a 250-seat Blended-Wing-Body (BWB) class aircraft [11], and the example relies on a set of assumptions listed in [10]. The discussion below presents two illustrations: a simulation run to demonstrate the decision rules arrived at by the optimizer and a connection to the Net Present Value technique shown by a plot of program value as a function of the initial forecast of annual demand.

5.1 Simulation Run

Figure 3 presents a simulation run, made after the algorithm itself was executed and an optimal solution was found. The simulation represents a sample path of demand through time—a hypothetical scenario constructed using a random number generator to approximate the stochastic behavior of demand. The upper half of the figure plots the random evolution of annual quantity demanded over time: this is a sample path of the underlying stochastic process. On the same plot is the optimizer's real-time strategy in response to the evolution of demand. This strategy consists of an annual selection of operating mode, based on the current year and the current demand level. Thus, at the beginning of the simulation, demand is at its baseline static forecast quantity, as calculated by the revenue model referenced above and described in [6]. This demand level, which happens to be 27.6 units per year, is insufficient for the firm to commit to developing the BWB (according to the optimal strategy). However, in year 3, demand increases as the result of a random fluctuation, and the choice is made to invest in non-recurring development for the aircraft. The investment choice is made because the new level of demand is greater than the threshold level corresponding to the "design" decision at time $t=3$ in the optimizer's solution. Recall that once the initial "design" operating mode is entered, the firm is committed to the first phase of the development process, until halfway through development, immediately before tooling. In this simulation run, demand falls immediately after design is started, but increases again when the halfway point is reached. Development is therefore continued, until time $t=9$, when the operating mode is 6 (end of development). At this point in time, demand is quite low, and the production decision is deferred.

However, one year later, at time $t=10$, demand increases past the threshold value for the "build" decision, and production is entered. The bottom half




Figure 3. Simulation run for BWB

of the figure shows the cashflows associated with the decisions made each year. Years 2 through 8 demonstrate the familiar bell-curve shape of a typical development effort. Year 10 shows why the optimizer chooses to wait at all before going into production: the switching cost to enter production is on the order of \$4B. This switching cost is the algorithm's way of handling the learning curve effect: the \$4B switching cost here is the present value of all the projected future costs in excess of long-run marginal cost for BWB production.

Once production is entered, after year 10, all units are produced at their long-run marginal cost⁵. The cash flows from production, in years 11 through 30, continue to fluctuate as a function of demand, and gradually creep upward with inflation. Returning to the upper half of the figure, the optimizer can be observed to respond to demand spikes in year 13 and then 17 by making incremental

investments in tooling to expand the capacity of the production line, first to a medium and then to a high level. In this simulation run, the high production capacity was put to good use only in year 17, as demand never reached that level again. However, the decision to enter high capacity production was optimal at that time, because the demand spike indicated a higher *expected* future demand.

The above simulation run is just one of millions of possible paths that can be taken by demand through time, but it effectively illustrates the decision-making element of the solution to the program valuation problem.

The actual expected program value corresponding to this solution is computed as \$2.26B. However, the magnitude of this value is not as important as the dynamics and approach illustrated by the valuation process. The number itself is strongly dependent upon the assumptions used in the underlying models, but the algorithm stands independently of the numerical results.

⁵ In reality, the \$4B would be distributed over the entire production run, with more weight on the early years.

5.2 Connection to Net Present Value

There is one primary conceptual difference between the dynamic programming approach used in this work and traditional project valuation approach of NPV: dynamic programming takes into account managerial flexibility, i.e. decision-making in real time. NPV analysis assumes a fixed schedule of actions and cash flows, and uncertainty regarding the magnitude of those cash flows is accounted for by appropriate selection of a discount rate. However, there is no uncertainty regarding which “operating mode” the firm using at any time—these decisions are made ex ante. Therefore, if the ability to make decisions is removed from the tool, it should reduce to a traditional NPV analysis. In other words, the switching costs between operating modes must be adjusted such that the optimizer has only one choice with a finite switching cost for any given operating mode. Referring to Figure 1, the only finite-cost path through the modes is now set as 0-1-2-3-10-11-12-15. This assumes an irreversible commitment, as of time 0, to design, tooling, and high capacity production.

Now, as the optimizer “solves” the problem, it is forced to make the same decisions regardless of the demand level. As a result, it is possible to generate negative program values, just as is it routine to find that a project has a negative NPV.

Figure 4 is a plot showing program value for the BWB as a function of the initial annual demand forecast. As detailed in [6], this demand value is a strong function of the characteristics of the aircraft—specifically, the range and seat count—but is also dependent upon the current condition of the market, and the resulting expectations and needs of the airlines. The plot thus considers the sensitivity of the program’s success to the current condition of the market. Demand is expressed as the number of aircraft per year that are demanded, and thus potentially sold, in year 1 of the analysis. This initial quantity is the starting point for the evolution of demand according to a stochastic over the time horizon of the problem. The simulation shown in Figure 3 above presents one possible sample path for this evolution, with the initial demand estimated at 28 aircraft per annum.

Figure 4 shows 2 plots of value on the same set of axes: “dynamic programming” and “Net Present Value”. The former shows the output of the algorithm as it finds the value of the program using dynamic programming to account for managerial flexibility. The latter is the NPV case described

above, where flexibility is removed from the program.

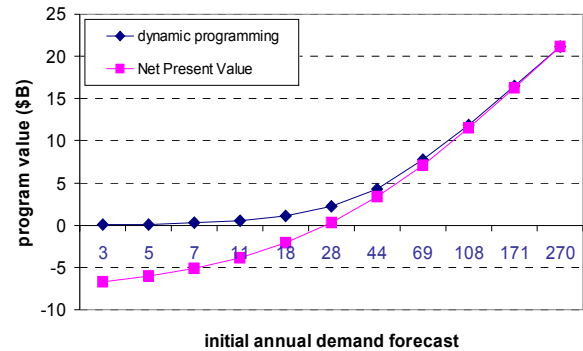


Figure 4. BWB program value.

As the initial demand index increases, the assumed baseline quantity of aircraft demanded per year increases. This quantity is the basis for the forecast of cash flows for the program. As the forecast increases, expected program value increases. If the forecast is very small, the value of the program with no-flexibility is negative—that is, the aircraft is developed, the non-recurring cost is incurred, but few if any units are sold. However, the value with flexibility for low demand indexes is zero—if no sales are expected, no investment is made in developing the aircraft.

Note that as the demand index increases, the no-flexibility program value quickly approaches value with flexibility. However, for small or marginal demand index numbers, there is a significant difference between the two valuations—one that may mean the difference between keeping a program and scrapping it. At the baseline initial demand, which happens to be 28 aircraft per year, the value with flexibility, \$2.26B, is almost seven times the value without flexibility, \$325M.

6 Conclusions

This paper provided an overview of an aircraft program valuation tool developed to conduct value-based trade studies and to gain insight into the value dynamics of aircraft design. The tool combines a performance model; a cost model; a revenue model; and a dynamic programming algorithm to measure the value of a set of aircraft designs to a firm. The value measurement is not based upon any technical characteristics per se, or any static forecast of cost and revenue, but on an analysis of an uncertain future, assuming that value-maximizing decisions

are made by the program's management as time goes on and uncertainty is resolved.

This approach captures the effect of flexibility, which has the potential of having a great impact on value. Flexibility is modeled and addressed by the dynamic programming "operating modes" formulation, which is an explicit method of formalizing and discretizing the decision-making process that is continuously ongoing for any project at any firm. The approach provides additional insight over traditional valuation techniques by its attempt to quantify the value created by flexibility.

If used with the right analytical models, the framework presented here can form the basis for a powerful tool that focuses the design process on the value of the entire system being created.

7 References

- [1] Reinhardt, 1973. "Break-Even Analysis of Lockheed's Tri Star: An Application of Financial Theory." *Journal of Finance*, Vol. 28, Issue 4, pp. 821-838, Sep. 1973.
- [2] Browning, Tyson R. "Modeling and Analyzing Cost, Schedule, and Performance in Complex System Product Development." PhD Thesis, MIT, 1998.
- [3] Zang, T. A., Green, L. L. "Multidisciplinary Design Optimization Techniques: Implications and Opportunities for Fluid Dynamics Research.." 30th *AIAA Fluid Dynamics Conference*. Norfolk, VA, AIAA Paper 99-3798, June 1999.
- [4] Kroo, I. "MDO Applications in Preliminary Design: Status and Directions," AIAA Paper 97-1408, 1997.
- [5] Wakayama, S., "Multidisciplinary Optimization of the Blended-Wing-Body." AIAA Paper 98-4938, Sep. 1998.
- [6] Markish, J., Willcox, K. "Multidisciplinary Techniques for Commercial Aircraft System Design." 9th *AIAA/ISSMO Symposium on Multidisciplinary Analysis and Optimization*. Atlanta, GA, AIAA paper 2002-5612, Sep. 2002.
- [7] Dixit, A. K. and Pindyck, R. S. *Investment Under Uncertainty*. Princeton University Press, 1994.
- [8] Trigeorgis, L. *Real Options: Managerial Flexibility and Strategy in Resource Allocation*. The MIT Press, 2000.
- [9] Kulatilaka, N. "Valuing the Flexibility of Flexible Manufacturing Systems." *IEEE Transactions on Engineering Management*, Vol. 35, No. 4, pp. 250-257, Nov. 1988.
- [10] Markish, Jacob. "Valuation Techniques for Commercial Aircraft Program Design." S.M. Thesis, MIT, 2002.
- [11] Liebeck, R.H. "Design of the Blended-Wing-Body Subsonic Transport." 2002 *Wright Brothers Lecture*, AIAA paper 2002-0002, Jan. 2002.