

Valuation Techniques for Commercial Aircraft Program Design

by

Jacob Markish

S.B., Aeronautics and Astronautics (2000)
Massachusetts Institute of Technology

Submitted to the Department of Aeronautics and Astronautics
in Partial Fulfillment of the Requirements for the Degree of
Master of Science in Aeronautics and Astronautics

at the

Massachusetts Institute of Technology

June 2002

© Massachusetts Institute of Technology, 2002
All rights reserved.

Signature of Author_____

Department of Aeronautics and Astronautics
May 10, 2002

Certified by_____

Dr. Karen Willcox
Assistant Professor, Department of Aeronautics and Astronautics
Thesis Advisor

Accepted by_____

Wallace E. Vander Velde
Professor of Aeronautics and Astronautics
Chair, Committee on Graduate Students

Valuation Techniques for Commercial Aircraft Program Design

by

Jacob Markish

Submitted to the Department of Aeronautics and Astronautics
in Partial Fulfillment of the Requirements for the Degree of
Master of Science in Aeronautics and Astronautics

ABSTRACT

This research considers the commercial aircraft design process from the perspective of program value. Whereas traditionally, the conceptual design of aircraft has often focused on minimum weight, or sometimes minimum cost, this approach demonstrates the feasibility and usefulness of design based on maximum value to the aircraft manufacturer.

A program valuation tool is developed and demonstrated that measures the overall program value associated with a set of either one or two new aircraft concepts. The tool is based on a combination of a performance model; a development and manufacturing cost model; a revenue model; and a dynamic programming-based algorithm that accounts for uncertainty in future market conditions and the program's ability to cope with such uncertainty through real-time decision-making.

The cost model, using a component-based representation of the aircraft, allows for the consideration of the effects of part commonality on development and production costs. The revenue model, based on an analysis of existing commercial aircraft, estimates a market price and demand forecast for a new aircraft based on several key characteristics. The dynamic programming algorithm, used to find program value, treats annual aircraft quantity demanded as a stochastic process, evolving unpredictably with time. The algorithm borrows from Real Options theory to discount future cash flows using risk-neutral expectations and models the aircraft program as an actively managed project with real-time decision-making to maximize expected program value.

Several examples are drawn from the Blended-Wing-Body aircraft concept to demonstrate the operation of the program valuation tool. The results suggest that the value of part commonality between aircraft may be strongly sensitive to the weight penalty and increased fuel burn resulting from a common derivative design. More generally, the example results illustrate the usefulness of the explicit consideration of flexibility in program valuation and the feasibility of a conceptual aircraft design tool based on the metric of program value.

Thesis Supervisor: Karen Willcox

Title: Assistant Professor, Aeronautics and Astronautics

Table of Contents

Chapter 1.	Introduction	13
Chapter 2.	Value-Based Design.....	17
2.1	Definition & Motivation	17
2.2	Finding and Using a Metric for Value	21
Chapter 3.	Uncertainty	25
3.1	Introduction	25
3.2	The Nature of Uncertainty	25
3.3	Modeling Uncertainty	26
3.3.1	Background.....	26
3.3.2	Weiner Process	27
3.3.3	Generalized Weiner Process.....	28
3.3.4	Ito Process	29
3.3.5	Random Walk	29
3.4	Conclusion.....	32
Chapter 4.	Flexibility	35
4.1	Introduction	35
4.2	The Nature of Flexibility.....	35
4.3	The Shortcomings of Net Present Value.....	37
4.4	Valuing Flexibility.....	39
4.4.1	Background.....	39
4.4.2	Black-Scholes Formula	40
4.4.3	Risk-Neutral Valuation	41
4.5	Conclusion.....	43

Chapter 5.	Program Valuation Tool: Sub-Components	45
5.1	Introduction	45
5.2	Performance Modeling.....	48
5.2.1	Overview	48
5.2.2	WingMOD	48
5.3	Cost Modeling	51
5.3.1	Overview	51
5.3.2	Development Cost.....	53
5.3.3	Manufacturing Cost.....	63
5.3.4	Summary	67
5.4	Revenue Modeling.....	68
5.4.1	Overview	68
5.4.2	Static Demand.....	70
5.4.3	Stochastic Demand.....	83
5.4.4	Summary	93
5.5	Conclusion.....	93
Chapter 6.	Program Valuation Tool: Synthesis	95
6.1	Introduction	95
6.2	Problem Formulation: Aircraft Program Design.....	95
6.3	Solution Approach: Dynamic Programming (DP)	97
6.3.1	General Theory	97
6.3.2	Specific Application: Operating Modes.....	99
6.4	Applying DP to the Aircraft Program Design Problem.....	101
6.4.1	Connection to General Theory.....	101

6.4.2	Connection to Operating Modes	102
6.4.3	Stochastic Process Dynamics and Risk-Neutral Expectations	107
6.4.4	Algorithm Review.....	112
6.5	Conclusion.....	113
Chapter 7.	Examples	115
7.1	Introduction	115
7.2	Background: BWB.....	117
7.3	Baseline Analysis	121
7.3.1	BWB-450.....	122
7.3.2	BWB-250C	125
7.3.3	BWB-250P	129
7.3.4	BWB-450 & BWB-250C	129
7.3.5	BWB-450 & BWB-250P.....	131
7.4	Connection to Discounted Cash Flow (NPV)	132
7.5	Sensitivity Analysis	134
7.5.1	CAROC	135
7.5.2	Demand Correlation.....	136
7.5.3	Volatility.....	136
7.6	Conclusion.....	138
Chapter 8.	Conclusions	141
References	147
Appendix:	Inputs Spreadsheet	151

List of Figures

Figure 1. Shareholder value as a driver of system design'	20
Figure 2. Random walk movement during time period Δt	30
Figure 3. Binomial tree representation of geometric Brownian motion.	31
Figure 4. Random walk representation of geometric Brownian motion.	32
Figure 5. Architecture of the program design and valuation tool	47
Figure 6. WingMOD representation of Blended-Wing-Body aircraft.	49
Figure 7. Non-recurring hours for a commercial aircraft project: the data	55
Figure 8. Non-recurring hours for a commercial aircraft project: the model	56
Figure 9. Cumulative non-recurring cost committed: the model	57
Figure 10. Linearized DAPCA IV model output: Total non-recurring cost	58
Figure 11. Estimated operating empty weight fraction breakdown for a typical commercial aircraft	60
Figure 12. Representative non-recurring cost breakdown by parts for a typical commercial aircraft	61
Figure 13. Representative recurring cost breakdown by parts	66
Figure 14. CAROC trend as a function of seat count: narrow body aircraft	74
Figure 15. CAROC trend as a function of seat count: wide body aircraft	74
Figure 16. Price model estimates compared to actual prices: narrow bodies	76
Figure 17. Price model estimates compared to actual prices: wide bodies	76
Figure 18. CAROC fuel cost fraction trend	79
Figure 19. Commercial aircraft forecasted deliveries: 2000-2019	81
Figure 20. Supply and demand: classical model	84
Figure 21. Narrow body aircraft historical prices	85

Figure 22. Wide body aircraft historical prices	86
Figure 23. Narrow body aircraft deliveries	86
Figure 24. Wide body aircraft deliveries.....	87
Figure 25. Price and quantity movements for wide body aircraft	88
Figure 26. Supply and demand: alternate model	89
Figure 27. Total deliveries for narrow bodies and wide bodies	91
Figure 28. Operating mode framework for a single aircraft.....	103
Figure 29. Learning curve effect and long-run marginal cost (LRMC).....	107
Figure 30. Blended-Wing-Body (BWB) aircraft	118
Figure 31. BWB example configurations.....	121
Figure 32. Decision rules for BWB-450, baseline	123
Figure 33. BWB-450 program value.....	124
Figure 34. Decision rules for BWB-250C, baseline	126
Figure 35. Simulation run for BWB-250C	127
Figure 36. Program value for BWB-450 & BWB-250C.....	130
Figure 37. No-flexibility program value for BWB-450, baseline assumptions.....	133
Figure 38. BWB-450 program value at time = 0	134
Figure 39. BWB-450 & BWB-250C program value: effect of correlation.....	136

List of Tables

Table 1. Key WingMOD Outputs.....	51
Table 2. Non-recurring effort process parameters	56
Table 3. Non-recurring part costs per pound.....	62
Table 4. Non-recurring cost reduction factors.....	63
Table 5. Manufacturing process parameters.....	64
Table 6. Recurring part costs per pound.....	67
Table 7. CAROC-neutral price regression parameters.....	72
Table 8. Lifecycle cost increment parameters	73
Table 9. CAROC fuel cost fraction calculations	78
Table 10. Seat category definitions and data.....	82
Table 11. Estimated parameters for stochastic behavior of demand	92
Table 12. Key input parameters used for all cases.....	116
Table 13. Example: “demand indexes”	117
Table 14. Selected features of BWB example designs.....	120
Table 15. Baseline analysis: key inputs.....	122
Table 16. Baseline analysis: key outputs.....	132
Table 17. Program value: results for CAROC sensitivity.....	135
Table 18. Reduced volatility: key inputs.....	137
Table 19. BWB-450 program value: effect of volatility	137
Table 20. BWB-450 and BWB-250C: effect of volatility.....	138

Acknowledgements

This work would have been impossible without the help and support of many people, some of whom I knew before I started my research, and many of whom I met through the research. I am indebted to all of them.

My research sponsors, NSF and LAI, provided an ideal level of support for the research while allowing me the freedom to explore wide-ranging topics. Furthermore, the research experience and quality of the work were both improved by the opportunities I had to continually interact with professors, lecturers, and researchers at MIT and to learn from their insights and experience. Fred Stahl, Al Haggerty, Joyce Warmkessel, Ian Waitz, Joo-Sung Lee, Dimitris Bertsimas, John Hansman, J. P. Clarke, Cindy Barnhart, Myles Walton, and Annalisa Weigel all gave me much-welcomed advice and feedback to help formulate my research framework.

Several individuals from outside of my immediate academic environment also contributed their knowledge and their thoughts to the development of this project: Ilan Kroo, Greg Bell, Nalin Kulatilaka, and particularly Dave Anderson, who expressed unabated interest in my thesis and greatly improved my understanding of the aircraft development process.

The foundation for this research, as well as much of its development, is owed to the Blended-Wing-Body group at Boeing Phantom Works. While the entire group deserves acknowledgement, I am especially thankful to Sean Wakayama, John Allen, Matt Wilks, Derrell Brown, Blaine Rawdon, Joshua Nelson, Jennifer Whitlock, Bill Vargo, Richi Gilmore, and Mary Pieper for giving me invaluable insights and a great experience at Boeing.

I was lucky enough to have four individuals who were all, to a greater or lesser extent, my thesis advisors. Each of them has amazed me with his or her knowledge, patience, and desire to make this work achieve its full potential. Earll Murman is the person without whom this thesis could not have become a reality. His untiring and relentless efforts to help were responsible for the conception of my research, making connections with many of the people above, maintaining a clear focus, and synthesizing a completed product. Karen Willcox continually surprised me with her approachability and friendliness, and the willingness and ease with which she engaged new and unfamiliar topics. Her intelligence and character made it a pleasure to work with her. I am very grateful to Stewart Myers for the time and energy he devoted to ensuring the soundness of this work from both a financial and a general academic standpoint. Finally, Bob Liebeck has been a stalwart proponent of my work as well as a great advisor, bringing to bear the breadth of his vast experience.

A special note of thanks goes to Michelle McVey, who has made the entire graduate research experience more enjoyable as a great friend, and who has made my thesis more focused as an esteemed colleague.

Finally, I wouldn't be in a position to write this work without the unfailing support and love of my mother and my sister, both of whom I thank for their caring and patience.

Chapter 1. Introduction

The design of a large commercial aircraft is a daunting task. It represents the synthesis of a staggering array of technologies, concepts, materials and subsystems into one functioning whole. A completed commercial aircraft is one of the most complex systems in operation today. Furthermore, its design is rendered all the more complex, because a commercial aircraft is not only an engineering system which performs a set of specified functions. It is also a value-creating product for its manufacturer, and a revenue generation instrument for its operator, the airline. A technically superb aircraft does not necessarily equate to a successful one. Thus, there is an imperative to design commercial aircraft not for maximum range or maximum speed or maximum payload, or even for minimum cost, but rather for maximum value.

The design for value imperative has been addressed in various forms by several authors. Reinhardt (1973) discusses a historical case study: the valuation techniques used for the Lockheed Tri Star commercial aircraft, and their effects on the program. Dickinson et al. (1999) discuss portfolio optimization—the problem of managing multiple interdependent development projects—and illustrate several aircraft development techniques used at Boeing to maximize value. Slack (1998) discusses the concept of value directly, and Browning (1998) addresses value through his consideration of the interactions between cost, schedule, and performance.

The field of technical aircraft design is very well-established and thoroughly documented. The problem of meeting a set of technical requirements with a system design has been solved hundreds of times by engineers developing new aircraft. Recently, there have been advances in performing these design tasks from a multidisciplinary standpoint, where several different types of analyses are combined into a simultaneous optimization process. Several examples of these multidisciplinary analysis and optimization techniques may be found in the work of Kroo (2000), Baker and Mavris (2001), and Wakayama (1994, 1996, 1998). The disciplines covered in the optimization frameworks

of the above examples include structural mechanics, aerodynamics, propulsion, stability and control, and even some cost measures.

While there exist numerous studies of value, and certainly a great body of knowledge on technical aircraft design, few have combined the two areas into a single design framework.

In general, the very mention of the word “value” begs two questions. First, value for whom? Second, how is value defined and how can it be measured? The answer to the first question is as follows. This study is conducted from the perspective of the aircraft designer, which is presumed to be the same entity as the manufacturer. Value is therefore considered strictly from the point of view of the firm developing the aircraft¹. Furthermore, value is defined in quantitative terms, based on economic principles. The answer to the second question, how to define and measure value, is more involved and subject to discussion. The definition of value is addressed in Chapter 2.

The measurement of value is clearly a crucial step in developing a maximum-value design. Such measurement must encompass not just the airframe, but the entire program. From design, tooling, and testing through production and sales, the entire process of bringing the aircraft to market must be considered to find the impact of the aircraft design on the firm. In order to fully capture this impact, two elements present in all major projects are introduced: uncertainty and flexibility.

Uncertainty, the result of the unpredictable nature of the future, cannot be ignored in designing or valuing any project. It is discussed in Chapter 3 with several mathematical representations of stochastic processes.

Flexibility is the ability of a firm or product to adapt and react to an uncertain future. The contribution of flexibility to a project’s value is discussed in Chapter 4. Borrowing from Real Options theory, a foundation is developed for measuring the value of a flexible project in the face of uncertainty. Real Options is a much-studied but still evolving field,

¹ Value to the customer also plays a crucial role in the success of any aircraft program. This value is usually coupled to the manufacturer’s value such that it is impossible to have one without the other.

with continuing theoretical research by authors such as Trigeorgis (2000), Kulatilaka (1998, 1993), Amram (1999), Dixit and Pindyck (1994), and others. Several efforts have also been made to apply Real Options to project valuation in practice (see Stonier, 1999, and Neely, 1998). However, there still remains much unexplored territory in the application of numerical techniques to valuation and project design.

Having defined value as a design metric and having identified the importance of uncertainty and flexibility in measuring value, an aircraft program valuation tool is constructed to demonstrate the value-based design approach. Chapter 5 outlines the tool as a combination of several standalone analytical models, capable of estimating the performance, cost, and revenue characteristics of a potential new aircraft design.

Chapter 6 describes the algorithm used to combine all three models and incorporate the concepts of uncertainty and flexibility introduced in Chapters 3 and 4. The algorithm linking the separate models into a complete program valuation tool uses dynamic programming to find an optimal strategy for bringing the potential new aircraft to market—that is, a strategy for managing the aircraft program. A set of assumptions is made to simulate the continuous decision-making process governing the aircraft program execution, and an optimal set of decision rules is found such that program value is maximized. Thus, a measure of overall program value is found for a given set of one or more new aircraft concepts.

To demonstrate the operation of the program valuation tool developed in Chapter 6, several new aircraft concept examples are presented and evaluated in Chapter 7. The examples are based upon a Blended-Wing-Body aircraft concept and illustrate a case of commonality across aircraft. The value of commonality in commercial aircraft has been explored both qualitatively and quantitatively by several authors, including Fujita et al. (1998), Nuffort (2001), and Willcox and Wakayama (2002).

Because the examples lack the high fidelity and accurate data available on a proprietary basis at aircraft manufacturers, they do not represent actual valuations. However, they provide a useful illustration of the program valuation tool and suggest several interesting insights into the dynamics of flexibility, commonality, and aircraft program value. Conclusions and opportunities for further research are summarized in Chapter 8.

Chapter 2. Value-Based Design

2.1 Definition & Motivation

In any problem-solving process, one of the most important steps is that of defining the problem and setting the scope. In a typical iteration of an engineering design cycle, this step is addressed by specifying a set of constraints, a set of design variables, and an objective function. The fundamental problem is to select the design variables such that all the constraints are satisfied and the objective function is optimized. For aircraft design, the constraints may include range and payload; the design variables may include vehicle geometry; and the objective function may be gross takeoff weight.

However, in a broader context, it is not immediately obvious what the appropriate objective function should be for designing a commercial aircraft. In practice, the formulation of the objective function is different for each stakeholder in the project, and the differences are often difficult to reconcile. To an aerodynamicist designing the wing planform, the objective function may be to maximize the vehicle lift-to-drag ratio. To a payloads engineer laying out the interior, it may be the payload capacity, or number of passenger seats. To the vehicle configurator, it may be to minimize weight, while to the cost estimator, it may be to minimize development or manufacturing cost. Finally, at a higher level, the objective function to the firm may be to maximize profit from the entire project, while to the customer (e.g., an airline), the objective function may be to maximize the net present value of the profit streams from the aircraft over its service lifetime.

For some of the above stakeholders, the constraints as well as the objective function are easily quantifiable—for example, range, fuel burn, drag, or gross weight. In these cases, the design process is relatively well established. There exists a set of effective techniques by which tradeoffs are conducted between the design variables, and an “optimal” design is reached which maximizes (or minimizes, as appropriate) the associated objective function. These techniques are sometimes executed by the designer, with the help of

discipline-specific tools; and they are sometimes integrated into a multidisciplinary optimization tool, which conducts the tradeoffs through an explicit set of optimization algorithms, often in a more automated fashion. Multidisciplinary optimization typically addresses its design problem from a broader, system-oriented scope, but it does so at the cost of lower fidelity compared to discipline-specific approaches.

For other stakeholders, like shareholders in the firm producing the aircraft, or the airline customers purchasing the aircraft, the constraints and objective are more difficult to quantify and characterize. Nevertheless, the associated tradeoffs represent a crucial element of commercial aircraft design. These are the tradeoffs between program-related parameters: cost, price, production quantity, and timing. Production quantity and timing are design variables in their own right, while cost and sometimes price may be thought of as characteristics of the aircraft that are determined by a combination of design variables. Clearly, these four factors are critical to the success of any commercial aircraft design. However, it would be incorrect to trade them in isolation. A proper evaluation of any such tradeoffs may only be made in the context of the entire system being considered—that is, the aircraft design itself and the associated program plan. The program plan represents the strategy and structured set of decisions by which the firm will bring a new aircraft to market, produce it, and sustain it, while the aircraft design determines the resources required to develop it, produce it, and operate it. All program-related parameters are dependent upon or constrained by technical parameters.

Traditionally, the program-related parameters of cost, price, quantity, and timing have been analyzed separately from the technical parameters. The analyses have often been conducted by different groups and at different times. Although the intent has typically been to complete a design by “closing the loop” for the whole system, the cycle time for this “loop” is fairly high, and the effect of uncoupling technical design and analysis from program-related design and analysis is a sub-optimization at each level. Such sub-optimization is the result of a failure to simultaneously consider the entire set of design variables, and sometimes a focus on the wrong objective altogether. For example, aircraft gross take-off weight is often used as an objective function in the conceptual design phase, with the goal of minimizing weight for a given mission. In this situation, weight is being used as a proxy for cost. However, such an approach suffers from several

inaccuracies. Cost is not perfectly correlated with weight; thus they should not be treated as a single parameter. Further, cost is not the only applicable system-level design metric—as mentioned above, the others include price, production quantity, and timing. Achieving optimality for the latter variables may require a sacrifice in cost. Ultimately, an aircraft design effort should not seek to minimize weight, or even minimize cost, but to maximize value.

If an appropriate method to measure “value” can be agreed upon by all the stakeholders involved in the design of an aircraft, design trades may be made in the context of the entire system, rather than a small portion thereof. However, there remains the question of how to define the term “value” in the first place.

Figure 1 illustrates a framework for two distinct kinds of value: customer value and shareholder value. Customer value, which is some measure of the utility a customer gets from a product, is derived from the product’s quality, the timeliness of its availability, and its price. In turn, shareholder value is derived from the firm’s net income, which is a function of the cost of goods sold and of the revenue, which is directly influenced by the customer value. Note that quality, cost, timing, and price are collectively referred to as “system design.”

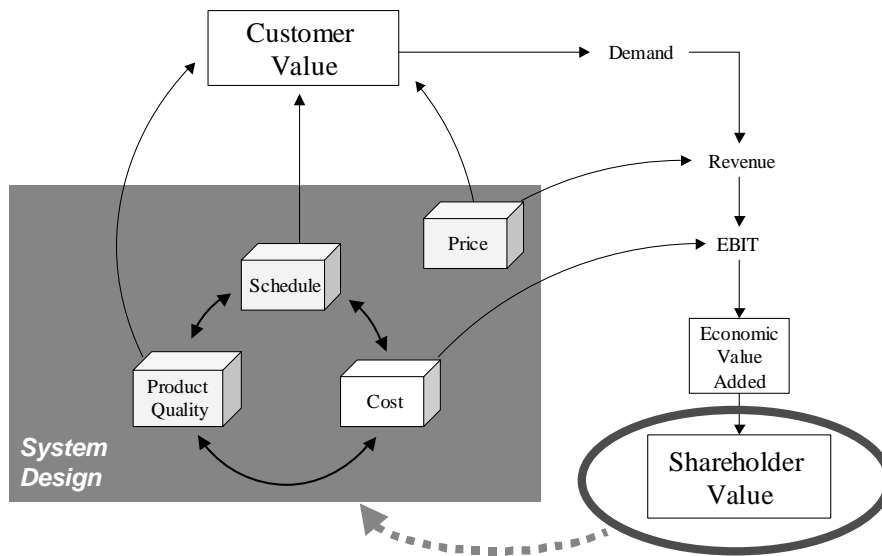


Figure 1. Shareholder value as a driver of system design^{2,3}

The assumption is made here that any commercial aircraft project has the objective of improving or sustaining the designing firm's financial value—that is, the shareholder value. Based on this assumption, it is reasonable to draw the conclusion that shareholder value should be the fundamental driving force behind system design.

Shareholder value, then, is the ultimate objective function from the perspective of the firm designing the aircraft. One may argue that such a narrow view is incorrect, and point out that it is a recipe for disaster for a firm to ignore value to other stakeholders—for example, customers, employees, and society. However, any solution maximizing shareholder value will necessarily need to address the remaining stakeholder values, *if* their effects are properly modeled. Clearly, if customers are unsatisfied, it is impossible for a firm to gain maximum benefit from selling its product. Similarly, if the firm's employees or its community are unsatisfied, the firm will face obstacles to growth and profitability, such as (in the extreme cases) strikes or legislative action.

To address all the dynamics of shareholder value is outside the scope of this work. In fact, if every element of value could be captured in a simple quantitative model, human

² Adapted from Slack, 1998.

³ EBIT = Earnings Before Interest and Taxes

management would be obsolete. A more modest approach is taken here by making a set of generous assumptions and limiting the scope of the problem.

Employee value and societal value, while very important to any firm, are not explicitly considered. Further, value to suppliers and partners of the firm is also not considered. Rather, the simplifying assumption is made that the same firm that designs the aircraft also develops, tests, and builds the entire vehicle, including all of its sub-components. While this is not true in practice, the trends resulting from analysis are presumed to be similar. Supply chain design and management is an important contributor to shareholder value, notably so in the aerospace industry, and to quantify its contribution is a separate topic, beyond the scope of this work. Given the above simplifying assumptions, the problem being addressed reduces to the following: find the aircraft, or set of aircraft, along with decision rules for bringing them to market, such that shareholder value is maximized.

2.2 Finding and Using a Metric for Value

Before any attempt can be made at finding a value-maximizing solution to the problem outlined in Section 2.1, a consistent and accurate method must be defined for measuring value. Several candidates exist for value metrics, each with its own set of rules and assumptions. While each of the candidates listed below has been used at one time or another for project valuation by firms, this analysis is built upon Net Present Value (NPV) as the baseline metric. The metric is later expanded to relax several of the assumptions of NPV valuation, but its foundation remains the same: the “value” of a project is defined as the market price the project would have today if it were a tradeable asset. A “project” is defined a set of one or more conceptual designs for new aircraft, complemented by (1) corresponding development and production strategies and (2) capital assets to bring the aircraft to market and generate sales.

Thus, the value of a project is the price a buyer would be willing to pay for the opportunity to invest in the project—that is, for the right to develop and build that aircraft⁴.

The following is a brief description of several alternate metrics considered for evaluation and comparison of projects⁵. NPV was the metric selected from this list as the most appropriate and useful measure of value.

Payback Period

This metric evaluates programs based upon the projected length of time before the entire initial investment is recovered through revenue. It suffers from several drawbacks, including an insensitivity to the order in which cash flows occur as the investment is being recovered—that is, an insensitivity to the time value of money—and a complete lack of consideration for cash flows occurring after the investment is recovered.

Net Present Value (NPV)

The Net Present Value (NPV) metric is based upon the existence and accuracy of a discount rate. This rate is used to discount all forecasted cash flows to reflect the opportunity cost of capital, and NPV is computed as the sum of all the discounted cash flows of a project. While the appropriate discount rate is often difficult to rigorously calculate and confirm, the NPV method is consistent across different projects and provides a good baseline for value measurement.

Internal Rate of Return (IRR)

The IRR metric, like NPV below, is based upon discounted cash flow (DCF) analysis, which reflects the time value of money as an opportunity cost of capital. The IRR is defined as the discount rate at which the sum of the discounted (present) values of the cash flows from a project is zero. If the IRR is greater than the opportunity cost of

⁴ One caveat to the simplicity of this definition is the need for a set of assumptions about existing capital and intellectual infrastructure already available to implement the project—design facilities and tools, engineering and other human resources, production facilities, etc.

⁵ The descriptions of the metrics presented here borrow from Brealey and Myers, 1996.

capital, traditional capital budgeting dictates that a project should be undertaken. However, IRR suffers from several difficulties of implementation, including inconsistent values for unusual cash flow sequences; occasional multiple IRR solutions for the same project; and difficulties in comparing projects of different magnitudes or durations.

Decision Tree Analysis (DTA)

The decision tree value metric adds a new dimension to valuation: it considers one or more decisions that may be made midway through the project based on some random event occurrence—for example, the market for widgets going up or down. Assuming the probabilities associated with this random future event are known, different cash flows are tied to different outcomes to reflect management decisions based on the random event, and a probability-weighted NPV is effectively obtained. However, finding accurate probability distributions to describe future events is nontrivial, and it becomes problematic to use the same discount rate as an NPV valuation. The discount rate, properly determined, is meant to account for all risk in a project, while the cash flows upon which it operates are meant to be expected values. When the cash flows are broken up according to decisions based on different future outcomes, it is no longer correct to use the same discount rate as before.

Despite its drawbacks, however, the DTA methodology offers a valuable insight: external uncertainties and the firm's ability to handle them can potentially have a great impact on the value of a project. This impact is very difficult to quantify with traditional NPV, because NPV does not explicitly consider flexibility. Flexibility represents the firm's ability to handle external uncertainties.

The topics of uncertainty and flexibility are worth exploring in more depth. In fact, as will be argued in the following two chapters, a thorough understanding of their effect is a requirement for a complete and accurate valuation of a program. Chapter 6 presents a tool that measures program value as an objective function as it searches for an optimal design solution. The tool combines traditional engineering analysis with cost and market analysis, and includes the consideration of uncertainty and flexibility.

Chapter 3. Uncertainty

3.1 Introduction

The future is unknown. This fact makes it very difficult to design airplanes. For that matter, uncertainty about the future makes it very difficult to design any long-lived project. Between the beginning of the project and the end, any number of events may occur that dramatically—or subtly—change the value of the project. As a result, aircraft manufacturers justifiably put enormous effort into deciding whether, what, and when to design and build. This chapter offers a classification of the primary types of uncertainty and surveys existing methods of quantitatively modeling uncertainty.

3.2 The Nature of Uncertainty

Uncertainty, otherwise known as risk, may be classified into two forms for the purposes of aircraft design. The first is technical risk, and the second is market risk.

Technical risk refers to the ability of the firm to develop the product to the desired performance specifications. It affects the product's feasibility from a technological standpoint. Managing technical risk is critical in the development of any complex system, especially when one or more of the technologies involved are not mature. There have been extensive studies by both industry and academia on technical risk management in aerospace and its relationship to cost and schedule (for example, see Browning, 1998). As technical risk represents uncertainty regarding the product's feasibility, it may be seen—from an academic standpoint—as entirely in the hands of the firm. Only the firm may resolve a project's technical risk, and it does so by putting resources into research and development (R&D). Technical risk is therefore completely internal to the project and firm—that is, it is idiosyncratic risk.

Market risk, on the other hand, is not idiosyncratic. It represents uncertainty about what will happen to the project given that the product is developed as planned. Much of this is beyond the firm's control, tied to external factors. Some external factors directly affect the project, while others affect not just one particular project but the entire market. These may include GDP growth, customer spending preferences, economic booms and recessions, regional conflicts, natural disasters, fuel prices, and numerous other processes which are, in general, unpredictable. Thus, a firm may take 7 years to develop what is believed at the start to be an exceptional aircraft filling an important gap in the market, but at the end of 5 years that gap may disappear or be filled by a competitor, and although the aircraft itself is flawless from a technical standpoint, the project is a failure. This is a fundamental challenge of aircraft design—the length of time required to bring a new product to market is comparable to the timescale of the expansions and recessions of the airline business cycle. Thus, a good market at program launch may or may not imply a good market at the time of sale.

While technical uncertainty has been researched in close connection with engineering design, there has been less explicit consideration of market uncertainty. This thesis focuses on market uncertainty and its effect on aircraft design.

3.3 Modeling Uncertainty

3.3.1 Background

In reality, market uncertainty in a project comes from innumerable sources. However, for the purposes of this study, from the perspective of the firm, all the unpredictable events and their effects are boiled down to cash flows from sales, or potential sales, of whatever product the firm provides. Either the price or the quantity of the product sold, or both, will fluctuate unpredictably as a result of global, regional, and local economic conditions. The question is how to observe, quantify, and model this unpredictable fluctuation. One familiar source for observation and quantification is the stock market. The stock prices of firms represent fluctuations in cash flows from sales, both current and forecasted. In fact, a historical analysis of stock prices is often used for traditional NPV valuation to

find the appropriate discount rate. The discount rate, or opportunity cost of capital, is one way of modeling the market uncertainty inherent in a project, and it is often calculated by comparing fluctuations in the firm's stock price to fluctuations of the entire market.⁶

The evolution of stock prices provides a useful starting point for studying uncertainty, because a stock price is a representation of cash flows from a group of projects being subjected to market risk. In general, a stock price may be thought of as a stochastic process—that is, a variable that changes over time in an uncertain way.⁷ In particular, stock prices are usually modeled as Markov processes—stochastic processes whose evolution depends only on the current value of the variable, and not on any of its past history. The discussion below introduces several specific types of Markov processes and their application to stock price modeling⁸.

3.3.2 Weiner Process

This is a process where the variable, z , evolves with an average growth of zero and an average variance rate of 1.0 per unit time. That is, during a small period of time Δt , $E[\Delta z] = 0$ and $\text{var}[\Delta z] = 1$. The change Δz during time Δt may be characterized as

$$\Delta z = \varepsilon \sqrt{\Delta t} \quad (3.1)$$

where ε is a random variable sampled from a Normal (Gaussian) distribution with mean 0 and variance 1, i.e. $\varepsilon \sim N(0,1)$. In continuous time, as Δt approaches zero,

$$dz = \lim_{\Delta t \rightarrow 0} \varepsilon \sqrt{\Delta t} \quad (3.2)$$

The Weiner process, also referred to as Brownian motion, is a basic building block that can be used to construct more complex stochastic processes.

⁶ This is strictly true for valuations where the project has the same risk characteristics as the firm's existing business, and the firm is entirely equity-financed.

⁷ Hull, 2000, p. 218.

⁸ More detailed and mathematically rigorous accounts may be found in Hull, 2000, and Dixit and Pindyck, 1994.

3.3.3 Generalized Wiener Process

The Wiener process described above had as its only component a random fluctuation based on a sampling from a Normal distribution. The generalized version of this process adds a non-random component—a drift rate—and relaxes the requirement that the variance rate equal unity. The generalized Wiener process may be expressed as

$$dx = a(dt) + b(dz) \quad (3.3)$$

where dz is the increment from a standard Wiener process as shown in equation (3.2). This type of generalized Wiener process is arithmetic—that is, the expected value of the change in x per unit time is equal to a constant, a ; and the variance of the change in x per unit time is equal to another constant, in this case b^2 . Each period, x is expected to increase by an amount a , with some random fluctuation around that increment. The discrete-time version of the arithmetic Wiener process is

$$\Delta x = a(\Delta t) + b(\varepsilon\sqrt{\Delta t}) \quad (3.4)$$

where $\varepsilon\sqrt{\Delta t} = \Delta z$, as in equation (3.1). An alternate model is the geometric Wiener process, where both the mean and variance of the change in x are proportional to x . In other words, the *percentage* change and *percentage* variance are constants, independent of the magnitude of x . This has been empirically observed to be a good approximation of the actual behavior of many stock prices. Replacing a with μ and b with σ to use standard notation,

$$dx = \mu x(dt) + \sigma x(dz) \quad (3.5)$$

or

$$dx / x = \mu dt + \sigma dz \quad (3.6)$$

where dz is the increment in a standard Wiener process as above. The discrete-time expression for the geometric Wiener process is

$$\Delta x / x = \mu \Delta t + \sigma \varepsilon \sqrt{\Delta t} \quad (3.7)$$

Here, μ represents an expected growth rate (rate of return) as a proportion of the magnitude of x , and σ represents the proportional standard deviation of returns per unit time—i.e., the volatility.

3.3.4 Ito Process

Taking the generalization process one step further, the Ito process is comprised of a random and a non-random component, both of which may in general be functions of time, t , and of the variable, x .

$$dx = a(x,t)dt + b(x,t)dz \quad (3.8)$$

This formulation is useful for extending the model to capture other dynamics of random behavior, such as mean reversion. A mean-reverting process is one where the variable is “pulled” toward some long-run average value, such that if it is abnormally high, its drift rate is negative, and if it is abnormally low, its drift rate is positive. Below is an example of geometric mean reversion:

$$dx/x = \lambda(\bar{x} - x)dt + \sigma dz \quad (3.9)$$

where λ is a parameter describing the strength of mean reversion. While mean reversion is not used in this study, it is useful for modeling industries or assets believed to be cyclical and may be a valuable extension for future research.

3.3.5 Random Walk

For purposes of numerical analysis, it is possible to discretize the geometric Brownian motion as a random walk. A random walk represents the process as binomial: each period of time, Δt , the variable in question may only do one of two things: increase or decrease, by a specified amount. For an arithmetic process, this amount is constant; for a geometric process the proportional amount is constant—that is, the fraction by which the variable may increase (or decrease) is constant. Each of the two possible events each period—increase or decrease—has an associated probability. For example, a standard Wiener process (no drift) can be represented as a random walk as follows. For a timestep Δt equal to 1.0, each period, the variable will either increase by 1 with probability 0.5, or decrease by 1 with probability 0.5. Because it is a Markov processes, the probabilities do

not change regardless of the value of the variable—whether it is zero or 100, the probability of a unit increase is 0.5. It can be confirmed that this binomial process has the same properties as the Weiner process—the mean of the change in the variable is zero, while the variance of the change⁹ is 1. Similarly, more complex processes such as the geometric Brownian motion can be expressed as a random walk. It is simply a matter of finding appropriate values for the probabilities of an increase and decrease, p and $1 - p$, and the associated amounts by which to increase or decrease the variable. As mentioned above, for a geometric process, the increase and decrease amounts are proportional. Let the increase factor equal u and the decrease factor equal d , such that given an initial value of x_0 , the next-period value will be x_0u with probability p , or x_0d with probability $1 - p$.

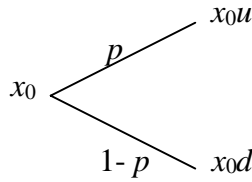


Figure 2. Random walk movement during time period Δt

If the geometric Brownian motion to be modeled as a random walk has parameters μ and σ as in equation (3.7), then the appropriate values for u , d , and p are¹⁰:

$$\begin{aligned}
 p &= \frac{e^{\mu\Delta t} - d}{u - d} \\
 u &= \eta \\
 d &= 1/\eta
 \end{aligned} \tag{3.10}$$

where

$$\eta = e^{\sigma\sqrt{\Delta t}} \tag{3.11}$$

As before, these parameters are fixed, regardless of the value of x . Because u and d are reciprocals, the random walk framework creates a recombining tree of possible values of

⁹ $\text{var}[\Delta x] = E[\Delta x^2] - E[\Delta x]^2 = 1 - 0 = 1$.

¹⁰ Cox, Ross, and Rubinstein, 1979, as quoted in Hull, 2000.

x . This binomial tree, illustrated in Figure 3, is a lattice representing the state space of the variable x as a function of time.

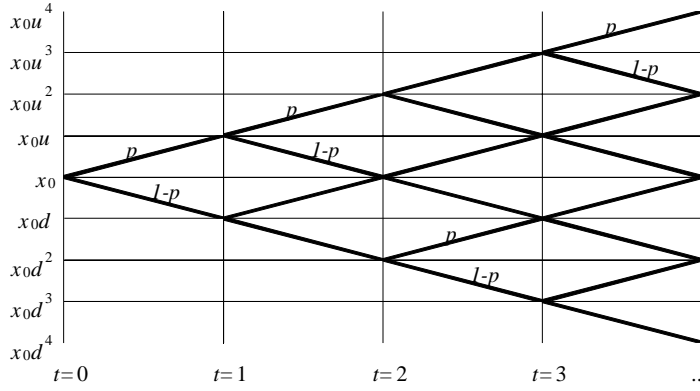


Figure 3. Binomial tree representation of geometric Brownian motion.

The space of the variable x has thus been discretized into a set of distinct values, separated by geometric (not arithmetic) intervals. The heavy lines represent possible transitions for x over time. At each time, regardless of the value of x , there is a probability p of an upward jump and a probability $1-p$ of a downward jump. As the time interval Δt is decreased, the upward and downward jumps u and d dictated by equation (3.10) also decrease, and in the limit as Δt approaches zero, it can be shown that the random walk representation approaches the geometric Brownian motion described by equation (3.5).

The above binomial tree is one representation of what is sometimes referred to as the “cone of uncertainty.” A particular asset or quantity, x , has a value that is currently known, but will evolve with some unpredictability in the future. In the general case, the tree does not have to be binomial and does not have to recombine, and the probabilities of upward and downward jumps do not have to remain constant with time or with the value of x ¹¹. However, the above representation is useful for its simplicity, and has been employed extensively in the valuation of assets that depend on the evolution of stock prices in particular and stochastic processes in general. Hence, the random walk forms

¹¹ For example: a mean reverting motion, with its instantaneous drift rate being a decreasing function of x , would also have p as a decreasing function of x .

the basis for uncertainty modeling in this study. Figure 4 shows several sample paths generated through a random walk representation to model a geometric Brownian motion. In this case, the volatility σ is set as 20% per annum; the initial value of the stochastic process is specified as 100; and the timestep is 1 month. The smooth line represents the expected growth path of the process, at 9% per annum. The three jagged lines are three different sample paths, constructed using a random number generator to approximate the random walk process.

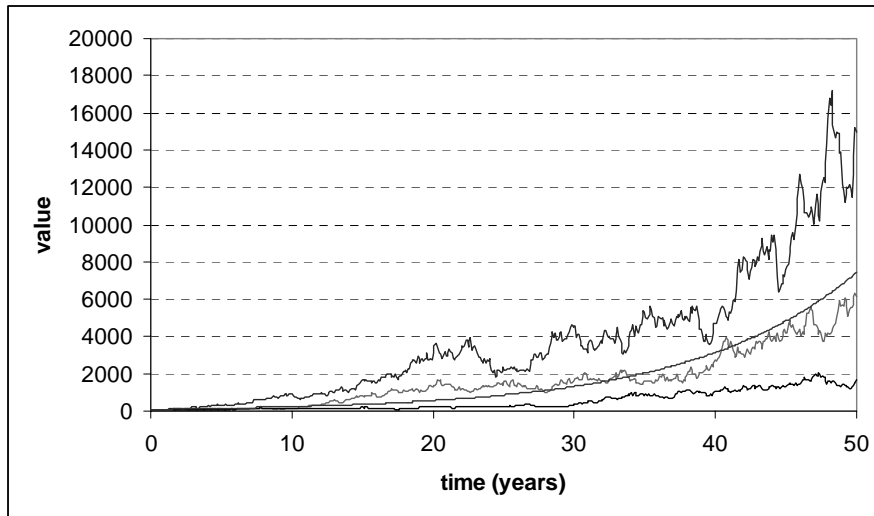


Figure 4. Random walk representation of geometric Brownian motion.

3.4 Conclusion

This chapter has introduced the topic of uncertainty. The unpredictability of future events is a fundamental challenge of capital budgeting. Some of this unpredictability can be addressed and mitigated by the firm's own actions—for the most part, this is referred to as technical risk. Other elements of unpredictability are beyond the firm's control—they are exogenous. These uncertain exogenous factors are modeled as stochastic processes, variables that evolve with a component of randomness. Some of the most familiar stochastic processes are found by observing the behavior of stock prices. These are most often approximated with geometric Brownian motion, although several other well-known models exist as well. A convenient way to numerically represent geometric

Brownian motion, as well as other stochastic processes, is as a discrete-time random walk, consisting of an upward or downward jump during each of numerous small time intervals. It is this representation that will be used in the remainder of this study to simulate unfolding market conditions and to enable the consideration of flexibility in an aircraft program.

Chapter 4. Flexibility

4.1 Introduction

The previous chapter discussed uncertainty—the unpredictable nature of the future. This chapter discusses flexibility—a firm’s ability to cope with uncertainty. The nature of flexibility is discussed below and the assertion is made that flexibility has value. The traditional NPV approach is shown to be deficient in capturing the value of flexibility, and the foundations of an alternate approach are presented to address this deficiency. The alternate approach draws on financial option pricing theory and its extension to project valuation, Real Options theory.

4.2 The Nature of Flexibility

All else equal, a project designed to accommodate an uncertain future is more valuable than an inflexible project. This is true because, in general, the future is impossible to predict with perfect accuracy. Therefore, an inflexible project is likely to be sub-optimal for the market conditions that actually occur. A flexible project, on the other hand, is by definition designed to adapt (to greater or lesser extent) to evolving conditions in real time to generate maximum value.

Flexibility may appear in multiple forms, some explicit and some implicit. Most projects have a degree of intrinsic, implicit flexibility, simply owing to the fact that there is a person responsible for the project outcome who is given the responsibility and authority to make decisions as the project continues. This person is known as a manager. Depending on the nature of the project, the manager has the discretion to adjust the workforce, schedule, expenditures, product mix, or other factors in response to unfolding events. Some forms of flexibility are explicit, such as an investment in a dual-fuel steam boiler that can switch from one fuel to another, depending on their relative prices

(Kulatilaka, 1993), or a production line designed to build any combination of several different aircraft. Flexibility is implicit in the simple decision to defer investment in a project until the market looks more favorable, or the decision to abandon a half-completed project if the market has evaporated.

Several examples of flexibility may be found in recent developments in aerospace. One example is the evolution of the Airbus A380, a “superjumbo” sized aircraft with capacity for 555 passengers. As Airbus moved toward committing to the A380 product, it passed through several gates, each time evaluating the maturity of the design and the perceived market for the aircraft. The final gate was a firm commitment from a sufficient number of customers to purchase the aircraft once it was designed and built. Had the threshold number of customers not been achieved, Airbus would not have invested in the detailed design effort and massive production facilities to bring the A380 the rest of the way to market. While Airbus was conceiving the A380, Boeing explored a stretched version of its existing jumbo aircraft, the 747, to compete with the new Airbus product. The 747 Stretch program was officially initiated and publicized, and design work progressed. However, approximately two years later, a sufficient customer base had failed to materialize for the 747 Stretch, despite the clear technical feasibility of the aircraft. Despite having invested significant resources into the program, Boeing accordingly halted its efforts¹². This action demonstrates rational decision-making on the part of Boeing management—decision-making that is difficult to explicitly include in a traditional NPV valuation. An in-depth study of the superjumbo development dynamics is presented by Esty and Ghemawat (2001).

Another example of flexibility in aerospace is the Airbus A320 series production line. The A320 narrow body aircraft product family consists of 4 aircraft—the A318, A319, A320, and A321—with a high degree of commonality across airframes. As a result, any of the 4 aircraft types can be built on the same production line, with comparable lead times. Therefore, an airline may place an order for an A319, to be delivered in several

¹² At the same time, Boeing announced a new development effort for a high-speed commercial aircraft, the Sonic Cruiser.

years¹³, with an option to switch to an A320 at a later time, much closer to the delivery date. This option provides significant value to the airline, which faces unpredictable demand fluctuations and consequently, unpredictable fleet evolution requirements. Consequently, Airbus may price such “switching options” accordingly, such that some of the additional value to the airline is passed on to the aircraft manufacturer.¹⁴

In summary, all forms of flexibility involve the ability to make decisions in the future that will affect the course of the project based on external events that also occur in the future, after the time of project valuation.

4.3 The Shortcomings of Net Present Value

The traditional Net Present Value approach suffers from an inability to rigorously quantify the value of flexibility. For example, consider a firm about to invest in a factory that produces widgets¹⁵. The factory takes 1 year to build, and once it is built produces widgets at the rate of 1 per year in perpetuity. The investment cost to build the factory is \$100; the risk-adjusted discount rate is 20% per annum; and the profit per widget, to be resolved 1 year from today, will be either \$25 or \$5 in perpetuity, with equal probability. To value this project using the traditional NPV method, one would consider the expected present value of the cash flows resulting from building the factory today. This value would equal the sum of -\$100 (investment in the factory) and the present value of the expected cash flows from widget sales starting next year and continuing in perpetuity. The expected profit per widget is \$15, since the profit will be \$25 or \$5 with equal probability. Thus the value of the project is calculated as

$$NPV = -100 + \sum_{t=1}^{\infty} \frac{15}{(1+0.2)^t} = -100 + \frac{15}{0.2} = -25 \quad (4.1)$$

¹³ Queue lengths on the order of several years are not uncommon for commercial aircraft.

¹⁴ The Airbus A320 production line example is taken from Stonier, 1999.

¹⁵ This example draws from content in Dixit and Pindyck, 1994.

The NPV is negative, and the project is rejected. However, this ignores the possibility of deferring the investment decision until the following year, once the profit per widget has been resolved. At that time, the firm may decide to build the factory only if the profit per widget is \$25. Assuming that this strategy is adopted, the firm would invest in the factory next year with 50% probability, corresponding to the event that the widget profit is \$25, and would in that case earn a perpetuity of \$25 starting in year 2. The NPV would be

$$NPV = 0.5 \left[\frac{-100}{1.2} + \sum_{t=2}^{\infty} \frac{25}{(1+0.2)^t} \right] = -83.3 + \frac{1}{1.2} \left(\frac{25}{0.2} \right) = 10.4 \quad (4.2)$$

In this case, the NPV is positive, because the firm exercises its flexibility by investing only if the conditions are right.

The above example shows the importance of flexibility, but equation (4.2) is not an accurate valuation. The explicit consideration of different possible future events is a way of accounting for uncertainty. However, traditional NPV uses the discount rate to account for all uncertainty. Therefore, the use of the original discount rate of 20% is flawed. More generally, the inclusion of flexibility in the analysis changes the risk characteristics of the problem, and thus the original risk-adjusted discount rate is no longer applicable.

To include the effect of flexibility in a valuation, then, it is necessary to (1) explicitly allow for decision-making in response to unfolding future events and (2) find an appropriate way to discount future contingent cash flows—that is, cash flows that depend on unpredictable events or decisions.

4.4 Valuing Flexibility

4.4.1 Background

The inclusion of real time decision-making in a project is quite context-specific and is best tailored to each situation individually. However, the discounting of contingent cash flows is a more general topic. It addresses the question of how to find the present value of future cash flows that depend on events driven by stochastic processes.

One approach suggested by a number of authors (Amram and Kulatilaka, 1999; Dixit and Pindyck, 1994; Trigeorgis, 2000) is to use Real Options Analysis. Real Options draws on financial option pricing theory, which provides a technique for finding the market value of a contingent claim on a financial asset. One example of a contingent claim is a call option on a stock: a piece of paper giving its owner the right, but not the obligation, to purchase the stock on (or before) a particular “exercise date” for a stated “exercise price.” If the stock price is high at expiration, owning the option will result in a large payoff, while if the stock price drops, the payoff will be zero—the option will go unexercised. Thus the cash flow from the call option is contingent upon the behavior of the stock price.

Just as the price of a stock, according to finance theory, reflects expectations about the value of holding the stock in the future, the price of an option is a function of its potential payoffs and their associated likelihoods. However, it is not obvious what this price should be due to the uncertainty of the option’s future payoffs, and their lack of exact correlation with the stock price. Rather, the payoffs are asymmetric, or downside-limited, because the minimum payoff to holding an option is always zero, regardless of how low the stock price falls. Thus, the problem of valuing a financial call option addresses the root of the topic presented at the beginning of this section: how to value contingent cash flows.

4.4.2 Black-Scholes Formula

Black and Scholes (1973) developed a method for pricing financial options, now famous as the Black-Scholes formula. The equation below gives the formula's result for the price of a European call option:

$$f = S_0 N(d_1) - X e^{-rT} N(d_2) \quad (4.3)$$

where

$$d_1 = \frac{\ln(S_0 / X) + (r + \sigma^2 / 2)T}{\sigma \sqrt{T}}$$
$$d_2 = d_1 - \sigma \sqrt{T}$$

The function $N(x)$ represents the cumulative probability distribution function for a normally distributed variable with mean 0 and variance 1, sometimes written as $\Phi(x)$. S_0 is the current price of the stock; X is the exercise price of the call option; T is the time to maturity of the call option; r is the risk-free rate; and σ is the volatility per unit time of the stock price. A thorough derivation of the formula may be found in Hull (2000), but for the purposes of this discussion it is important to note simply that the value of a call option is a function of the current stock price; the exercise price; the risk-free rate; the time to maturity; and the stock volatility, but not a function of the expected growth rate of the stock price, shown as μ in equation (3.5). According to the Capital Asset Pricing Model¹⁶, the equilibrium growth rate μ is the rate required by investors in the market to compensate for the risk inherent in that particular stock. The riskier the stock, the greater is the equilibrium growth rate. The fact that μ is not present in the Black-Scholes formula indicates that the value of a call option—or in general, a contingent claim—does not depend upon the premium that is required by investors to compensate for risk. The price of an option does not depend upon investors' risk preferences. This insight is explained in more detail below as the principle of risk-neutral valuation.

¹⁶ For a description of this well-known model, consult a finance text such as Brealey and Myers, 1996.

4.4.3 Risk-Neutral Valuation

The basis of risk-neutral valuation is the observation that the price of a contingent claim is independent of risk preferences—that is, independent of the required return, μ , for a certain level of risk¹⁷. No matter what assumptions are made about the required premium for risk, the price (value) of the contingent claim will be the same. If one set of assumed risk preferences yields the same result as another, one might as well use the simplest set of risk-preferences: all investors are risk-neutral. This means that all assets earn the risk-free rate, regardless of how risky they are. The only applicable discount rate in this risk-neutral world is also the risk-free rate. To find the price of a contingent claim, then, it is sufficient to find its expected future value in a risk-neutral world—assuming that all assets earn the risk-free rate—and discount it to the present at the risk-free rate. This assertion is demonstrated in the simple example below.

Example: Call Option Pricing. Consider a call option on a stock whose price is \$100 today. The option will expire in 1 year, and has an exercise price of \$100. If the stock price is greater than \$100 in a year, the option will be “in the money” and the option holder will exercise it to get a payoff. Suppose the stock has risk characteristics such that its expected growth rate μ is 10% per annum, while its volatility σ is 20% per annum. Suppose further that the risk-free rate is 6% per annum. For simplicity, the stock price is modeled as a random walk with a timestep of 1 year. Therefore, using equation (3.10), the stock price in 1 year will either be \$122.14, with a probability p of 71.1%, or it will be \$81.87, with a probability of $1-p = 28.9\%$. These possible outcomes result in an expected price of \$110, which corresponds to $\mu = 10\%$. At the end of the year, the payoff from the option will therefore be either $\$122.14 - \$100 = \$22.14$, or 0 if the stock price falls. Thus, the problem is as follows: given that the option will be worth either \$22.14 or \$0 in 1 year, what is its value today? Two solutions are given below. The first constructs a portfolio of assets, other than the option, that results in identical payoffs in 1

¹⁷ “Risk” may be defined as the beta of the stock price—a parameter measuring the stock’s sensitivity to the market. The higher the sensitivity, the higher the beta and the level of risk. For a complete discussion of risk and return, see Brealey and Myers, 1996.

year. The present value of that portfolio must be equal to the present value of the option. The second solution uses risk-neutral valuation: it assumes that the world is risk-neutral and that the stock's expected growth equals the risk-free rate. Adjusting the random walk model parameters accordingly, the expected value of the option payoff is found and then discounted back to the present by 1 year. The two solutions result in identical option values.

Solution 1: Replicating Portfolio. Consider purchasing 0.55 shares of the stock¹⁸ and borrowing \$42.40 for 1 year at the risk-free interest rate. After 1 year, the stock is sold and the loan is paid off. The resulting payoffs, depending on what happens the stock price, are either $0.55(\$122.14) - e^{0.06}(42.40) = \22.14 , or $0.55(\$81.87) - e^{0.06}(\$42.40) = \$0$. Note that these payoffs exactly replicate the payoffs from holding the option¹⁹. The portfolio of 0.55 shares and a loan of \$42.40 creates a synthetic call option that behaves identically to the real thing. Since the payoffs of the portfolio and the option in 1 year are the same, the present values must also be the same. The present value of the portfolio is simply the cost of constructing it—the amount paid for the 0.55 shares, less the \$42.40 loan: $0.55(\$100) - \$42.40 = \$12.60$. This must be the value of the call option.

Solution 2: Risk-Neutral Valuation. Assuming that the world is risk-neutral, the random walk model of the stock price is modified by setting μ equal to the risk-free rate of 6%. The new random walk calculations yield unchanged future values of the stock price (still \$122.14 and \$81.87), but the probability p of an upward jump is now 60.4%, while the probability of a downward jump is $1-p = 39.6\%$. The possible option payoffs are still \$22.14 and \$0, respectively. The expected value of the payoffs is simply their probability-weighted average: $60.4\%(\$22.14) + 39.6\%(\$0) = \$13.37$. This is the expected future value of the option in a risk-neutral world. To find the present value, the number is discounted back at the risk-free rate by 1 year: $e^{-0.06}(\$13.37) = \12.60 . Note

¹⁸ In practice, one would not purchase fractional shares, but multiple shares in proportion to the number of options to be purchased.

¹⁹ Calculations assume continuous compounding of interest. Thus borrowing \$42.40 requires paying back $e^{0.06}(\$42.40)$ one year later.

that this solution, which assumed a risk-neutral world, yielded the same answer as Solution 1 above.

Summary: Risk-Neutral Expectations. The method presented in Solution 2 is applicable not only to call options, but to any contingent claims (derivatives). Any contingent claim may be valued by replacing the growth of its underlying asset with the risk-free rate, taking the resulting risk-neutral expectations (probability-weighted averages), and discounting them at the risk-free rate. Note that the probabilities calculated as part of the risk-neutral valuation process are not true probabilities. They are artificial constructs that represent hypothetical probability values in a risk-neutral world; hence they are referred to as “risk-neutral probabilities.”

4.5 Conclusion

This chapter has presented a brief overview of the topic of flexibility. Through several historical and hypothetical examples, it has been shown that flexibility plays an important role in almost any project in general, and in aircraft development projects in particular. Flexibility may be explicitly built into a product or into the project framework; and it may be implicit in the way a project is actively managed. Both forms of flexibility can potentially make significant contributions to the value gained by a firm from a project, but neither is well represented in traditional NPV-based valuation techniques. One method that seeks to identify and quantify the impact and value of flexibility is Real Options. It borrows from financial option pricing theory to draw analogies between derivatives of financial assets and decisions relating to real assets and projects. The Real Options approach has two primary components: one is the creation of a framework to model the project as a collection of decision points analogous to financial options, and the other is a scheme for discounting contingent cash flows. While the former is better explained in the context of the specific application for which it is needed, in Chapter 6, the latter has been explained in this chapter as background material. Given a stochastic process which forms the basis for possible future decisions, the resulting contingent cash flows may be valued by using the method of risk-neutral expectations, which modifies

the stochastic process into its hypothetical equivalent in a risk-neutral world. The expected cash flows from the risk-neutral process may then be discounted at the risk-free rate.

Chapter 5. Program Valuation

Tool: Sub-Components

5.1 Introduction

This chapter summarizes and details the sub-components of the valuation tool developed to enable the design and optimization of an aircraft program. The following chapter links the sub-components to synthesize the entire tool. This tool does not have the capability to either perform aircraft design, or to find project market values, at an industry-level fidelity. Such a capability is not the goal of this research. Rather, the goal is to model the salient dynamics of the design problem; to gain insight into the interactions between technical and program design; and to demonstrate the utility—or lack thereof—of this particular approach.

The following are the three distinguishing characteristics of the approach outlined in this chapter and the following one:

1. Engineering analysis is combined with economic analysis. For example, a single design iteration may involve the following “chain reaction” of effects, starting with a decision to set an increased cruise speed:

- a. Modify wing sweep to optimize aerodynamic performance at new cruise speed.
- b. Modify wing and fuselage structural layout for new aerodynamic loads.
- c. Generate a new vehicle weight estimate based on modified structural layout.
- d. Modify development and production costs based on new weight breakdown.
- e. Generate a new fuel efficiency based on speed, weight, and aerodynamics.
- f. Find a new baseline market price based on change in fuel efficiency.
- g. Modify optimal timing of design and production based on new cost & price.
- h. Find new program value.

2. Program management decisions may be deferred and made in “real time.”

Flexibility is modeled as a set of decisions available to management throughout the lifetime of the program. Thus, the program schedule is not fixed at the time of valuation—program value is found under the assumption that management will actively control the project as future events unfold. For example, the maximum rate at which aircraft may be produced (the factory capacity) is not known or specified a priori; rather it is a decision that is made midway through the design process, several years into the project, depending upon market conditions at that time. The rules for making such decisions optimally are determined as part of the analysis

3. Demand is modeled as a stochastic process. While traditional valuation techniques rely on static forecasts of demand and account for uncertainty through a risk-adjusted discount rate, this tool explicitly models the uncertainty inherent in the market for commercial aircraft. The explicit consideration of uncertainty is useful, as described in Chapters 3 and 4, because (1) it facilitates the inclusion of flexibility in the analysis, and (2) it forces the user of the tool to give serious and concrete consideration to the volatility of the market and its effect on the program.

The tool is synthesized by creating three free-standing analytical models and linking them to compute a measure of program value. The free-standing analytical models, described in this chapter, are *static* estimators of (1) performance, (2) cost, and (3) market characteristics. The linking process involves a *stochastic* market model—one which captures the market’s time-varying and uncertain nature—and a method (described in the following chapter) for evaluating the complete aircraft program as a sequence of decisions and resulting cash flows occurring over a period of time. The final result is a measure of program value.

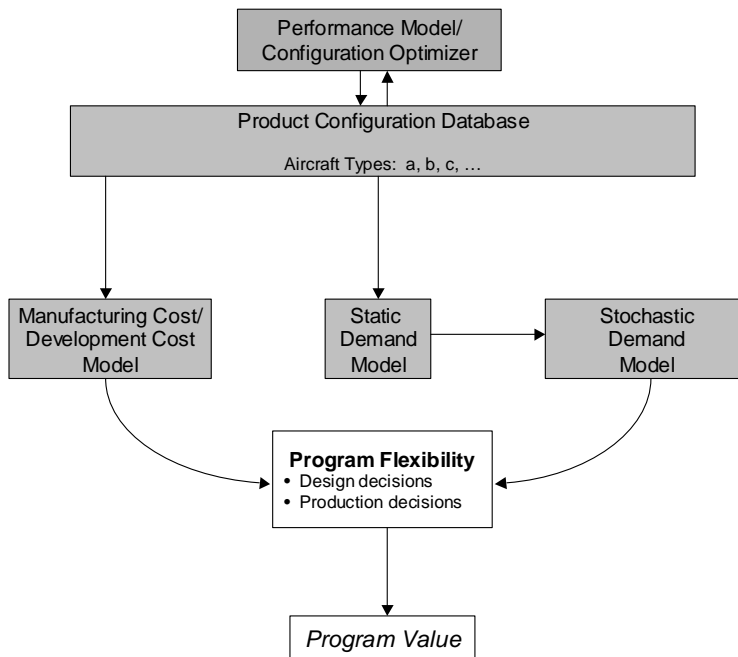


Figure 5. Architecture of the program design and valuation tool

Figure 5 illustrates the conceptual framework unifying the models noted above. Each of the models is described in detail in the following sections. The performance model is a sizing and configuration tool, which closes the engineering design loop between technical parameters and performance metrics²⁰ (range, capacity, fuel burn). The cost model generates estimates of manufacturing and development costs given the technical parameters specified in the performance model. The static revenue model captures the behavior of the market for commercial aircraft, and characterizes airline demand for the aircraft in question, given the performance predicted by the performance model. Finally, the stochastic market model (described in this chapter) and the valuation algorithm itself (described in the following chapter) link the performance, cost, and revenue analyses with the concepts of uncertainty and flexibility to generate a measure of program value using a dynamic programming approach.

²⁰ As used here, “performance metrics” represent a quantifiable characteristic of the aircraft that relates to its operation. Range and capacity were mentioned previously as examples of constraints. In that context, they would be constraints on the aircraft’s performance.

5.2 Performance Modeling

5.2.1 Overview

The performance model serves as the front end of the program valuation tool. It provides as inputs to the rest of the tool the functional and technical characteristics of the aircraft, or set of aircraft, to be developed. These inputs are calculated using a set of engineering analysis and estimation techniques that translate a set of requirements and constraints into a baseline airframe. The resulting design is optimized with respect to gross weight—a frequently used objective function in aircraft design. This analysis and optimization is performed by the WingMOD aircraft design tool (see Section 5.2.2).

However, there is no guarantee that the optimal performance-based vehicle design (i.e., minimum-weight design) can or should be a part of the optimal program design. Therefore, the performance-based design is input into the other parts of the program valuation tool—cost modeling, revenue modeling, and the dynamic programming framework for linking all the parts—and an evaluation is made of the overall program. Based on the insights gained from this evaluation, changes may be proposed to the design constraints imposed on the performance model. Such changes may in turn result in an altered performance model output, and another iteration of the program valuation process may be executed.

5.2.2 WingMOD

As described by Wakayama (1994) and Wakayama and Kroo (1995), WingMOD is a multidisciplinary design optimization (MDO) code that optimizes aircraft lifting surfaces subject to numerous practical constraints. The original application for WingMOD was the re-design of a wing for a variant of the MD-90 narrowbody aircraft (Wakayama, Page, and Liebeck, 1996). After undergoing further development and modification at Boeing Phantom Works, the code was applied to the design of the Blended-Wing-Body (BWB) aircraft concept (Wakayama & Kroo, 1998; Wakayama, 1998).

WingMOD uses several intermediate-fidelity analysis techniques to analyze a set of constraints specified for an aircraft at each of over twenty flight conditions. The

constraints address numerous technical and regulatory requirements, as well as mission requirements. The analyses involved include performance, aerodynamics, loads, weights, balance, and stability & control. Because the analysis techniques are conducted at an intermediate-fidelity level, computation cost is relatively low, allowing for an extensive optimization with over 100 design variables. The optimization uses a clearly defined technical objective function: minimize gross takeoff weight. There is no explicit consideration of cost or revenue; the focus is on engineering design and its effect on performance. Below is a brief description of the mechanics of WingMOD.

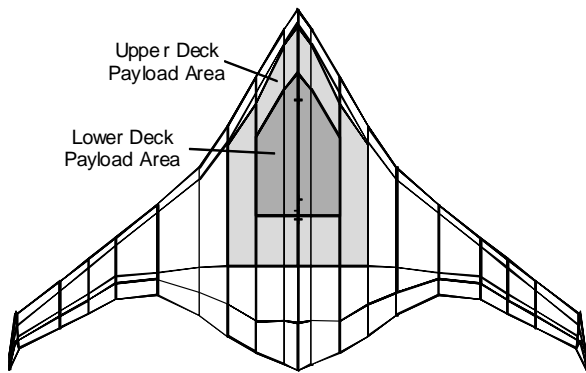


Figure 6. WingMOD representation of Blended-Wing-Body aircraft.

The optimizer, which uses the “Genie” computational framework (Wakayama & Kroo, 1998), treats the Blended-Wing-Body (BWB) planform as a series of spanwise elements. This representation of the aircraft is illustrated in Figure 6. The planform is modeled as an aerodynamic surface with a modified vortex-lattice code and a monocoque beam structural analysis. These analyses are coupled to generate aeroelastic loads. Drag characteristics are generated by trimming the vehicle at several flight conditions to find loads and induced drag; and by using empirical relations to find profile and compressibility drag at each station across the planform span based on the lift coefficients from the vortex lattice code. Structural weight is generated by finding the maximum expected elastic loads over all the specified flight conditions, such as maneuver, vertical gust, and lateral gust; and by sizing the structural elements based on bending strength and buckling stability. The maximum lift coefficient for the entire planform is found by a critical section method, which assumes that the planform has reached its maximum usable lift at the point when any section has reached its maximum section lift coefficient,

which is found from empirical data. A balance analysis is conducted over the duration of the mission (Wakayama, 2000).

In summary, the WingMOD optimization framework takes a set of constraints representing mission requirements (range, payload capacity, cruise speed, approach speed, balance, etc.); and finds an optimal aerodynamic and structural configuration such that the resulting aircraft satisfies the constraints. In this context, “optimal” means minimum gross takeoff weight (GTOW). In addition to a detailed description of each the constraints at each of the flight conditions noted above, the outputs of the optimization include general weight characteristics and planform geometry of the optimized vehicle.

While WingMOD was developed as and still is used as a standalone technical design code, it has been integrated into the program valuation tool to the extent that its outputs are directly used by the other analytical models described below to complete the design loop around the whole system being designed (the aircraft program), as opposed to just the airframe. This study does not attempt to close the design loop by feeding program value back to WingMOD as an objective function. A closed feedback loop based on value would be a logical extension to the research.

Several WingMOD outputs and constraints are used directly by the other models described below. For the purposes of the models used here, these outputs and constraints completely characterize the physical concept for the aircraft. They are listed in Table 1.

Table 1. Key WingMOD Outputs

Output	Description
MTOW (maximum takeoff weight)	Gross vehicle weight at start of max. range mission. Equivalent to GTOW.
DLW (design landing weight)	Gross vehicle weight at end of max. range mission.
Design range	Maximum range. Specified as a constraint.
Seats	Number of passenger seats. Specified as a constraint.
Weight breakdown	Estimates of weight for each primary component of the aircraft. The weight categories below sum to Operating Empty Weight (OEW).
<ul style="list-style-type: none">• Structure<ul style="list-style-type: none">○ Fuselage bay 1○ Fuselage bay 2○ Inner wing○ Outer wing○ Etc...	Structural elements (spars, ribs, skin, etc.). The structural weight is broken down into a set of large “parts” which together comprise the airframe.
<ul style="list-style-type: none">• Propulsion	Engines, nacelles, and supporting equipment.
<ul style="list-style-type: none">• Systems	Onboard systems, including avionics, fuel, flight control, hydraulics, and others.
<ul style="list-style-type: none">• Payloads	Seats, bag racks, cargo equipment, etc.

5.3 Cost Modeling

5.3.1 Overview

The goal of the cost model is to determine the cost characteristics of an aircraft given its technical parameters—that is, its physical characteristics. There are two primary categories of cost in aircraft programs: development cost and manufacturing cost. The former is mostly non-recurring, while the latter is mostly recurring. Both are addressed

in detail below; however, the baseline cost model described in this section must be modified and simplified for integration into the dynamic programming algorithm (see Chapter 6), which adds the dimensions of uncertainty and flexibility to the overall valuation tool.

Importantly, the cost characteristics found by the cost model include the effect of commonality between several different airframes. In general, then, both the development cost and manufacturing cost of a new aircraft will depend upon (1) the aircraft's technical parameters; and (2) the technical parameters of other aircraft types that have already been designed.

Detailed descriptions of each cost estimation process (development and manufacturing) follow. Both estimation algorithms are based upon the decomposition of the aircraft into a set of components, or "parts." In this context, the word "parts" will be used to refer to fairly large sub-assemblies, such as wing sections or fuselage bays; or to refer to entire groups of components, such as the "payloads" group or the "systems" group, or even to processes, such as "final assembly." Thus, any aircraft can be explicitly defined as a set of "parts." Both development cost and manufacturing cost are estimated on a *per part* basis: each part has its own set of cost characteristics, which determine the formula used to compute its contribution to cost. Existing cost models and statistical data are used for calibration, resulting in a baseline *cost per pound* for each part. This cost per pound is converted to a total cost through the technical parameters of the aircraft (i.e., part weight) generated by WingMOD, and further modified based on other factors, such as other aircraft types already designed. The total manufacturing, or development, cost of an aircraft is then computed as the sum of the costs calculated for each of its parts. Clearly, this approach implies an assumed direct relationship between weight and cost. However, a minimum-weight design does not necessarily equal a minimum-cost design. One example of such a situation is an aircraft shares multiple common parts with other, existing aircraft. In this case, the cost benefits from part commonality may balance the cost drawbacks of a higher weight and result in a lower total cost.

The parts-based approach facilitates the consideration of commonality between multiple aircraft. For example, if a family of two aircraft is being considered for design and

production with a common wing, the cost model will recognize that the same wing part is used on both airframes, even though the rest of the aircraft may be completely different. Since the cost for each part is calculated independently, the effect of commonality on the wing will be explicitly captured. The mechanics of adjusting part cost for commonality are discussed in more detail below.

5.3.2 Development Cost

Development cost is the non-recurring effort required to bring an aircraft concept to production. It includes preliminary design, detail design, tooling, testing, and certification. As noted above, the total development cost estimate is built up as the sum of the costs for each of the parts comprising the aircraft.

The development cost model was created in seven basic steps, which are described sequentially in the remainder of this section.

1. The non-recurring effort is broken into several processes (engineering, tool design, etc.) for each aircraft part. These processes are characterized in terms of relative time and cost based on non-dimensional industry data for typical commercial aircraft.
2. An estimate is made of total weight and non-recurring cost for a modern wide-body aircraft: the 777-200. The estimate is based on public domain data and tools.
3. An estimate is made of total duration of the non-recurring effort for typical commercial aircraft, based on expert testimony.
4. An estimate is made of the fractional weight breakdown, by parts, of typical commercial aircraft. The estimate is based on non-dimensionalized industry data.
5. An estimate is made of the fractional cost breakdown, by parts, of typical commercial aircraft. Again, the estimate is based on non-dimensionalized industry data.
6. Fractional cost and weight breakdowns for typical aircraft are combined with cost and weight estimates for the 777-200 to yield cost per pound estimates for each

part in the 777-200, and for the contribution of each process from step (1) to each part²¹. The resulting cost per pound values are assumed to apply to all hypothetical new commercial aircraft.

7. The effect of commonality is approximated through a set of factors modifying the cost per pound values.

Non-recurring processes.

In addition to nominal dollar cost, the model estimates the profile of cash flows as a function of time. A typical new commercial aircraft development process lasts on the order of 5 years, and the associated non-recurring expenses are spread out over that duration. Figure 7 shows a non-dimensionalized plot of the non-recurring person-hours used for a representative aircraft development project at a major aerospace company. The plot abscissa is a timeline roughly corresponding to the period between program launch and the delivery of the first product. The non-recurring hours are broken down not by aircraft parts, but by separate processes comprising the development effort: engineering; manufacturing engineering; tool design; tool fabrication; and two others, which can be classified as “support.” Note that the time profile for the entire process, as well as the individual processes which comprise the total, has roughly the shape of a Bell curve—that is, a normal distribution.

²¹ Recall that a “part” is used here to represent some subcomponent of the aircraft, e.g. a wing.

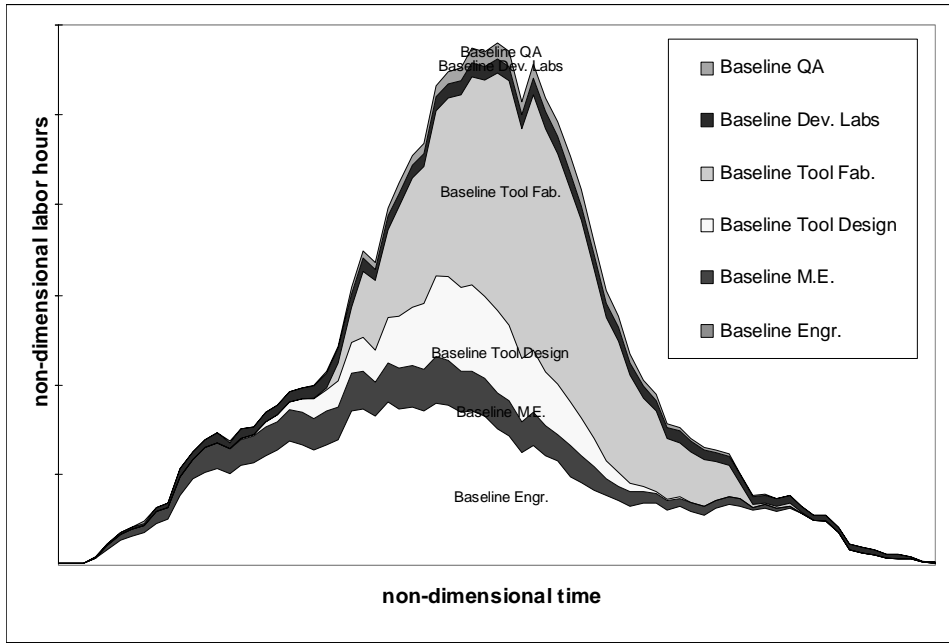


Figure 7. Non-recurring hours for a commercial aircraft project: the data²²

To approximate the process illustrated in Figure 7, a simple model was constructed: each of the 5 major processes represented was assigned a normalized weighting (i.e., percentage of total cost); a normalized start time and duration; and two parameters— α and β —defining the shape of the cost profile for that particular process. The cost profile for each process was defined by

$$c(t) = Kt^{\alpha-1}(1-t)^{\beta-1} \quad (5.1)$$

where t is non-dimensionalized time, $c(t)$ is non-dimensionalized process cost (as represented by person-hours), and K is a constant. The parameters α , β , and K for each process; as well as the weights, start times, and durations, were all varied until the resulting “synthetic” cost profile—the sum of the profiles from equation (5.1)—approximated the shape of the profile in Figure 7. The resulting set of parameters is summarized in Table 2, and the synthetic non-recurring cost profile thus generated is shown in Figure 8.

²² D. Anderson, Boeing Commercial Aircraft, email communication, 4/17/2001

Table 2. Non-recurring effort process parameters

	Engineering	Mfg. Engr.	Tool Design	Tool Fab.	Support
Weighting (%total cost)	0.400	0.100	0.105	0.348	0.047
Start time	0	0	0.22	0.27	0
Duration	1.00	0.85	0.45	0.50	1.00
α	2.2	2.5	3.5	3.0	1.5
β	3.0	3.0	3.0	3.0	1.5

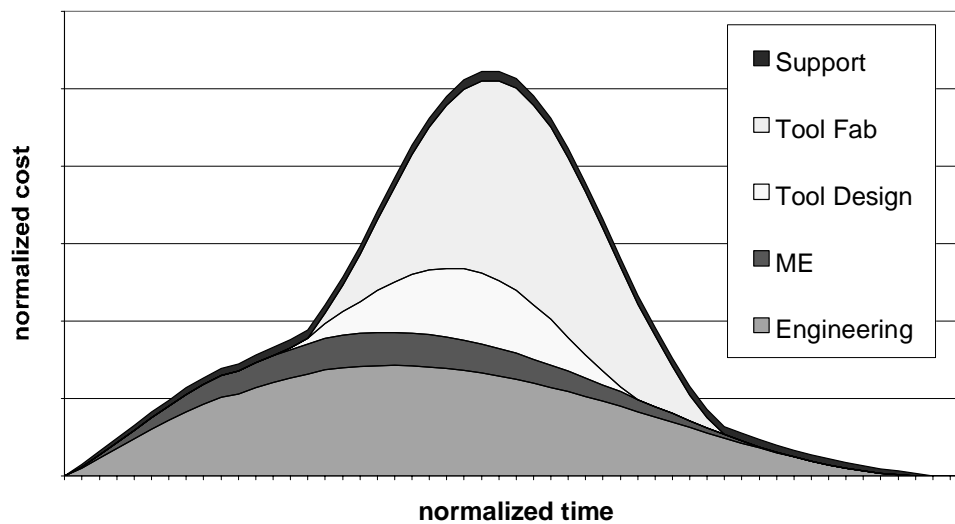


Figure 8. Non-recurring hours for a commercial aircraft project: the model

To confirm the general validity of the cost profile modeled above, a plot was made showing cumulative non-recurring cost committed as a function of time. This plot, shown in Figure 9, displays the familiar “S” shape characteristic of large-scale aircraft development projects.

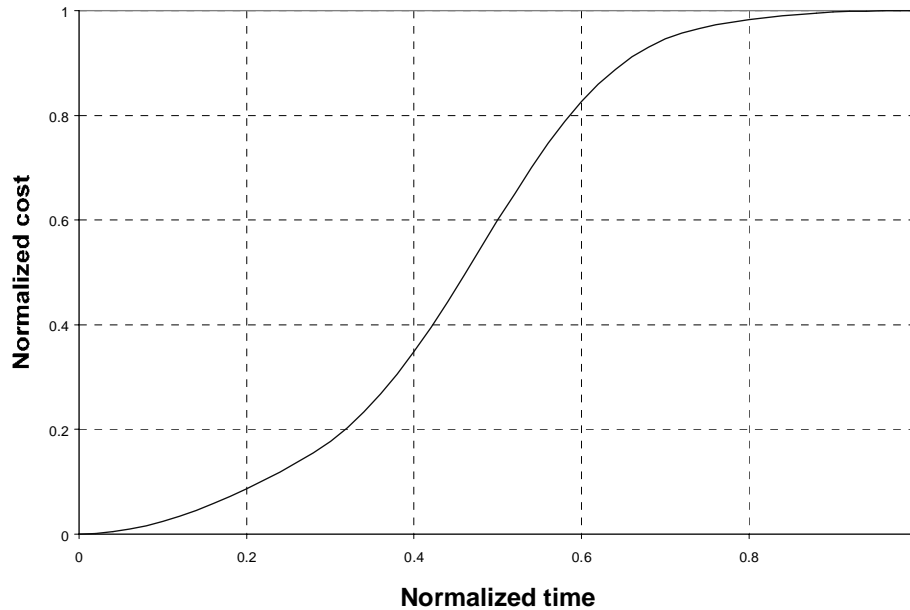


Figure 9. Cumulative non-recurring cost committed: the model

Total non-recurring cost.

To build a complete development cost model, it was necessary to find information in the public domain that could be used to estimate the actual cost magnitudes involved, and how such magnitudes relate to aircraft weight. For simplicity, both the development model and the manufacturing model assume that cost scales directly with operating empty weight (OEW), commonality effects notwithstanding. This assumption breaks down for parts of significantly different complexities or significantly different materials, but given the scope of this research, it is a reasonable simplification.

The necessary baseline set of cost and weight information was found by using an existing parametrics-based model to estimate the development cost of a modern wide body aircraft, for which the OEW is available in the public domain (*Jane's All the World's Aircraft*, 1991). The aircraft selected was the 777-200. The model used was Raymer's (1999) DAPCA IV cost model²³, which estimates RDT&E (research, development,

²³ Development And Procurement Cost of Aircraft

testing and evaluation) costs and manufacturing costs for both military and commercial aircraft by breaking down the cost structure into categories such as engineering hours, tooling hours, quality control hours, development support costs, and flight test cost. Each cost category has an associated cost-estimating relationship (CER) based upon a regression fit of a number of existing aircraft for which the data is available.

For estimation of total non-recurring cost, the DAPCA IV model was linearized with respect to weight about the data point for the 777-200—that is, the assumption was made that for aircraft in the same weight class as the 777-200, cost scales linearly with weight. This relationship is shown in Figure 10, where the 777-200 OEW of 305,000 lb corresponds to a non-recurring cost of \$6.08B. No claim is made here that this number is representative of the actual non-recurring cost of the 777-200. This is simply an estimate based on existing public domain cost-modeling technology. Its purpose here is to provide first-order accuracy and to demonstrate a realistic sensitivity to aircraft weight.

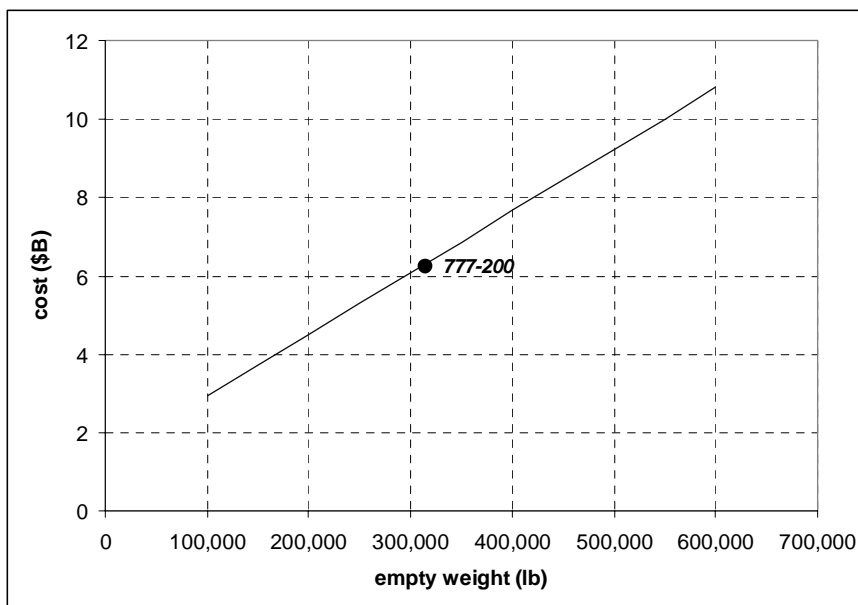


Figure 10. Linearized DAPCA IV model output: Total non-recurring cost

The information presented so far provides a relationship between total non-recurring cost magnitude and total vehicle OEW, as well as a rough estimate for the breakdown of the total cost by processes and over time. To complete the model, it was necessary to (1) find a baseline time duration for the development effort; (2) estimate a weight breakdown of typical commercial aircraft by parts, corresponding to the part types described by

WingMOD; (3) estimate a cost breakdown of typical commercial aircraft by parts, corresponding to the same part types; and (4) using the total cost estimated for the 777-200, find cost per pound values for each type of part and for each process (engineering, manufacturing engineering, tool design, etc.) for the 777-200. The next step was to assuming that these cost per pound values are independent of the actual aircraft, such that a wing part costs the same, per pound, on a 777-200 as on an A340 or 747 aircraft.

Total non-recurring duration.

The baseline time duration of a representative aircraft development process is set at 66 months, based upon an expert estimate.²⁴ The development activities represented by this duration include preliminary design, program launch, firm configuration freeze, detailed design, tooling, flight test, and certification. This represents a significant simplification, as the time duration of aircraft development is in general dependent upon the characteristics of the aircraft being designed. A useful extension of the research would be to model this dependency.

Fractional weight breakdown.

To estimate the fractional weight breakdown of typical commercial aircraft, historical data was collected on weight characteristics of a set of existing aircraft (Raymer, 1999, and Roskam, 1989). The results—average values for weight fractions of major parts—are shown in Figure 11. The data available for the analysis was for the following aircraft: MD-80, DC-10-30, 737-200, 747-100, and A300-B2.

²⁴ Conversation with D. Anderson, Boeing Commercial Aircraft, 3/5/2001

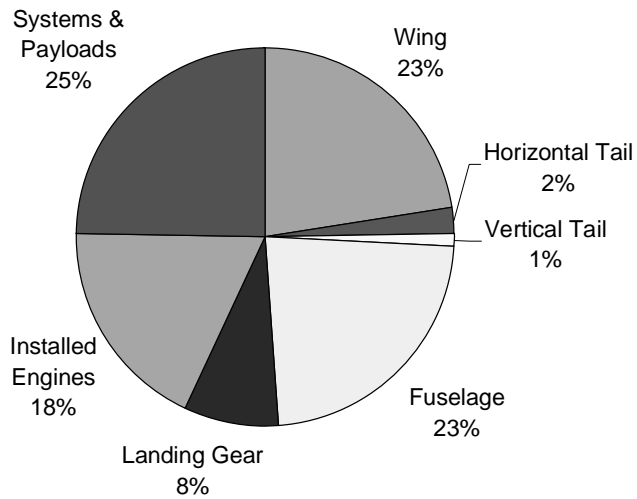


Figure 11. Estimated operating empty weight fraction breakdown for a typical commercial aircraft

Fractional cost breakdown.

To estimate a cost breakdown of a typical commercial aircraft by parts, a set of assumptions was made, based on non-dimensionalized data provided by industry sources.²⁵ The resulting assumptions describing non-recurring cost fraction by parts are illustrated in Figure 12.

²⁵ Data from R. Liebeck, Boeing Phantom Works, 9/2001.

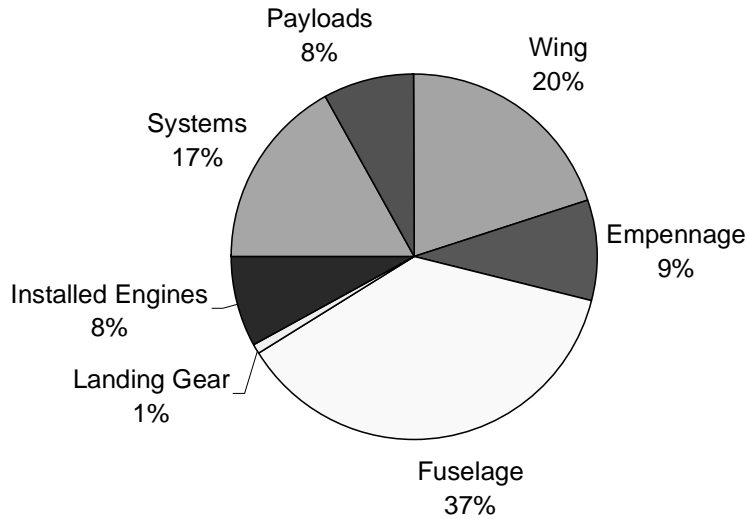


Figure 12. Representative non-recurring cost breakdown by parts for a typical commercial aircraft
Cost per pound values.

Having developed estimates of both weight fractions and non-recurring cost fractions, it is possible to apply these to the baseline vehicle, the 777-200, for which estimates have been made above for both total weight (OEW) and total non-recurring cost. The result is a set of estimated values for non-recurring cost and weight for each of the major parts of a 777-200 (recall that the term “parts,” as used here, refers to major subassemblies, or “building blocks”). Taking the ratio of one set of data to the other, a normalized non-recurring cost per pound is obtained for each part. Each of these costs is further split according to the total cost fractions of the non-recurring processes identified in Table 2. The resulting data is shown in Table 3: normalized non-recurring costs by aircraft part and by non-recurring process. These normalized costs, along with the time profiles of the processes (from Table 2) may now be applied to each part of a hypothetical new aircraft design, assuming that all the components of the new aircraft are classified into one of the existing part categories and assuming that a weight estimate exists for each part. Thus, each part of a new aircraft is assigned its own non-recurring cost and duration profile for each non-recurring process.

Table 3. Non-recurring part costs per pound

	Engineering	ME	Tool Design	Tool Fab	Support	Totals
	40.0%	10.0%	10.5%	34.8%	4.7%	100.0%
Wing	\$7,093	\$1,773	\$1,862	\$6,171	\$833	\$17,731
Empennage	\$20,862	\$5,216	\$5,476	\$18,150	\$2,451	\$52,156
Fuselage	\$12,837	\$3,209	\$3,370	\$11,169	\$1,508	\$32,093
Landing Gear	\$999	\$250	\$262	\$869	\$117	\$2,499
Installed Engines	\$3,477	\$869	\$913	\$3,025	\$408	\$8,691
Systems	\$13,723	\$3,431	\$3,602	\$11,939	\$1,612	\$34,307
Payloads	\$4,305	\$1,076	\$1,130	\$3,746	\$506	\$10,763

The complete non-recurring cost and its time profile for a given aircraft design is simply the sum of the costs and profiles of each of the aircraft's component parts. For example, an aircraft may consist of two wing parts (inner and outer wing); 4 fuselage parts (cockpit, forward/mid/rear fuselage); 3 empennage parts (winglets, horizontal stabilizer, vertical stabilizer); a landing gear part; an "installed engines" part; a systems part; and a payloads part; each with its own weight.

Effect of commonality.

One complication requires a modification of the non-recurring cost model: the effect of commonality. Consider an aircraft about to be developed that is a derivative of an existing design—one that has already been developed and built. The hypothetical derivative may have a high degree of commonality with the aircraft already in production—for example, it may use the same wing, or the same fuselage section, or even both. In such a case, the non-recurring cost for those components of the aircraft that already exist should be significantly lowered from the baseline case of an all-new design.

Unfortunately, very little information is readily available to quantify the design cost savings associated with commonality. Accordingly, a set of assumptions was made to model such savings, based on reasonable first-order estimates. Table 4 presents these assumptions as cost reduction factors associated with previously designed parts. For

every part of a new aircraft design that has already been developed on a different aircraft, the non-recurring cost is multiplied by the corresponding factor shown in the table.

Table 4. Non-recurring cost reduction factors

	Engineering	ME	Tool Design	Tool Fab	Support
Wing	20%	50%	5%	5%	50%
Empennage	20%	50%	5%	5%	50%
Fuselage	20%	50%	5%	5%	50%
Landing Gear	20%	50%	5%	5%	50%
Installed Engines	20%	50%	5%	5%	50%
Systems	20%	50%	5%	5%	50%
Payloads	20%	50%	5%	5%	50%

The above estimates for reduction factors assume that the greatest design cost savings will be in tooling design and fabrication--unnecessary for parts which have already been developed, because tooling equipment for those parts should already exist. However, this assumption breaks down when the desired production rate requires additional tooling to be built.

The development (non-recurring) cost model described above is a static one, producing a time profile of cashflows necessary to bring an aircraft design to production readiness. This model is later modified, as explained in Section 6.4, to fit into the dynamic framework of the program valuation tool such that the start and end times of the development process are no longer firmly fixed, but flexible.

5.3.3 Manufacturing Cost

Like the development cost model, the manufacturing cost model is centered around the assumption of a linear relationship between cost and weight. Also, the manufacturing cost model uses the same parts-based approach as development cost—the cost to build an aircraft is computed as the sum of its part costs. However, unlike the development cost, the nature of manufacturing cost is recurring—that is, the cost is incurred repeatedly for

each unit, and as such, it is subject to a learning curve effect. This well-known phenomenon is characterized in aircraft production by a significant decrease in unit cost as additional aircraft are built. The decrease is most noticeable early in the production run, and eventually decays to a negligible level, when unit cost remains roughly constant. As described below, the learning curve effect is captured in the manufacturing cost model. However, as described in Chapter 6, it must be modified for incorporation into the overall program valuation tool due to the dynamic nature of the framework.

As in the development cost model, manufacturing cost for any aircraft part is split into several processes. The two most significant processes contributing to aircraft manufacturing cost are labor and materials. There are also several others, which for the purposes of this model are grouped in the “other” category: quality assurance, recurring engineering, and recurring tooling. The resulting three processes are summarized in Table 5, with their associated typical fractional contributions to recurring cost and typical learning curve slopes.

Table 5. Manufacturing process parameters

	Labor	Materials	Other
% of total recurring cost ²⁶	41%	33%	26%
Learning curve slope	85%	95%	95%

The learning curve “slope” parameter describes the magnitude of the learning curve effect mentioned above. A slope of 100% indicates no learning—the initial unit cost remains constant throughout the production run. As the value decreases, the learning effect grows stronger. As described by Fabrycky (1991), the effect is quantitatively modeled as

$$MC = TFU \times Q^{\ln(s)/\ln(2)} \quad (5.2)$$

where MC is marginal (unit) cost; TFU is theoretical first unit cost; Q is quantity built to date; and s is the learning curve slope parameter. Thus, when the number of units built

²⁶ R. Liebeck, 9/2001, non-dimensionalized production cost data.

doubles, the marginal cost of producing one additional unit is s percent of its original value.

The learning curve slope values used for the model, as shown in Table 5, are quite high for the “materials” and “other” processes, while the value for labor is lower, implying a greater impact. This reflects the reality that most production efficiency improvements, implemented as a result of experience, are realized through reductions in labor hours required to build the aircraft.

The cost model uses the above process assumptions to vary the manufacturing cost per pound for any given part as a function of the number of such parts that have already been built. However, this requires knowledge of the first-unit cost of each part—that is, the *TFU* term in equation (2).

The first-unit, or baseline, recurring costs for each part are found through a process similar to the one used for the development cost model: a baseline aircraft, the 777-200, is analyzed based on public domain data, and estimates are made for the baseline recurring cost of each of its parts. Those estimates, combined with the 777-200 part weight information introduced in Section 5.3.2, yield values for baseline cost per pound for each part, split up by process (labor, materials, or other).

The baseline recurring costs by part for the 777-200 are estimated in two steps: first, the total recurring cost of the aircraft is estimated using Raymer’s (1999) DAPCA IV cost model; and second, fractions of the total cost are assigned to the component parts according to non-dimensionalized recurring cost data for typical commercial aircraft, shown in Figure 13.

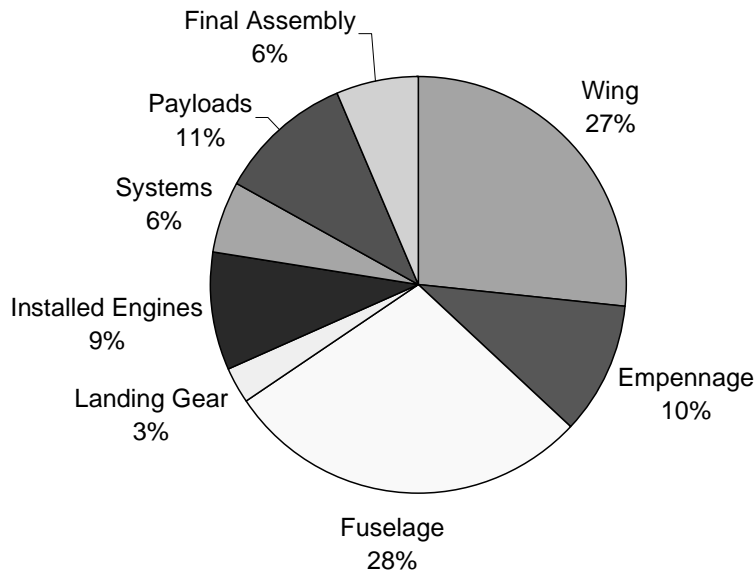


Figure 13. Representative recurring cost breakdown by parts²⁷

Note that the recurring cost breakdown includes a “part” not listed in the development cost model—final assembly. Although this is not technically a component of the aircraft, final assembly is treated as a separate part. It is assigned a distinct cost per pound value, as well as a “weight” value equal to the OEW of the aircraft.

To use the DAPCA IV cost model for 777-200 recurring cost estimation, a production run of 100 units was assumed. The output was used to find an estimate of the unit cost for the 100th aircraft to be built, and the resulting breakdown of cost by parts. The cost values were then converted to first-unit costs, using the learning curve assumptions given in Table 5. Finally, the first-unit cost for each part was normalized by its respective weight, and broken up by recurring process, to yield the data in Table 6.

²⁷ Non-dimensionalized data from R. Liebeck, 9/2001

Table 6. Recurring part costs per pound

	Labor	Materials	Other	Total
Wing	\$609	\$204	\$88	\$900
Empennage	\$1,614	\$484	\$233	\$2,331
Fuselage	\$679	\$190	\$98	\$967
Landing Gear	\$107	\$98	\$16	\$221
Installed Engines	\$248	\$91	\$36	\$374
Systems	\$315	\$91	\$46	\$452
Payloads	\$405	\$100	\$59	\$564
Final Assembly	\$58	\$4	\$3	\$65

The estimated normalized cost data in Table 6 enables the calculation of the unit cost for any given aircraft, given (1) the set of its component parts, each with a well-defined weight, and (2) data on how many units of each part have already been built to date. Thus, the learning curve effect is applied to each part individually, and the aircraft unit cost is the sum of the costs of its component parts.

A useful feature of the above approach is its ability to account for derivative aircraft already in production, which may have some parts in common with the aircraft in question (see Willcox and Wakayama, 2002). As long as a record is kept of the number of units of each part produced, regardless of what aircraft that part was used for, the recurring cost for new aircraft will be computed accurately, correctly including any learning curve effects “inherited” for any components shared by the aircraft with others in production.

5.3.4 Summary

As any cost estimation expert will agree, calculating aircraft part cost—recurring or non-recurring—as a linear function of its weight is a crude approximation at best. However, there are several reasons why this approach was chosen over a more rigorous one. First, only the baseline values for part costs are calculated as a linear function of their weight.

Actual cost values also take into account learning curve (for recurring costs) and commonality effects. Second, a higher-fidelity model, while conceivable, would be impractical for the purposes of this study. Specifically, a more accurate model would consider the geometric complexity of the parts and the materials used; it would use nonlinear cost-estimating relationships (CERs); and it would, by necessity, split the aircraft into a much finer (more numerous) set of parts. Because this study focuses on the conceptual design stage, many of these cost modeling techniques are very difficult to implement. The aircraft may not be well-defined enough to break down into smaller parts and classify all of them. Further, a detailed bottoms-up parts-based cost buildup would be impractical for conducting trade studies because such a buildup is not readily automated.

The cost model, as presented above, is based upon a collection of data of varying fidelity available in the public domain or provided for public release. It is plausible that some of the assumptions made are inaccurate, and the development and manufacturing cost estimates that are outputs of the model may not match with the actual development and manufacturing costs incurred for real aircraft. However, a perfect match is not the goal of this study—rather, it is to demonstrate a methodology, part of which is the above cost model; and to gain insight into program value dynamics. To that end, the above model captures several key trends in aircraft cost, such as a near-linear relation with weight; the learning curve effect; the bell-shaped profile of non-recurring expenses; and the effect on development and manufacturing of inter-aircraft commonality.

5.4 Revenue Modeling

5.4.1 Overview

Demand for commercial aircraft is difficult to model, because commercial aircraft are not a commodity. They are sold in relatively low volumes, to a relatively small number of buyers, by an even smaller number of sellers. The industry is in general an oligopoly, and in some cases practically a duopoly: Boeing and Airbus dominate the market for aircraft bigger than regional transports. There are several airlines and several other

customers (leasing companies) who have a significant degree of buyer power, and regularly negotiate prices at deep discounts compared to other, less powerful buyers. To further complicate matters, international politics sometimes affect aircraft purchases and sometimes affect new aircraft development programs.

Bearing in mind the above peculiarities of the commercial aircraft market, it is necessary to construct a revenue model sophisticated enough to characterize the potential cash flows from the sales of new aircraft, and simple enough to be readily understood and readily used as part of the program valuation tool in this study. Therefore, several assumptions are made.

First, up to a certain maximum production capacity, exactly as many aircraft are produced in each period of the analysis as there are aircraft demanded. The reason for this assumption is that once a production line is in place for a particular aircraft, sales in any given period are almost exclusively determined by demand. That is, commercial aircraft are very rarely manufactured without a buyer already committed to the aircraft.

Second, the effect of competition is not explicitly considered. To do so would be an interesting exercise, but outside the scope of this work. Rather, the firm designing the aircraft in question is treated as a price taker—that is, demand is completely exogenous, and buyers are a non-differentiated group with identical preferences. The implications of this assumption are made clear in the following sections, 5.4.2 and 5.4.3.

The third significant assumption concerns the process by which demand is translated into sales. Specifically, all aircraft demanded in any given period from the producing firm are sold in the same period—that is, the queue length (lag between order time and delivery time) is assumed to be zero. In fact, this queue time is often on the order of a year or longer, but several factors make the assumption of zero queue length not entirely unreasonable. First, while orders are sometimes canceled or changed, they are for the most part tracked well by deliveries. Second, while payment schedules vary from customer to customer, the majority of the price of an aircraft is typically paid upon delivery. Therefore, the only significant breach between the zero queue length assumption and reality is an offset of approximately a year in production schedule, and thus production costs incurred.

Characterizing the quantity of aircraft demanded per unit of time—first, its magnitude, and second, its variability—is one goal of the revenue model. The second goal is to understand the behavior of the price, as well. In fact, classical economics states that the two are closely linked, and in fact it is impossible to consider one without the other. Aircraft prices are therefore also analyzed and modeled here.

Both price and quantity are treated in each of two separate sections—one focusing on a static analysis, and the other shifting to a dynamic analysis. The first section, 5.4.2, is concerned with generating an estimate of the instantaneous values for market price and annual quantity demanded for a given new aircraft design. No consideration is given to market dynamics, price elasticity, or future uncertainty. The second section, 5.4.3, searches for a relationship between price and quantity and studies the time-varying and stochastic (i.e., random) nature of both. The final result is a set of estimating relationships that takes as inputs several basic characteristics of a new aircraft, and provides as outputs an initial market price and annual quantity demanded, as well as a characterization of their future behavior.

5.4.2 Static Demand

The purpose of the static demand model is to generate an estimate of a market price and initial quantity demanded per unit time given a new aircraft concept described by a set of well-defined parameters. Rather than attempt to find a relationship between price and quantity, this section treats their initial, or baseline, values separately. Price is estimated with a regression model based on existing aircraft—their characteristics and corresponding prices—and quantity is estimated by averaging and filtering a set of existing market forecasts for commercial aircraft fleet evolution.

Price

Price is modeled as a function of several variables representing an aircraft's value to its operator. Several sets of variables and several functional forms were tested by applying

candidate price models to 23 existing aircraft: 11 narrowbodies and 12 widebodies²⁸. The outputs generated by the price model were compared to best estimates available in the public domain for the actual sale prices for each of the aircraft²⁹. For each functional form tested, the function parameters were adjusted to minimize the mean squared error of estimated price.

The resulting function and its variables are shown below. Note that speed (or Mach number) is not one of the variables. No significant statistical relationship between price and speed was found in the range of available data.

$$Price = f(Seats, Range, CAROC) \quad (5.3)$$

where

$$\begin{aligned} Seats &= \text{Passenger capacity} \\ Range &= \text{Design range} \\ CAROC &= \text{Cash Airplane-Related Operating Costs} \\ &= \text{Total operating costs less ownership costs} \\ &\quad \text{per available seat-mile} \end{aligned}$$

$$Price = \underbrace{\left[k_1 \left(\frac{Seats}{Seats_ref} \right)^\alpha + k_2 \left(\frac{Range}{Range_ref} \right) \right]}_{\text{“CAROC-neutral” price}} Price_ref - \Delta(LC) \quad (5.4)$$

where

$Seats, Range$ = Aircraft seat count and range (nm); input variables.

$Seats_ref$ = Reference value used to normalize seat count

²⁸ The author is indebted to J. Allen and J. Nelson, Boeing Phantom Works, for their assistance and guidance in developing the price estimation algorithm.

²⁹ Estimated price data taken from *Aircraft Value News, The Airline Monitor* (2001).

$Range_{ref}$ = Reference value used to normalize range

$Price_{ref}$ = Reference value used to normalize price

k_1, k_2, α = Model parameters, selected to minimize mean squared error of estimated prices.

$\Delta(LC)$ = Increment in lifecycle cost due to off-nominal CAROC

The price estimate given by equation (5.4) is the difference between a certain “CAROC-neutral” price and an adjustment based on CAROC—the “increment in lifecycle cost.” The reasoning behind this structure is explained below, while the parameter values identified by the regression analysis for the price estimating function are shown in Table 7.

Table 7. CAROC-neutral price regression parameters

	Narrow body aircraft	Wide body aircraft
k_1	0.735	0.508
k_2	0.427	0.697
α	1.910	2.760

It was clear that the price estimation algorithm must exhibit a sensitivity to operating cost: all else equal, an aircraft that is less expensive to operate is more valuable to a customer. At the same time, the other two variables identified as key determiners of price were seat count and range—aircraft that fly farther and carry more passengers are more valuable than shorter-range, smaller vehicles. While a regression based on seats and range produced reasonable, statistically significant results, an attempt to regress the price based on seats, range and CAROC produced coefficients that were not statistically significant. This result occurred because there is a significant correlation between CAROC and aircraft size—larger aircraft tend to be less expensive to operate per available seat mile (ASM). Therefore, the regression algorithm had difficulty distinguishing the effect of CAROC on price from the effect of seat count or range on price.

The solution to this issue involved making an assumption about the value of CAROC to customers. Given a particular new aircraft design, let “nominal CAROC” denote the expected operating cost of the aircraft based on the average for existing aircraft of its size. As mentioned above, the existing CAROC trend is a decreasing function of seat count (see Figure 14 and Figure 15). The customer is willing to pay a certain price for the aircraft, given its seat count and range, assuming its CAROC is at the nominal value. However, if such is not the case and the aircraft in fact has a higher (or lower) CAROC than expected, the customer will see an extra loss (benefit) from purchasing the aircraft. This loss (benefit) can be quantified as the expected present value of the additional expense (savings) from operating the aircraft over its lifetime. Theoretically, the customer should be willing to pay to the producer an amount less (greater) than the “CAROC-neutral” price by the expected present value of the loss (benefit). This amount is referred to as the lifecycle cost increment, $\Delta(LC)$, in equation (5.4). The lifecycle cost increment is calculated as follows:

$$\Delta(LC) = [\Delta(CAROC) \times (Seats) \times (Annual\ Utilization) \times (Trip\ Length)] \rightarrow \text{discounted over lifetime} \quad (5.5)$$

$$\left[\frac{\$}{\text{seats} \times \text{nm}} \times \text{seats} \times \frac{\text{trips}}{\text{year}} \times \frac{\text{nm}}{\text{trip}} \right] \times \text{years}$$

where $\Delta(CAROC)$ and *Seats* are functions of the new aircraft design, and the other parameters are constant, as given in Table 8.

Table 8. Lifecycle cost increment parameters

	Narrow body aircraft	Wide body aircraft
Annual utilization (trips/year)	2120	590
Average trip length ³⁰ (nm/trip)	500	3000
Service lifetime ³¹ (years)	27	27
Operator discount rate	12%	12%

³⁰ Utilization and trip length both from *Boeing Quick Reference Handbook* non-proprietary data.

³¹ *The Airline Monitor*, May 2001, p. 24: Average Age of 2000 retirements

Thus, the term in brackets in equation (5.5) is discounted over the service lifetime of the aircraft with a discount factor for an annuity with a duration of 27 years and a discount rate of 12%. $\Delta(\text{CAROC})$ is computed as the difference between the “nominal CAROC” for the aircraft (based on its seat count, according to the regression equations shown in Figure 14 or Figure 15—for narrow bodies or wide bodies), and the actual CAROC of the aircraft.

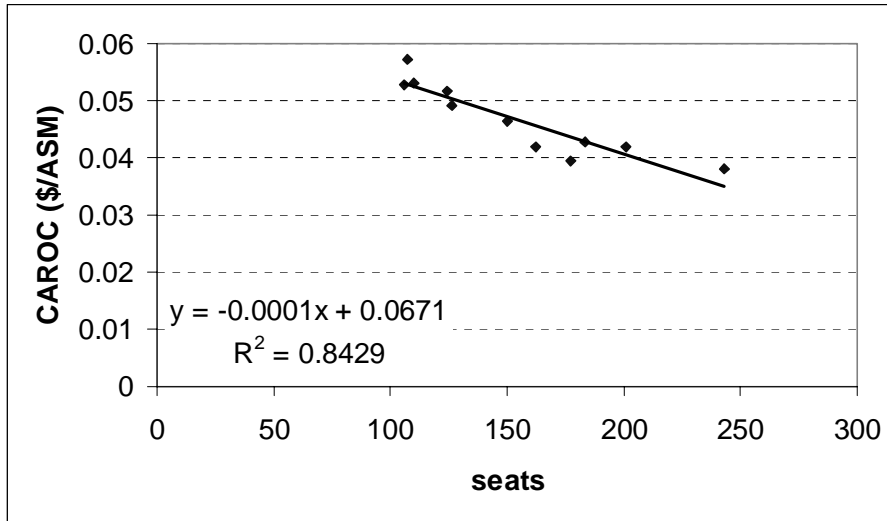


Figure 14. CAROC trend as a function of seat count: narrow body aircraft³²

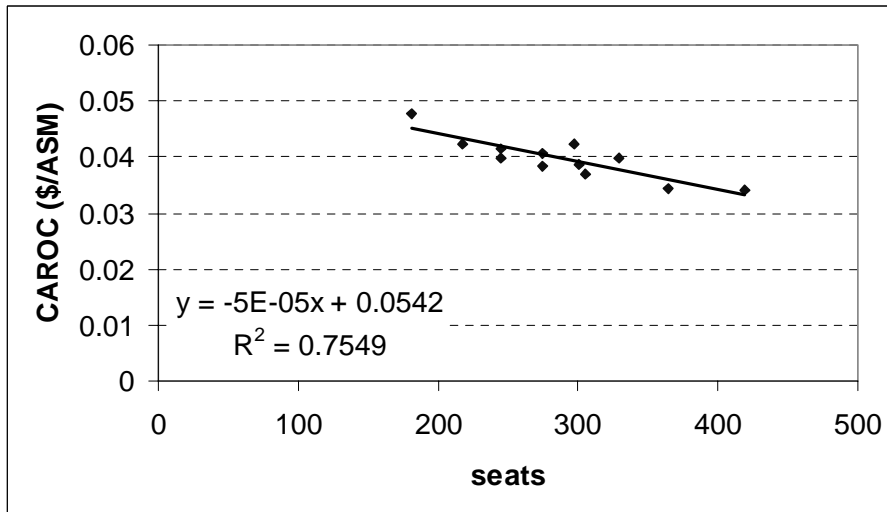


Figure 15. CAROC trend as a function of seat count: wide body aircraft³³

³² Boeing Quick Looks public domain data.

Recalling the overall expression for a price estimate in equation (5.4), it is the difference between the “CAROC-neutral” price and $\Delta(LC)$, as found in equation (5.5). Note that the CAROC-neutral price is a linear function of range, and a non-linear function of seat count. It was found that this functional form produced a reasonable regression fit, while constraining the function to be linear in both range and seat count was thought to be too restrictive and just as arbitrary.

To execute the regression analysis and find best values for the parameters k_1 , k_2 , and α , the lifecycle cost increment $\Delta(LC)$ was computed for each aircraft in the existing aircraft data set and subtracted from the aircraft’s known (best available estimate for) price³⁴. The result was a set of CAROC-neutral prices for the existing aircraft. These modified prices were then used as the target values which the CAROC-neutral price function approximates. An iterative numerical solution method was used where the seats exponent, α , was varied while the linear coefficients, k_1 and k_2 , were continuously computed with a linear least-squares regression. The optimum value of α was selected to minimize the mean squared error of the estimated CAROC-neutral prices. Figure 16 and Figure 17 compare the total aircraft price estimates (CAROC-neutral price - $\Delta(LC)$) generated by the price model to the actual prices for the existing aircraft data set. The estimates do not coincide exactly with the actual values, but incorporating the major price dynamics above into the model is more important than obtaining a perfect match.

³³ *Boeing Quick Looks* public domain data.

³⁴ Best available estimates for actual prices obtained from *Aircraft Value News* and *Airline Monitor*, May 2001.

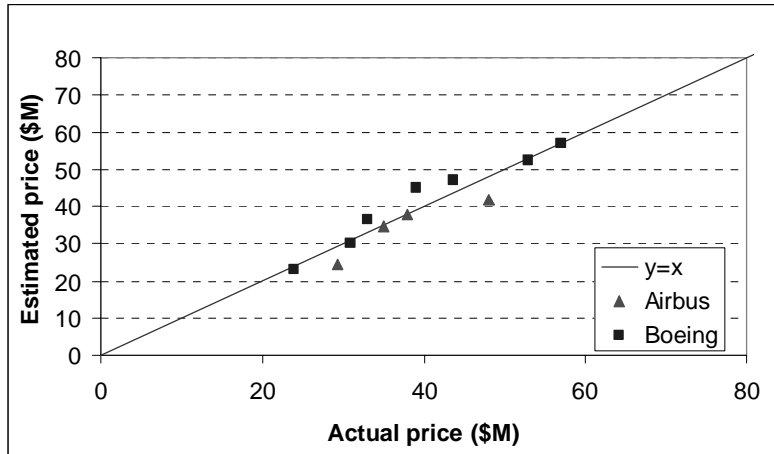


Figure 16. Price model estimates compared to actual prices: narrow bodies

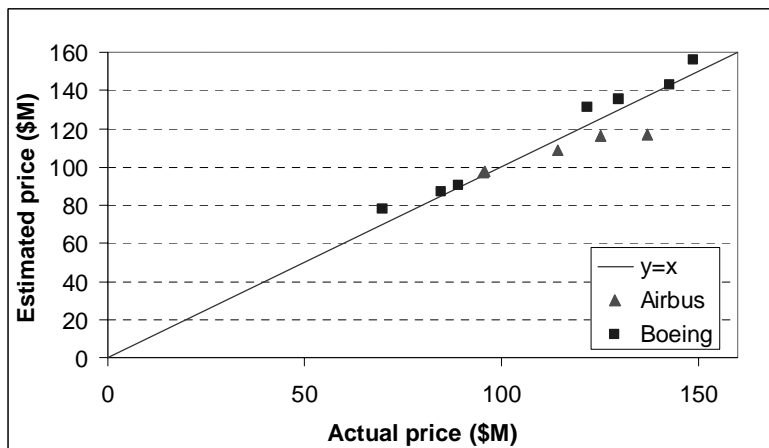


Figure 17. Price model estimates compared to actual prices: wide bodies

In review, the static price model takes as inputs the seat count, range, and operating cost of a new aircraft and outputs a baseline market price. Higher seat counts and higher ranges, and reduced operating costs all result in higher prices. The former two inputs (seat count and range) are readily available for any new aircraft design—in fact, they comprise two of the most important constraints given as inputs to the WingMOD-based performance model in Section 5.2. However, the third input—CAROC—requires some computation. CAROC is not an output of WingMOD; rather, it must be estimated with a separate model.

CAROC, or Cash Airplane-Related Operating Cost, has several primary components: fuel cost, crew cost, maintenance cost, and other costs (ground handling, communications, landing fees). The first three components dominate CAROC, and of

those, the one most easily quantifiable and most directly attributable to aircraft properties is the fuel cost. Fuel cost is simply the volume of fuel burned on a typical flight³⁵, multiplied by the price of aviation fuel. Fuel burn is strongly affected by the aircraft's aerodynamic, structural, and propulsion efficiency. The other primary components of CAROC also depend upon the aircraft, but those relationships are harder to quantify and are often customer-dependent. For example, different airlines have different pay structures for their pilots, as well as different maintenance infrastructures for their aircraft. Therefore, for the purposes of this model, the assumption is made that for a given aircraft size, fuel cost is a constant percentage of CAROC. To check and quantify this assumption, several existing aircraft were sampled, and a calculation of the fuel cost percentage was made for each. The calculations are summarized in Table 9³⁶.

³⁵ Reference missions used here for operating cost calculations are 3000 nm for wide bodies; 500 nm for narrow bodies. (*Boeing Quick Reference Handbook*)

³⁶ Because the remainder of this study focuses on wide body aircraft as examples, the CAROC analysis presented here is shown only for wide bodies, but it is equally applicable to narrow body aircraft.

Table 9. CAROC fuel cost fraction calculations**Assumptions**

ref. range	3000
(nm)	
fuel cost	0.65
(\$/gal)	
fuel density	6.7
(lb/gal)	

Inputs³⁷

aircraft	767-200ER	767-300ER	A330-200	A330-300	A340-300	777-200ER	747-400
seats	181	218	245	275	275	305	416
powerplant	PW4062	PW4062	PW4068A	PW4168A	CFM56	Trent 895	PW4062
MTOW (lb)	395,000	412,000	507,050	507,050	597,450	656,000	875,000
MLW ³⁸ (lb)	300,000	320,000	396,825	407,850	418,875	460,000	652,000
max range	6,555	6,075	6,450	5,600	7,200	7,595	7,335
(nm)							

Outputs

fuel burn	51,355	55,704	63,232	69,358	79,732	82,548	104,169
(lb)							
fuel cost	0.0080	0.0072	0.0073	0.0071	0.0082	0.0076	0.0070
(\$/ASM)							
CAROC ³⁹	0.0478	0.0423	0.0415	0.0385	0.0406	0.0369	0.0341
(\$/ASM)							
fuel/CAROC	16.69%	16.99%	17.49%	18.42%	20.08%	20.63%	20.65%

The key output in the above fuel cost fraction calculations is the fuel burn for the reference range mission. This number was found using the Breguet range equation:

$$R = V \left(\frac{1}{TSFC} \right) \frac{L}{D} \ln \left(\frac{W_0}{W_1} \right) \quad (5.6)$$

³⁷ Aircraft input information (seats, powerplant, weights, range) from *Jane's All the World's Aircraft*, 2001.

³⁸ MTOW = max takeoff weight; MLW = max landing weight. MLW is assumed to equal 110% of design landing weight (DLW)

³⁹ *Boeing Quick Looks* non-proprietary data.

where R is range; V is flight velocity; $TSFC$ is thrust specific fuel consumption (pounds of fuel burned per hour per pound of thrust generated); L/D is lift to drag ratio, a measure of aerodynamic efficiency; W_0 is initial (takeoff) weight; and W_1 is final (landing) weight. Range, W_0 and W_1 are all known for the maximum range mission. The remaining parameters in the equation are constant for all missions. Thus, knowing the reference range and design landing weight, W_1 (assumed to equal 90% of MLW), W_0 may be found for the reference range mission. The weight of fuel burned equals the difference between W_0 and W_1 .

The calculations suggest a weakly increasing trend in fuel cost fraction with seat count. A regression fit is found for the trend in Figure 18. This completes the CAROC model: for a given aircraft, fuel cost fraction is available as a function of seat count, and fuel cost itself for any range mission may be calculated from the MTOW (max takeoff weight), DLW (design landing weight), and maximum range WingMOD outputs using the Breguet range equation.

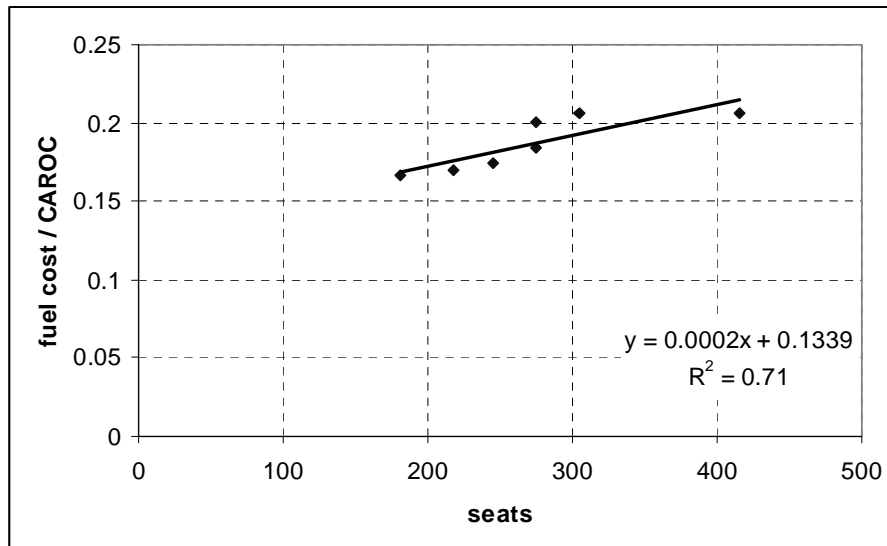


Figure 18. CAROC fuel cost fraction trend

Once the fuel cost is found, it is divided by the estimated fuel cost fraction to generate a value for CAROC. CAROC, in turn, determines the increment in lifecycle cost to the operator compared to a “nominal” aircraft, and adjusts the CAROC-neutral price as in equation (5.4).

Thus, given a new aircraft's seat count, range, maximum landing weight, and design landing weight, the above procedure finds the expected market price of the aircraft.

Quantity

Baseline quantity demanded for a given new aircraft design is estimated based on the aircraft's characteristics and a set of available market forecasts for commercial aircraft. The "baseline quantity" represents the expected number of units that would be sold in the initial period⁴⁰ of the analysis if production were instantaneous and costless. While this definition may seem somewhat odd, it is founded on the initial assumption that orders and deliveries are simultaneous. Section 5.4.3, which models the time-varying and uncertain nature of demand, uses the baseline quantity as a starting point.

Quantity data is based on 3 distinct forecasts of the commercial aircraft market outlook for the next 20 years. The sources of the forecasts are Boeing, Airbus, and The Airline Monitor, an aviation consulting service.⁴¹ Each forecast describes the future commercial aircraft market in terms of projected aircraft deliveries over the course of 20 years. The projections are based on current and expected future fleet growth trends; aircraft retirement and replacement rates; and observations and predictions regarding the evolution of fleet composition—that is, the fractional breakdown of the global airline fleet by aircraft type. The primary distinguishing characteristic for aircraft type is seat count, followed by range. In general, the global airline fleet may be broken down into a number of groups, each comprising a range of seat counts, e.g. 90-110 seats. Such a classification generally obviates a breakdown by range, as aircraft of similar seat counts typically have similar range classes. The two major range categories coincide with the two major aircraft classes: narrow bodies, which typically carry less than 200 passengers, have ranges of less than 3500 nm, and are typically used for transcontinental

⁴⁰ While the example cases in this study all use a period of one year, the methodology is extendable to other timesteps, e.g. monthly or quarterly.

⁴¹ *Boeing Current Market Outlook*, 2000; *Airbus Global Market Forecast*, 2000; *Airline Monitor*, July 2000.

or shorter flights; and wide bodies, which typically carry more than 150 passengers, have ranges of over 5500 nm, and may be used for intercontinental flights.

Each of the three forecasts breaks down the global airline fleet into groups, as above, but each defines its groups differently. Therefore, all 3 forecasts were recast into a single, consistent set of “seat categories” based on aircraft class (narrow body or wide body) and seat count. Forecasted deliveries are assumed equivalent with quantities demanded at current market prices. The results are shown in Figure 19.

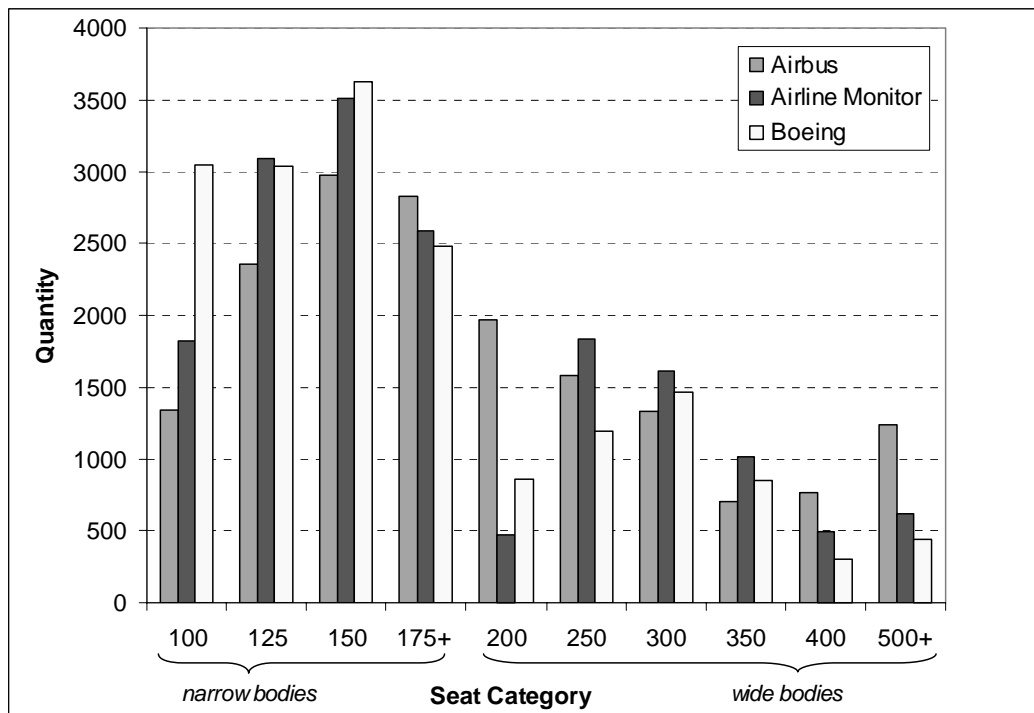


Figure 19. Commercial aircraft forecasted deliveries: 2000-2019

Depending on seat category, there is considerable variance between the 3 forecasts. This reflects differences in the forecasters’ assumptions; methodology; and to some extent, corporate strategy. Further, the high variance reflects the high degree of uncertainty regarding future revenue cashflows.

For a given new aircraft design, estimation of baseline quantity demanded proceeds as follows. First, the aircraft is assigned to a seat category. This is a straightforward decision if its seat count falls within a range specifically defining a seat category—in that case, demand for the aircraft is derived solely from that category. However, if the seat

count falls between the ranges for two adjacent seat categories, demand for the aircraft is fractionally assigned to both categories. For definitions of the seat categories as seat count ranges, refer to Table 10. The seat count ranges defining the seat categories were chosen based on groups in which existing aircraft are observed to compete. Note that the seat count ranges for the 175+ and the 200 seat categories overlap. This is because the former category is a narrow body group; while the latter is for wide bodies. For the purposes of this analysis, an aircraft is strictly a narrow body if its range is less than 3500 nm, and it is strictly a wide body if its range is greater than 5500 nm. If the aircraft does not fall into one of these range groups, it is assigned fractionally to each of the two borderline seat categories—high-end narrow bodies (175+) and low-end wide bodies (200). If an aircraft has a smaller seat count than the lower bound of the 100 category, or greater than the upper bound of the 500+ category, it is either competing in a market not analyzed in this study (regional transports), or it is assumed to be competing in a market which does not yet exist (very large aircraft), and is assigned a fraction of the demand from the minimum or maximum seat category.

Once a new aircraft is assigned the demand from one or more seat categories, the demand level is quantified as the average of the three market forecasts illustrated in Figure 19 multiplied by a market share assumption. Both of these values are given in Table 10 for each seat category.

Table 10. Seat category definitions and data

Seat category		Narrow bodies				Wide bodies					
		100	125	150	175+	200	250	300	350	400	500+
lower count	seat	90	120	150	175	180	240	275	340	410	530
upper count	seat	110	130	160	250	220	250	305	370	440	550
average forecast quantity, 20 yrs		2071	2828	3370	2630	1100	1534	1472	855	520	767
market share		50%	50%	50%	50%	50%	50%	50%	50%	50%	50%

In the event that there is more than one new aircraft design being considered for simultaneous production, and if two or more such designs fall in the same seat category, the model assumes that the demand for that category is split between the aircraft such that the total is still equal to the market share times the average forecast quantity.

Thus, given a set of new aircraft, represented by their seat counts and range, the above procedure finds a baseline value for expected quantity demanded over a period of 20 years. This number is then adjusted to an initial demand per period of time, depending upon the definition of the period length. In this study, time is discretized by years, so an initial annual demand is found such that the sum over 20 years of the annual demand growing at its expected rate equals the average forecast gross quantity. Expected demand growth rate was not addressed here, as this section is concerned with static analysis; however it is discussed and quantified in the following section as part of the examination of price and quantity dynamics.

5.4.3 Stochastic Demand

While the previous section estimated instantaneous market price and quantity demanded for a new aircraft, this section investigates the time-varying properties of the two and their relationship, if any. The purpose of this investigation is to gather the knowledge required for a complete program valuation algorithm based on dynamic programming, as described in Chapter 6. As will be explained, the dynamic programming valuation approach is founded upon one or more state variables which evolve continuously and randomly, and program management decisions are continuously made, period by period, as the state variable uncertainty is resolved. The uncertainty that drives the decisions, and that is responsible for much of the risk in any aircraft program, is the future condition of demand. If demand were known with certainty, aircraft program design would be straightforward—all the future revenue cash flows would be known at program launch and the program would always be a success. In reality, the best available information about future revenue cash flows is a forecast, which, strictly speaking, will never turn out exactly right. It is therefore important to quantify the amount of uncertainty present in the continuously evolving variable that is aircraft demand.

In Section 5.4.2, demand was considered as a composite of two parts: price and quantity. The assumption that price and quantity are uncoupled is relaxed in this section and a brief investigation is presented into the relationship between the two.

Figure 20 shows a classical microeconomics perspective on price and quantity with a demand curve (D), representing customer preferences, a supply curve (S), representing producers' marginal costs of production, and an equilibrium price and quantity at the curves' intersection. In this model, all else equal, customers are willing to buy more units (q) at a lesser price (p), and vice versa. Therefore, if a producer lowers the price of the good, quantity demanded will rise. This relationship is referred to as the price elasticity of demand. Furthermore, if demand increases—that is, customers are willing to buy more quantity for each possible price—the demand curve will shift to the right, and equilibrium price as well as quantity will rise. In other words, an increase in demand results in an increase in price.

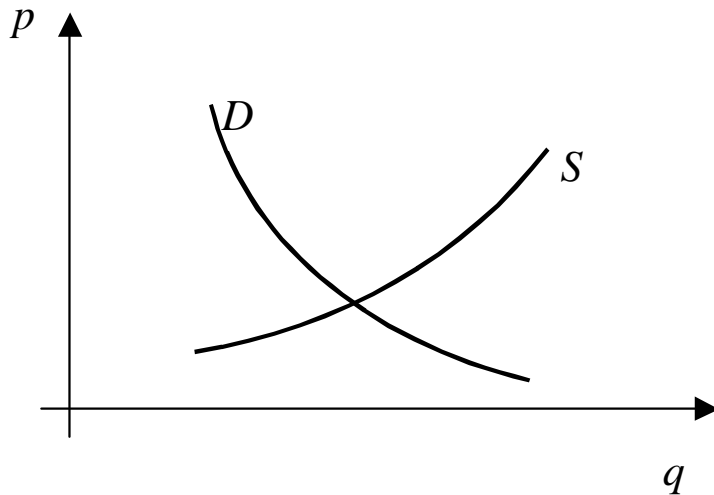


Figure 20. Supply and demand: classical model

The market for commercial aircraft, like most other markets, is subject to unpredictable fluctuations in the demand level. These fluctuations, as expressed through aircraft orders (which become deliveries), are caused by a variety of factors, such as fuel prices, the public's willingness to travel, airline profitability, gross domestic product, interest rates, and general economic conditions. All of the above factors exhibit, to a greater or lesser extent, randomly fluctuating behavior. The question to answer is how such fluctuations

in demand are reflected in quantity and price. If the classical microeconomics framework shown in Figure 20 holds, then changes in quantity demanded should be correlated with changes in price. This hypothesis is explored by comparing historical data for deliveries and prices of commercial aircraft over the past 15 years.

Figure 21 and Figure 22 show estimates of historical prices for new narrow body and wide body aircraft, not adjusted for inflation.⁴² The data shows that, on average, aircraft prices grow at a rate of approximately 1.2% per year. Note that this is below the rate of inflation, which implies that real aircraft prices are falling. This makes sense: as a design becomes older, new aircraft of that design become less valuable.

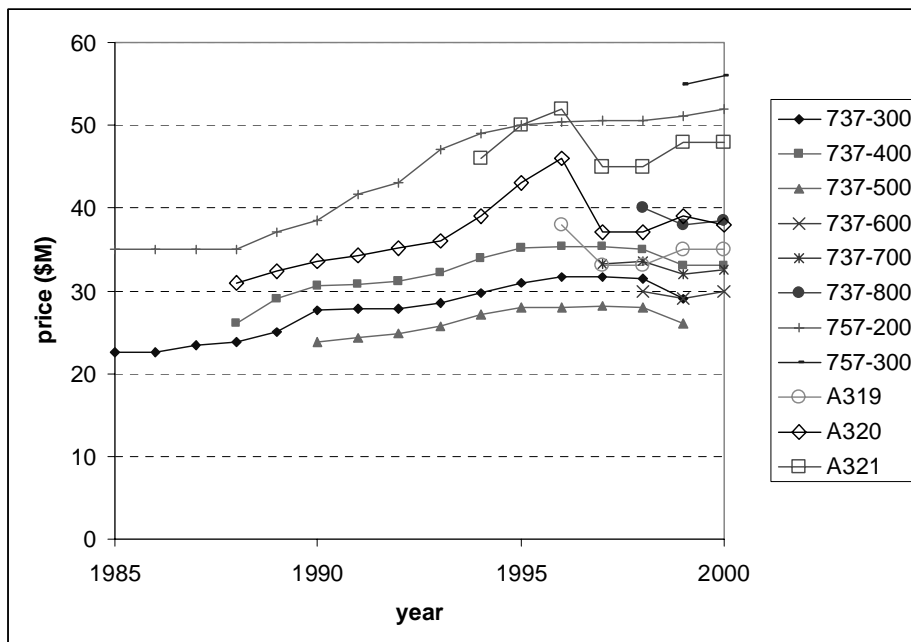


Figure 21. Narrow body aircraft historical prices

⁴² *The Airline Monitor*, May 2001.

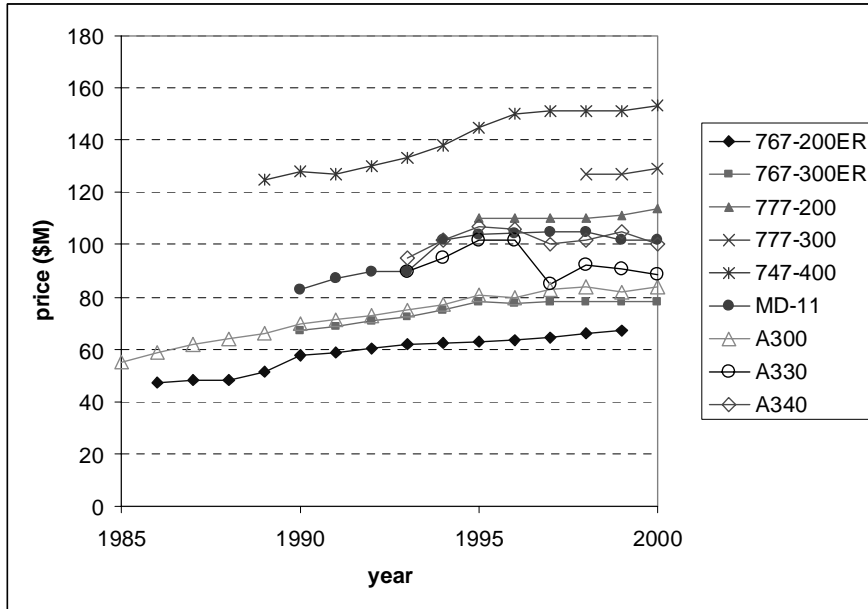


Figure 22. Wide body aircraft historical prices

Figure 23 and Figure 24 show annual deliveries for roughly the same aircraft.⁴³ As can be seen, quantities delivered are considerably more volatile than prices. Whereas prices exhibited a fairly uniform, if weak, upward trend, quantities delivered fluctuate significantly.

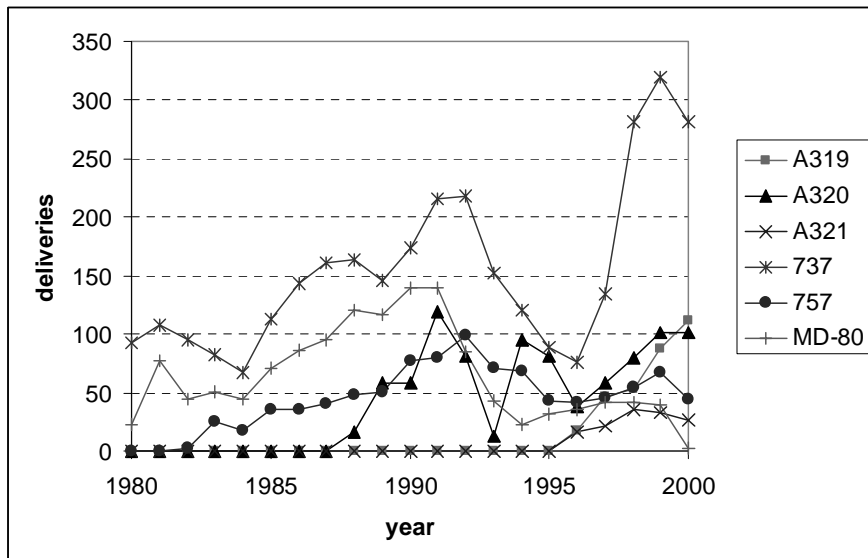


Figure 23. Narrow body aircraft deliveries

⁴³ Data from Morgan Stanley.

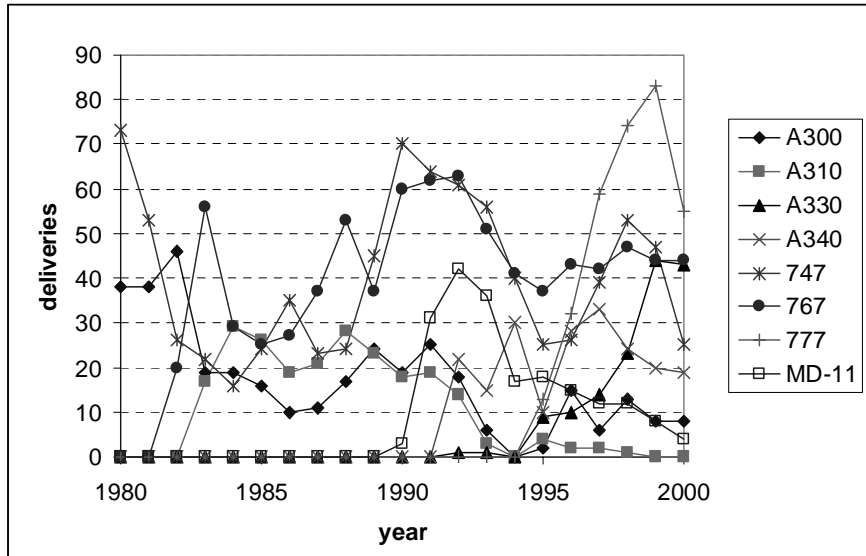


Figure 24. Wide body aircraft deliveries

The next step is to check the two sets of data, price and quantity, for any correlation, possibly with a lead or lag. After a considerable number of comparisons, however, it was concluded that no significant correlation exists between price and quantity. For an example of one of the comparisons made, refer to Figure 25. Here, proportional changes in total quantity delivered for wide bodies are compared to proportional changes in average price charged for wide bodies. Similar analyses were done for narrow body aircraft in the aggregate, all aircraft types in the aggregate, as well as numerous comparisons of quantity and price movements for individual aircraft. No consistent correlation was found, positive or negative. This result is surprising. It seems quite reasonable to expect prices lowered in times of slow demand, and prices increased in boom cycles. That no correlation was found in the data may be a function of the quality of the data and the assumptions of the analysis more than it is a representation of reality. Specifically, the price data is a set of estimates rather than firsthand information. The only publicly available price data is in the form of list prices published by the manufacturers, which are routinely on the order of 10% to 30% above actual sale prices, depending upon the customer's (airline's) clout and negotiating ability. As the actual discounts are held in secret, the list prices are all but useless in tracking the actual price dynamics.

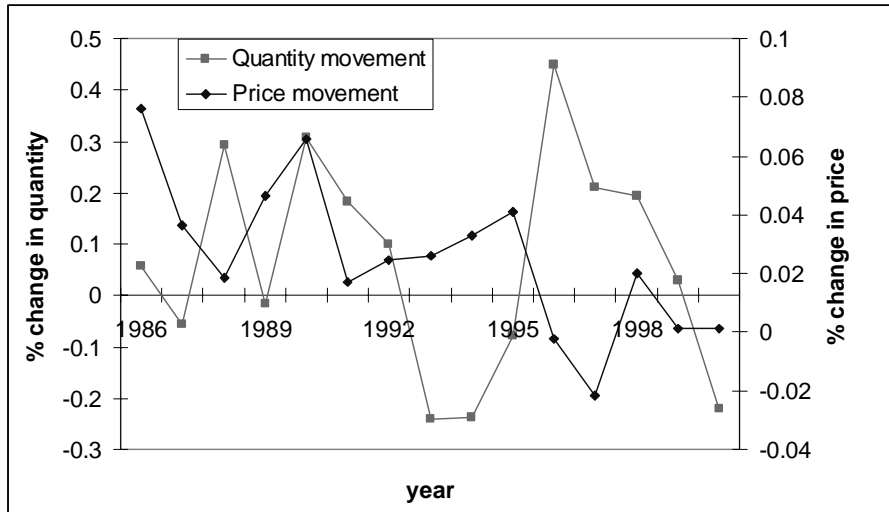


Figure 25. Price and quantity movements for wide body aircraft

In summary, although it is probable that a relationship exists between quantity and price, the available data does not provide sufficient information to quantify such a relationship. Therefore, for the purposes of this study, the simplifying assumption is made that quantity and price evolve independently. Based on the relatively insignificant volatility of prices compared to quantities, it is assumed that quantity demanded is the only source of randomness in the market, whereas price evolves steadily through an annual growth rate. Note that modeling price as a steadily increasing, non-random variable rests upon an assumption made in Section 5.4.1—that buyers are a non-differentiated group with identical preferences. In reality, prices may fluctuate dramatically from customer to customer, especially if there is a large order at stake and the competition is trying to provide a more attractive package. Furthermore, prices may in fact fluctuate with quantities demanded, but the available data does not provide enough information to quantify such a relationship. A more in-depth study of the relationship of price and quantity would be a useful extension to this research.

Figure 26 offers a possible explanation of the price behavior observed above. In this alternate economic model, the demand curve is not gradually downward-sloping, but is a step function. For a given aircraft, a producer may sell any amount, up to a point, but always at the same price—that is, demand is perfectly elastic. The producer is unable to raise the price because of competitive forces and buyer power. However, past a certain critical quantity, there is no price low enough to induce additional sales, because no more

aircraft are required by the customer. In fact, an excessive number of aircraft would hurt the airlines' operations with increased overhead costs and little extra revenue due to congestion. S_1 , S_2 , and S_3 represent several possible short-run supply (marginal cost) curves for the producing firm, which would be driven to invest in capital such that its supply curve is S_2 . Whether this alternate economic model is an accurate representation of the market is not known with certainty—it is simply one theoretical representation of the trends observed in the available data.

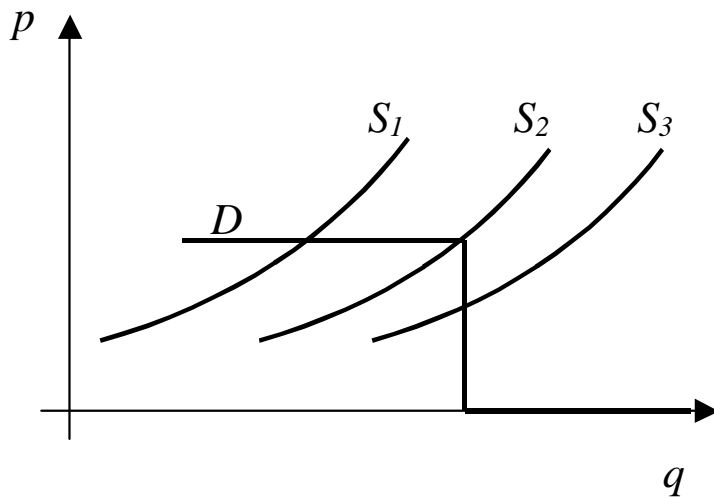


Figure 26. Supply and demand: alternate model

Having made the simplifying assumption that quantity evolves independently of price, it remains to characterize the behavior of quantity demanded. Specifically, the dynamic programming algorithm used for program valuation, which is described in Chapter 6, requires several parameters describing the random variable upon which the analysis is based—in this case, quantity demanded for some given new aircraft design. The required parameters to describe a random variable are average volatility per unit time; average growth rate per unit time; and the correlation coefficient between the random variable and the market portfolio.

The volatility is a measure of the degree of uncertainty regarding future values of the random variable. It defined here as the square root of the variance per unit time of the proportional changes in the variable.

The growth rate is the basis for the current forecast of the variable's expected future value. It may be estimated as the arithmetic average of returns (proportional changes) over multiple periods. On average, the variable is expected to increase at this rate in the future.

The correlation with the market portfolio is a measure of the systematic risk inherent in the variable—that is, a measure of how closely the random variable is linked to the fortunes of the market. A correlation coefficient of 1 (-1) implies that the variable moves in lockstep with (exactly opposite of) the market portfolio, while 0 implies that it evolves completely independently of the market. This measure is necessary to ensure an appropriate discounting scheme to account for the time value of money, as explained in Section 6.4.3.

To find an estimate for volatility, recall the data shown in Figure 23 and Figure 24. From the available data, it is very difficult to conceive a relationship between the characteristics of an individual aircraft and demand volatility for that aircraft. Rather, a representative volatility value is found for all wide body aircraft and all narrow body aircraft, and it is assumed that any aircraft in one of those two classes (narrow or wide) faces the corresponding volatility for that class. Volatility is found as the average of individual aircraft volatilities, weighted by the number of units delivered of each aircraft (see Table 11).

A similar assumption is made for growth rate: this parameter is only a function of aircraft class—wide body or narrow body. However, the growth rate estimation process is somewhat different from volatility. The data in Figure 23 and Figure 24 is fairly non-homogeneous: multiple aircraft have extended periods of zero deliveries. Several of these are new aircraft that did not enter production until recently, and have experienced a ramp-up associated with product introduction. If an average of growth rates were taken over individual aircraft, these newly introduced aircraft would bias the growth rate upward. Therefore, an alternate approach is taken: average growth rates are observed from total deliveries of narrow bodies and wide bodies (see Table 11). In the long run, these should equal the average growth rate for any given individual aircraft. Note that such is not the case for volatilities: while the average volatility for an individual aircraft

is quite high, the volatility of total narrow or wide body deliveries is significantly lower. This effect can be observed in Figure 27, which plots total deliveries for wide bodies and narrow bodies. There is noticeably less fluctuation of these aggregate deliveries than for individual aircraft. This is due to the effect of competition among firms and sometimes cannibalization of products within a firm—that is, one aircraft may often see a large jump in demand at the expense of another. Figure 27 also superimposes best-fit exponential regressions on the delivery data, illustrating the annual growth rates. However, it is important to point out that this regression should not be used for growth rate estimation—the correct growth rate estimate is an arithmetic average of annual growth rates.⁴⁴

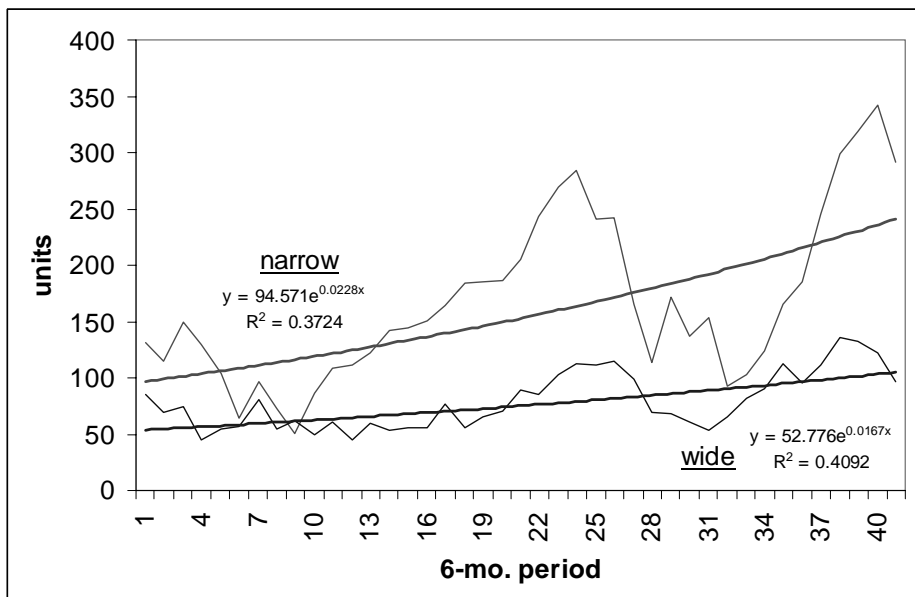


Figure 27. Total deliveries for narrow bodies and wide bodies

To estimate correlation with the market, the assumption is made that the market portfolio is well represented by the Standard & Poor's S&P 500 index. When aircraft deliveries are compared to the S&P 500, several interesting results are found.⁴⁵ First, the aggregate deliveries of wide bodies and narrow bodies do exhibit a nonzero, positive correlation with the S&P 500. However, a sampling of correlations between the S&P 500 and individual aircraft deliveries yields a wide spectrum of results, ranging from negative to

⁴⁴ Brealey & Myers, 1996, p. 147.

⁴⁵ S&P 500 historical data from Datastream Database, available on-line.

positive to insignificant⁴⁶. Obviously, because the correlation is positive for the aggregate deliveries for both narrow bodies and widebodies, the average of the individual correlations should also be positive. However, this value was so small and the data was so sparse that it was decided to simply use a value of zero as a lower limit. The correlation coefficient under this assumption (using individual aircraft delivery correlations with the S&P) is therefore set to zero. Like the phenomenon observed with volatilities, this indicates that while the overall market behaves in one pattern, the pattern is mostly dissolved by the effect of market dynamics and competition on individual aircraft. These results are also summarized in Table 11.

Table 11. Estimated parameters for stochastic behavior of demand

	Narrow bodies	Wide bodies
Average volatility, σ , per annum	42.73%	45.57%
Average growth rate, α , per annum	9.23%	4.43%
<u>Correlation coefficient, ρ</u>		
Total deliveries v. S&P 500	0.269	0.328
Individual deliveries v. S&P 500	0	0

Several points are worth noting about the stochastic demand model. In searching for a correlation between quantity demanded and the market metric of the S&P 500, the assumption is used that quantity demanded represents the present value of an asset, or set of assets, that could potentially be traded in the market. This assumption relies in turn on the assumption that price evolves with certainty at a steady growth rate, independent of quantity demanded. Furthermore, given that price is fixed, quantity demanded does not represent the present value of the profit streams from aircraft production, but the revenue streams. Strictly speaking, it is the profit streams that should be correlated with the market. Finally, note the extremely high values of annual volatility for quantity demanded. While they are based on observed data, they are used together with a random

⁴⁶ Several lead and lag correlations were also attempted, producing similarly inconclusive results.

walk representation of the behavior of quantity demanded. In fact, the demand behavior may be closer to a mean-reverting process, sometimes referred to as a cyclical process. Mean reversion is not modeled in this study, but its effect is to reduce the amount of uncertainty that accumulates over time. It thus lowers the effect of volatility on expected cash flows and on project value. Because of the likely presence of mean reversion in the commercial aircraft market and because of the effect of mean reversion, the high volatilities in Table 11 may be seen as upper bounds. Accordingly, the sensitivity analysis presented in Section 7.5 considers the effect of using a significantly lower volatility.

5.4.4 Summary

To model any market is a non-trivial task, and the commercial aircraft market is no exception. The above analysis contains a great deal of assumptions and simplifications to facilitate a first cut at the problem. Numerous opportunities exist for improvement and increased sophistication, both in estimating baseline market prices and quantities; and in characterizing their dynamics. One interesting expansion of the model would be to explicitly consider the effect of competition on the dynamics of price and quantity. This would involve possible applications of game theory for a duopoly or oligopoly industry structure.

However, it is noteworthy that the goal of the above demand model is not to find the most accurate possible representation of the market, but to provide inputs to the dynamic programming algorithm described in the following section. Specifically, those inputs are a baseline price; an initial quantity demanded per unit time; volatility of demand; expected growth rate of demand; and correlation of demand with the market portfolio. For the purposes of the dynamic programming valuation algorithm as it is implemented in Chapter 6, these inputs completely describe the market for a given new aircraft.

5.5 Conclusion

The preceding chapter described three distinct models: a performance model, a cost model, and a revenue model. Each ties one or more physical characteristics of a potential

new aircraft concept to a functional description: what is the aircraft capable of doing, what will it take to design and built it, and what kind of sales will the aircraft generate? Each of the models is based on information available in the public domain as opposed to higher-fidelity proprietary data that would be available to an aircraft manufacturer. Therefore, while the models are unlikely to provide accurate numerical results, they do exhibit realistic trends, and thus form a useful foundation for further development. The following chapter links the three models described above with a framework based on dynamic programming to find a measure of value for the entire program associated with a potential new aircraft concept.

Chapter 6. Program Valuation

Tool: Synthesis

6.1 Introduction

This chapter presents the approach used to link the three models above—performance, cost, and revenue—and to find program value for a given set of new aircraft designs. Given the performance, cost, and revenue inputs, program value is found by identifying the optimal program strategy in the face of an uncertain future.

This chapter is organized as follows. First, in 6.2, a simplified review is given of the problem to be solved—the aircraft program design problem. Second, in 6.3, a stochastic dynamic programming (DP) approach is introduced. It is described in general, theoretical terms, as applicable to a wide variety of problems, and also described in the context of a variation of DP, referred to as the “operating mode” framework. Next, in 6.4, the DP approach is applied specifically to the solution of the problem described in 6.2. Several variations and extensions are made to the basic DP approach, and the complete valuation algorithm is synthesized. The synthesis includes the adaptation of the performance, cost, and revenue model outputs to the algorithm structure.

6.2 Problem Formulation: Aircraft Program Design

The basic problem to be solved by using the dynamic programming algorithm 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.

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. These conceptual designs are subject to changes by the firm. 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.

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 firm⁴⁷. 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.

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.

⁴⁷ The *baseline*, or initial, price and quantity demanded are both functions of the aircraft’s characteristics.

6.3 Solution Approach: *Dynamic Programming (DP)*

6.3.1 General Theory

Dynamic programming is a well-known and well-explored optimization technique. While there exist numerous ways to outline and define the DP method, the summary presented here is just one perspective, chosen for its simplicity and conciseness⁴⁸.

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

1. **State variables.** These variables evolve with time, or distance, or any other parameter analogous to time that represents a conceptual sequence of events. The set of state variables completely defines the problem at any point in time: all relevant information about the problem is continuously present in the values of the state variables.
2. **Control variables.** These variables are values that are set at any point in time by the decision-maker, that is, the firm. They are the means by which decisions are implemented, and, depending on their value, they have an impact on 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. This process must be quantitatively characterized to solve the problem.
4. **Profit function.** In this context, the goal of the dynamic programming method is to maximize some objective function⁴⁹, in this case the program value. The value will, in general, be a function of certain “profits” incurred every period. These profits are functions of the state variables.
5. **Dynamics.** Dynamics represent the set of rules that govern the evolution of the state variables. These may be a combination of the effects of randomness, the effects of control variables, and any other relationships.

⁴⁸ This view of dynamic programming was suggested by D. Bertsimas, MIT Sloan School of Management.

⁴⁹ The goal could just as well be one of minimization; the formulation is unchanged.

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:

$$F_t(s_t) = \max_{u_t} \left\{ \pi_t(s_t, u_t) + \frac{1}{1+r} E_t[F_{t+1}(s_{t+1})] \right\} \quad (6.1)$$

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 the “Specific Application” discussion 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.

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.” (Dixit and Pindyck, 1994). 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 (6.1) reduces to

$$F_T(s_T) = \max_{u_T} \{ \pi_T(s_T, u_T) \} \quad (6.2)$$

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 (6.1) 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.

6.3.2 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 (p. 171, Trigeorgis, 2000; and Kulatilaka, 1988). 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:

$$F_t(x_t, m_t) = \max_{u_t} \left\{ \pi_t(x_t, u_t) - I(m_t, u_t) + \frac{1}{1+r} E_t[F_{t+1}(x_{t+1}, u_t)] \right\} \quad (6.3)$$

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 (6.1) 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 (6.3) 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. If the risk-neutral valuation technique introduced in Chapter 4 is properly applied, (this implies altering the probability distribution of changes in x to their risk-neutral equivalent) then the appropriate discount rate is simply the risk-free rate—that is, the long-term interest rate on government bonds.

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.

6.4 Applying DP to the Aircraft Program Design Problem

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 6.3; second, a further connection is made to the specific application of dynamic programming to an “operating mode framework”; third, the dynamics of the underlying stochastic process are reviewed in the context of risk-neutral expectations; and finally, the entire algorithm is summarized.

6.4.1 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, as described in the “General Theory” discussion in 6.3.

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. The operating mode for an aircraft is defined in more detail in the discussion below.
2. **Control variables.** For each new aircraft design being simultaneously considered, one control variable exists: the choice of operating mode for the current period. Again, a precise definition of operating mode follows below.
3. **Randomness.** For each new aircraft design being simultaneously considered, one state variable exists with random characteristics: the quantity demanded. It

evolves from a given initial value as a stochastic process, which is discussed in “Stochastic Process Dynamics,” below.

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. The profit functions are specified in more detail in the “Connection to Operating Modes” discussion below.
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 specific characteristics of the evolution of quantity demanded are addressed in the “Stochastic Process Dynamics” discussion below.

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 Section 2.1: 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.

6.4.2 Connection to Operating Modes

Whereas the operating modes introduced in 6.3 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 entail 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 28. 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.

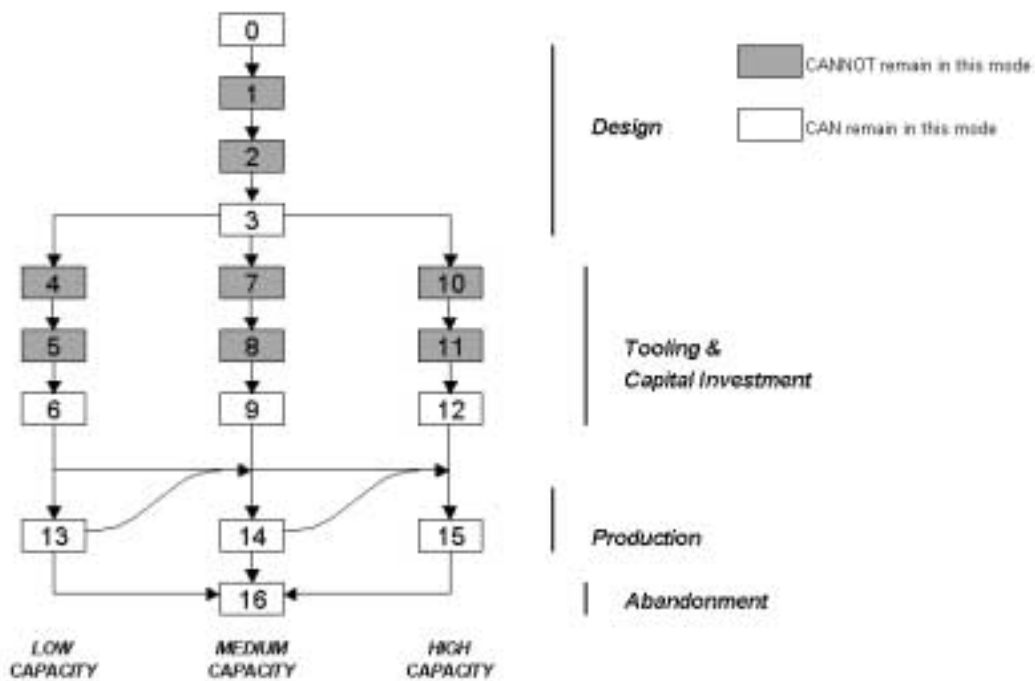


Figure 28. Operating mode framework for a single aircraft

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 market⁵⁰. 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, and they are presented for each example in Chapter 7. 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 cost model described in Section 5.3—both the development cost model and the manufacturing cost model are adapted to the operating mode framework.

Development cost. As explained in 5.3.2, 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

⁵⁰ In fact, the 6-year baseline duration may be altered, as discussed below, depending on previous design experience.

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, as described in 5.3.2. 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). The capacity scaling factors are shown as inputs for the examples presented in Chapter 7.

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 factor⁵¹ 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 29). 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 vector⁵². 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 29. 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*.

⁵¹ 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.

⁵² One possibility would be to include units built to date as an additional state variable, but computation time would suffer drastically as a result.

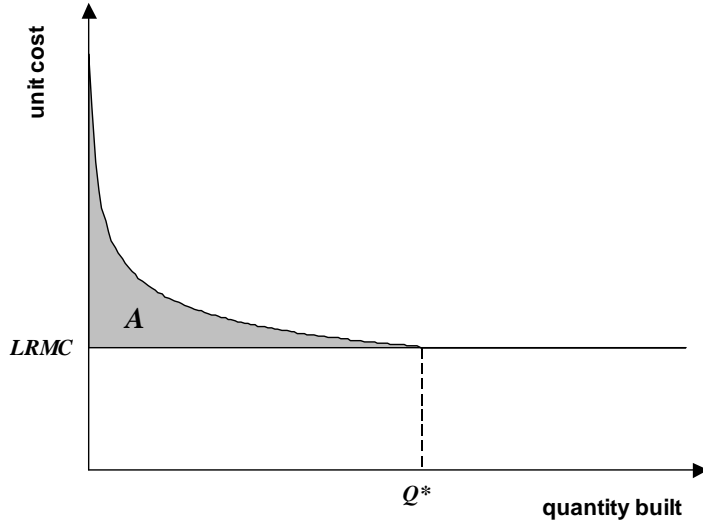


Figure 29. Learning curve effect and long-run marginal cost (LRMC)

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 6.3, “Operating Modes”, 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 $[u_1, u_2]$ is simply equal to the sum of $I(m_1, u_1 | m_2)$ and $I(m_2, u_2 | m_1)$.

6.4.3 Stochastic Process Dynamics and Risk-Neutral Expectations

As outlined in Chapter 4, the risk-neutral expectations method effectively accounts for the uncertainty in an underlying asset by transforming its dynamics into their equivalents in a “risk-neutral world.” The stochastic process followed by the underlying asset is represented as

$$\frac{dx}{x} = \alpha dt + \sigma dz \quad (6.4)$$

where α is the drift rate; σ^2 is the variance rate; and dz is the increment in a Weiner process. The behavior of x may be discretized using a binomial tree approach, where in any given period, the asset, x , will either increase in value to ηx or fall in value to $(1/\eta)x$. In a risk-neutral world, discretizing with time intervals of Δt , the expected rate of return

on x will be the risk-free rate, r_f , and the variance of the expected return must equal $\sigma^2 \Delta t$. A common approach, proposed by Cox, Ross, and Rubinstein (1979), to find acceptable values for η and the risk-neutral probability of a rise, p , is as follows⁵³:

$$\eta = e^{\sigma \sqrt{\Delta t}} \quad (6.5)$$

$$p = \frac{e^{r_f \Delta t} - 1/\eta}{\eta - 1/\eta} \quad (6.6)$$

A binomial tree may thus be constructed with a lattice of possible values of x (x_0 , $x_0\eta$, $x_0\eta^2$, $x_0\eta^{-1}$, ...) linked by the risk-neutral probabilities of rises and falls. This lattice can be used as the basis for calculating the future value expectation in equation (6.1) or (6.3) as a probability-weighted average of the two possible future values at any time: a higher one and a lower one.

However, the above approach is completely valid only if the underlying asset, x , represents a present value that is freely traded in the market. This condition is true for most financial assets, but is not true for many real assets. In particular, this application uses the annual quantity of aircraft demanded as the stochastic underlying “asset” that determines the value of the project. This quantity is by no means traded on the market, and there is therefore no guarantee that its expected growth rate is commensurate with its risk—in other words, its expected growth rate may not equal the equilibrium rate of return it would have if it were a tradeable asset (Kulatilaka, 1988).

If the rate of return of the “asset” (in this case, aircraft quantity demanded) is less than its hypothetical equilibrium rate of return by δ , then in a risk-neutral world, the asset would generate a rate of return of δ less than the risk-free rate⁵⁴ (Amram and Kulatilaka, 1999). In this case, the risk-free rate, r_f , in the exponent of the expression for p in equation (6.6) should be replaced by $(r_f - \delta)$, such that:

⁵³ Hull, 2000.

⁵⁴ δ is sometimes referred to as the “convenience yield.”

$$p = \frac{e^{(r_f - \delta)\Delta t} - 1/\eta}{\eta - 1/\eta} \quad (6.7)$$

The above equation correctly sets the risk-neutral probability of a rise in the underlying asset when the asset is non-traded, which is the case for aircraft quantity demanded.

At this point, it is possible to construct a binomial lattice representing the possible movements and values of the random variable representing quantity demanded: volatility is an input generated by the stochastic demand model in Section 5.4.3; the risk-free rate is observed to be 5.5%⁵⁵, and Δt has been defined as 1 year. The last input to be found is δ , the shortfall between the equilibrium return, denoted here as α^* , and actual return, α , for the asset x . While α is observed as described in Section 5.4.3, α^* must be estimated using the CAPM (Capital Asset Pricing Model), as follows:

$$\alpha^* = r_f + \beta_x(r_m - r_f) \quad (6.8)$$

where β_x is the “asset beta” of x , defined as the ratio of the covariance of the returns of x and the market to the variance of the return of the market,

$$\beta_x = \frac{\sigma_{xm}}{\sigma_m^2} \quad (6.9)$$

and $(r_m - r_f)$, the difference between the market return and the risk-free return, is known as the market risk premium. A commonly used historical average for this quantity⁵⁶ is 8.4%.

Asset beta, as defined above, may be calculated by using the covariance of the changes in aircraft quantity demanded and the returns of the S&P 500, as found in Section 5.4.3. Alternately, asset beta may be set to zero, based on evidence that individual aircraft demand, as opposed to aggregate aircraft demand, is very poorly correlated (if at all) with the market. Thus, a value is found for the shortfall δ :

⁵⁵ Average monthly value for 30-year constant maturity treasury bonds for 2001, from <http://www.economagic.com/em-cgi/data.exe/fedbog/tcm30y>.

⁵⁶ Brealey and Myers, 1996, p. 146.

$$\delta = \alpha^* - \alpha = r_f + \frac{\sigma_{xm}}{\sigma_m}(r_m - r_f) - \alpha \quad (6.10)$$

To summarize, the above discussion defines a scheme for forming a risk-neutral binomial tree to model the dynamics of aircraft quantity demanded as it affects program value. The scheme consists of a binomial tree formed with the risk-neutral probability of a rise⁵⁷ as defined by equation (6.7), and with the proportional amount of the rise (inverse of a fall) defined by equation (6.5).

However, this framework addresses just one aircraft demand quantity. The algorithm needs to model two simultaneously evolving random variables describing the quantities demanded for two different aircraft. In general, it is reasonable to assume that these two variables will have a nonzero correlation⁵⁸. The two-variable risk-neutral formulation is constructed as follows. The stochastic process for one aircraft quantity (“aircraft 0”) is assumed to be unchanged: the values of p (risk-neutral probability of rise) and η_0 (the “rise” factor) are calculated as before. However, the process for the other aircraft quantity (“aircraft 1”) is defined as a function of the change in aircraft 0. That is, the risk-neutral probability of a rise in aircraft 1 depends upon whether aircraft 0 rose or fell. Let v be the risk-neutral probability of a rise in aircraft 1 *given* that aircraft 0 rose; and let w be the risk-neutral probability of a rise in aircraft 1 *given* that aircraft 0 fell. Finally, let η_1 be the “rise” factor for aircraft 1, independently of the aircraft 0 movement.

The problem is to find values for v , w , and η_1 such that the expected return of aircraft 1 is $(r_f - \delta_1)^{59}$; the variance per unit time of aircraft 1 is σ_1 ; and the covariance of the returns of aircraft 0 and 1 is equal to some specified value, σ_{01} . The following equations satisfy the above conditions:

⁵⁷ The risk-neutral probability of a fall is equal to $1-p$.

⁵⁸ However, from the analysis in Section 5.4.3, it is difficult to predict a value for this correlation.

⁵⁹ Recall that δ_1 is the shortfall between the equilibrium rate of return, α^* , and the actual rate of return, α , for the “aircraft 1” process.

$$\eta_1 = e^{\sigma_1 \sqrt{\Delta t}} \quad (6.11)$$

$$v = \frac{A_1 \eta_1 - 1}{\eta_1^2 - 1} + \frac{\eta_0 \eta_1 \sigma_{01}}{p(\eta_0^2 - 1)(\eta_1^2 - 1)} \quad (6.12)$$

$$w = \frac{A_1 \eta_1 - 1}{\eta_1^2 - 1} - \frac{\eta_0 \eta_1 \sigma_{01}}{(1 - p)(\eta_0^2 - 1)(\eta_1^2 - 1)} \quad (6.13)$$

where $A_1 = e^{(r_f - \delta_1) \Delta t}$. Note that η_1 is calculated in the same way as η_0 . As a check on the accuracy of the equations, it can be observed that as the covariance, σ_{01} , increases, v increases and w decreases. Also, when $s_{01} = 0$, that is, when the two processes are independent, the expressions for the (conditional) risk-neutral probabilities v and w reduce to

$$v = w = \frac{A_1 \eta_1 - 1}{\eta_1^2 - 1} = \frac{e^{(r_f - \delta_1) \Delta t} - 1/\eta_1}{\eta_1 - 1/\eta_1} \quad (6.14)$$

This is the same expression as equation (6.7): the risk-neutral probability p for an independently evolving random variable.

A transition probability matrix can now be constructed linking through a risk-neutral probability each possible pair of demand quantities to 4 distinct outcomes for the future period (1. up, up; 2. up, down; 3. down, up; 4. down, down). For example, the risk-neutral probability of an (up, down) outcome equals the product of the probability of a rise for aircraft 0 (p) and the conditional probability of a fall for aircraft 1 given that aircraft 0 rose ($1 - v$). The transition probability matrix (TPM) is simply a record of each such product. While the above formulation results in risk-neutral probabilities that are independent of the current state, the TPM structure allows for variations in the risk-neutral probabilities depending upon the pair of demand quantities from which the up/down jumps are about to occur. This can be used to model more complex processes than the random walk (Brownian motion) assumed above—for example, mean reversion, where the shortfall δ will be a function of the state.

The transition probability matrix is used to compute the expected future value of the project in the backwards-iterative solution process, as prescribed by equation (6.3).

6.4.4 Algorithm Review

The above section, 6.4, has described the dynamic programming algorithm as it is applied in this work to solving the problem of optimal decision-making and valuation of an aircraft program. To summarize the algorithm, it is presented here as a sequence of steps and iterations.

1. Perform preliminary calculations.
 - a. Cost model \rightarrow construct switching cost matrix.
 - b. Static demand model \rightarrow baseline price, quantity.
 - c. Stochastic demand model \rightarrow transition probability matrix.
2. Start program valuation calculation at $t = T$ (final time period).
3. Cycle over all values of random variables (aircraft demand quantities).
 - a. Cycle over all values of operating mode from last period.
 - i. Cycle over all values of control variables.
(operating mode decisions)
 1. Find profit function.
 2. Find discounted expected future value.
 3. Compare their sum to current “maximum value.”
 - ii. Record “maximum value” for this aircraft demand, operating mode, and time.
 - iii. Record control variable (operating mode choice) associated with “maximum value.”
4. If $t > 0$,
 - a. Step backwards 1 time period.
 - b. Repeat step 3.
5. If $t = 0$, “maximum value” for baseline quantity demanded equals program value.

For a summary of all the input parameters used for the program valuation tool, refer to the Appendix, which shows a printout of the Excel spreadsheet used as the input interface for the tool.

6.5 Conclusion

This chapter introduced and developed an algorithm, based on dynamic programming theory, to find the value of an aircraft program, defined as a portfolio of aircraft designs accompanied by the plans and necessary resources to bring them to market. The dynamic programming approach 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.

Together, the preceding two chapters 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 product and program flexibility, which has the potential of having a great impact on value. Program flexibility is modeled and addressed directly by the dynamic programming “operating modes” formulation presented in 6.4, while product flexibility is addressed indirectly, through the inclusion of

two aircraft designs in the valuation tool. The simultaneous program valuation of two designs allows the optimization to identify and exploit any synergies they may have—for example, part commonality—and quantify the value of such synergies. Conversely, it may find that the synergy value is insignificant. In either case, the valuation tool performs a useful service.

The following chapter provides several concrete examples of this tool in use, demonstrating the nature of the outputs, the tradeoffs that may be modeled, and the insights that may be gained from the tool.

Chapter 7. Examples

7.1 Introduction

The following chapter demonstrates the program valuation tool through several examples and associated trade studies. As the tool itself and the data in the examples are based on public domain information, and a reasonably modest level of fidelity where cost and price estimation are concerned, the examples should in no way be seen as actual valuations. Rather, they are meant to illustrate the operation of the valuation tool and to suggest some trends and phenomena present in the dynamics of aircraft program value.

The chapter is organized into several sections. First, some background is given on the aircraft design used in the examples. Three versions of the aircraft are introduced: one high-capacity configuration, and two alternate lower-capacity configurations. Second, the operation of the program valuation tool is demonstrated on each of the aircraft versions—first separately, then in parallel. Third, a connection is made to traditional project valuation techniques to highlight the distinguishing features of the program valuation tool. Finally, a sensitivity analysis is conducted to investigate the effects on program value of CAROC (Cash Airplane-Related Operating Cost), demand volatility, and correlation between quantities demanded for different aircraft.

Each of the sections below summarizes the key inputs and assumptions used for its particular calculations, but several parameters that remain unchanged throughout the analysis are shown in Table 12. Obviously, a complete list of model parameters is impractical here, as it would involve all the inputs used in the cost and revenue models from Sections 5.3 and 5.4.

Table 12. Key input parameters used for all cases

Number of periods	30
Timestep per period	1 year
Number of demand values	60
Initial demand index	30
Risk-free rate, r_f	5.5%
Aircraft price inflation ⁶⁰	1.2%
Maximum annual production rates:	
low capacity	48 units/year
medium capacity	96 units/year
high capacity	192 units/year
Non-recurring cost scaling factors ⁶¹ :	
low capacity	0.75
high capacity	1.66
Tooling disruption scaling factor ⁶²	1.25

At this point, a note is necessary regarding notation. Table 12 refers to the number of demand values. This corresponds to the number of distinct values into which the underlying asset (annual aircraft quantity demanded) has been discretized. In this case, there are 60 possible values, corresponding to a time period of 30 years and a timestep of 1 year: starting from a given level in year 0 and jumping up or down each year for 30 years, demand may evolve to equal any one of 60 possible values at the end of the time

⁶⁰ This inflation factor, taken from *The Airline Monitor* (2000) is also applied to cash flows from non-recurring and recurring costs: it is assumed that the real prices charged by suppliers to the firm track the real prices of aircraft.

⁶¹ Used to calculate switching costs corresponding to tooling for low- or high-capacity production as a function of the baseline (medium capacity) non-recurring costs. See Section 6.4.2 (p. 102) for more detail.

⁶² Used to calculate switching costs to increase capacity once a production line is already in place.

horizon. Recall that this random walk process, as described in Section 5.4.3, is based upon proportional (not arithmetic) jumps every period—either increasing by a factor of η , or falling by a factor of $1/\eta$. Therefore, the level of aircraft demand in any given period may be referred to in one of two ways: either explicitly, by the number of aircraft units demanded, or by referring to a demand index representing the “location” in the lattice of possible demand values. Equivalently, the index may be thought of as the logarithm of the aircraft demand. For example, if the baseline quantity demanded is 100 aircraft per year; the “initial demand index” corresponding to the baseline quantity is 30; and the volatility is such that $\eta = 1.1$; Table 13 lists several “demand index” values and their corresponding demand quantities. Note that the baseline quantity of 100 was chosen arbitrarily for this example. For actual cases, the baseline quantity is calculated by the quantity estimator of the static demand model, as described in Section 5.4.2.

Table 13. Example: “demand indexes”

Demand index	27	28	29	30	31	32	33
Demand quantity (units/year)	...	82.64	90.91	100	110	121	...

7.2 Background: BWB

The Blended-Wing-Body (BWB) concept, presented here as an example aircraft, was originally developed at McDonnell Douglas, under funding from NASA⁶³. The BWB is a proposed new concept for air transport, combining the wings and fuselage into one integrated structure. Figure 30 illustrates one variant of the BWB design.

⁶³ Liebeck, 2002.

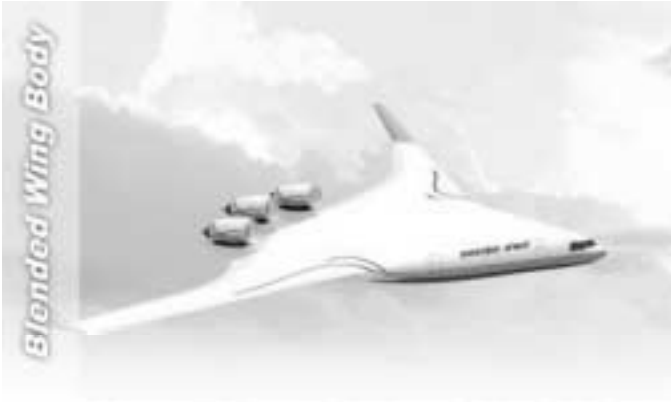


Figure 30. Blended-Wing-Body (BWB) aircraft⁶⁴

The BWB offers several potential advantages over conventional air transport aircraft. Its design holds the promise of improved structural efficiency due to span-loading, as well as improved aerodynamic performance due to reduced wetted area. These improvements, if realized, would result in a lower weight and a lower fuel burn for a given range and passenger capacity. The BWB design is well-documented by Liebeck (2002), Liebeck et al. (1998), and the “Blended-Wing-Body Technology Study” (1997).

Among its other virtues, the BWB design is well amenable to stretching—that is, to incremental changes in its design to increase (or decrease) passenger capacity. As explained by Liebeck (2002), it is conceivable to create an entire family of BWB aircraft, ranging from 200 seats to 450 seats, with a common set of wings and a varying number and size of fuselage bays used to make up the centerbody between the wings.

The design and construction of a number of highly common aircraft offers several potential benefits, such as savings in non-recurring development cost and an accelerated progression along the learning curve in manufacturing. On the other hand, enforcing commonality between several different aircraft designs typically results in weight penalties, as commonality imposes additional constraints on the designer. Weight penalties, in general, translate into additional cost and/or fuel burn. In this context, the following question arises: how much is commonality worth?

⁶⁴ <http://www.geocities.com/witewings/bwb/>

The examples below use three hypothetical BWB variants, designated here as BWB-450, BWB-250C (“Common”), and BWB-250P (“Point design”). The variants are a 450-passenger class aircraft and two 250-passenger class aircraft, respectively. The BWB-250C is assumed to share a number of components with the BWB-450; specifically, the wings, winglets, and fuselage bays. The BWB-250P is assumed to be designed from the ground up without any commonality considerations—it is optimized as a point design based on minimum gross takeoff weight. All three designs are processed in WingMOD, as described in Section 5.2, and the resulting outputs are provided to the program valuation tool. Several primary characteristics of the designs—highlights of the WingMOD outputs—are summarized in Table 14, and the vehicles are illustrated in Figure 31.

Table 14. Selected features of BWB example designs

	BWB-450	BWB-250C	BWB-250P
Seat count	475	272	272
Maximum range	8550	8550	8550
<u>Weight breakdown⁶⁵</u>			
Cockpit ⁶⁶		0.0115	
Fus. bay 1 (x2)	0.0347		
Fus. bay 2 (x2)	0.0187	0.0187	0.0237
Fus. bay 3 (x2)	0.0138	0.0138	0.0127
Fuselage total	0.1344	0.0766	0.0728
Inner wing	0.0451	0.0451	0.0319
Outer wing	0.0111	0.0111	0.0065
Winglets	0.0012	0.0012	0.0008
Wing total	0.0574	0.0574	0.0392
Structures total	0.1918	0.1339	0.1120
Landing gear	0.0420	0.0317	0.0262
Propulsion	0.0677	0.0451	0.0451
Systems	0.0525	0.0359	0.0415
Payloads	0.1039	0.0615	0.0615
Total: Operating Empty Weight	0.4579	0.3082	0.2863
Maximum Takeoff Weight	1.0000	0.7557	0.6244
Design Landing Weight	0.6080	0.4600	0.3750

Note that several parts in the weight breakdown have identical weights for the BWB-450 and BWB-250C: Fuselage bays 2 and 3; the inner wing; the outer wing; and the winglets. These components are assumed to be identical for both airframes, and thus define the commonality built into the designs.

⁶⁵ All weights given in this table are normalized by the BWB-450 maximum takeoff weight.

⁶⁶ The cockpit is counted as part of fus. bay 1 and fus. bay 2, respectively, in the BWB-450 and the BWB-250P; while it is accounted for explicitly in the BWB-250C.

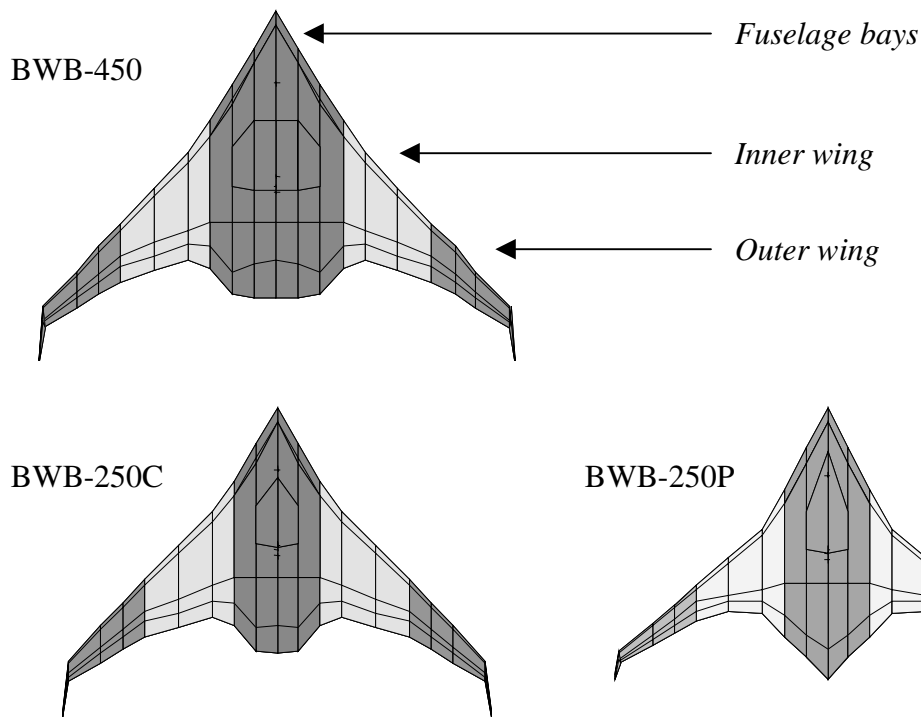


Figure 31. BWB example configurations

The examples below do not attempt to answer the question of the value of commonality. The models and assumptions used in this work do not have the required accuracy. However, the examples illustrate techniques which are applicable to the program value question in general, and to the commonality question in particular. If properly applied, these techniques will yield insight.

7.3 Baseline Analysis

Each of the examples is evaluated using a set of baseline assumptions for the program valuation tool. Some of the assumptions are relaxed or varied in sections below. First, each of the three designs is input into the program valuation tool individually, such that there is only one aircraft available in the firm’s “portfolio.” Then, the tool is given two aircraft designs simultaneously—the BWB-450 paired with one of the BWB-250 variants. Table 15 lists the key inputs used for all three designs throughout this baseline analysis.

Table 15. Baseline analysis: key inputs

Aircraft demand volatility, σ	45.6%
Equilibrium return shortfall, δ	2.19%
(fuel cost / CAROC) fraction	20%

The baseline volatility is quite high. As discussed in Section 5.4.3, it represents an average measured value for the historical sales records of several commercial aircraft. Also as discussed in Section 5.4.3, the baseline shortfall, δ , corresponding to this volatility, is used to adjust the risk-neutral probability to reflect the fact that aircraft demand is a non-tradeable asset. The fuel cost fraction represents an assumption that fuel costs represent a constant percentage of CAROC—in this case, 20%. This is a simplification from the CAROC model presented in the “Price” discussion in Section 5.4.2, and as such it is relaxed in Section 7.5 below, which considers the model’s sensitivity to CAROC and other calculations.

7.3.1 BWB-450

This calculation assumes that the only design available to the firm is that of the BWB-450. In this case, the solution found by the optimizer to the dynamic programming problem may be partially illustrated in a plot, as follows. Figure 32 shows, as a function of time, the *minimum* value of the demand index for which it is optimal to make certain program-related decisions—namely: design (enter mode 1 from mode 0); build (enter mode 13, 14, or 15 from mode 6, 9, or 12); switch from low- to medium-capacity production (enter 14 mode 13); and switch from low- to high-capacity production (enter 15 from 13)⁶⁷. Recall that the demand index is a logarithmic measure of the annual quantity demanded of the BWB-450. A demand index value of 30 represents the static forecast of annual demand—that is, whatever quantity is calculated by the static demand estimator in Section 5.4.2. Referring to Table 16 at the end of this Baseline Analysis (Section 7.3), this quantity happens to be 16.7 units per year for the BWB-450.

⁶⁷ Refer to Section 6.4.2 for mode definitions.

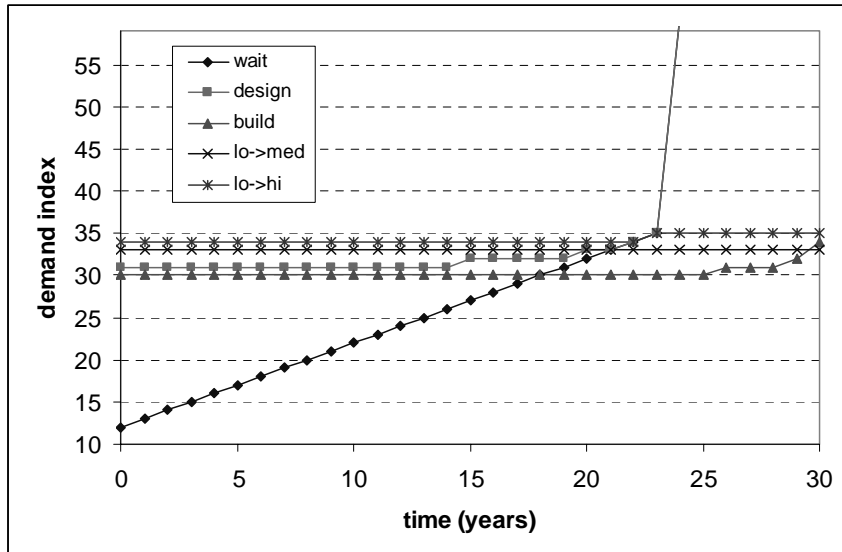


Figure 32. Decision rules for BWB-450, baseline

The third decision illustrated on the plot is the “decision” to wait—not really a decision, because no expenditure is required to remain in mode 0 (do nothing). The “wait” decision curve is simply the lowest demand index, as a function of time, for which the expected value of the program is nonzero. Note that this condition is necessary, but not sufficient, for the firm to commit to designing the aircraft. For some demand indexes at some times, the program has a positive expected value but the optimal decision is to continue waiting until demand reaches a sufficiently high level. The intuition behind this is that demand may go up, but it may also go down, in which case the firm is better off not investing at all rather than sinking money into a nonexistent market. In this case, the “wait” decision curve is linear for most of the time horizon, because the binomial tree assumption prevents the demand index from ever jumping up by more than one unit. Note that in year 24—6 years before the end of the time horizon—the “wait” decision curve jumps off the scale. If by year 24 the firm has still not begun design of an aircraft, there will not be enough time to make any sales, regardless of the condition of the market. Recall that the baseline assumption for the duration of the non-recurring development process is set at 66 months (Section 5.3.2), which is rounded to 6 years.

The above decision rules provide a glimpse of the overall solution, which specifies the optimal decision (mode selection) for every possible initial mode, for every time, and for every state. The maximized objective function accompanying this solution yields a

program value of \$5.95B. This number, along with several other relevant outputs of the valuation tool, is shown in Table 16 at the end of Section 7.3.

Another interesting representation of the solution is shown in Figure 33. This is a plot of program value as a function of demand index and time, assuming that the firm is in mode 0—that is, assuming that the firm has done nothing to date.

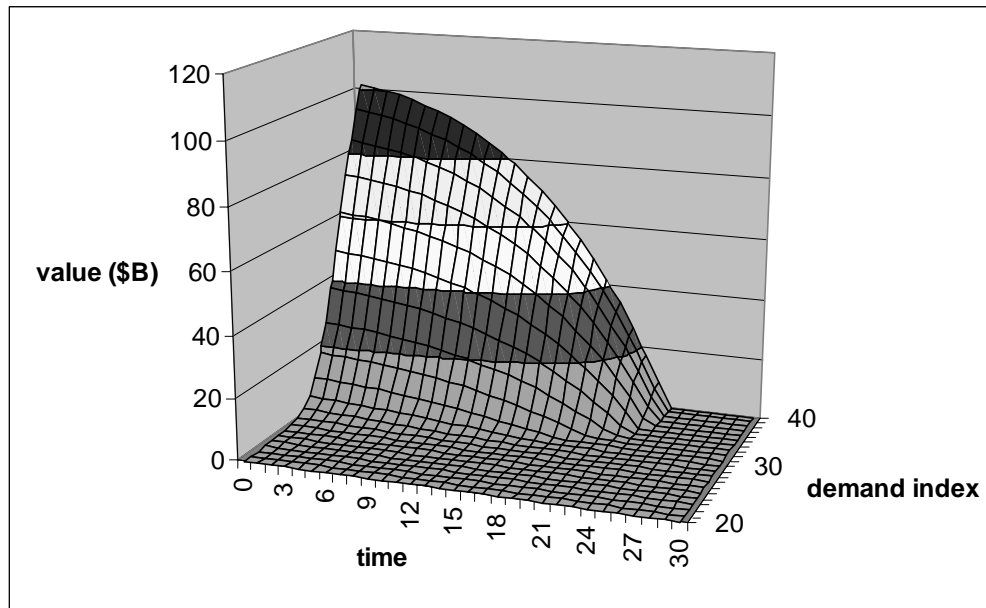


Figure 33. BWB-450 program value

One surprising feature of the plot is the extreme range of possible values—the peak is over \$100B. This number seems quite unrealistic, if not absurd. The value of the program is nowhere near \$100B. Rather, that data point represents the hypothetical situation where the initial demand index is 40. Recall that the initial demand index is in fact 30, and it would be impossible—by the assumptions of the binomial tree model—for the index to reach 40 earlier than 10 years into the time horizon. Furthermore, a demand index of 40 corresponds to a quantity demanded of 1592 aircraft per year. This points to a weakness of the model: the assumption of a random walk, where the proportional jumps η are an increasing function of the observed demand volatility. The high volatility used in this calculation, 45.6%, results in a high value for η and consequently very large possible outcomes for demand quantity. The only way to reconcile such high demand quantities is to recognize that they are theoretically possible, but extremely unlikely.

Thus, if the firm does nothing for 10 years (not an optimal strategy), while demand continuously increases (jumps up every period), then according to the plot the value of the program to the firm in year 10 will be \$84.8B. However, this number has little physical significance. In general, the plot shows that given a certain demand index for which the program value is positive, waiting will decrease program value. This is a fairly obvious result, given the assumption that time runs out after 30 years. The other trend clearly demonstrated is the nature of program value as a function of demand index, given a particular time. Value is never negative, because a negative expected value at a particular time and demand index would prompt the optimizer to select to do nothing (remain in mode 0) instead of investing. This does not mean that a firm following the optimal decision rules will never lose money. Rather, it demonstrates the model's acknowledgment of management's ability to defer investment if the *expected* value at that time is negative.

7.3.2 BWB-250C

This calculation focuses on a 250-passenger class BWB. The BWB-250C shares several components with the BWB-450: its 4 fuselage bays, its inner wing, outer wing, and winglets. As a result, this aircraft is heavier than the BWB-250P, its point-designed counterpart.

The dynamic programming problem is solved in the same way as described for the BWB-450 above. The decision rules demonstrate similar trends as the ones found for the -450, as shown in Figure 34.

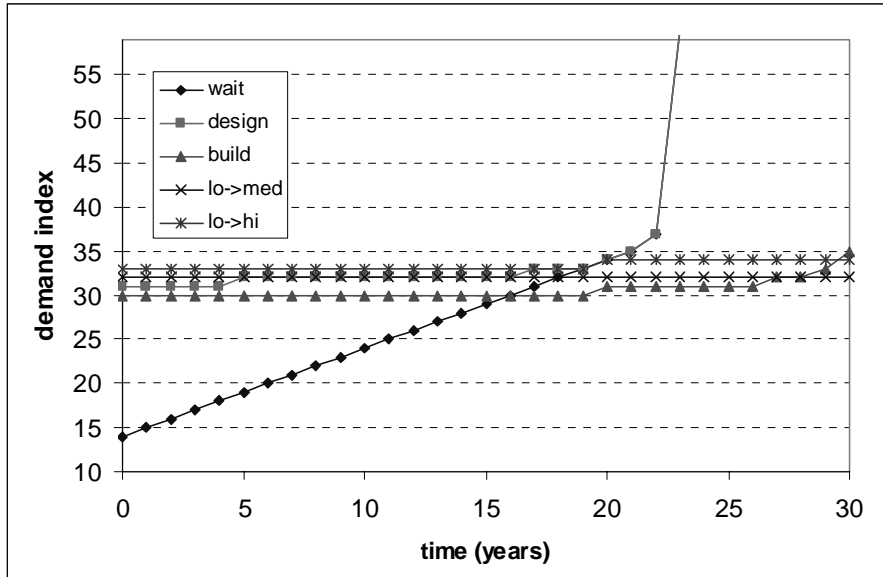


Figure 34. Decision rules for BWB-250C, baseline

To better illustrate the meaning of the decision rules, Figure 35 presents a simulation run. 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 as described in Section 6.4. 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 static demand quantity estimator in Section 5.4.2. This demand level, which happens to be 27.6 units per year (see Table 16), is insufficient for the firm to commit to developing the BWB-250C (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 Figure 34⁶⁸. Recall that once the initial “design” operating mode is entered, the firm is committed to the first phase of the development process⁶⁹, 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.

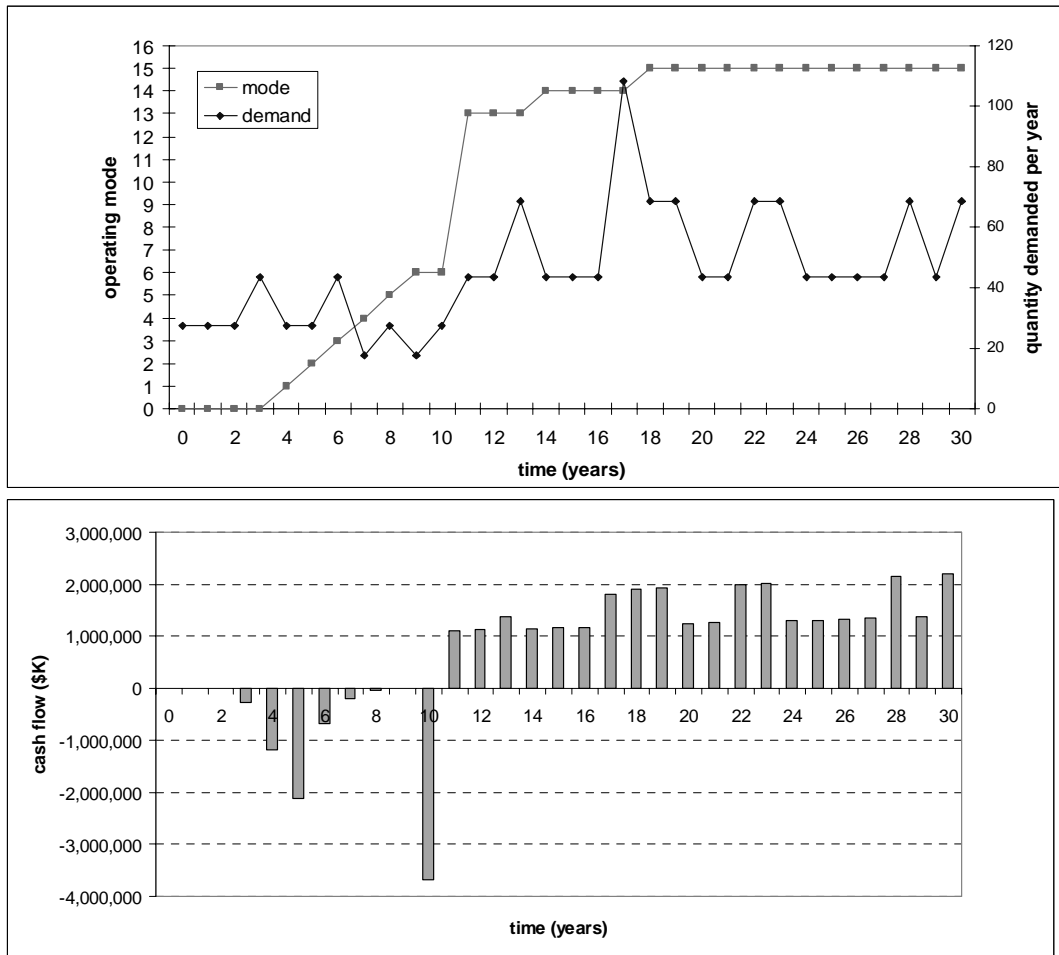


Figure 35. Simulation run for BWB-250C

⁶⁸ Note that Figure 34 refers to demand level by index, while Figure 35 refers to demand level by actual quantity.

⁶⁹ See Section 6.4.2.

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 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-250C production. Once production is entered, after year 10, all units are produced at their long-run marginal cost⁷⁰. 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 general solution for the overall program value of the BWB-250C, as a function of time and demand index, behaves similarly to the BWB-450 value plot of Figure 33, only with a different magnitude. The result for program value for the baseline BWB-250C is \$2.26B.

Table 16, at the end of this section (7.3), shows several other relevant outputs: long-run marginal cost, baseline price, and baseline quantity demanded. The latter, an output of

⁷⁰ In reality, the \$4B would be distributed over the entire production run, with more weight on the early years.

the quantity estimator in the static demand model from Section 5.4.2, corresponds to the initial demand index of 30. This is the “root” of the binomial tree lattice representing the possible movements in demand.

7.3.3 BWB-250P

The point-designed version of the 250-passenger class BWB is significantly lighter in empty weight (OEW) than the common version. Because the cost model is fundamentally based on weight, this translates into a somewhat reduced cost—both from the development and manufacturing perspective. In the context of another aircraft, such as the BWB-450, these advantages may be outweighed by the synergies of the –250C through commonality. However, this part of the analysis treats the –250 versions as standalone projects, and as such the –250P is clearly superior in terms of low cost. Furthermore, in addition to its cost advantage, the –250P also gains a price advantage through its lower weight. Its slightly lower empty weight results in a significantly lower maximum takeoff weight, which results in a significantly lower fuel burn for a given range. Fuel burn is converted to fuel cost and then to CAROC according to the fuel cost fraction presented in Table 15; and CAROC is incorporated into the baseline price calculation as described in Section 5.4.2. As a result, the BWB-250P commands a much higher price than its all-common counterpart, as shown in Table 16. In fact, comparing the price to the long-run marginal cost, the –250P enjoys a profit margin of ~67%, compared to ~23% for the –250C⁷¹.

The lower cost and higher price—maybe unrealistically so—of the BWB-250P, as modeled, result in a much higher standalone program value than the BWB-250C: \$14.62B.

7.3.4 BWB-450 & BWB-250C

This calculation determines the optimal decision strategy and program value for a portfolio of designs that includes the baseline BWB-450 and the derivative BWB-250C, which shares fuselage bays, wings, and winglets with the larger aircraft. The baseline

⁷¹ Corporate taxes are not included in this model.

assumption is made that the coefficient of correlation between the demand quantities for the two aircraft equals 0.5. As mentioned in the “Stochastic Process Dynamics” discussion of Section 6.4, it is difficult to observe a well-founded estimate for this parameter, therefore it is arbitrarily set to 0.5. As long as the two aircraft designs aren’t in direct competition, it is reasonable to assume that there will be some degree of positive correlation in their demand, driven by global economic factors affecting the entire air travel industry.

It is difficult to plot a representative set of decision rules as in Figure 32, because decisions now depend on two random variables evolving simultaneously. One representation of the results of the optimization is shown in Figure 36: a plot of program value at time 0, as a function of the initial demand indexes for both aircraft. The number corresponding to the demand index pair (30,30) is the actual program value: \$8.95B. Note that the program value is greater than the sum of the values of the standalone programs for the BWB-450 and BWB-250C: $\$5.95\text{B} + \$2.26\text{B} = \$8.21\text{B}$. The premium in value due to commonality is $(8.95/8.21 - 1)$, or 9%.

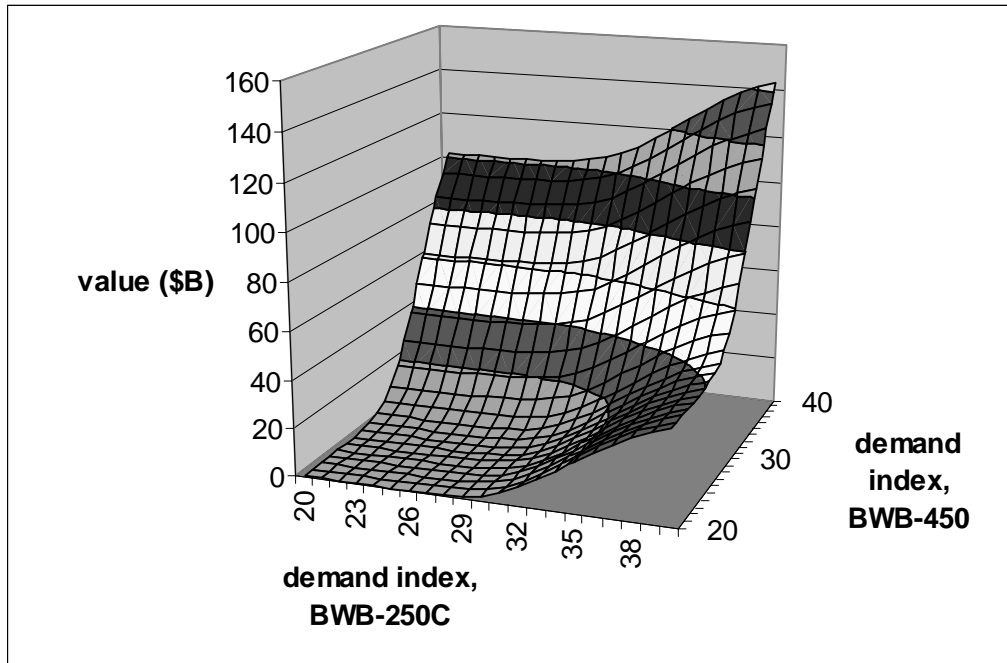


Figure 36. Program value for BWB-450 & BWB-250C

The commonality premium, 9%, is surprisingly low. In fact, the premium is not large enough to be obvious in the above plot, where it seems like the value for the demand index pair (40,40) is very close to the sum of the values for (40,0) and (0,40). Therefore, given the baseline assumptions, the value of commonality is in this case rather small. This may be partly explained by the following observation. In a world with no uncertainty, the development and production schedules of the two aircraft would be planned out to maximize the savings due to commonality. First, one aircraft would be developed, then the other, for a reduced cost. This cost savings would be guaranteed. However, in a world with uncertainty, conditions may change after the first aircraft is developed and the second aircraft may be deferred indefinitely. The benefits of commonality are only realized in some of many possible scenarios. The high volatility of demand used in the baseline assumptions exacerbates this effect and weakens the impact of commonality.

7.3.5 BWB-450 & BWB-250P

This calculation finds the optimal set of decisions and associated program value for a “design portfolio” consisting of the BWB-450 and the BWB-250P. These are two completely distinct aircraft with no commonality. As expected, the resulting program value is exactly equal to the sum of the values of the standalone programs for the –450 and –250P separately: \$20.57B. Note that this observation is unaffected by the correlation in demand between the –450 and the –250P, because there is no coupling: the two aircraft programs have no effect on each other.

Table 16. Baseline analysis: key outputs

	Baseline quantity demanded (units/yr)	Long-run marginal cost (\$M)	Baseline price (\$M)	Program value (\$B)
BWB-450	16.7	139.0	195.0	5.95
BWB-250C	27.6	93.8	116.1	2.26
BWB-250P	27.6	84.9	142.2	14.62
BWB-450 & BWB-250C				8.95
BWB-450 & BWB-250P				20.57

7.4 Connection to Discounted Cash Flow (NPV)

As mentioned before, 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.

It should be noted that nowhere in the model is a risk-adjusted discount rate explicitly calculated. The dynamic programming algorithm uses the risk-free rate to discount cash flows, which are probability-weighted averages calculated using risk-neutral expectations, as described in Chapter 4 and Section 6.4.

Given the assumptions of the model, if the risk-neutral expectations and risk-free rate are applied to a pre-determined set of decisions, the resulting program value should be equivalent with traditional NPV. 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 28 (p. 103), the only finite-cost path through the modes is now 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 37 is a plot analogous to Figure 33, showing program value for the BWB-450 as a function of time and demand index. In this case, program value is negative for numerous combinations of time and demand.

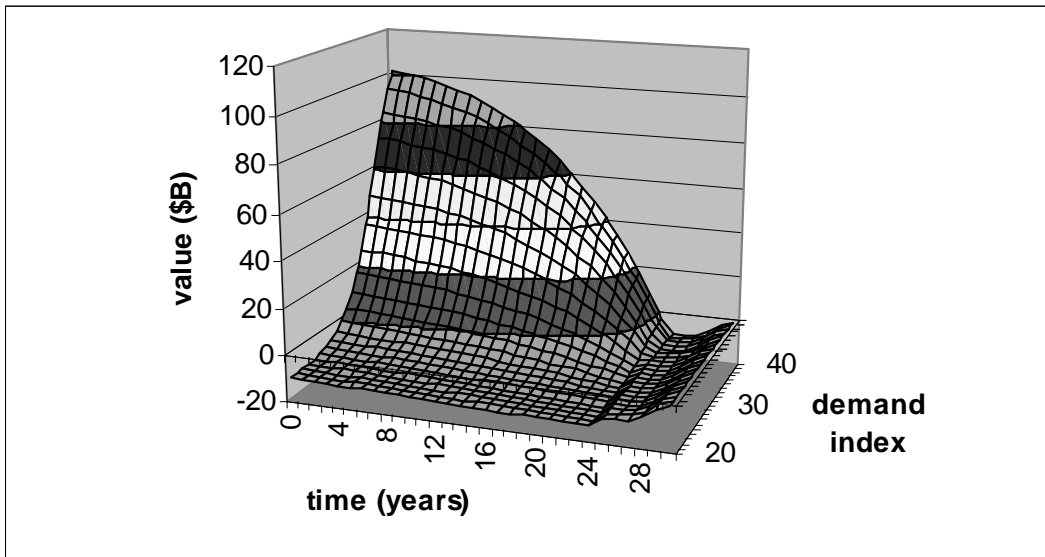


Figure 37. No-flexibility program value for BWB-450, baseline assumptions

Consider a cut through the above plot, made at the plane defined by (time = 0). The result gives program value as a function of demand index as of time 0. This function is shown in Figure 38, along with its equivalent from Figure 33: program value with flexibility.

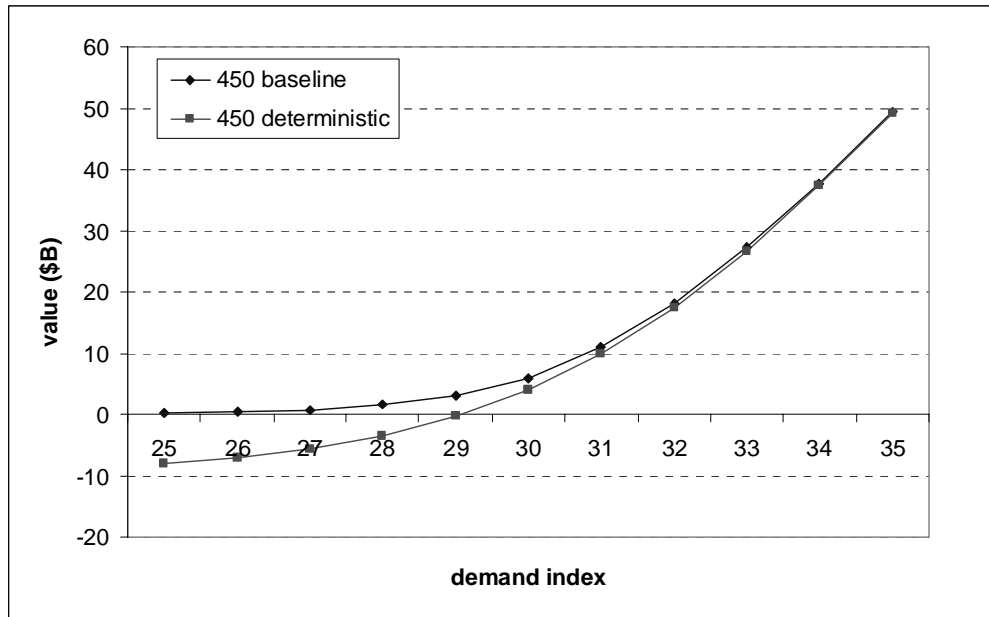


Figure 38. BWB-450 program value at time = 0

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 index, the value with flexibility, \$5.95B, is 48% greater than the value without flexibility, \$4.03B.

7.5 Sensitivity Analysis

This section investigates the effect of varying several of the baseline assumptions made in Section 7.3. First, model sensitivity to CAROC (Cash Airplane-Related Operating

Costs) is addressed by comparing the baseline to results using the CAROC model developed in Section 5.4.2. Next, the effect of demand correlation is considered in finding the value of two common aircraft. Next, the effect of volatility is demonstrated.

7.5.1 CAROC

In Section 5.4.2, a CAROC model was developed by finding a regression equation for the fraction of CAROC that is accounted for by fuel costs. This fraction was estimated to be a linear function of aircraft size—that is, seat count. In the baseline analysis of Section 7.3, the simplifying assumption was made that the fuel cost fraction is constant at 20%. If this assumption is relaxed, larger aircraft will tend to have decreased CAROC values, and smaller aircraft will have increased CAROC values. Table 17 presents the results of re-solving the problem for each of the cases presented in 7.3.

Table 17. Program value: results for CAROC sensitivity

	Fuel cost fraction = 20%			Use CAROC model		
	CAROC	Baseline price	Program value	CAROC	Baseline price	Program value
	(\$/ASM)	(\$M)	(\$B)	(\$/ASM)	(\$M)	(\$B)
BWB-450	0.0323	195.0	5.95	0.0285	219.9	11.58
BWB-250C	0.0425	116.1	2.26	0.0455	104.6	0.24
BWB-250P	0.0357	142.2	14.62	0.0382	132.5	11.11
BWB-450 & BWB-250P			8.95			12.44

Clearly, the assumptions made about CAROC play a significant role in determining program value. Using the CAROC model instead of the assumption of 20% fuel cost fraction nearly doubled the value of the BWB-450, because its price rose by over 20% as a result of a lowered CAROC estimate. Conversely, the value of the BWB-250C shrank significantly from the drop in price precipitated by a higher CAROC.

7.5.2 Demand Correlation

As described in Section 6.4, the model makes provision for a correlation coefficient between the two processes describing the demand quantities associated with two aircraft designs being considered simultaneously. While the correlation is irrelevant in considering a program with two aircraft unrelated to each other, it does have an effect on program value if the aircraft have commonality. The effect is quantified for several values of the correlation coefficient, ρ_{01} , in Figure 39.

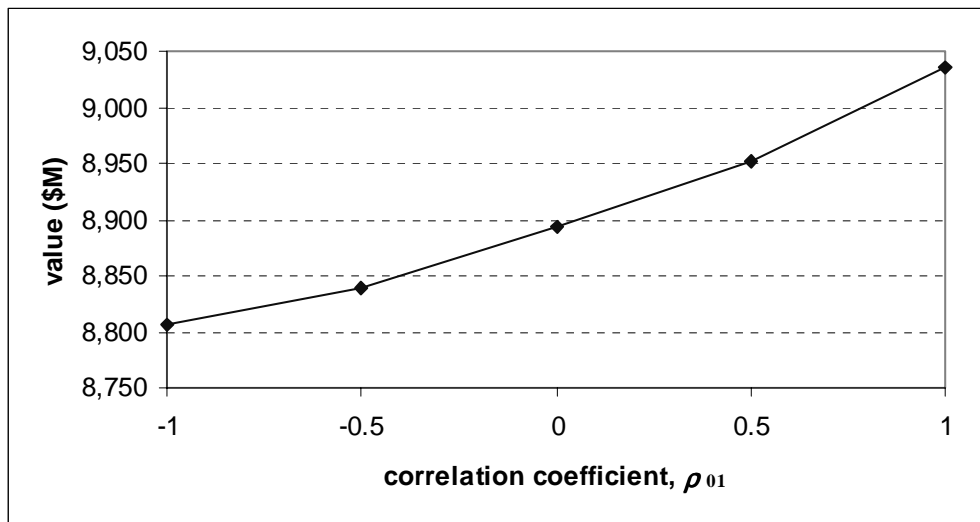


Figure 39. BWB-450 & BWB-250C program value: effect of correlation

The stronger the correlation between demand for the two aircraft, the more valuable is the commonality of the BWB-250C. If a good market for the BWB-450 guarantees a good market for the BWB-250C, investing in the larger aircraft implies a higher probability of also investing in the smaller aircraft, and vice versa. There is thus less likelihood that the savings from commonality will go unused.

7.5.3 Volatility

To investigate the effect of volatility, the BWB-450 problem was re-solved with the following revised set of assumptions. For a review of the equilibrium return shortfall, δ , refer to the “Stochastic Process Dynamics” discussion in Section 6.4. The exact definition of δ is given in equation (6.10) in that section.

Table 18. Reduced volatility: key inputs

Aircraft demand volatility, σ	19.6%
Equilibrium return shortfall, δ	3.68%

The revised value for σ was taken from the average of the volatility of aggregate historic wide body deliveries. The aggregate tends to be less volatile than the average individual aircraft delivery profile. Hence, since the revised assumption here is based on an aggregate, it is likely not an accurate indicator of the actual demand volatility. However, the point of the sensitivity analysis is simply to observe the response of the algorithm to a change in the volatility input parameter. Table 19 compares the results to the baseline runs, for both the program values with flexibility, and program values without flexibility, as described in Section 7.4.

Table 19. BWB-450 program value: effect of volatility

		Reduced volatility	Baseline volatility
Volatility, σ	(%)	19.6%	45.6%
Shortfall, δ	(%)	3.68%	2.19%
Program Value			
(1) with flexibility	(\$B)	6.65	5.95
(2) without flexibility	(\$B)	5.48	4.03
(1) – (2)	(\$B)	1.17	1.92

Note that the incremental value of including flexibility in the program analysis increases with increasing volatility. This is consistent with Real Options literature—the Real Options approach becomes more valuable and more necessary as uncertainty increases. The program value, including flexibility, is 21% greater than the value excluding flexibility for the low-volatility scenario, while the program value gains an additional 48% through flexibility for the high-volatility scenario.

To evaluate the effect of reduced volatility on the value of commonality, the above assumptions for σ and δ were applied to the BWB-250C and to the two-aircraft scenario

involving BWB-450 and BWB-250C. Table 20 compares the results to the baseline analysis.

Table 20. BWB-450 and BWB-250C: effect of volatility

		Reduced volatility	Baseline volatility
Volatility, σ	(%)	19.6%	45.6%
Shortfall, δ	(%)	3.68%	2.19%
Program Value			
(1) BWB-450	(\$B)	6.65	5.95
(2) BWB-250C	(\$B)	2.90	2.26
(1) + (2)	(\$B)	9.55	8.21
BWB-450 & BWB-250C	(\$B)	11.10	8.95
Commonality premium		16.2%	9%

The last entry in the table, commonality premium, is the percent increase in value from the sum of the individual program values to the value of the combined program with both aircraft. The reduced volatility, almost half of its original value, results in almost a doubling of the commonality premium, from 9% to 16.2%. The greater the amount of certainty regarding the evolution of demand, the more valuable is commonality.

7.6 Conclusion

This chapter has presented several applications of the program valuation tool to a hypothetical aircraft design project. It demonstrated the tool's functionality for a standalone single-aircraft program evaluation, as well as for the simultaneous evaluation of two aircraft designs, which may or may not be designed and/or built concurrently. Each test case took approximately 10 minutes to solve using the dynamic programming algorithm presented in Chapter 6, implemented in C on a Pentium III laptop.

A solution consists of the set of decision rules, as a function of time and state, that maximizes the value of the program. In this case, the term "state" indicates the demand level and the current operating mode, while the term "decision rule" refers to the choice of operating mode for next period.

The numerical results of the example runs in this chapter demonstrate the critical dependence of the accuracy of any model upon its assumptions. Since the dynamics of cost, price, and demand are not described by exact physical sciences, any economic model is an approximation. The model used here is no exception, and the assumptions forming its basis are subject to change upon further study. For example, the BWB-250P aircraft generates a phenomenal business case as presented in this chapter. It creates so much value because of a calculated profit margin of almost 70%. Based on the historical performance of aircraft programs in the industry, it is hard to believe that such a margin is attainable in practice. More likely, the pricing model has built in an excessive sensitivity to CAROC, which resulted in an unrealistically high price estimated for the BWB-250P. Given more extensive data, it would be possible to adjust the dependence of price on CAROC in a way that more accurately represents reality.

However, even if it is decided that the numerical result generated for the BWB-250P is highly unrealistic, the example still serves a useful purpose. In this case, it highlights the considerable sacrifice made in the 250-passenger class aircraft design to accommodate commonality. A modest increase in empty weight translated to a medium increase in takeoff weight, which translated to a significant difference in fuel burn and operating cost, and an even greater difference in market price. Regardless of the accuracy of the model, this snowballing phenomenon underscores the importance of considering the downstream effects of a design change on program value. Of course, the other side of the coin is the value benefit gained by commonality: a savings in development and manufacturing costs.

Within the framework of flexibility and decision-making introduced by the dynamic programming algorithm, the choice to use commonality may be framed using real options. When the firm develops the BWB-450, it acquires an option to develop the BWB-250C for a reduced cost and at a time of its choosing. The penalty paid—i.e., the price of the option—is the present value of additional profits the firm would receive had it instead developed the BWB-250P as a point design to maximize its performance. From a program flexibility standpoint, the firm still has an option to develop a second aircraft even if there is no commonality—in such a case, the exercise price of the option is simply higher by the amount of cost savings from commonality.

The conclusion of this analysis, therefore, is not that commonality isn't justifiable. Rather, for commonality to be justifiable, the benefits must outweigh the costs. The benefits include the development and manufacturing cost savings gained *if* the derivative aircraft is in fact built. The costs include any additional design or manufacturing costs as a result of commonality, but most importantly, any resulting performance penalty on the aircraft. This performance penalty must be translated into an opportunity cost: the revenues foregone by *not* selling a higher-performance aircraft. The set of aircraft designs used in this example, with the baseline parameters specified, did not indicate a higher program value for commonality, because the opportunity cost of lost revenues was very high.⁷²

Despite the small value of commonality demonstrated in this chapter, several additional observations can be made on this topic. The volatility analysis in Section 7.5.3 showed that the value of commonality increases as volatility decreases, and the demand correlation analysis in Section 7.5.2 showed that commonality value increases with correlation. This results reflect the fact that designing a high degree of commonality into two aircraft ties their fortunes together to some degree. That is, if one is built, it is more attractive to develop the other, and vice versa. The greater the uncertainty in future demand, the less the likelihood that both will be developed and built in a sequence that maximizes their synergies. If two aircraft are being designed independently, their total program value will be independent of the relative timing and magnitude of the realized demand levels for each aircraft. Thus, whatever the disadvantages of having distinct designs, this approach entails less risk. Whether it also entails more value depends on the degree and nature of the commonality—its costs and benefits, as mentioned above.

⁷² Note that the example aircraft do not represent actual current Blended-Wing-Body designs.

Chapter 8. Conclusions

The motivation for this work lay in sharpening the focus of the aircraft design process. It centered on the idea that an aircraft should not be designed as a single machine with a set of cost and performance characteristics. Rather, an aircraft should be designed as part of its program—the process by which the firm brings the product to market and extracts value from it. At an even broader level, an aircraft should be designed in the context of the firm's entire product line, and that of its competitors.

This modified perspective requires a modified objective function—a metric to evaluate the aircraft and to enable a feedback loop in the design process to find an optimal product. The ideal objective function for aircraft design was identified as net present value. Value was defined as the (equilibrium) amount a potential buyer would be willing to pay for the aircraft project, if the project were an asset being freely traded on the market.

It is also important to recognize two factors that have a significant effect on value, but are sometimes overlooked in aircraft design in particular and project valuation in general. These are uncertainty and flexibility. The former refers to the continuously evolving and unknown future state of the world, while the latter refers to a firm's ability to react to this random process. In general, an accurate consideration of uncertainty and flexibility is very difficult to include in a program valuation. One method of constructing a model to approximate their effect borrows from the field of derivative pricing in finance. The pricing of options and other derivatives involves well-understood notions of both uncertainty—the behavior of securities—and flexibility—the type of derivative (call option, put option, swap, etc.).

The adaptation of option pricing theory to the valuation of projects, known as Real Options, offers an interesting alternative to traditional NPV valuation. By using risk-neutral expectations and discounting at the risk-free rate, the approach tries to relate the risk inherent in the project to the risk and return characteristics of the market, and thus

find a market value for the project. The advantage of Real Options is that the approach lends itself to explicit modeling of managerial flexibility in a project. The drawback is that, in practice, linking the project risk to the characteristics of the market can be very difficult, because in general projects are not traded financial assets. Thus, it is difficult to apply a market risk-pricing model to an asset that is not on the market.

While the real options approach should by no means be regarded as a replacement for human insight, it is a useful tool to complement traditional capital budgeting decision methods. It broadens the scope of numerical analysis by quantifying some of the elements of strategy that are not explicitly addressed by the traditional methods. Real options provides no explicit set of quantitative rules by which to select and manage projects—if that were so, corporate management would add no value to a firm and companies would be run exclusively by computers. Real Options does provide a more complete framework by which to evaluate and plan projects, and in the context of this framework, better decisions can be made.

This idea was demonstrated by implementing a Real Options-based valuation approach through the use of a dynamic programming algorithm, and modeling the aircraft program as a series of “modes” in which management can choose to operate. To support the dynamic programming framework, a series of models was developed for the estimation of development cost, manufacturing cost, baseline cost and price, and future demand evolution for any given new aircraft design. Together with the dynamic programming algorithm and a performance and sizing code, these models synthesized a tool for aircraft program valuation. While they lack a high degree of fidelity, the models comprise a useful framework upon which to build more complex studies. The models may readily be modified with more accurate data and incorporated into alternate valuation algorithms.

In particular, the cost model captured some possible effects of part commonality with its treatment of aircraft as a collection of “parts”, or major components. It accounted for learning curve effects in different categories of recurring cost, and for varying cost per pound of structure for different types of major aircraft components. It also simulated the simplifying effect of commonality on the design of a derivative aircraft by shortening the development time and reducing the development cost. However, accurate quantification

of the time and cost savings was not available, and would be a valuable area for further research. In addition, while part commonality was addressed through the above approach, process commonality was not. Multiple examples exist of aircraft that share few identical parts, but many elements of a production line—tooling, staff, assembly procedures—and thus generate significant savings in production cost. A consideration of the effects of process commonality would be an interesting extension to the cost model. Finally, another area for cost model refinement is the supply chain, which was reduced to just one firm in this study. It may be worthwhile to study the effect of procuring many of the aircraft components from suppliers outside the firm, as is typically the case.

The baseline price and demand quantity estimators developed for the revenue model are also a useful foundation for further study. Price was estimated as a function of seat count, range, and CAROC (Cash Airplane-Related Operating Cost), while the static forecast of quantity demanded was found by classifying the aircraft based on its size and assigning it to a market category in an existing set of forecasts for aircraft deliveries over the next 20 years. As new or better data becomes available, the pricing model may easily be adjusted for more realism or varying degrees of sensitivity to its input parameters, and as new demand forecasts are created, the quantity estimator may be refined as well. One salient feature of the revenue model was the significant sensitivity of price to operating cost, which is driven by fuel burn. The high sensitivity was the result of an assumption that an abnormal aircraft operating cost will change the amount the customer is willing to pay for the aircraft by the present value of the lifecycle savings (or additional cost) to the customer due to the abnormally low (or high) operating cost. It may be useful to alter this assumption by specifying that the customer will only be willing to pay some fraction of the lifecycle savings in the case of abnormally low operating cost—in other words, if the firm develops an unusually efficient aircraft, it must split the additional value it creates with the customer. This would probably be a more realistic representation of aircraft pricing, but it necessitates determining exactly how the additional value is split up. Often, that question is answered on a customer-by-customer basis.

To demonstrate the operation of the program valuation tool, several aircraft designs were evaluated, using the Blended-Wing-Body concept as an example. A highlight was placed on the idea of commonality between different aircraft.

The results of the example cases provide several insights.

First, the downstream effect of commonality on operating cost and consequently price may be quite significant, and is critically important to consider in making a commonality tradeoff. While the most immediate impact of commonality between aircraft may be a weight penalty, the most significant value penalty may come not from the associated increase in cost due to weight (development or manufacturing)⁷³, but from the increase in fuel burn that is tied to increased weight. Higher fuel burn results in a higher operating cost, which dictates a smaller sale price and less profitability.

Second, while commonality was introduced as a type of flexibility, the results suggested that increasing uncertainty reduced the value of commonality. Real Options literature frequently mentions the increasing value of flexibility as uncertainty increases. That observation is difficult to apply in this case because of the nature of the flexibility provided through commonality. It is not a switching option which provides the ability to select the maximum of two assets, but rather a compound option: investing in one aircraft design provides the ability to invest in another one later for less. However, the question is whether the option to save money on a future design is worth the penalty in performance and therefore price that accompanies commonality. If not, the flexibility afforded by an independent design may be more valuable.

Third, the dynamic programming algorithm demonstrated some potential drawbacks of traditional NPV valuation for commercial aircraft programs. Especially in situations where the business case for the program is marginal, the dynamic programming (DP) approach demonstrates a significantly greater value than the NPV approach because of the consideration of flexibility afforded by DP. This observation does not suggest that this particular DP approach is the ideal one, just that the consideration of flexibility in a

⁷³ Such costs may be offset by the benefits of commonality—e.g., faster learning, economies of scale, development cost and time reductions.

program offers an extra dimension over traditional valuation techniques—a dimension that may make the difference between accepting a design and rejecting it.

As mentioned above, numerous assumptions were necessarily made, in the interest of time and simplicity, to develop the program valuation tool. Accordingly, several opportunities exist for extension of the model to yield increased sophistication and accuracy.

One of the problems with traditional NPV approaches for long-lived projects such as aircraft programs is the assumption that the underlying stochastic process follows a random walk, not unlike typical models of stock price behavior. The variance of the returns grows linearly with time, and the discount rate used, based on the Capital Asset Pricing Model, is the same regardless of the timing of the cash flows. One possible alternative approach is to model the randomness in an aircraft program as a mean reverting process. This would involve modifying the risk-neutral probability calculation used in the binomial tree to be a function of the demand index value. If implemented successfully, this approach would place more emphasis on cash flows far in the future than does the traditional NPV approach, which renders insignificant any cash flow past roughly 10 years from the present.

Another limitation of the valuation tool is its treatment of only two aircraft. In practice, it would be very useful to evaluate portfolios of multiple aircraft designs—for example, 3, 4, or 5 aircraft. The value of commonality may be significantly affected by the inclusion of additional aircraft, and any interactions between the aircraft in the market, such as cannibalization, would be interesting to model. The result of this extension would be an optimal strategy for aircraft portfolio development for the firm.

A third area for extension is competition. The present-day market for commercial aircraft is a fairly good example of a duopoly, and the inclusion of competitor dynamics and game theory in the program valuation model would provide a significant added level of realism which the present model does not attain.

The above opportunities for further research suggest that the concept of value-based design is a useful one. As both a quantitative approach and a qualitative philosophy, designing an aircraft to maximize value as part of a much larger and more complex system is an idea with much potential.

References

Amram, M., Kulatilaka, N. *Real Options: Managing Strategic Investment in an Uncertain World*. Boston, MA: Harvard Business School Press, 1999.

Airbus Global Market Forecast, 2000-2019. Appendix G, Detailed passenger fleet results, p. 74.

Aircraft Value News Aviation Newsletter. www.aviationtoday.com/catalog.htm

The Airline Monitor, ESG Aviation Services, July 2000. Projected Jet Airplane Deliveries—2000 to 2020, pp. 10-11.

The Airline Monitor, ESG Aviation Services, May 2001.

Baker, A. P., Mavris, D. N. “Assessing the Simultaneous Impact of Requirements, Vehicle Characteristics, and Technologies During Aircraft Design.” AIAA paper 2001-0533, Jan. 2001.

Black, F., Scholes, M. “The Pricing of Options and Corporate Liabilities,” *Journal of Political Economy*, No. 81, pp. 637-59, May-June 1973.

“Blended-Wing-Body Technology Study,” Final Report, NASA Contract NAS1-20275, Boeing Report CRAD-9405-TR-3780, Oct. 1997.

Boeing Current Market Outlook, 2000. Appendix B, p. 42.

Boeing Quick Looks online aircraft data, 2000. (non-proprietary)

Boeing Quick Reference Handbook, Economic Comparison—US rules, 2000. Company publication with basic economic analysis assumptions. (non-proprietary)

Brealey, R. A., Myers, S. C. *Principles of Corporate Finance*. New York, NY: McGraw-Hill, 1996.

Browning, Tyson R. “Modeling and Analyzing Cost, Schedule, and Performance in Complex System Product Development.” PhD Thesis; Technology, Management, and Policy. MIT, 1998.

Cox, J., Ross, S., Rubinstein, M. “Option Pricing: A Simplified Approach,” *Journal of Financial Economics*, No. 7, October 1979.

Dickinson, M. W., Graves, S., Thornton, A. C. “Technology Portfolio Management: Optimizing Interdependent Projects Over Multiple Time Periods.” MIT working paper, 1999.

Dixit, A. K., Pindyck, R. S. *Investment Under Uncertainty*. Princeton, NJ: Princeton University Press, 1994.

Esty, B., Ghemawat, P. “Airbus vs. Boeing in Superjumbos: Credibility and Preemption.” Draft paper, Harvard Business School. Aug. 2001.

- Fujita, K., Akagi, S., Yoneda, T., Ishikawa, M. "Simultaneous Optimization of Product Family Sharing System Architecture and Configuration." *Proceedings of DETC '98, ASME Design Engineering Technical Conferences*. DETC98/DFM-5722, Sep. 1998.
- Hull, J. C. *Options, Futures, and Derivatives*. 4th ed. Upper Saddle River, NJ: Prentice-Hall, 2000.
- Jane's All the World's Aircraft*. London : Sampson Low, Marston & Co., 2001.
- Fabrycky, W. J. *Life-Cycle Cost and Economic Analysis*. Englewood Cliffs, NJ: Prentice Hall, 1991.
- Kroo, I., Manning, V. "Collaborative Optimization: Status and Directions." AIAA paper 2000-4721, Sep. 2000.
- Kulatilaka, N. "Valuing the Flexibility of Flexible Manufacturing Systems." *IEEE Transactions on Engineering Management*, Vol. 35, No. 4, pp. 250-257, Nov. 1988.
- Kulatilaka, N. "The Value of Flexibility: The Case of a Dual-Fuel Industrial Steam Boiler." *Financial Management*, Vol. 22, No. 3, pp. 271-279, 1993.
- Liebeck, R. H. "Design of the Blended-Wing-Body Subsonic Transport." *2002 Wright Brothers Lecture*, AIAA paper 2002-0002, Jan. 2002.
- Liebeck, R. H., Page, M. A., Rawdon, B. K., "Blended-Wing-Body Subsonic Commercial Transport," AIAA paper 1998-0438, Jan. 1998.
- Neely, J. E. III. *Improving the Valuation of Research and Development: A Composite of Real Options, Decision Analysis and Benefit Valuation Frameworks*. PhD Thesis, MIT, May 1998.
- Nuffort, M. R. *Managing Subsystem Commonality*. S. M. Thesis, MIT, Feb. 2001.
- Press, W., Teukolsky, S., Vetterling, W., Flannery, B. *Numerical Recipes in C: The Art of Scientific Computing*. 2nd edition. New York, NY: Cambridge University Press, 1992.
- Raymer, D. P. *Aircraft Design: A Conceptual Approach*. Reston, VA: American Institute of Aeronautics and Astronautics, 1999.
- Reinhardt, 1973. "Break-Even Analysis of Lockheed's Tri Star: An Application of Financial Theory." *Journal of Finance*, Vol. 28, Issue 4, Sep. 1973, pp. 821-838.
- Roskam, J. *Airplane Design. Part V: Component Weight Estimation*. Ottawa, KS: Roskam Aviation and Engineering Corporation, 1989.
- Slack, R. "The Application of Lean Principles to the Military Aerospace Product Development Process." S.M. Thesis, MIT, Dec. 1998.
- Stonier, J. "What is an Aircraft Purchase Option Worth? Quantifying Asset Flexibility Created through Manufacturer Lead-Time Reductions and Product Commonality." *Handbook of Airline Finance*, 1st edition. Executive editors G. Butler and M. Keller. Aviation Week: 1999.
- Trigeorgis, L. *Real Options: Managerial Flexibility and Strategy in Resource Allocation*. Cambridge, MA: The MIT Press, 2000.

Wakayama, S., Kroo, I., "Subsonic Wing Planform Design Using Multidisciplinary Optimization." *Journal of Aircraft*, Vol. 32, No. 4, Jul.-Aug. 1995, pp. 746-753.

Wakayama, S., *Lifting Surface Design Using Multidisciplinary Optimization*. Ph.D. Thesis, Stanford University, Dec. 1994.

Wakayama, S., Kroo, I., "The Challenge and Promise of Blended-Wing-Body Optimization." AIAA Paper 98-4736, Sep. 1998.

Wakayama, S., Page, M., Liebeck, R., "Multidisciplinary Optimization on an Advanced Composite Wing." AIAA Paper 96-4003, Sep. 1996.

Wakayama, S., "Multidisciplinary Optimization of the Blended-Wing-Body." AIAA Paper 98-4938, Sep. 1998.

Wakayama, S., "Blended-Wing-Body Optimization Problem Setup." AIAA Paper 2000-4740, Sep 2000.

Willcox, K., Wakayama, S. "Simultaneous Optimization of a Multiple-Aircraft Family." AIAA Paper 2002-1423, Apr. 2002.

Appendix: Inputs Spreadsheet

NOTE: Example aircraft part weights have been replaced by "xx".

Summary of Key Inputs

Case Description: Two aircraft design: BWB3450x600 and BWB3250x600

Input	Value	units	Notes
-------	-------	-------	-------

DYNAMIC PROGRAMMING DATA

timestep per period	1	year	
number of periods	30		in Constants.k
number of modes	17	per aircraft	in Constants.k
number of state values	80	per aircraft	in Constants.k
low capacity	40	units per period	in Constants.k
med capacity	90	units per period	in Constants.k
high capacity	132	units per period	in Constants.k
MR_SPANTIME	6		
MR_TOOLMODE	3		
MR_BUILDMODE	13		
MR_TOOL_L0	0.75		
MR_TOOL_H4	1.90		
TOOL_DISRUPT	1.20		
risk-free rate	0.065	per period	in r_f.xls
correlation coefficient	0	per period	(not used)
air cost inflation	0.012157883	per period	
air price inflation	0.012157883	per period	
initial mode	all 0		
initial state index	all 0		
state volatility	0.455737278	per period	from stochastic demand.xls (averaged)
state growth rate	0.044261369	per period	used for MR random path simulation
convergence yield	0.021037388	per period	from stochastic demand.xls

DEMAND DATA

Expected growth rates per period in quantity demanded

Narrow body	0.045796695	per period	from Quantiles analysis.xls
Wide body	0.054032093	per period	from Quantiles analysis.xls

Gross Quantity demanded over 20 YEARS

	(narrowbody)	(narrowbody)	(narrowbody)	(narrowbody)	(widebody)	(widebody)	(widebody)	(widebody)	(widebody)	(widebody)
seat category	100	125	150	175+	200	250	300	350	400	500+
seat category index #	0	1	2	3	4	5	6	7	8	9
seat lower limit	90	120	150	175	190	240	275	340	410	530
seat upper limit	110	130	160	200	220	290	305	370	440	550
quantity demanded	2071.199	2826.131	3369.636	2638.186	1089.896	1634.173	1472.055	865.037	620.371	758.907

quantity in period 0	64.81630818	88.61367631	105.4457252	82.30554785	39.27542132	54.78610393	52.56789832	30.53980938	18.582715	27.3854988
growing	at	1 year	per period	per period	per period	per period	per period	per period	per period	per period
growth rates	at	1 year	per period	per period	per period	per period	per period	per period	per period	per period

Market share	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5
--------------	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----

Max. narrowbody index	3
Max. narrowbody range	6500
Max. widebody range	3500

COST DATA

Recurring part costs per lb

NOTE: these costs are in dollars; they are linked to Parameters' as THOUSANDS of dollars

	Wing	Engine	Fuselage	Landing	Geared	Engines	Systems	Payloads	Final Assembly
Labor	608.57	1614.38	679.2	107.38	247.87	315.46	405.47	58.5	
Materials	203.75	483.79	190.14	98.47	80.96	90.97	100.43	3.65	
Other	87.91	233.23	98.11	16.62	36.0	45.65	58.67	3.33	

Learning curves

	Labor	Materials	Other
LC slope parameter	0.95	0.95	0.95

Non recurring part costs per lb (not adjusted for learning/commonality effects)

NOTE: these costs are in dollars; they are linked to Parameters' as THOUSANDS of dollars

	Wing	Engine	Fuselage	Landing	Geared	Engines	Systems	Payloads	Final Assembly
Engineering	7092.57	20962.47	12937.38	999.42	3476.52	13722.85	4305.21	0	
ME	1773.14	5215.62	3209.34	249.86	868.13	3430.71	1075.38	0	
Test Design	1861.88	5476.48	3368.81	262.36	912.89	3602.26	1138.12	0	
Test Fab	8170.54	18188.35	11188.52	889.58	3024.58	11938.89	3745.53	0	
Support	833.38	2461.34	1508.39	117.43	438.49	1612.44	506.88	0	

Non recurring part cost reduction factors (for parts which have already been designed)

	Wing	Engine	Fuselage	Landing	Geared	Engines	Systems	Payloads	Final Assembly
Engineering	0.2	0.2	0.2	0.2	0.2	0.2	0.2	0.2	1
ME	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	1
Test Design	0.05	0.05	0.05	0.05	0.05	0.05	0.05	0.05	1
Test Fab	0.05	0.05	0.05	0.05	0.05	0.05	0.05	0.05	1
Support	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	1

Baseline non recurring process parameters (i.e., no learning/commonality effects)

	Engineering	ME	Test Design	Test Fab	Support
Start Time	0	0	0.22	0.27	0
Total Duration	1	0.85	0.45	0.5	1
Alpha	2.2	2.5	3.5	3	1.5
Beta	3	3	3	3	1.5

Input Value units Notes

CAROC DATA

Reference mission range	3000	nm	from CAROC Model.xls
Fuel price	0.65	\$/gal	from CAROC Model.xls
Fuel density	6.7	lb/gal	from CAROC Model.xls
CAROC regression:			
slope, m	0.000194215		from CAROC Model.xls
y-intercept, b	0.158831633		from CAROC Model.xls

AIRCRAFT DATA

Aircraft 1

Name	BWB-3-250-x01	
source	BWB3250-x00.xls	using WMDD 98

Performance:

seats	272	
MTOW	xx	lb
CLW	xx	lb
range	2250	nm

Weight:

Structure	xx	lb
Propulsion	xx	lb
Systems	xx	lb
Furnishings & Op Items	xx	lb
Ballast	xx	lb
OCW	xx	lb
Max Payload	xx	lb
Partial Payload	xx	lb

Design Payload	xx	lb
Design ZFW	xx	lb
Max Payload	xx	lb
MZFW	xx	lb
Partial Payload Fuel	xx	lb
Design Fuel	xx	lb
Max Payload Fuel	xx	lb
MTOW	xx	lb
Structural Weight:		
	Total	Bay 1
	xx	xx
		Bay 2
		xx
		Inner Wing
		xx
		Outer Wing
		xx
		Winglet
		xx

parts list:				NOTE: This data must be entered manually into "Aircraft" worksheet.	
part name	part weight	part quantity			
cockpit 1	xx	1	lb		
fus bay 2	xx	2	lb		
fus bay 3	xx	2	lb		
inner wing 1	xx	1	lb		
outer wing 1	xx	1	lb		
winglet 1	xx	1	lb		
landing gear 1	xx	1	lb		
propulsion 1	xx	1	lb		assume LG weight = 0.042* TOGW
systems 1	xx	1	lb		assume LG weight is included in "Systems";
payloads 1	xx	1	lb		subtract calculated LG weight from Systems
					assume = furnishings & op items
Total (DEW check)	xx		lb		less ballast of xx

Aircraft 2

Name	BWB-3450-a00	never optimal to build
source	BWB0450-a00.xls	

Performance:		
seats	475	
MTOW	xx	lb
CLW	xx	lb
range	6550	nm

Weight:		
Structure	xx	lb
Propulsion	xx	lb
Systems	xx	lb
Furnishings & Op items	xx	lb
Ballast	xx	lb
OEW	xx	lb
Max Payload	xx	lb
Partial Payload	xx	lb
Design Payload	xx	lb
Design ZFW	xx	lb
Max Payload	xx	lb
MZFW	xx	lb
Partial Payload Fuel	xx	lb
Design Fuel	xx	lb
Max Payload Fuel	xx	lb
MTOW	xx	lb
Structural Weight:		
	Total	Bay 1
	xx	xx
		Bay 2
		xx
		Bay 3
		xx
		Inner Wing
		xx
		Outer Wing
		xx
		Winglet
		xx

parts list:				NOTE: This data must be entered manually into "Aircraft" worksheet.	
part name	part weight	part quantity			
fus bay 1	xx	2	lb		
fus bay 2	xx	2	lb		
fus bay 3	xx	2	lb		
inner wing 1	xx	1	lb		
outer wing 1	xx	1	lb		
winglet 1	xx	1	lb		
landing gear 1	xx	1	lb		
propulsion 1	xx	1	lb		assume LG weight = 0.042* TOGW
systems 1	xx	1	lb		assume LG weight is included in "Systems";
payloads 1	xx	1	lb		subtract calculated LG weight from Systems
					assume = furnishings & op items
Total (DEW check)	xx		lb		