Learning-Derived Cost Evolution in Materials Selection

by

Trisha M. Montalbo

S.B. Physics
S.B. Materials Science & Engineering
Massachusetts Institute of Technology, 2002

S.M. Materials Science & Engineering
Massachusetts Institute of Technology, 2004

Submitted to the Department of Materials Science and Engineering
in partial fulfillment of the requirements for the degree of

Doctor of Science in Materials Science and Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2010

Author . . . .
Department of Materials Science and Engineering
May 18, 2010

Certified by . . . .
Randolph E. Kirchain
Associate Professor of Materials Science and Engineering
Thesis Supervisor

Accepted by . . . .
Christine Ortiz
Chair, Department Committee on Graduate Students

Learning-Derived Cost Evolution in Materials Selection

by

Trisha M. Montalbo

Submitted to the Department of Materials Science and Engineering
on May 18, 2010, in partial fulfillment of the
requirements for the degree of
Doctor of Science in Materials Science and Engineering

Abstract

Materials selection is a complex, but important, problem for manufacturing firms. Poor material choices can negatively affect the firm's market share or profits. In the face of this complexity, most selection methods make a number of simplifications, including limiting problem scope to selection for a single product or application, and assuming material properties and design criteria are constant over the problem's time horizon. Such assumptions, however, do not always apply, especially when material preference is based on the materials' "emergent properties," the values of which are context-dependent. Consequently, these properties can evolve with changes in context and potentially alter the preferred material identified by the selection method.

This thesis investigates the impact of considering cost evolution on a firm's materials selection decision, and seeks to identify strategies the firm can adopt when introducing new materials to its products. To that end, a framework for incorporating cost evolution, specifically from learning, into the materials selection process is proposed and demonstrated using single-product and multi-product automotive case studies. In the single-product method, material options are ranked by their respective manufacturing costs. The multi-product problem is more complex and requires an analytical framework that combines an integer linear program and a genetic algorithm to select materials for any number of products over a specified time horizon.

Case study results indicate that when selection problem scope is limited to a single product, accounting for learning in the decision process has minimal impact on the preferred material. When several products are included in the problem scope, however, the firm is able to leverage "shared learning" so that experience gained from manufacturing one product can be applied to lower the costs of other products that share a common resource, such as a manufacturing process line, with the initial product. Not only does the consideration of shared learning impact the preferred materials that are suggested by the selection framework, it also helps to better characterize the circumstances under which the firm should introduce a new material on a test bed. Additionally, the case study results emphasize the use of one material across multiple applications and indicate that this approach helps the firm cope with uncertainty in selection criteria.

Thesis Supervisor: Randolph E. Kirchain
Title: Associate Professor of Materials Science and Engineering

Acknowledgments

A big thank you Prof. Randy Kirchain and Dr. Rich Roth, two of the best advisors I could have ever asked for. Their insight and guidance have been invaluable to me over the years, as has their patience and their willingness to make time and talk about whatever research problem I'm facing. Thank you also to my committee members, Dr. Frank Field, Prof. Sam Allen, and Prof. David Roylance. To Frank especially, who is always willing to ask the hard questions of me—and everyone else—and who, over the last few months, has helped me frame my story and give it meaning. Thanks to everyone at MSL, in particular to Ashish, Catarina, Elisa, Jeff, Lynette, Nate, and Rob for all your input during our subgroup meetings; and to everyone else, including Jeremy, Elsa, Terra, Joel, Anna, Boma, Catarina, Gabby, Sup, Tracey, Tommy, YingXia and others. You've made the experience truly enjoyable in and out of work and I will miss all our random conversations. Thank you to everyone at General Motors, in particular Theresa Lee, for teaching me the world of vehicle design and manufacturing and for funding my research but giving me the freedom do my own thing. Thank you to my family and friends, who provided distractions from life in front of a computer. And finally, there is no thank you big enough for Justin—for all the love and support he has given me over the years and for always believing in me.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Materials selection is a common, but important, step in product design. Any manufacturing firm—or product designer—looking to develop a new product or modify a current will have to identify an appropriate material for that product, given design criteria that place constraints on the material's performance. Choosing an inappropriate material can be a costly mistake for the firm and can potentially lead to product, or even firm, failure. The extensive literature available on the topic of materials selection is indicative of the importance of finding a satisfactory material for a product, as well as the difficulties the firm faces when doing so.

Identifying an appropriate material, however, is a complex problem because materials are chosen for their properties, their ease of processing, or for the product capabilities they make possible, rather than because the designer likes "bronze" or "wood." Consumers, though, are likely to only care about product attributes enabled by the materials and not about processing or specific properties. For example, a consumer purchasing a winter coat will choose based on the coat's fit, style, and warmth rather than for the specific type of thread the manufacturer uses to sew the coat together. It is therefore left to the product designer or the firm to assess how the selection of any given material will influence a product's attributes and whether the changes in those attributes are desirable from a cost or consumer preference standpoint.

One of the fundamental concepts in modern materials science states that a connection exists between a material's properties and its performance within a product. Equally important are the material's structure and the processing it undergoes during product manufacture, both of which have an effect not only on the material's performance, but also on its properties, and

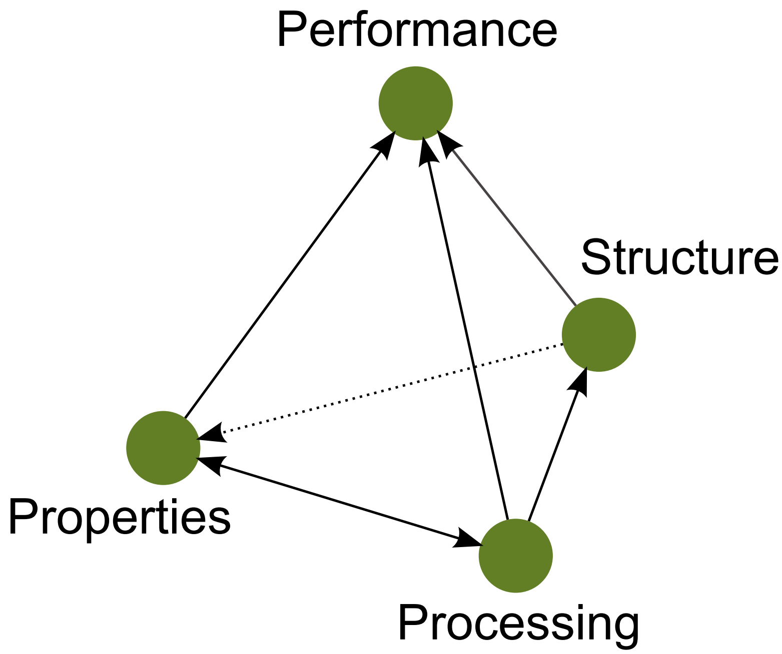

Figure 1-1: Materials science tetrahedron illustrating the inter-relationship between performance, properties, processing, and structure [21, 58].

are themselves inter-related [21, 58]. The inter-relationships among these four qualities is often depicted with a tetrahedron; arrows in Figure 1-1 indicate direction of influence.

Because of these relationships, a product designer cannot simply look up material properties in a database and find the one material that best satisfies design criteria. Many properties are context-dependent because of their relationship to structure and processing. For example, the environmental impact of a material will vary according to that material's mining location, the method of transportation to the manufacturing plant, the product being manufactured, the manufacturing process flow, the firm's energy supplier, the operational environment of the firm, the market conditions the firm faces, and so forth. Secondly, there are usually several materials that satisfy the design criteria, but likely no single "best" option due to the presence of trade-offs among the different criteria. These trade-offs force the firm to make sacrifices in some criteria in order to obtain improvements in others. One common example is the cost-mass trade-off in which a firm that wishes to reduce its product's mass will have to invest in more expensive, but lighter weight, materials such as titanium or carbon fiber composites. The selection problem then becomes a question of how much is the firm willing to pay for a reduction in mass, or is shifted to other properties that differentiate the materials. Thus, when selecting a material, a firm has to consider not only the product attributes it is looking to achieve, but also the problem's context so it can accurately assess the material's properties and the trade-offs among different material options.

Given the difficulty and cost of identifying the right material for a product, firms are understandably reluctant to switch to a new material without a compelling reason, especially once they have established that their current choice works—that is, it satisfies design criteria and the resulting product is appealing to consumers. Firms, however, will still shift to a new material if there exists a good reason driving the switch—for instance, if the new material is a clear improvement over the current one. This was the case in the 1970s when economic factors and technology improvements led to the beverage industry's decision to adopt plastic for soda bottles in place of glass [55]. Likewise, firms will switch to a new material if they can compensate for the cost of doing so. For example, certain market segments are often willing to pay a premium for product attributes enabled by new materials, as often happens with sporting equipment such as tennis racquets, golf clubs, and bicycle frames [17]. It is worth noting, though, that in some instances, the attribute consumers are paying for is the sexiness of a new material rather than any notable gains in performance brought on by the new material.

In general, it is often difficult for a firm to prove conclusively that the benefits of using a new material outweigh the costs of switching to that material. The absence of a compelling reason to switch materials, combined with the firm's conservative behavior, means that it is difficult for new materials, or materials the firm is not accustomed to using, to gain market share. Nevertheless, firms may still be forced to consider alternative materials due to rapid changes in its technology options or operational environment, such as updates in a competitor's product, or changes in the market or regulations. These changes force the firm outside its "comfort zone," where it will have to move more aggressively, especially with regards to materials selection. Under these circumstances, the firm will therefore have to redesign its products, and in the process, identify the best material to use given the altered operational environment.

Partly because of its complexity, materials selection is a well-researched problem and even a brief literature survey yields a number of selection methods that have been developed to inform the materials selection decision. These methods are components of a design process that strives to provide a systematic means to sift through and evaluate potential materials options and identify

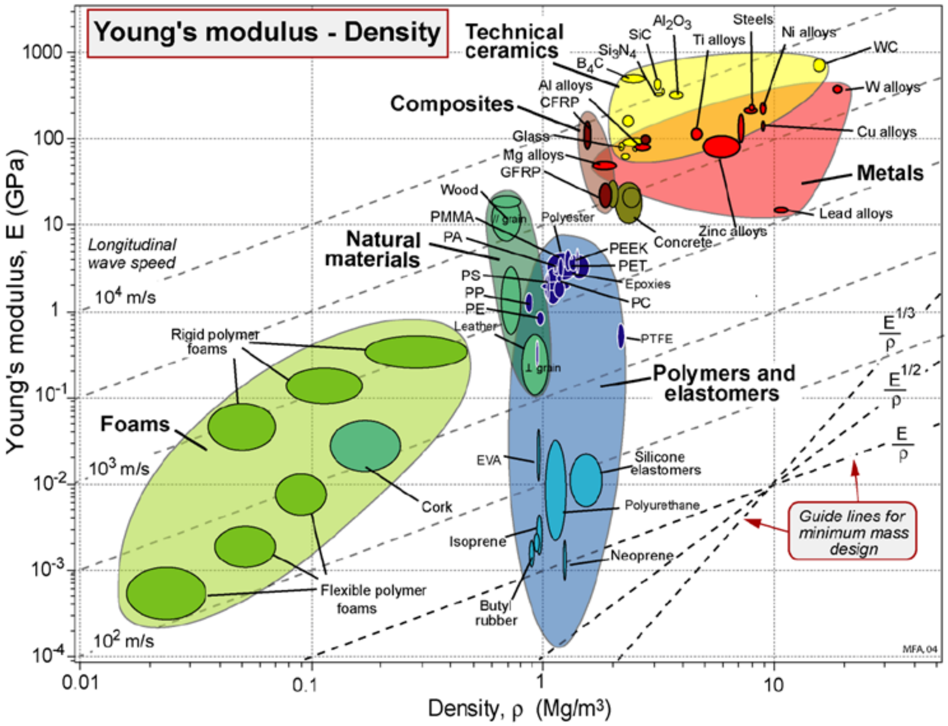

Figure 1-2: Example of a material property chart for Young's modulus and density [15].

satisfactory options based on design criteria as dictated by desired product attributes. One of the better-known methods, screening and ranking, has been developed and refined through a number of studies, most notably Ashby et al. [14]. In this approach, a designer first screens materials to weed out unsuitable candidates (i.e. materials that are infeasible in the application); the remaining options are then ranked according to performance indices, which are defined based on design criteria. For example, a designer seeking to minimize the mass of a beam while maintaining a specified stiffness will use the index $\sqrt{E}/\rho$ (where E is the material's modulus and ρ its density) to rank appropriate materials. Plots such as the one in Figure 1-2 can be generated to aid the designer. In the last step of Ashby's method, additional information such as case studies, known applications, and supplier details are gathered on the highest-ranked materials and used to make a final decision.

Other selection tools include out-ranking approaches such as ELECTRE [66, 69] and TOPSIS [42], which set preference and indifference thresholds for materials [19]; and heuristics like Sapuan's Knowledge-Based System [67] that define a set of rules or steps a designer can walk through to identify the preferred material. All of these methods, however, must make assumptions in order to simplify the selection process and arrive at a decision without overwhelming the designer. Without any simplifications, materials selection would be an even more burdensome and inefficient process; adding simplifications, on the other hand, makes the process more manageable and, in most cases, does so without compromising the results. One of these simplifications restricts the problem's scope to the selection of a single material for use in a single application within a product. Other products or applications that may also require a new material are typically not considered in the selection process of the primary application. Another common simplifying assumption is that design requirements and material properties do not vary over the time horizon of the decision, but instead remain fixed at a single value. While this is valid for most cases, it is not necessarily true for new or unfamiliar materials a firm is forced to work with: properties of these materials such as manufacturing cost or product performance can evolve over time as the firm becomes more comfortable using them and refines the product's design and manufacturing process.

Parameter evolution over time is particularly relevant in selection processes involving new materials for what Field et al. [29] term to be *emergent properties*. These properties are named thus because they are inherently context-dependent and "emerge" only when the details surrounding a material's production and use have been characterized. The context necessary to determine the value of an emergent property varies according to the property. For some properties, only the material's application within the product is necessary; others require more information such as industry-wide regulations, market prices for raw materials, or consumer behavior. An example of an emergent property is a product's safety, which depends on the product's design as well as factors external to the firm such as consumer usage. Manufacturing cost is another such property as it is a function of process parameters like cycle time and reject rate, along with raw material price, labor wage, and electricity cost.

The context-dependent nature of emergent properties means that, if the context changes over the decision's time horizon, so too do the properties and, potentially, the preferred material. A firm whose selection criteria are based on a material's emergent properties will therefore require a selection framework that can account for the evolution over time in these properties. It is impractical, though, for the framework to consider *all* the ways material properties and the selection problem's context can evolve over time. Variations in some parameters will not affect the selection decision, and can be removed in others—for example, through contracts with energy or raw material suppliers.

Consequently, property evolution will only be considered if it results from *learning by doing*. Learning or learning by doing refers to the process of gaining experience through the repetition of a task—in the case of a manufacturing firm, through designing and manufacturing its products. This research focuses on learning as the basis for the evolution of emerging properties because the nature of the process is such that the firm can act to influence the evolution by deliberately gaining experience. As the firm learns, it applies its experience toward improving its products and manufacturing processes, which in turn alters the context under which emergent properties are defined and thereby the values of the properties.

Given that a firm often identifies preferred materials based on material properties that are affected by learning, such as manufacturing cost and product performance, it is worthwhile to assess the impact of considering learning on a firm's materials selection decision. Focusing on learning as a basis for property evolution has additional benefits in that learning is widely accepted and has been observed in a number of industries and products (e.g. [37, 43, 52]). Consequently, it is associated with an established framework for evaluating property evolution over time through its use of cumulative number of units produced as a proxy for the firm's experience. The rate and extent of change in an emergent property can then be predicted as a function of the total number of units a firm has produced by a given point in time. Incorporating this framework into materials selection can provide a means to account for the evolution over time of the emergent properties a manufacturing firm uses in its selection process.

Most firms, however, do not explicitly consider learning—let alone other reasons for property evolution—when selecting a new material. Not only does this incorrectly represent emerging properties used in materials selection, it also limits a firm's ability to systematically evaluate strategies for introducing new materials to their products. One such strategy is the use of a test bed—typically a low-volume, high-end product which a firm uses to experiment with new materials or features before expanding their use to other product lines. Experimentation with test beds enables the firm to learn, either by gaining experience with, or by obtaining information (e.g. market research) about, the new material. Both help the firm improve its products as well as its profits, but only the former implies that the firm believes material properties will evolve: if properties were invariant, as traditional selection methods assume, the firm would not use a test bed to gain experience because there would be no benefit to doing so.

The use of test beds in industry indicates that firms recognize that material properties evolve. Instead of explicitly accounting for this evolution, though, firms are more likely to rely on expert judgment to compensate for a traditional selection method's assumption that material properties are invariant. Consequently, there is a need for a systematic means to address the evolution of emergent material properties and the use of test beds in materials selection for product design and manufacturing.

This thesis therefore proposes to account for property evolution, specifically through the consideration of learning, in the materials selection process. The goal is to better inform the selection decision and to identify strategies, such as the use of a test bed, a firm can adopt to introduce new materials to its products. This is accomplished by using the learning framework to develop a more rigorous treatment of the evolution of emergent properties—in particular, manufacturing cost—within materials selection. The consideration of learning in the selection process enables a more accurate representation of each material's cost and therefore, a better-informed comparison of current and alternative options. It also allows the assessment of conditions under which learning will impact the selection decision, as well as strategies a firm can adopt when introducing new materials and when such strategies are applicable.

1.1 Thesis Outline

The remainder of this thesis presents materials selection methods, illustrated with case studies, that are aimed at better understanding when the consideration of cost evolution due to learning impacts the selection decision. Chapter 2 reviews literature concerning traditional materials selection methods, which assume time-invariant design requirements and material properties over the decision's time horizon; a background on learning is also covered in this chapter. The following chapter presents a framework for incorporating manufacturing cost learning into a traditional selection method, which is then illustrated with a case study concerning the selection of a material for the body-in-white of a midsize sedan.

Limiting the selection problem scope to a single application, however, cannot fully capture all the benefits a firm may realize from learning. It also restricts the potential strategies a firm can consider when introducing new materials to its products.

Using a test bed, for instance, would make more sense if the firm were able to apply the experience it gained to reduce not only the cost of the product used as the test bed, but also that of other, similar products that share the same resources and or use the same materials as the first product. Chapter 4 uses a stylized exercise to explore the use of a test bed in more depth and to motivate scope expansion of a selection problem to encompass multiple products and explicitly account for the decision's time horizon. A more formal framework for a multiple-product selection method that accounts for manufacturing cost evolution through learning is proposed in Chapter 5. This new selection method is illustrated with two case studies: first, the stylized exercise from Chapter 4 and second, materials selection for several applications within an automaker's fleet. Finally, Chapter 7 summarizes selection approaches and conclusions and presents potential next steps for this research. Case study details and code for materials selection algorithms are contained in the appendices.

Chapter 2

Background

In order to evaluate whether property evolution has any effect on a firm's preferred materials, it first helps to understand both the traditional materials selection methods and the learning framework as a basis for property evolution in manufacturing industries. This in turn enables the development of a selection framework that incorporates learning and thereby, the assessment of conditions under which learning will impact a firm's selection decision. The following chapter reviews materials selection methods proposed in literature and provides an overview of the basics of learning.

2.1 Materials Selection

The goal of materials selection tools is to inform a product's selection decision given design criteria that place constraints on the material's or product's performance. Since the material is selected for its role in the product's design, criteria are therefore set according to design requirements, which in turn are determined by consumer demand, manufacturing limitations, regulatory requirements, and so forth. Although its goal is straightforward, materials selection is, in reality, a complex problem. Part of this complexity arises from the number of properties or attributes a firm or product designer may have to consider, as well as the number of criteria dictating acceptable values of said properties. Other problem challenges include the inter-relationship among properties, processing, structure, and performance (see Figure 1-1); the potentially large number of materials a designer has to evaluate and select from; and uncertain or unknown properties, or ones that cannot be easily quantified [14].

Between this complexity and the importance of materials selection in product design, it is no surprise that there exists a large body of literature devoted to the topic. Numerous methods have been proposed for applying selection criteria to potential materials and identifying, if not *the* best solution, at least acceptable options. These methods include Ashby's performance indices [14] as well as other ranking methods such as TOPSIS [42], ELECTRE [66, 69], GRA [19], and TIES [46, 63]. Multi-attribute methods like linear programming [49] and utility analysis [40, 64] can also be used to identify preferred materials based on a designer's preferences, as can heuristic methods like CBR [12] and KBS [67]. Other optimization methods that have been applied to the materials selection problem include mixed integer programming [71], stochastic programming [36], and genetic algorithms [45, 63].

2.1.1 Literature Gaps

In addition to proposing materials selection methods, available literature also describes strategies for coping with some of the challenges a firm or designer faces when forced to choose a new material for its product. For example, some studies describe the application of fuzzy logic to quantify qualitative properties [20, 77]. Variations of certain methods, such as ELECTRE III, provide the designer with the necessary tools for handling incomplete or inaccurate material property data [65]. Even so, many of the methods in literature still have to make assumptions, such as those presented in the introduction, to simplify the selection process. One of these simplifications limits the decision scope to the identification of a single material for a single application or product; the majority of the methods reviewed are formulated around this premise. A few, however, such as the studies by DeCicco et al. [23, 24], Khajavirad et al. [45], and Roth et al. [63] are capable of selecting several technologies for implementation on a single product. Other papers, including work by Owen [59] and the Volpe model [68], all expand the scope even further to include other products manufactured by the firm or the industry.

Another simplifying assumption made by most selection methods is that material properties, product attributes, and design criteria are all invariant over the decision's time horizon. While this assumption is valid in most cases, is not necessarily true for emergent properties. These properties are, by definition, context-dependent, so if the context is evolving, the properties are as well. While there are many factors that can lead to changes in context over time, this research focuses on learning by doing as the reason for the change. Unlike other driving forces behind context—and thereby material property—evolution, the extent to which a firm learns is directly influenced by that firm's actions. Consequently, the firm has some control over the process and, when forced to work with new or unfamiliar materials, can choose to deliberately gain experience with them through repetition and cause their emergent properties to evolve.

In order to incorporate the evolution of properties, specifically through learning, into the materials selection process, the selection decision has to be evaluated over the firm's time horizon. Most selection methods proposed in literature do not account for the passage of time in the selection process, let alone address the evolution of material properties or the impact of considering learning on the preferred materials. Of the few studies that explicitly include a decision time horizon, none of the selection methods they propose are well-suited to accommodating learning in the selection process because, with learning, the results are path-dependent in that past actions can influence future decisions. This path-dependence is due to the learning framework's use of cumulative production volume to predict the evolved values of material properties. If these properties improve with increased cumulative volume, a firm will likely—but not necessarily—favor a material it has used in the past over a material it has no experience working with because it is further down the "learning curve" for the former.

Multi-stage programs, such as those developed by Gupta [36] and Li [49], mark the passage of time by identifying preferred materials based on current parameters (Stage 1), perturbing the system, and then re-evaluating material preference according to the system's new state (Stage 2). While this method can evaluate what happens given one particular path, it does not provide a means to easily and systematically evaluate and compare different paths (i.e. different perturbations) the system might follow when learning is present. Other publications, like those by Dutta [26] and Owen [59], use multi-period linear programs to address decision-making over a time horizon. Linear programs are capable finding an optimal path for a firm, even when problem parameters vary over the time horizon. Exactly how the parameters vary in any given context, however, has to be an input to the problem and therefore a known quantity. This is not the case with learning, when future decisions depend on previous actions due to its use of cumulative production volume to predict material property values—a non-linear calculation.

Both gaps discussed above are addressed in the materials selection methods proposed in this thesis. First, cost evolution due to learning is explicitly incorporated into a traditional, single-product selection method. The need for a multi-product method that also considers learning is later motivated and such a method developed. The resulting methods are used to analyze the impact of considering learning on a firm's material preferences.

2.2 Learning

A firm's ability to learn by doing is one of the many reasons a material's emergent properties evolve over time and is the focus of this research. Not only has learning been observed in a wide variety of industries, it also has an already established framework that can be incorporated into a materials selection method. Additionally, many of the emergent properties that can be affected by learning, including manufacturing cost and product performance, often factor into a firm's materials selection process. This thesis focuses specifically on the evolution of manufacturing cost due to learning.

In order to discuss learning within materials selection, it helps to first understand what learning involves and how it can be incorporated into the selection framework. Learning by doing refers to the observation that as firms—or people—perform a single task repeatedly, they gain experience at that task and thereby become more proficient at it. Consequently, the cost of performing that task—whether measured in units of time, money, material usage, etc.—decreases each successive time that task is carried out. This cost can be plotted, often as a function of the cumulative number of times the task is performed, to create what is known as a *learning curve*. A firm that starts using a new manufacturing process therefore begins at the top of the learning curve; as it gains experience through the production of each additional unit that uses the new process, it requires fewer resources to produce each subsequent unit and thus moves down the curve.

An article by Wright [79] concerning learning in the manufacture of aircraft is often cited as the first publication to frame the concept of learning. Since then, learning has been documented in various industries and for a number of products, including automotive [43], chemical processing [50], semiconductor manufacturing [37, 41], metal products [25], and energy technologies [52]. Learning curves in these studies plot unit cost, direct labor hours per product, yield, or

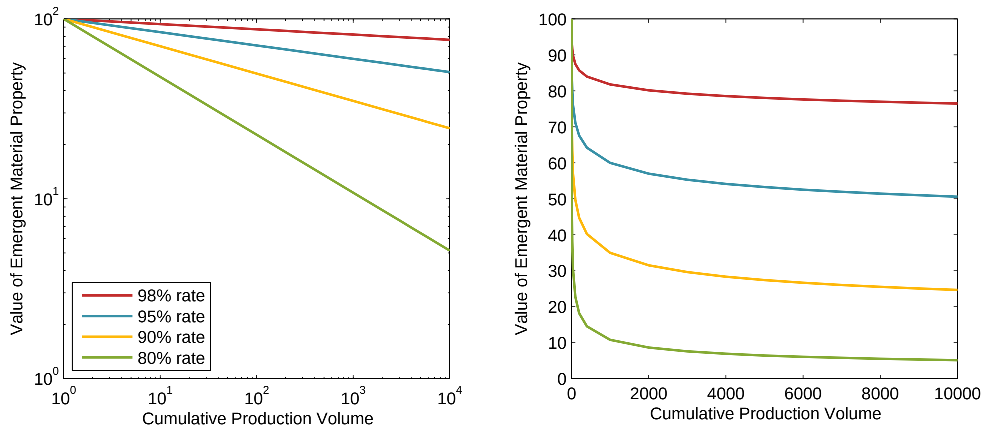

Figure 2-1: Log-linear learning curves for a generic emergent property in log-log and linear space.

other parameters as a functions of the cumulative number of units produced. A log-linear representation is one of the preferred functional forms of the learning curve [18, 81] and has the equation:

$$y = y_o \cdot x^b, \quad b = \frac{\log r}{\log 2}$$

where y_o and y are the initial and learned values of a material property, x is the cumulative volume, and r is the learning rate. Given this functional form, property y decreases by $(1 - r)$ percent every time cumulative volume doubles. Figure 2-1 plots log-linear learning curves for learning rates of 98%, 95%, 90%, and 80%. S-curves and other functional forms are also used in learning literature (e.g. [4, 43, 56, 68, 81]).

Often, a curve is simply fitted to cost or labor data and the learning rate (r) estimated from there. Some literature takes the extra step of investigating the driving forces behind the process improvements since in order to realize lower costs or other manufacturing process improvements, a firm has to make changes to that process [51, 57]. For example, improvements in manufacturing cost can be attributed to changes in process parameters such as cycle time, reject rate, and so forth, as well as to more efficient design [43]. Some changes are consciously made, whereas others take place simply because workers perform their tasks more efficiently [37].

Regardless of whether the underlying reasons for the evolution of emergent properties are known, learning curves are a compact way to model property evolution over time and ought to be accounted for in the materials selection process to more accurately represent properties like manufacturing cost that are known to be influenced by learning. The following chapter presents an approach for incorporating learning curves for unit manufacturing cost into a traditional materials selection method.

Chapter 3

Materials Selection for a Single Product

This chapter presents a possible approach for considering learning in the materials selection process for a single product or application. A selection methodology is proposed and subsequently demonstrated with a case study built around selecting the material for the body-in-white of a midsize car. A sensitivity analysis of the case study is conducted to determine conditions under which learning alters the preferred material, and thereby when firms should be concerned about considering learning.

3.1 Traditional Selection Methodology

The goal of this chapter is to investigate whether the consideration of evolution in manufacturing cost due to learning impacts a firm's materials selection decision. This is accomplished by modifying a traditional selection method to account for learning in the selection process. Regardless of the method's specifics, however, it will have to perform the basic functions of a materials selection tool and identify the best material for a product given design criteria.

To keep things straightforward, a basic ranking approach was chosen for the selection method. These methods rank materials according to a metric, which in turn is chosen by the firm based on design criteria. The metric can be something as simple as density or elastic modulus, or can combine a number of material attributes into a single measure. Typically, ranking methods also include an initial screening step to eliminate the material options that are clearly unsuitable for the product. This step, however, is assumed to have already been performed for the analyses in this research so that the firm's options have already been limited to a small set of candidate materials. The metric is then evaluated for each of these remaining options and the alternatives ranked according to metric value. Finally, the material associated with the optimal (i.e. highest or lowest, depending on design criteria) metric value is designated as the preferred option.

3.1.1 Adding Learning

The selection method described above only dictates how the materials are ranked and the preferred material identified. In order to determine whether the consideration of learning impacts the ordering of the material options, learning first has to be incorporated into the selection process. In the case of the above selection method, this is done through the metric by choosing a property that is affected by learning and is a concern of all manufacturing firms—namely, manufacturing cost. In this analysis, manufacturing cost is evaluated in three different ways, two of which are short-term and long-term unit cost. Short-term or *unlearned* cost represents the initial cost of a new material or the cost a manufacturer pays when it first starts working with that material. Conversely, long-term or *learned* cost is the cost after the firm has done all it could to improve the manufacturing process and has reached some physical or other limitation that prevents further evolution in cost. Depending on the process and the firm's annual production volume, it may take years or even decades before a firm's manufacturing cost reaches its long-term value and ceases to evolve.

Neither short-term nor long-term unit costs, though, represent the true cost to a firm because analyses that use these values still make the assumption that cost does not evolve over the decision's time horizon. That is, the manufacturer pays either each material's short-term unit cost, or its long-term unit cost for each unit of product manufactured over the time horizon. The third way of evaluating manufacturing cost, on the other hand, accounts for cost evolution due to learning by calculating the total manufacturing cost to the firm over the time horizon. Total manufacturing cost adds the unit cost of each unit of product manufactured by the firm within the time horizon. The product's unit cost is determined by a material's learning curve, which defines the path the evolving cost takes from when the material is first used until it becomes an established technology. Each additional unit the firm produces adds to its experience and, thus, to its familiarity with a new material—represented by a step along the learning curve. To-

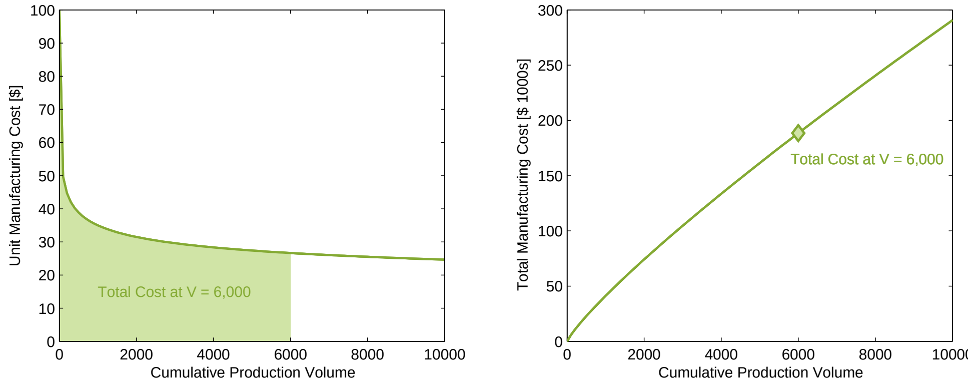

Figure 3-1: Total manufacturing cost for 90% log-linear learning curve.

tal (undiscounted) manufacturing cost from when a firm first introduces a new material until the end of the time horizon is therefore the area under the learning curve from when the curve begins at zero units to the total number of units the firm has produced by the end of its time horizon. This total cost given cumulative volume is plotted in Figure 3-1 for the log-linear curve with a 90% learning rate (see the yellow curve in Figure 2-1) and a cumulative volume of 6,000 by the end of the time horizon. The impact of the consideration of learning on a firm's material preference is then assessed by comparing the ordering of the material options when ranked according to their short- or long-term manufacturing cost to their ordering according to total manufacturing cost as calculated from each option's learning curve.

3.2 Case Study: Body-in-White of a Midsize Car

A case study is employed to illustrate the selection method described in the preceding section and to test whether the consideration of learning can influence a materials selection decision. The study focuses on an automaker seeking to improve the fuel economy of one of its midsize cars by reducing the weight of the car's body-in-white (Figure 3-2) through the use of alternative materials. An automotive case study is used to provide a real-world example in which a manufacturing firm is faced with the need to adopt new or unfamiliar materials for its products, whether because of consumer demand or tighter government regulations that place constraints on the vehicle's fuel economy.

This case study deals with an automaker's desire to find an alternative, lightweight material

Figure 3-2: A car body-in-white [75].

for the body-in-white of a midsize car that it is planning to manufacture at an annual production volume of 200,000 units per year over a period of five years. The preferred material is selected based on its total manufacturing cost over the decision's time horizon, with lower-cost alternatives considered more desirable. Since all material options are assumed to satisfy the minimum performance requirements for the body, only their costs matter to the automaker, who prefers the least expensive option.

It would be more economically accurate, however, to base the decision upon firm profitability rather than production cost as the two approaches are anticipated to yield identical results in this case study, provided the automaker is a price taker for its midsize car. The cost minimization approach also has the added benefit that it does not require knowledge of the vehicle's price elasticity of demand for its calculations. Although if the automaker does have knowledge of the demand curve, it can use that to its advantage and better inform the selection decision by optimizing for profit. Profit-based analyses are feasible within this framework (see Appendix A for a discussion of an approach), but are generally not illuminating in the case of materials selection for a single application.

In order to investigate whether considering learning has any impact on the preferred material, the unit manufacturing cost of each alternative is evaluated using the three approaches described in Section 3.1.1: 1) short-term unit cost, before any cost evolution takes place; 2) long-term unit cost, after costs have ceased to evolve; and 3) total (evolving) cost over the five-year time frame. Only the final option explicitly considers cost evolution and, thereby, learning in the decision metric. In order to be directly comparable to the short- and long-term units costs, the total cost of each material is scaled by a constant equal to the annual production volume (APV) multiplied by the sum of the discount factors over the five years or

$$\begin{aligned} \text{Average cost factor} &= APV \cdot \sum_{y=1}^{5} \frac{1}{(1+r)^{y-1}} \\ &= 200,000 \cdot 4.31 = 862,000 \end{aligned} \tag{3.1}$$

to arrive at an "average" unit cost. The calculation in Equation (3.1) assumes the discount rate, $r = 8\%$. Once the manufacturing costs have been evaluated, the material options are then ranked according to each cost calculation approach and the orderings compared. Learning is deemed to affect the selection decision when the ordering by average cost produces different results from orderings by either short-term or long-term unit cost.

3.2.1 Case Study Inputs

In this case study, the automaker has a choice of four feasible alternative materials for the body-in-white of the midsize car: high strength steel, aluminum, glass fiber composite, and carbon fiber composite. Mild steel is also included in the analysis as it represents the baseline material the automaker is currently using for the body. The material options are then ranked according to their total manufacturing costs, and the results analyzed to answer whether the consideration of learning can help the automaker better inform its selection decision.

Before the options can be ranked, the manufacturing cost of each has to be obtained. Estimating the manufacturing cost of a vehicle body produced from any of the above materials, however, is not as simple as multiplying the raw material's unit price by the body's mass. Because cost is an emergent property, its calculation requires additional information beyond a material's specifications, including body designs, manufacturing process parameters, plant operational conditions, and so forth. To this end, data describing body designs for each of the material options were obtained from various sources and scaled to ensure the final results are comparable. Table 3.1 shows a few high-level attributes for each body design. This data was then entered into process-

Table 3.1: Lightweight body designs for a midsize car.

Strategy ID	Primary Material	Manufacturing Process	Mass [kg]	Mass [% of MS]	Number of Components	Source
MS	Mild steel	Stamping	322	100%	143	[33]
HSS	High-strength steel	Stamping	243	75%	60	[74]
AL	Aluminum	Stamping, extrusion	193	60%	111	[44]
GF	Glass fiber composite	SRIM	219	68%	62	[35]
CF	Carbon fiber composite	SRIM	138	43%	62	[35]

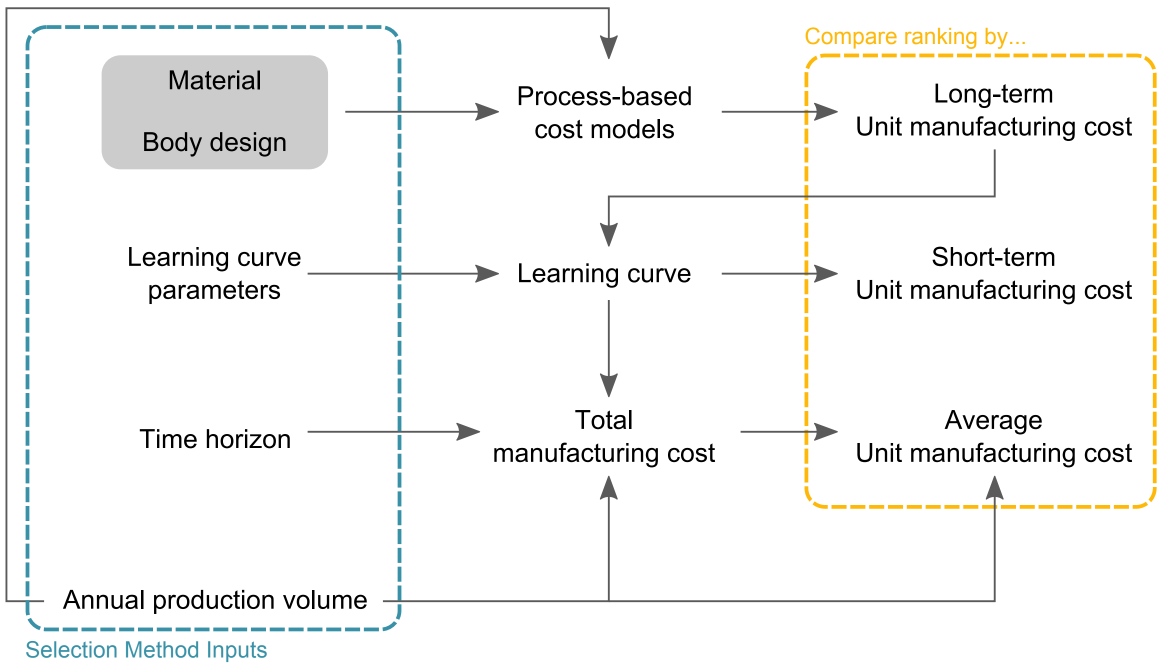

Figure 3-3: Schematic of models and other parameters required for calculating manufacturing cost of car bodies.

based cost models (discussed in more depth in the following section) to predict long-term unit cost. Short-term unit cost and average unit cost over the five-year time frame are then calculated based on the material's long-term cost and learning curve parameters. Figure 3-3 summarizes the models, learning curves, and other parameters required to calculate manufacturing cost.

PBCM Calculation of Long-Term Manufacturing Cost

Process-based cost models (PBCMs) were employed to predict the long-term unit manufacturing cost of vehicle bodies at the specified annual production volume of 200,000 units per year. These process-specific models use detailed data that describe the individual parts of each body design to estimate the cost of forming and joining these parts to produce a vehicle body. With the exception of the assembly (joining) cost model, the process-based cost models take a part to be manufactured and from there estimate process requirements such as press tonnage and number of dies necessary to produce a given volume. These requirements, in turn, are combined with operational conditions for the plant (e.g. number of shifts and downtime) and factor costs (e.g. raw material price and labor wage) to calculate the unit cost of forming the part. Each part in a body's design is cost modeled and the resulting unit costs summed to generate the unit cost of forming the whole body. The bodies used in this case study required models for the following manufacturing processes:

- Die casting
- Extrusion
- Stamping
- Structural reaction injection molding (SRIM)
- Tailor-welded blanking (TWB)

Because process-based cost models strive to be predictive, they are very useful in comparing emerging technologies or design concepts that are not yet standard for a firm. These models are also valuable for predicting the unit cost of a part over different production volumes. Cost calculations throughout the models and in their final outputs are broken down into fixed and variable costs. Fixed costs are amortized over the lifetime of the product, equipment, or building in order to assign a portion of the cost to each unit or part. If the equipment is non-dedicated, its amortized cost is spread over multiple parts according to the fraction of its up-time it spends manufacturing any given part. Cost breakdown within the cost models is as follows:

- Variable costs
    - Materials – raw materials that constitute the part itself
    - Process materials – materials such as lubricants required by the process
    - Labor – wages paid to factory workers (can include benefits)
    - Energy – energy required to power the equipment
- Fixed costs
    - Main and auxiliary equipment – equipment such as presses

Table 3.2: Exogenous cost model assumptions.

Parameter	Value
Working days per year	235 days
Labor wage	\$35 / hr
Energy cost	\$0.07 / kWh
Building unit cost	\$1,500 / sq m
Interest	8%
Working capital period	3 months
Product life	5 years
Equipment life	13 years
Building life	40 years
Indirect workers per direct worker	0.25
Indirect workers per line	1
Idle space	25%
Capacity utilization	100%

◇ Tooling – tools or dies specific to the part

◇ Building – cost for the factory

◇ Maintenance – upkeep of equipment, tools, and the building

◇ Overhead – overhead labor

◇ Working capital cost – invest first then earn

The usefulness of the cost model results depends on the model's assumptions. Variations in factors such as forming reject rate, labor wage, material formability, and raw material price can all potentially affect cost. Tables 3.2 and 3.3, respectively, list exogenous model assumptions and material prices used in this study. These numbers, along with other process-specific inputs, take a long-term perspective on the manufacturing process, meaning that they assume the automaker is already experienced with the process and does not expect to realize any further process improvements (and thereby reductions in cost). For instance, in addition to the numbers listed in Table 3.2, the stamping process-based cost model assumes a low reject rate, fast cycle time, minimal downtime for equipment maintenance or repair, and so forth. Consequently, model results represent long-term unit manufacturing cost. Material prices, though, are set to reflect market conditions during the first half of 2009. More information on these cost models can be found in [47].

Table 3.3: Raw material and scrap prices.

Material	Price [$ / kg]	Scrap price [$ / kg]
Mild steel sheet[a]	$0.87	$0.26
High strength steel sheet (210 MPa)	$0.92	$0.26
High strength steel sheet (280 MPa)	$0.97	$0.26
Aluminum ingot[b]	$1.50	$1.01
Aluminum blanks – 5182	$3.48	$1.52
Aluminum blanks – 6111	$3.48	$1.52
Aluminum blanks – 5754	$3.35	$1.52
Aluminum sheet – 5083, BN lubricated[c]	$4.23	$1.99
Magnesium ingot[d]	$2.92	$1.46
Magnesium sheet – AZ31[e]	$5.00	$2.25
Magnesium sheet – AZ31, BN lubricated[c,e]	$5.82	$1.65

[a] Average of MEPS cold rolled coil price between January and April 2009 [2].
[b] Estimated from Platt's Metals Week prices between January and June 2009, for example, [7].
[c] Boron nitride lubricant cost is $0.55 per kilogram.
[d] Taken from Platt's but divided by two for long-term contract estimate [7].
[e] Sheet fabrication prices estimated from [38].

Assembly Cost Model

The assembly model is another process-based cost model, but the process line it represents is constructed differently from the forming process lines in that it scales by increasing line length rather than by adding parallel lines. Assembly lines are composed of stations that are responsible for performing the necessary process steps. The number of stations in each line—and therefore the equipment and tooling required—depends on annual production volume, which also determines the line rate. At higher volumes, finished vehicles have to leave the line at a faster rate. The rate, in turn, limits the amount of time a vehicle can spend at any given station. Consequently, multiple stations may be required to complete a single joining process step in order to keep the line moving at the required rate. In contrast, line rate and the number of stations are fixed parameters in the "classic" process-based cost models; higher volumes are accommodated by adding parallel lines rather than additional stations to the existing line.

The assembly model predicts the cost of joining any number of units given the joining process for any two parts or subassemblies and the "quantity" of that process. Because of the structure of the line, as discussed above, the assembly process does not have a fixed process flow. Rather, the cost to assemble a multiple-component subassembly, such as a door, is evaluated for a number

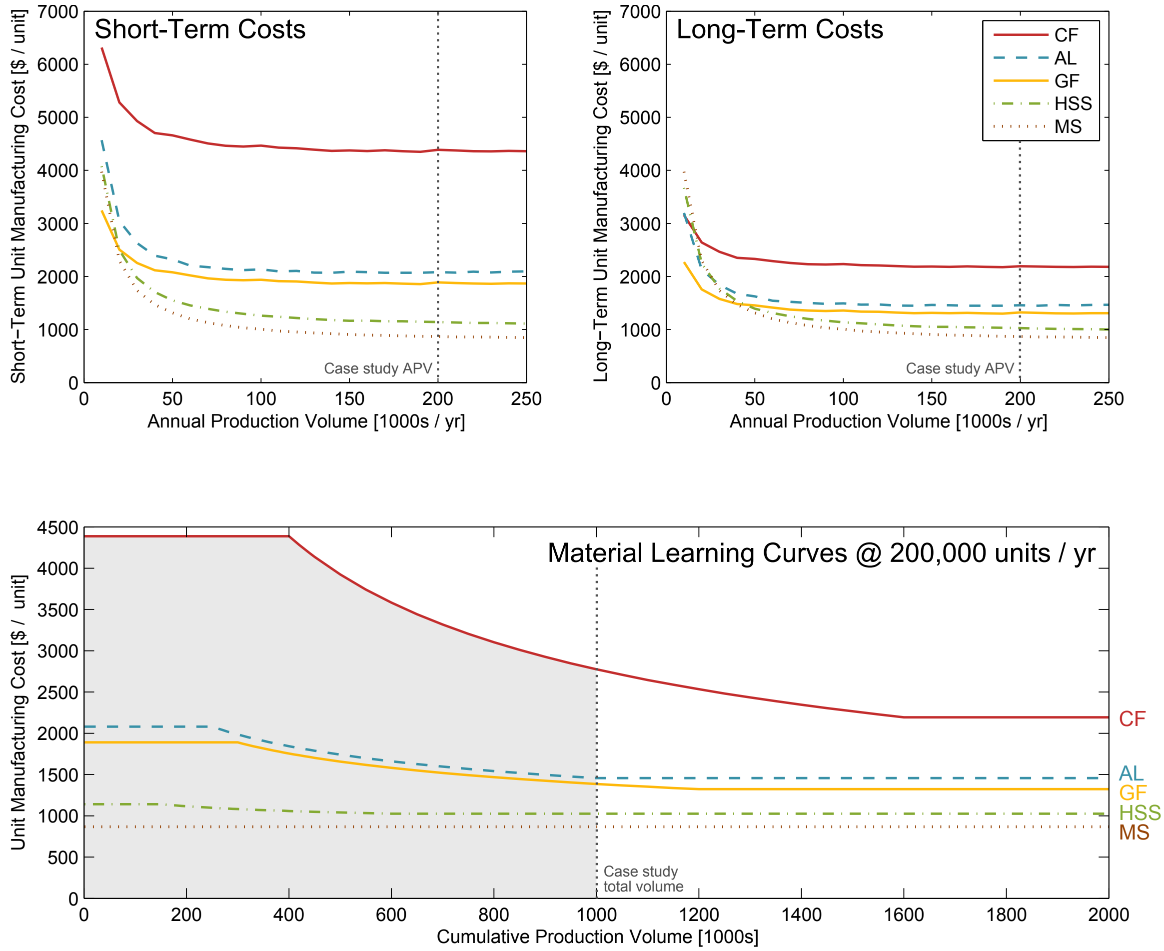

Figure 3-4: Short- and long-term unit costs, and learning curves, of each material option considered in the body-in-white case study.

of proposed serial lines generated by the model from the inputted joining processes and with varying levels of automation. The model then attaches a cost to the final, most cost-effective iteration. Varying production volume causes the model to adjust line rate and generate a revised assessment of the line design and thus a new cost estimate [80].

Short-Term and Average Costs

The chart in the upper right of Figure 3-4 plots long-term unit manufacturing costs, estimated by the process-based cost models, of each material option as functions of the annual production volume of the midsize body. The production volume of 200,000 units per year is also denoted on the plot. Once the long-term unit manufacturing costs of the material options are generated, they are used to construct the learning curves and to calculate the short-term unit manufacturing costs (also shown in Figure 3-4), as well as the average unit costs over the automaker's five-year time frame. Rather than use a log-linear functional form, which exhibits a steep drop in parameter value over the first few units manufactured by the firm, the learning curves in this case study follow S-shaped paths, in which cost saturates at both low and high cumulative volumes, but evolves in a log-linear fashion (i.e. each doubling of cumulative volume leads to the same percent decrease in cost) at intermediate volumes. This curve is the same as the one employed by the Volpe Model, which is used by NHTSA (National Highway Traffic Safety Administration) to set fuel economy targets for CAFE (Corporate Average Fuel Economy) regulations [4, 68]. This case study adopts a slightly modified form of the equation:

$$y(V) = y_o \left(\sqrt{1-\sigma}\right)^{\max(0,\log_2(\min(V,V_{hi})/V_{th}))} \tag{3.2}$$

where y is the unit manufacturing cost (or more generally, the emergent property), V the cumulative volume, V_{th} the threshold volume at which y first starts changing, V_{hi} the maximum volume at which y saturates, and σ the learning scope. For this curve, learning scope represents the extent to which manufacturing cost has evolved by the time cumulative volume reaches V_{hi} so a scope equal to 0 corresponds to no cost evolution and a scope approaching 1, to the property y evolving to its maximum extent (that is, $\lim_{\sigma \to 1} y(V > V_{hi}) \to 0$). The learning rate, ρ, or how fast cost decreases each time production volume doubles, is determined by σ according to

$$\rho = 1 - \sqrt{1-\sigma}$$

provided $V_{hi} = 4V_{th}$. Higher values of σ reflect faster rates or higher values of ρ because cost evolves further over the same number of units (as defined by V_{th} and V_{hi}). Values of the S-curve parameters for each material are shown in Table 3.4; they are loosely based on the numbers used by NHTSA for its 2011 CAFE target analysis [4] with $V_{hi} = 4V_{th}$ for all material options. For instance, the learning scopes of 30% to 50% are in line with NHTSA's assumption that technology cost decreases by 20% ($= \rho$) with each doubling of cumulative production volume. Such numbers are realistic for the automotive industry—and may even be on the low side—as evidenced by data collected by Kar [43] and Nadeau [56]. Figure 3-5 illustrates the relationship between short- and

Table 3.4: Learning curve parameters for each body-in-white material.

Strategy ID	V_{th}	V_{hi}	σ	ρ
MS	150,000	600,000	0.0	0.0
HSS	150,000	600,000	0.0	0.0
AL	250,000	1,000,000	0.3	0.16
GF	300,000	1,200,000	0.3	0.16
CF	400,000	1,600,000	0.5	0.29

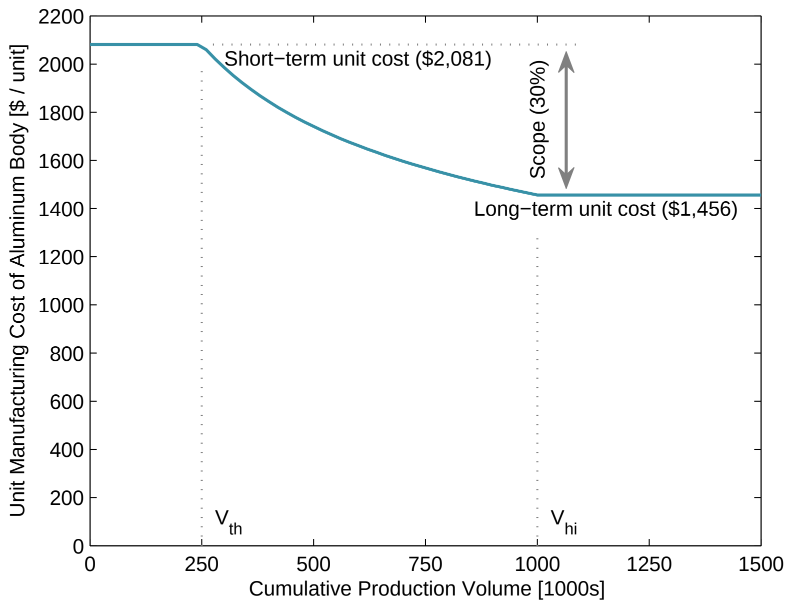

Figure 3-5: Learning curve for aluminum body.

long-term unit cost using the learning curve for the aluminum body, which is produced using Equation (3.2) and the material's parameters from Table 3.4.

The final step in analyzing manufacturing cost for the selection method is to compute the average unit cost of each material option. Assuming that the cost of each option follows its respective learning curve exactly as it evolves, the average cost is the total cost to the automaker over that five-year period, divided by the aforementioned factor (Equation (3.1)).

The total cost of all bodies produced given a material option is related to the total area under the learning curve between zero and the cumulative volume after five years or one million units

Table 3.5: High volume body-in-white manufacturing costs.

Strategy ID	Unit Cost Short-Term	Unit Cost Long-Term	Unit Cost Average	Total Cost
MS	\$868	\$868	\$868	\$750 M
HSS	\$1,140	\$1,026	\$1,067	\$920 M
AL	\$2,081	\$1,456	\$1,805	\$1,560 M
GF	\$1,890	\$1,322	\$1,689	\$1,460 M
CF	\$4,389	\$2,194	\$3,856	\$3,330 M

(200,000 units per year × 5 years). The area is calculated by integrating the learning curve:

$$\begin{aligned}\text{Average unit cost} &= \text{Average cost factor} \cdot \text{Total cost} \\ &= \text{Average cost factor} \cdot \int_0^{10^6} y(V)dV\end{aligned}$$

where V is the cumulative production volume over the five-year time frame and $y(V)$ is the learning curve in Equation (3.2). It should be noted that the above equation assumes cost is not discounted (which affects the factor in Equation (3.1). Adding a discount factor requires piecewise integration with the integration range broken down into blocks of 200,000 units (the annual production volume). Each of these integrals is then multiplied by its appropriate discount factor, and these values summed to calculate discounted average unit cost. The resulting unit cost associated with each material is shown in Table 3.5, along with long-term unit cost from the process-based cost models and short-term unit cost.

3.2.2 Results and Sensitivity Analysis

The manufacturing costs calculated above are next used to rank the different materials for the midsize car's body-in-white. These rankings are shown in Table 3.6 and indicate that the different approaches to evaluating manufacturing cost lead not only to the same preferred material—mild steel—but also to the same ordering of the options. Even if mild steel is removed from consideration in the selection process, high-strength steel is identified as the preferred alternative material regardless of how unit manufacturing cost is calculated. It therefore appears that cost evolution has no impact on the preferred materials. A comparison of the learning curves for the material options (bottom plot in Figure 3-4) confirms this conclusion: the learning curves never

Table 3.6: Ranking of material options according to high volume unit manufacturing cost.

Rank	Short-Term	Long-Term	Average
1	MS	MS	MS
2	HSS	HSS	HSS
3	GF	GF	GF
4	AL	AL	AL
5	CF	CF	CF

cross. Consequently, there is no reason to expect that the ordering of the materials will change between cost evaluation approaches, even if the decision's time horizon is shortened or extended.

This conclusion, however, is specific to the conditions outlined above and does not imply that the consideration of learning never impacts the materials selection decision. On the contrary, manufacturing cost is context-dependent so changes to material parameters or to the problem's context can potentially lead to a different outcome. One possible context change is to lower the vehicle's annual production volume, which, in turn, affects the unit manufacturing cost of each material because the firm will no longer be able to leverage economies of scale. Since the material options have different fixed and variable costs, their unit costs will respond differently to a change in annual production volume. Another possible change is to alter a material's learning curve parameters; such a change may be warranted, especially if there is some uncertainty in the original values. The remainder of this section explores whether changes to the vehicle's annual production volume and material learning curve parameters impact the preferred material or the rank of the material options.

In the first sensitivity analysis, annual production volume is reduced to 20,000 units per year, down from 200,000 units per year. Twenty thousand units per year, in the automotive industry, is considered to be low volume production because at this level, fixed costs dominate total manufacturing cost and increasing the volume by the marginal vehicle has a significant affect on the unit cost. In contrast, 200,000 units per year is considered to be high volume production because variable costs dominate total manufacturing cost and producing an additional vehicle has minimal impact on unit cost.

Lowering the volume, however, means that unit manufacturing costs have to be recalculated. As before, the long-term cost is predicted using process-based cost models, with short-term and average costs calculated as described above. The resulting (discounted) costs are shown in

Table 3.7: Low volume body-in-white manufacturing cost.

	Unit Cost			
Strategy ID	Short-Term	Long-Term	Average	Total Cost
MS	\$2,297	\$2,297	\$2,297	\$200 M
HSS	\$2,497	\$2,247	\$2,497	\$220 M
AL	\$3,047	\$2,133	\$3,047	\$260 M
GF	\$2,508	\$1,755	\$2,507	\$216 M
CF	\$5,282	\$2,641	\$5,281	\$455 M

Table 3.8: Ranking of material options according to low volume unit manufacturing cost.

Rank	Short-Term	Long-Term	Average
1	MS	GF	MS
1	HSS	AL	HSS
2	GF	HSS	GF
3	AL	MS	AL
4	CF	CF	CF

Table 3.7. These costs are clearly higher than those in Table 3.5 because tooling and other fixed costs are spread over fewer units.

Table 3.8 shows the materials ranked according their to low volume manufacturing costs. Looking at the table, it is immediately apparent that the short- and long-term costs result in different material orderings. This is partly due to the variation in learning parameters among the different options, partly to the approach used to derive short-term cost from long-term cost. (The ratio between short-term cost and long-term cost is fixed; therefore, at high volumes when car bodies are cheaper, the difference between short-term and long-term cost is small; in contrast, at low volumes when bodies are more expensive, the difference between the two costs is larger. For the above case, this difference is sufficient for both aluminum and glass fiber to trade places with high-strength steel and mild steel.) The ordering of the material options by average cost, however, is no different from that of short-term unit cost because the production volume is so low that the automaker does not reach V_{th} by the end of the five-year time frame and therefore does not experience any cost evolution. In other words, the automaker does not produce enough vehicles to gain sufficient experience to improve its manufacturing process and reduce its cost. This is supported by the data in Figure 3-6, which plots the learning curves of the material options at an annual production volume of 20,000 units per year. Although the curves eventually cross—which leads to a different ordering when long-term unit costs are considered—the automaker,

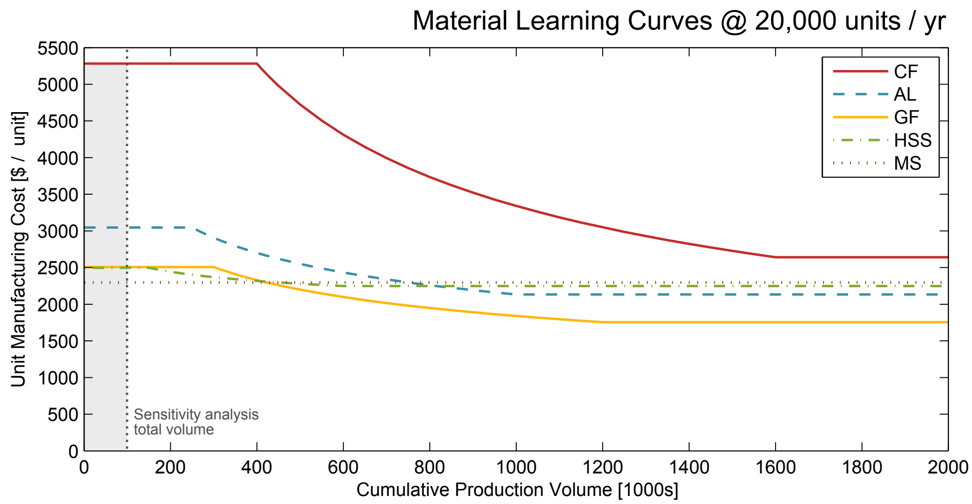

Figure 3-6: Learning curves of material options for an annual production volume of 20,000 units per year.

by the end of the five-year time frame, is nowhere near that point and has only produced 100,000 total vehicles. In order to see a difference, the automaker would have to manufacture almost 400,000 vehicles; at an annual production volume of 20,000 per year, this would take nearly 20 years. Consequently, considering learning at low volumes does not provide any new information because the ordering of material options for average cost remains the same as that for short-term cost.

The second sensitivity analysis for this case study uses the initial high-volume scenario to investigate whether changes to the parameters of a material's learning curve can affect the ordering of the material options. Specifically, it considers a faster learning rate (i.e. lower values of V_{th} and V_{hi}) for the aluminum body. This faster rate reduces the average unit cost of aluminum but leaves the average costs of all other materials, as well as all short- and long-term unit costs, untouched. Figure 3-7 plots the average unit cost of aluminum as a function of its threshold volume (assuming $V_{hi} = 4V_{th}$), along with the average costs for the other materials. As can be seen from the figure, the cost of aluminum is below that of glass fiber for threshold volumes below 160,000 units—indicating that at faster learning rates, aluminum is preferred over glass fiber for a midsize car's body. In contrast, the short- and long-term results suggest that glass fiber is preferred over aluminum regardless of the learning rate (see Table 3.6).

The change in order at faster learning rates is because the learning curve of aluminum is

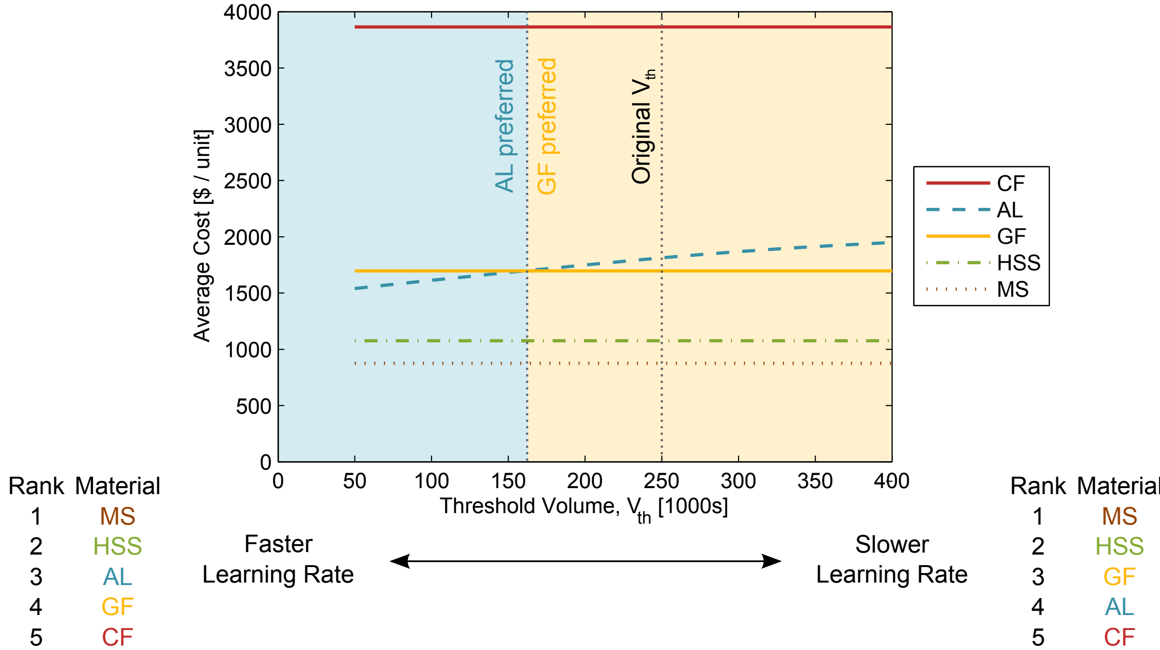

Figure 3-7: Material average unit cost versus the learning rate for aluminum, as set by threshold volume.

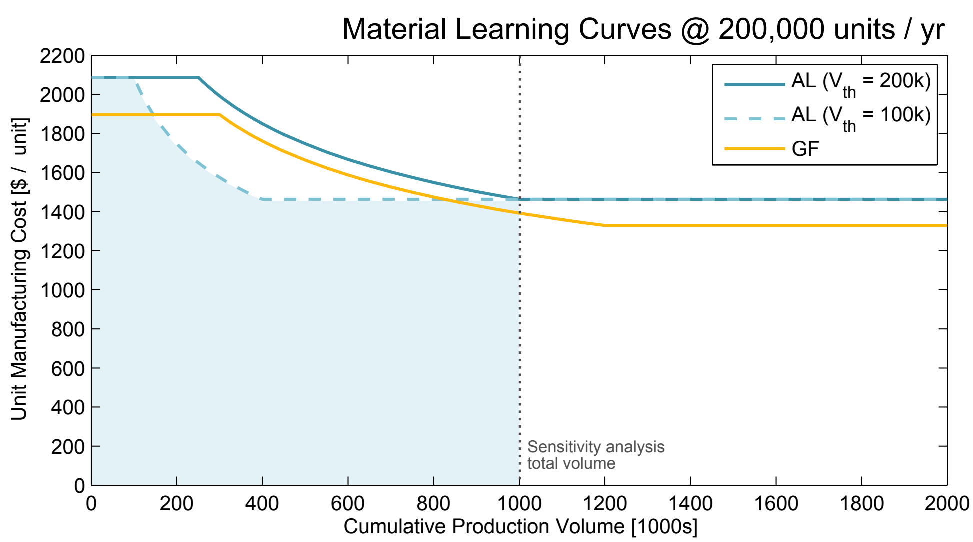

Figure 3-8: Modified aluminum learning curve with $V_{th}^{AL} = 100,000$.

able to undercut that of glass fiber over a range of cumulative volumes, and thus, reduce that material's average cost by reducing the area under its learning curve (Figure 3-8). If the time frame is long enough, however, the average cost of the glass fiber body will eventually "catch up" to that of the aluminum body and the former will prove to be the cheaper solution because of its lower long-term cost. Changing the threshold volume of aluminum, however, does not change the fact that mild steel remains the least expensive option. Regardless, under certain conditions—in this case, faster learning rates for aluminum—considering learning can make a difference in material rank. A higher scope can also impact the materials' rankings, but is more complex to analyze because changes in scope affect both the short-term and the average costs; consequently, material order according to average cost may be different from that of the long-term cost, but tell the same story as short-term cost.

3.2.3 Summary

The above case study investigated whether the consideration of learning in the manufacturing cost calculation could impact the material selection decision for the body-in-white of a midsize car. To accomplish this, materials options were ranked using three approaches to estimating manufacturing cost: 1) short-term unit cost, 2) long-term unit cost, and 3) average cost over the five-year time frame. The resulting orderings were then compared to determine whether ranking the materials according to average cost, the only cost calculation approach that considers cost evolution through learning, would lead to a different preferred material or, at the very least, a different ordering of the material options. Case study results showed that:

- ◆ At high volume (200,000 units per year)...
  - ◇ Mild steel, the baseline material, has the lowest cost regardless of how manufacturing cost is calculated and is always the preferred option.
  - ◇ Learning curves do not cross so the ordering of the material options does not change regardless of cost calculation approach.
- ◆ At low volume (200,000 units per year)...
  - ◇ Mild steel has the lowest short-term unit manufacturing cost and the lowest average cost; however, glass fiber composite has the lowest long-term unit cost.

- ◇ Glass fiber composite has the lowest long-term unit cost; however, it would take the automaker around 60 years (given input parameter values) to gain sufficient experience to reach long-term costs.
- ◇ Unit manufacturing costs are not as spread out as at high volume so learning curves cross and the ordering of the material options changes between cost calculation approaches.

- ◆ Learning curves can also cross, depending on material learning rates and scopes.
  - ◇ Increasing the learning rate of aluminum reduces the average cost of the aluminum body so that at high rates, it is eventually preferred over the glass fiber composite body.
  - ◇ If the material's learning scope is not large enough, however, a higher learning rate will not make a difference.

Exactly which material is preferred, however, will ultimately depend on the automaker's time horizon and the long-term costs of the material options. In the end, the results are case-specific: different materials or products can lead to different conditions under which considering learning has an impact so each situation will have to be analyzed separately. However, the midsize car case study at least proves it is possible for considering learning to lead to a different conclusion; the automaker can thereby better inform its selection decision by accounting for learning in its cost calculations.

Chapter 4

Rethinking Learning in Materials Selection

The preceding chapter investigated the impact of incorporating cost evolution due to learning by doing into a traditional materials selection method. The modified selection method was demonstrated with a case study of materials selection for the body-in-white of a midsize car. Results from the case study illustrated that, under certain circumstances, the consideration of cost evolution in the materials selection process could impact the ranking of the material options when ordered according to manufacturing cost. Mild steel, however, was the preferred material throughout the case study, even if less preferred materials, such as aluminum and glass fiber composite were reordered between different scenarios. (The one exception to this observation was when the materials were ranked according to their long-term costs at low production volume: in this case, glass fiber composite was preferred. This scenario, however, can be disregarded as not viable because of the low annual production volume and consequently, the extremely long time frame required for the automaker to gain enough experience and realize long-term manufacturing costs.) Therefore, one possible conclusion from the case study is that an automaker will never introduce an alternative material for use on a vehicle unless it is forced to do so, whether because of regulation requirements, consumer demand, or other constraining factors.

This statement, though, does not match reality: it is a known fact that automakers—and other manufacturing firms—adopt new materials, even when those materials are not essential to the product's design and despite their higher costs. For example, automotive manufacturers such as

Audi and General Motors currently sell vehicles with aluminum body structures or carbon fiber panels (e.g. [53, 61]); likewise, Airbus and Boeing have incorporated composites into their commercial aircraft designs (e.g. [16, 31, 62]). The beverage and sporting goods industries have also, in the past, adopted new materials for their products [17, 55]. Many of these firms, however, still manufacture and sell products made from the "original" materials they used prior to the introduction of the new material: as a case in point, the automotive and aircraft industries continue to use steel and aluminum, respectively, for their products. These observations indicate that manufacturing firms, including automakers, introduce alternative materials to their products, even when they are not forced to use such materials—which directly contradicts the conclusion drawn from the case study's results that firms will not implement alternative materials unless those materials are absolutely necessary.

One potential reason for the discrepancy in these conclusions is that firms behave irrationally by adopting new materials when they do not have to. More likely, though, is that the proposed materials selection framework is missing a key component and consequently, is limited in its comprehension of why a firm might be motivated to switch to an alternative material, especially when that material is more expensive than the original one. More to the point, the framework, as presented in Chapter 3, will only identify an alternative material as the preferred option if that material is cheaper than the baseline material currently in use. For the body-in-white case study, this means that mild steel, the baseline material and the least expensive option, is always preferred. However, it was established above that this conclusion contradicts observations of firms' actions; therefore, there must be other reasons why a firm might adopt an alternative material that the selection framework, in its current state, cannot evaluate.

The use of a *test bed* is one potential reason why a firm might choose to adopt a new, more expensive, material before that material is necessitated by design constraints. By employing one of its products as a test bed, the firm can deliberately gain experience working with the new material and reduce that material's unit cost so that it is cheaper by the time the firm requires it to satisfy constraints. Test beds are a well-known industry practice: examples of use include military aircraft and Formula One cars [31, 32, 72].

Learning is essential, however, for a firm to even consider introducing an alternative material to its product line via a test bed: the firm has to believe that, at the very least, cost reduction due to learning will happen and that it will realize economic benefits in exchange for the early adoption of a new, costly material for one of its products. Without the conviction that cost will evolve, the firm will have no logical reason to use a test bed because doing so would mean implementing a more costly material without any foreseeable economic benefits. Regardless, the decision to use a test bed can lead a firm to adopt an alternative material, even when that material is not strictly necessary for its products.

4.1 Expanding Selection Framework Scope

4.1.1 Adding Multiple Periods

Any firm planning to use a test bed with the goal of gaining experience and thereby reducing a material's unit cost will first have to answer 1) whether future constraints will require the use of a specific alternative material on that firm's products and 2) whether that material is suitable for use on a test bed—that is, whether using it on a test bed will enable the firm to lower its overall manufacturing cost. The proposed materials selection framework is on the right track in that it incorporates learning into the decision-making process; however, it is unable to analyze whether a material is suitable for use on a test bed due to its restricted scope in which it makes a single selection decision that is limited to the choice of a single material for a single product. Consequently, it lacks the ability to make a selection *now* given some future constraint that necessitates the use of a specific alternative material, and analyze any potential benefits, which come in the form of reduced manufacturing cost, the firm may realize in the future.

The selection framework is therefore revised so that its scope encompasses a longer time horizon, which is divided into two (or more) periods. In the second period, the firm is forced to use a specific alternative material because that material's properties enable the firm to satisfy a design constraint that the baseline material cannot. In the first period, though, the firm confronts the choice of whether to continue using the baseline (current) material until the alternative material is required in the second period, or introduce an alternative material in that period. These options are detailed by the two scenarios presented in Figure 4-1.

If material costs do not evolve, the choice for the first period is trivial: the material with the lowest unit manufacturing cost is identified as the preferred option, regardless of which alterna-

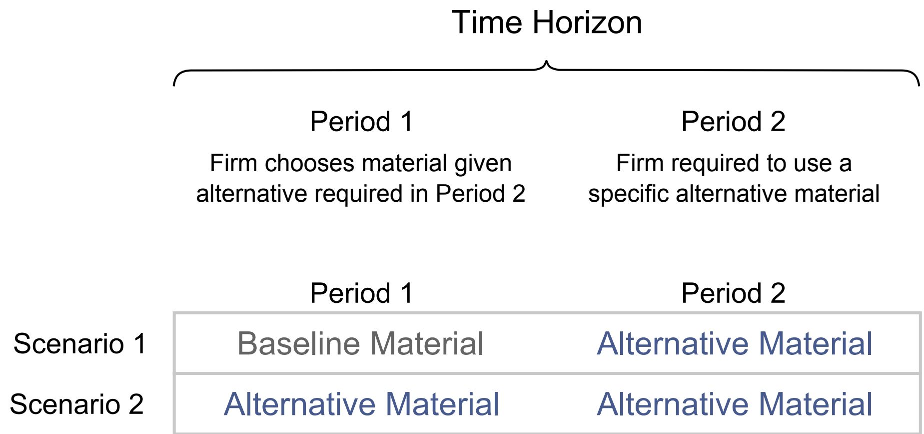

Figure 4-1: Two scenarios for a single-product selection problem with a two-period time horizon.

tive material is required in the second period. On the other hand, if unit manufacturing costs evolve due to learning, the preferred material is not as obvious. While the firm can still choose the material associated with the lowest *total* manufacturing cost (to account for cost evolution in the calculations) over the first period, it may not lead to the lowest-cost decision once the second period is taken into account. For instance, the firm may be able to further reduce total cost over the entire time frame by selecting the same alternative material required in the second period for use in the first period (Scenario *2*); in this way, the firm can gain experience working with that material and thus reduce its cost in the second period, at least relative to the amount it would have paid had it used the alternative material in only the second period (Scenario *1*). Whether or not using the same alternative material in both periods will reduce overall cost relative to choosing a different material in the first period will ultimately depend on the additional amount the firm pays in the first period relative to the savings it realizes in the second period.

Figure 4-2 presents a simple two-period, two-material illustration that compares the undiscounted costs of two possible scenarios. On the left, the firm is assumed to use its baseline material in the first period and an alternative material in the second (Scenario *1*); on the right, the alternative material is used in both periods (Scenario *2*). Consequently, in the right-hand scenario, the experience the firm gains in the first period enables it to lower the material's unit cost in the second period. Only the alternative material's unit manufacturing cost evolves; the baseline's does not because the firm is already familiar working with it. The alternative material's cost is also assumed to be higher than that of the baseline material—at least when the firm

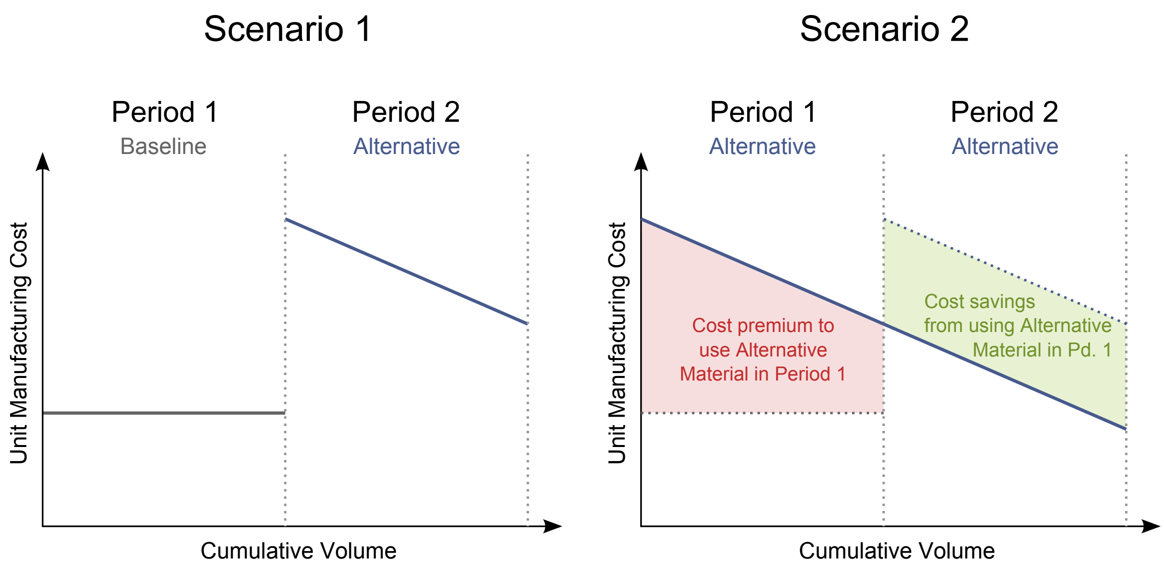

Figure 4-2: Scenario cost comparison for single-product, two-period selection problem.

first implements it. Additionally, the firm is assumed to produce the same number of products in each period.

The preferred scenario will depend on which is cheaper according to the total manufacturing cost of each. Total cost, in turn, is a function of the baseline material's unit manufacturing cost, as well as that of the alternative material and how fast the latter evolves with each additional unit produced. As with the learning curves in Chapter 3, total cost over the entire time horizon is represented by the area under the curves; therefore, the costs of the two scenarios can be compared by comparing the areas under their unit manufacturing cost curves (as shown in Figure 4-2).

If the discount rate is 0%, the firm will pay a cost premium, denoted by the light red region in Figure 4-2, to use the alternative material in the first period, before that material is absolutely necessary. Since this early adoption of the alternative material also leads to reduced unit cost in the second period, the firm will realize savings in that period, denoted by the light green region. The preference for either scenario will then depend on the relative area of these two shaded regions: if the green region has the larger area, Scenario *2* will be preferred because savings will exceed the cost premium associated with early adoption. On the other hand, if the red region has the larger area, Scenario *1* will be preferred and the firm will choose the baseline material for the first period.

Another way of looking at this is to compare the area under the baseline material's curve in the first period of Scenario *1* with the area under the alternative material's curve in the second period of Scenario *2*. The other two periods that use the alternative material can be ignored because their areas cancel out—assuming a 0% discount rate. If the area under the baseline material's curve is less than the area under the alternative material's curve (as is the case in Figure 4-2), the firm will prefer to use the baseline material in the first period; otherwise, the firm will prefer the alternative material.

Increasing the discount rate to a non-zero value will shift the firm's preference toward the use of the baseline material in the first period. A higher discount rate will devalue the second period's cash flow more than that of the first period and thus, shrink the area of the green region relative to the area of the red region. The firm will, therefore, have to realize even more savings with the alternative material—either through a faster learning rate, larger scope, or a combination of both—in order for Scenario *2* to be preferred over Scenario *1*.

Ultimately, which scenario the firm prefers will depend on a number of parameters, including the length of the time horizon, production volume in each period, cost of both materials, the functional form and parameters of the alternative material's learning curve, and so forth. Regardless, the results of this brief exercise indicate that including a second period, which requires the use of an alternative material, in the selection problem's scope can potentially lead the firm to adopt that alternative material in the first period, before that material is required by design constraints. Adopting the alternative material in the first period, however, is only preferred if doing so enables the firm to lower its total manufacturing cost over its time horizon. This, in turn, requires that the cost savings in the second period exceed the cost premium in the first period. A non-zero discount rate will only favor the use of the baseline material over the (more expensive) alternative material in the first period.

The exercise in Figure 4-2 illustrates that although it is possible for the early adoption of an alternative material to take place, the preference for the alternative material in the first period requires that the average unit manufacturing of that material in the second period of Scenario *2* be lower than the unit cost of the baseline material in the first period of Scenario *1*. While satisfying this condition might be feasible for some industries and material options, it, unfortunately, is not a realistic requirement for the material options in the automotive industry, especially if automakers produce vehicles at high annual production volumes in order to leverage economies of scale. At high volumes, the long-term unit costs of all alternative materials are above that of the mild steel baseline (see Figure 3-4). These alternative materials, consequently, will never be less expensive than the baseline and Scenario *1* in Figure 4-2 will always be favored. Therefore, the selection framework, even with a scope that has been expanded to include multiple periods, still cannot account for why an automaker or other manufacturing firm might adopt a new material that it is not required for use in the short-term.

4.1.2 Including Multiple Products

The conclusion of the preceding section indicates that something is still missing from the analysis since clearly, expanding problem scope time-wise is not sufficient for selection results to accurately portray a firm's actions. Even when two periods are considered in the selection process, a firm will still avoid adopting a new material—even for a test bed—unless that material will eventually be cheaper than the baseline option. The selection problem's scope is therefore expanded again, this time to include the multiple products, or applications within a single product, manufactured by a single firm.

Broadening the selection framework's scope to include multiple products is a more accurate representation of a firm's operations because the selection process will no longer assume any one product is designed and manufactured independently of a firm's other products. Most firms, anyway, have manufacturing portfolios that consist of several product lines, each line targeted at a different market segment or serving a distinct purpose. Thus, the selection scope is expanded from selecting materials for a single product over some time horizon, to selecting materials for multiple products over the same time horizon (Figure 4-3). As proposed in the preceding section, the time horizon is divided into periods so that one material is chosen for each product in any given period, with the potential for the preferred material for any one product to change between periods as material properties or problem context evolves. Preferred materials, though, are still identified according to manufacturing cost, only this time, the total cost of all products over the time horizon is included in the metric calculation.

In addition to providing a more accurate portrayal of a firm's operations, a multi-product scope is anticipated to be necessary for the materials selection framework to be able to compre-

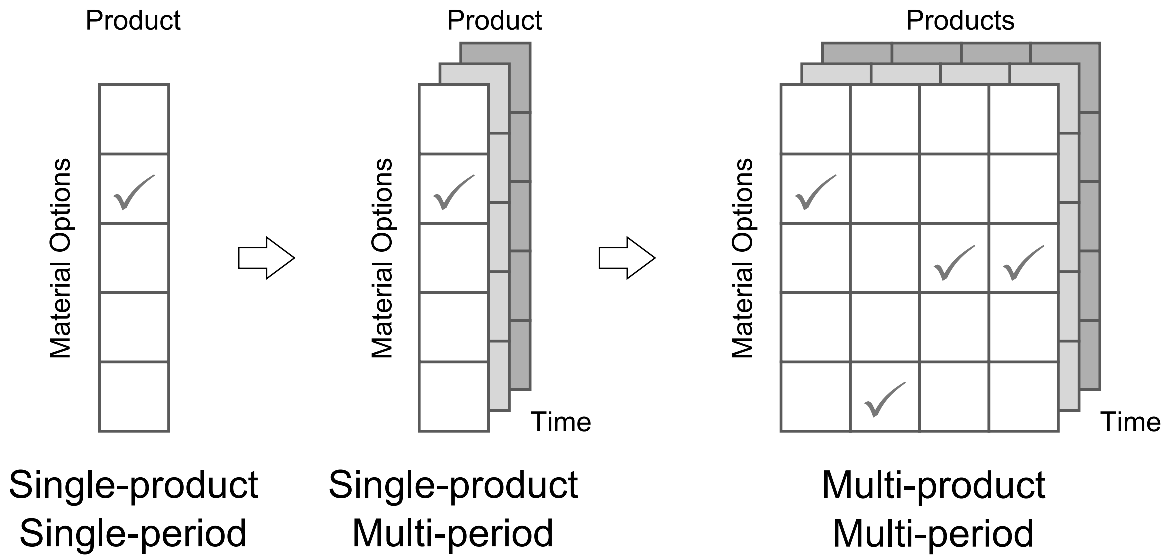

Figure 4-3: Contrast of single-product and multi-product selection method scope.

hend a firm's decision to use a test bed. In the beginning of this chapter, it was noted that the purpose of a test bed is so a firm can deliberately gain experience working with a new material and reduce that material's cost in preparation for future use. The experience the firm gains, however, should be applicable not only on the product that first implemented the new material, but also on the other products that eventually adopt the new material after the firm is more comfortable working with it. If these products benefit from the firm's experience with the initial product, they should be included in the selection framework in order to account for that benefit. By expanding the framework's scope to include these products, the firm can then fully assess the benefits derived from expanding the new material's use to other products, as well as which products should adopt the new material in the first place.

4.1.3 Additional Considerations

There are consequences, however, for expanding the scope of the materials selection problem. Unlike the single product analysis, which relies on manufacturing cost to rank material options, minimizing the total manufacturing cost of several products is not necessarily equivalent to maximizing total profit. Not only is each product likely to have a different production volume and profit margin, the firm can vary both these values (within reason), which in turn affects the returns on each product. Consequently, using total profit rather than total cost as the decision metric may make more sense for a multi-product problem. Doing so, however, is not without challenges: calculating profit requires knowledge of profit margins of the products, as well as price elasticity of demand to determine the extent demand—and therefore profit—will be affected upon changes to production volume or price. Consequently, the materials selection analyses in this study are based on cost minimization, although selection according to profit maximization is briefly addressed in Appendix A.

Secondly, learning is assumed to be *shared* among products, rather than to take place independently on each product in a firm's portfolio. Shared learning occurs when experience a firm gains while working with one product can be transferred to other similar products that share a common resource (e.g. a manufacturing process line for a specific material) as the first product [13, 22]. In such cases, all these products can contribute to the firm's cumulative experience—as measured by the cumulative number of units of product manufactured—with that resource and facilitate the firm's progress down the learning curve. The firm can therefore reduce not only the cost of the product it used to gain its initial experience, but also the costs of other products that share the common resource. For example, in the case of automotive manufacturing, shared learning can arise when products that are designed using the same material and share the same production line; as the automaker learns and improves line efficiency, all products benefit.

Learning has been observed in the automotive industry, in both manufacturing [43] and general assembly [56]. Shared learning has also been documented, specifically for the transfer of knowledge between shifts and over time [27]. More broadly, other studies have examined learning transfer: for example, between plants owned by a single firm [10], as well as between different firms [41].

Shared learning is also anticipated to be necessary for a firm to decide to use a test bed for the introduction of a new material. The exercise in Section 4.1.1 illustrated that if the product that implements the new material is the only one that benefits as the firm gains experience, the new material will be implemented early (i.e. in the first period) only if its average unit cost in the second period of Scenario *2* is lower than the unit cost of the baseline material. If the alternative material's unit cost, even in the long-term, is not expected to drop below that of the baseline material (as is the case for most of the alternative materials used in the automotive industry), the alternative material will never be adopted unless the firm is forced to do so. On the other hand, the ability to transfer knowledge between the test bed product and other products that share a production line will enable the firm to realize additional economic benefits that, in turn, will facilitate its ability to offset the cost premium of a test bed. Given this information, it is now conceivable that a firm will deliberately introduce a new, more expensive material in order to gain experience working with that material, even if the material is not necessitated by design constraints.

4.2 Next Steps

Although the single-product selection framework proposed in Chapter 3 accounts for cost evolution, it was shown to be incomplete because it is unable to account for any sharing of learning among products due to its limited scope. This was demonstrated as problematic because its conclusions did not match firm behavior: the case study's results suggested that firms will not use alternative materials unless the materials are cheaper than the baseline, or unless the firms are forced to do so. In contrast, the actions of manufacturing firms indicate that the firms introduce alternative materials even when such materials are not required to satisfy design constraints. Therefore, the scope of the selection framework was expanded to account for multiple periods and products in the selection process (Figure 4-3), as well as shared learning among products that use similar materials and manufacturing processes so that when a firm gains experience with one product, other products benefit as well.

Despite its challenges, expanding the scope of a selection problem enables a more accurate representation of a manufacturing firm's product design process and is likely to better inform the selection decision than a selection method that is limited to a single product. A method with a broader scope not only enables the evaluation of shared learning, but also provides the ability to identify and systematically analyze strategies a firm can adopt for introducing new materials to its products. *How* a firm should introduce new materials to its products cannot be decoupled from materials selection because factors like learning not only favor certain materials over others, but also favor strategies the firm can use to cost effectively introduce the new materials—such as early adoption as in the case of a test bed. In contrast, traditional selection methods do not address the introduction of the preferred materials to the firm's products, but instead simply assume that once the materials are identified, they are applied whenever the firm gets around to doing so. The traditional selection framework will therefore have to be substantially modified in order for it to accommodate the expanded scope and comprehend the associated consequences of the broader problem.

The following chapters present and demonstrate a multi-product, multi-period materials selection framework. This framework is designed to identify appropriate materials that enable the firm to satisfy its design constraints at minimum total manufacturing cost. Shared learning is incorporated into the manufacturing cost calculation, which accounts for the cost of each material selected for each product in each period included the selection problem's scope. The firm is permitted to switch materials between periods because both material properties and problem context can evolve throughout the time horizon and potentially alter the preferred materials.

This proposed selection framework is used to investigate whether the consideration of learning—specifically, shared learning—can impact the preferred materials for a firm's many products. On a more practical level, it is employed to identify appropriate materials for the firm, as well as strategies the firm can adopt for introducing those materials to its product lines and the contexts under which such strategies are financially beneficial to the firm. The use of a test bed is one example of an introduction strategy and the primary focus of this study. By deliberately choosing to introduce a new material on a test bed, a firm can gain experience with that material before it is required to satisfy constraints or before adopting it for other products. Test beds were discussed earlier in this chapter, but in the context of the need to expand selection problem scope to encompass multiple products and periods.

Now that the need for a broader scope has been established, the next logical question is which materials should the firm choose for introduction on a test bed and which products should be used as that test bed. The selection framework can be employed to perform this analysis and identify both the preferred material and the product, as well as the conditions under with the firm will financially benefit from using a test bed. Ideally, the product selected as a test bed will be one whose consumers are willing to pay for the added benefits of the new material, or will have a low price elasticity of demand so the firm can recoup the test bed's cost without losing sales. If it turns out the firm does not need to adopt new materials, the question of whether to use a test bed may be irrelevant. Either way, a firm will introduce a new material on a test bed only if it expects to realize sufficient economic benefits in exchange for the early adoption of that material. Since the benefits can be realized on both the original product used as a test bed and on other products that eventually adopt the "not-so-new" material, the consideration of multiple products, as well as shared learning, in the selection framework are thereby essential to systematically analyze the use of a test bed (and other strategies for introducing materials) and when that strategy will be worthwhile to a firm.

This chapter and the ones following it analyze a firm's decision to deliberately gain experience by introducing a new material on a test bed, as well as the more general question of whether the consideration of shared learning impacts a firm's preferred materials in a multi-product selection problem. A stylized test bed exercise is used to motivate first, the incorporation of shared learning into the materials selection process, and second, learning's role in a firm's decision to deliberately gain experience with a new material by introducing that material on a test bed. The stylized exercise shows that, under certain (controlled) circumstances, a firm can benefit financially from deliberate learning with a test bed. Chapter 5 builds upon this simpler analysis and presents a multi-product selection framework for the purposes of informing the materials selection decision for two or more of a firm's products, and targeting the more general question regarding the consideration of shared learning in the selection process. The proposed framework is designed to account for evolution in each product's manufacturing cost through shared learning and is demonstrated with two case studies, the first of which revisits the stylized exercise and the second, evaluates the materials selection decision of an automaker seeking to improve the fuel economy of its fleet through use of alternative lightweight materials. The second case study is also used to systematically investigate potential strategies, in addition to the use of a test bed, a firm could adopt when introducing new materials.

4.3 Stylized Test Bed Exercise

The use of a test bed is one of the many strategies available to a firm for introducing new materials to its products. Firms can use test beds to deliberately learn how to better work with the new material so that they can later apply that knowledge to other products that share a common resource, and hopefully realize benefits in both the initial and the additional products [22]. This section develops a stylized exercise to take a closer look at that strategy; in particular, to assess whether there is merit in investigating it before building a more formal model. More

Table 4.1: Scenarios for stylized test bed exercise.

	Scenario *1*		**Scenario *2***	
Year	Product *A*	Product *B*	Product *A*	Product *B*
1	BL	BL	ALT	BL
2	BL	BL	ALT	BL
3	BL	BL	ALT	BL
4	ALT	ALT	ALT	BL
5	ALT	ALT	ALT	BL
6	ALT	ALT	ALT	BL

specifically, it evaluates the financial consequences to the firm of using a test bed and identifies the conditions under which doing so will lead to a reduction in the firm's manufacturing cost.

To accomplish this, this exercise evaluates and compares two scenarios, each involving the selection of two materials—a baseline material and an alternative—for two products over a six-year time frame. Both scenarios assume that the alternative material is required—whether due to government regulations, consumer demand, or simply because the baseline material is no longer available—on both products in the final three years of the time frame because it enables necessary product capabilities the baseline material does not. The initial three years, however, differ in that the second scenario uses one of the products as a test bed for the alternative material, whereas the first scenario assumes both products are manufactured from the baseline material. These scenarios are detailed in Table 4.1.

The total manufacturing cost of both products is chosen as the metric for comparing the scenarios and for assessing the financial consequences of introducing the alternative material on a test bed (Scenario *2*) versus waiting until it is absolutely necessary in both products (Scenario *1*). Total costs are calculated under three different circumstances for the cost of the alternative material: short-term cost, long-term cost, and evolving cost. In the first two, costs are invariant over time so the firm pays either the short- or long-term unit cost for each product it manufactures with the alternative material. In the third, the firm applies whatever experience it gains to improving its process line and, over time, reduces the unit manufacturing costs of both products. Although firm profitability would be a better choice than manufacturing cost, calculating profit requires additional information such as price elasticity of demand for each product. Profit and its affect on a firm's decision to use a test bed is thus limited to a qualitative discussion for this stylized exercise, with Appendix A containing more detail on the incorporation of profit into the

Table 4.2: Inputs for the stylized test bed exercise.

Parameter	Variable	Value
Annual Production Volume		
Product A	V_A	20,000 units / year
Product B	V_B	300,000 units / year
Baseline Material		
Fixed Cost		\$10,000,000 / year
Variable Cost		\$300 / unit
Alternative Material		
Fixed Cost		\$6,000,000 / year
Variable Cost		\$375 / unit
Learning Scope	σ	20%
Threshold Volume	V_{th}	300,000 units
High Volume	V_{hi}	1,200,000 units
Other Inputs		
Discount Rate	r	8%
Time Frame		6 years

Table 4.3: Unit costs of products.

	Variable	Product A	Product B
Baseline Material	C_{BL}	\$800	\$333
Alternative Material	C_{ALT}		
Short-Term	C_i	\$844	\$494
Long-Term	C_f	\$675	\$395

selection framework. Finally, a sensitivity analysis is used to investigate the conditions under which the second scenario reduces the total manufacturing cost for the firm.

4.3.1 Inputs

The preference for either scenario is based on the total manufacturing cost of each, which, in turn, depends on the annual production volume of the products and the unit cost of using either material. Table 4.2 provides a summary of the inputs. In the analysis, the annual production volume of Product A (V_A) is 20,000 units per year, while the volume of Product B (V_B) is 300,000 units per year. Consumers are assumed to purchase all these units and are indifferent to the firm's choice of material.

To simplify the math, both products are assumed to have the same fixed and variable material costs, but because of their different annual production volumes, different unit costs. If the fixed

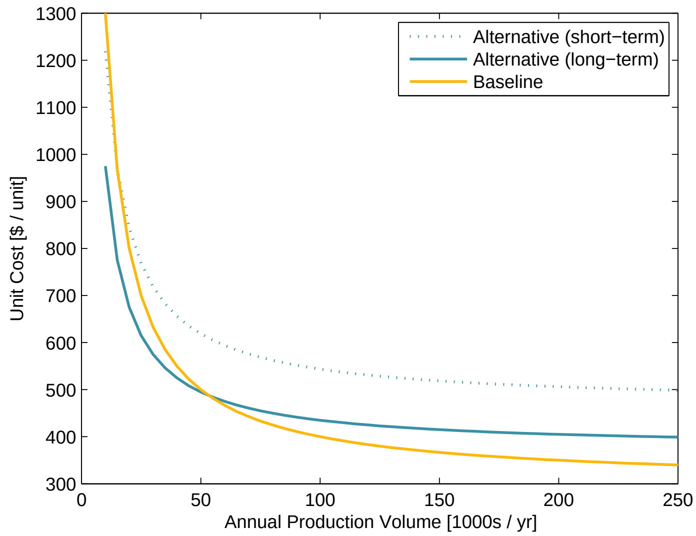

Figure 4-4: Unit manufacturing cost of the alternative and baseline materials as a function of annual production volume.

costs listed in Table 4.2 are paid every year, the unit cost for a product with annual production volume APV is

$$\text{Unit Cost} = \frac{\text{Fixed Cost}}{APV} + \text{Variable Cost}.$$

The unit costs for both materials are plotted as functions of annual production volume in Figure 4-4 and summarized in Table 4.3 for the volumes given in Table 4.2.

As with the process-based cost model results from the body-in-white case study (see Section 3.2.1), the alternative material's unit cost as calculated directly from the fixed and variable costs in Table 4.2 is assumed to represent the long-term cost (C_f) of using the material. A short-term unit cost (C_i) and a corresponding learning curve for the alternative material are constructed from the long-term cost plus the parameters in Table 4.2. No learning curve is required for the baseline material because its cost does not evolve. The short-term cost is related to the long-term cost through learning scope, $\sigma = 1 - C_f / C_i$. Higher values of σ imply greater learning and therefore a larger difference between short- and long-term costs. Between these two extremes, the unit cost of the alternative material follows a learning curve similar to the one in Figure 3-5, but with a linear rather than log-linear decay to simplify the math. Cost evolution starts at $V_{th} = 300,000$ units and ceases at $V_{hi} = 1,200,000$ units—the same values as those used by NHTSA in its 2011

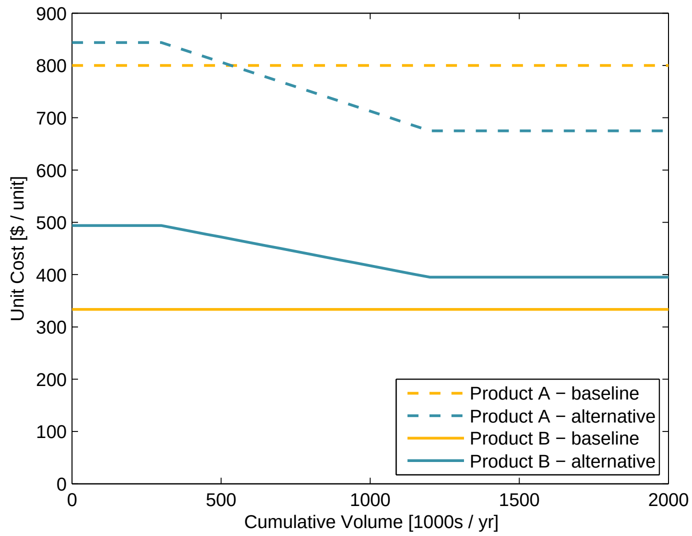

Figure 4-5: Contrast of learning curves for both products in stylized exercise.

CAFE target analysis [4]. The slope of the curve is then calculated from V_{th}, V_{hi}, and σ. The learning curve equation for the alternative material in either product has the form

$$C_{ALT}(V) = \begin{cases} C_i & V < V_{th} \\ \frac{-\sigma}{V_{hi}-V_{th}} V + \left(C_i - V_{th} \frac{-\sigma}{V_{hi}-V_{th}}\right) & V_{th} < V < V_{hi} \\ C_i \left(1 - \sigma\right) & V_{hi} < V \end{cases} \tag{4.1}$$

where V is the cumulative production volume of the alternative material over both products. As before, the total undiscounted cost of manufacturing any number of products using the alternative material is equal to the area under that material's learning curve (see Figure 3-1). Figure 4-5 contrasts the learning curves for both products given their respective production volumes.

The inputs from Tables 4.2 and 4.3 are used to calculate the total manufacturing cost of each scenario, which consists of the discounted cost of each product in each year. In these calculations, learning is assumed to be perfectly shared between products so that a unit of either product counts equally toward the cumulative production volume of the alternative material. Assuming the unit cost of the alternative material follows the learning curve exactly (i.e. the firm continually improves its process to stay on the curve), the undiscounted cost of using that material on either product over any given year is therefore the area under the learning curve

Table 4.4: Total manufacturing costs of each scenario in the stylized exercise.

Cost Evaluation Approach	Scenario *1*	Scenario *2*
Short-Term	**$687 M**	$690 M
Long-Term	$615 M	**$608 M**
Evolving	$670 M	**$669 M**

between the cumulative volume at the beginning of the year and that volume plus the product's annual production volume.

When both products use the alternative material in the same year, the products are assumed to be manufactured in sequence with the entire annual volume of Product *A* manufactured first, and the volume of Product *B* following. Consequently, by the time it manufactures Product *B* in a given year, the firm has additional experience obtained from manufacturing Product *A* in that year and is further down the learning curve. Product *B* therefore "sees" a higher cumulative volume that includes the production volume of Product *A*. The costs of both scenarios are calculated and the results compared to assess whether introducing the alternative material on a test bed enables a firm to reduce its total manufacturing cost. The following section presents the manufacturing costs of the two scenarios as computed using the above inputs, as well as a sensitivity analysis to identify the conditions under which deliberate learning financially benefits the firm.

4.3.2 Results and Sensitivity Analysis

The total manufacturing costs for both scenarios are shown in Table 4.4. These results indicate that the second scenario is favored by both the long-term and evolving manufacturing cost calculations. For the long-term cost calculation, this preference is driven by the lower cost of the alternative material relative to the baseline material on Product *A* ($675 versus $800) rather than by the firm's desire to use a test bed. To reach this cost point for the alternative material, however, the firm would have to somehow gain experience and learn. The short-term cost calculation provides a better point of reference because it represents the cost to the firm should the firm have chosen to not learn or improve the alternative material's manufacturing process. Under these conditions, there is no point to using a test bed because there is nothing to be gained. On the other hand, if the firm did choose to learn and to improve its manufacturing process, then the

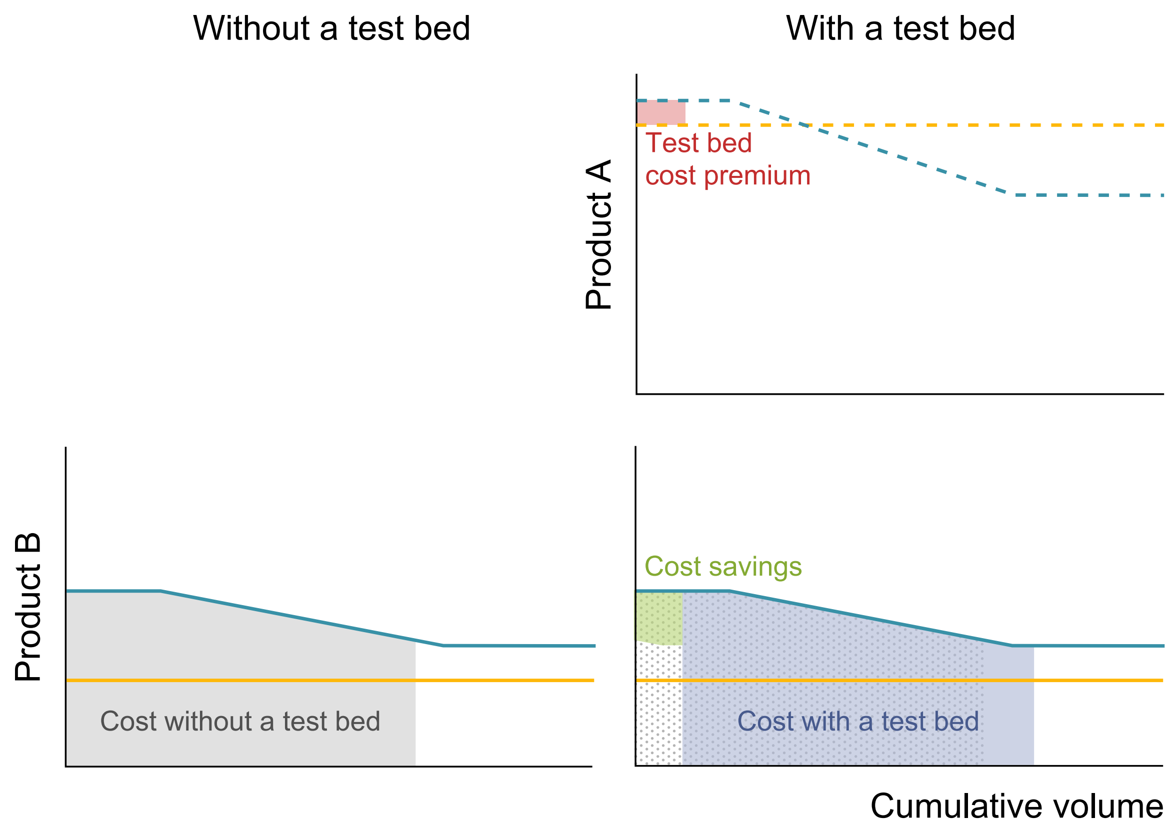

Figure 4-6: Scenario cost comparison depicted with learning curves.

use of a test bed would be preferable—as the results for the evolving cost calculation indicate. Since the cost of the second scenario is less than that of the first by approximately one million dollars, it can be concluded that under the above conditions, introducing the alternative material on Product *A* over the first three years is financially beneficial to the firm.

Before moving on to a sensitivity analysis, it is worth investigating the above situation in order to better understand why spending money upfront for a test bed leads to future savings once the alternative material is required on both products. When the alternative material's short-term cost is greater than the baseline material's cost, the firm will have to pay a cost premium to use a test bed. This cost premium is shown by the light red region in the graph at the top-right of Figure 4-6. Adding a test bed to the picture means that any products manufactured afterwards with the alternative material will be shifted to the right along the learning curve's x-axis, as illustrated in the bottom half of the figure. This shift reduces the area of the gray region on the left to that of the lavender region on the right, in effect by reducing the number of products manufactured at the short-term unit cost and increasing the number manufactured at the long-term unit cost. The difference between these two areas (green) represents the savings the firm realizes from using a test bed. If these savings are greater than the cost premium of the test bed—that is, if the green region has a larger area than the light red region—the firm can benefit financially from introducing the alternative material on a test bed.

Since total manufacturing cost is an emergent property and therefore context-dependent, a sensitivity analysis is used to better understand when changes to the context and to the inputs affect the firm's decision to use a test bed (or not). As illustrated in Figure 4-6, this decision will ultimately depend on how the upfront costs compare to the future savings in the firm's total manufacturing cost. The remainder of this section uses the stylized exercise framework to explore whether variations in 1) the annual production volumes of both products and 2) the timing of the constraint that necessitates the use of the alternative material in the fourth year impact the costs and savings associated with a test bed—and thereby the firm's decision to use one. All calculations are performed assuming costs evolve.

In the analysis, Product *A* is used as the test bed in the second scenario because, compared to Product *B*, it has lower upfront costs due to its lower annual production volume and lower unit cost premium. The latter is a direct result from the product's annual production volume and arises from the difference between the fixed and variable costs of the two materials. Therefore, changing the product's volume will impact the added cost of using it as a test bed; it will also influence the future savings the firm realizes by affecting the amount of experience the firm gains working with the alternative material.

Figure 4-7 plots both added test bed cost, from the first three years of the time frame, and future savings from the last three years as a function of the annual production volume of Product *A*. Net savings, the difference between the costs and savings associated with a test bed, is also plotted; this quantity is equivalent to the cost difference between the two scenarios so that negative values correspond to a preference for Scenario *1*, while positive values to Scenario *2*—the test bed scenario. Both costs and savings decrease with decreasing production volume of Product *A*; savings, though decreases at a slower rate than cost so net savings increase. Lower volumes of Product *A* are associated with lower cost premiums because the unit costs of the baseline and alternative materials are nearly equal (see Figure 4-4) and the manufacturer produces

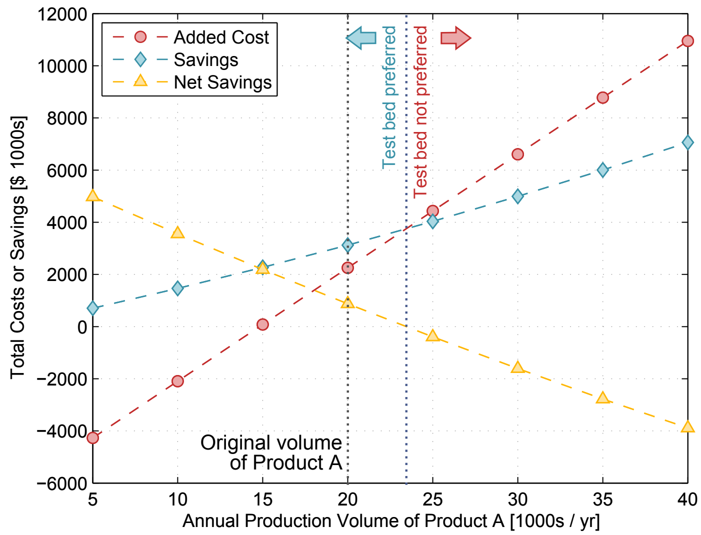

Figure 4-7: Added cost, savings, and net savings as functions of the annual production volume of Product *A*.

fewer products with the alternative material. These volumes, however, are also associated with lower savings because the manufacturer does not gain as much experience.

At very low volumes (below 15,000 units per year), the added cost of using a test bed drops below zero because the alternative material is cheaper than the baseline material, even before any cost evolution occurs; consequently, the firm should favor it whether or not it has plans to use a test bed. Increasing the test bed product's annual production volume, on the other hand, increases the short-term unit cost of the alternative material relative to that of the baseline material (short-term unit cost is 120% of long-term unit cost) and therefore cost premium. The savings increase as well, but at a slower rate. Eventually—just above a volume of 23,000 units per year—costs exceed savings so net savings drop below zero and Scenario *1* will be preferred over Scenario *2*.

Future savings also depend on the annual production volume of Product *B*, which accounts for the majority of the alternative material's total volume over the six-year time frame. These savings are shown in Figure 4-8, along with the added cost of a test bed and net savings. At low volumes of Product *B*, there are no savings because, although learning takes place, the firm does not gain enough experience—or manufacture enough units—to pass the threshold volume (V_{th}) and see any cost evolution. (In other words, the gray and lavender regions in Figure 4-6 both

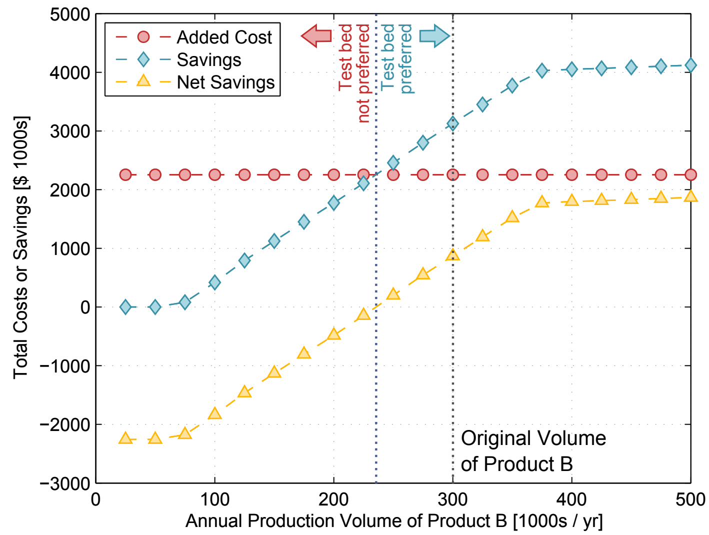

Figure 4-8: Added cost, savings, and net savings as functions of the annual production volume of Product *B*.

end before the first "knee" in the learning curve.) Consequently, the firm will prefer Scenario *1* over Scenario *2* because implementing a test bed will increase its manufacturing cost without providing any benefits. As the annual production volume of Product *B* increases, so too do savings because the firm is able to move further down the learning curve, which increases the extent to which unit cost evolves. Eventually, around a volume of 230,000 units per year, the savings are high enough to offset the added cost of a test bed and lead the firm to prefer the second scenario over the first. At high volumes of Product *B*, the savings plateau again because of the saturation of the learning curve above V_{hi}. The use of a test bed continues to be preferred, but the firm no longer realizes any additional benefit from increasing the volume of Product *B*.

Ultimately, a firm's preference for a test bed will depend not only on the production volume of the product used for the test bed, but also on the volumes of other products that adopt the alternative material. If the combined volume of these products is too low, there is no motivation for a firm to deliberately learn on a test bed because it will not be able to gain enough experience to offset the added cost of doing so. Likewise, at high volumes, using a test bed will be favorable, but the firm will not necessarily realize any additional benefits from increasing the number of Product *B* it produces (market considerations aside).

Exactly whether a production volume is "low" or "high" will depend on the learning curve

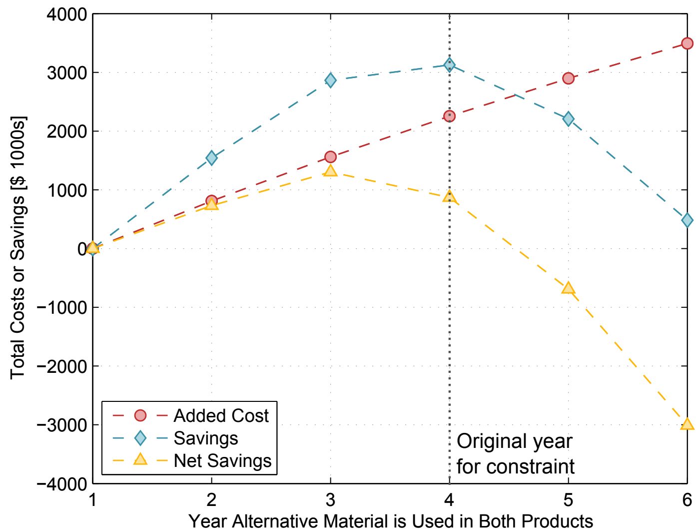

Figure 4-9: Added cost, savings, and net savings as functions of constraint timing.

and the values of its parameters. For instance, if instead V_{th} equaled 1 million units for the alternative material, an annual production volume of 300,000 units per year for Product B would be considered too low for a test bed to be favorable. Consequently, learning curve parameters can also influence a firm's decision to use a test bed. A firm will be reluctant to invest in a test bed if it believes learning will take place very slowly or if the annual production volume of Product B is low *compared to* V_{th} *and* V_{hi}. Under these conditions, the savings may be insufficient to offset the added cost, or the firm may not see any savings at all. Similarly, the firm will not use a test bed if it recognizes that the learning scope (σ) is very small so, again, savings would never be able to compensate for the added cost of using a test bed (e.g. if $C^B_{i,ALT} - C^B_{f,ALT} < C^A_{i,ALT} - C^A_{BL}$).

The second part of this sensitivity analysis explores the consequences of altering the year in which both products start use of the alternative material on the firm's decision to use a test bed. Both scenarios are adjusted accordingly to reflect this change in the constraint. The analysis is then simplified by assuming the firm always implements the test bed starting in the first year, regardless of when the constraint requiring the alternative material comes into play. Varying the year the constraint takes effect impacts the cumulative volume of the alternative material, as well as that of the test bed, and therefore the amount of experience the firm can gain before the constraint tightens.

The added cost of the test bed, future savings, and net savings (i.e. the cost difference between the two scenarios) are plotted as functions of constraint timing in Figure 4-9. If the constraint becomes a factor early on in the time frame, a firm has minimal time to learn with a test bed. However, the exercise indicates that the use of a test bed is still preferable because the total volume of Product *B* that is manufactured with the alternative material is high enough that the firm is able to move completely down the learning curve and offset the test bed's added cost. As the constraint is pushed further into the future, the number of years the test bed is used increases, as does its added cost. The savings, though, also continue to increase, despite the decreasing number of Product *B* that is manufactured with the alternative material, because the higher cumulative volume of the test bed shifts Product *B* further to the right along the x-axis of the learning curve (i.e. the green rectangle in Figure 4-6 becomes wider). This shift decreases the number of units of Product *B* manufactured at short-term cost and increases the number manufactured at long-term cost. Eventually, however, the total volume of Product *B* manufactured with the alternative material becomes too low for the firm to learn enough and compensate for the added cost. Under these conditions, introducing the alternative material on a test bed is no longer desirable and the firm will prefer to follow the first scenario.

The above conclusions are, of course, subject to the learning curve's parameters. Preference for introducing the alternative material on a test bed will ultimately depend on how much experience a firm can gain—as measured via cumulative production volume—and how fast its unit cost is expected to evolve with that experience. When the cumulative volume, by the end of the time frame, is too low *for the firm to see any cost evolution*, the firm will prefer to not use a test bed—which is exactly what happens when the constraint was pushed back to Year 5 or 6. If, instead, cost evolution were to happen extremely fast over relatively few units, the firm may still opt to use a test bed even if the constraint is far in the future.

4.3.3 Discussion

The stylized exercise above shows that, under certain conditions, introducing a new material on a test bed can lower a firm's total manufacturing cost over a given time frame. The test bed works by allowing the firm to deliberately gain experience with an unfamiliar material, and later, apply that knowledge to a larger number of products—hopefully at reduced cost. Only when the future savings are greater than the cost premium of using a test bed will the firm prefer to introduce a new material in this manner.

In this exercise, Product A was chosen as the test bed because of its low unit cost premium and low annual production volume, both of which enable the firm to minimize the added cost of using a test bed. Increasing Product A's annual production volume, as the sensitivity analysis shows, simply increases the test bed's cost premium, not only because the alternative material is more expensive than the baseline material at higher volumes, but also because the firm would have to use the alternative material on more units.

There are other strategies, however, not presented by the stylized exercise, that a firm can adopt to reduce the cost premium of using a test bed when low-volume production is not an option. For instance, the firm can choose a product with inelastic demand for its test bed. Since, by definition, the demand for these products is not responsive to changes in price, the firm can offset the cost of the test bed by raising product price with minimal losses in sales, but still realize benefits from cost evolution. Luxury cars such as the Corvette often serve this purpose for an automaker seeking to introduce a new materials or features to its fleet [72].

Even when product demand is elastic, a firm can still offset some or all of the added cost of the alternative material, particularly when the material improves one or more of a product's attributes, and the firm can find consumers who place a high value on those improvements and are willing to pay for them. The firm can then select as its test bed the product targeted at those consumers and recover its costs by increasing that product's price to capture consumer value. If the firm is smart about changing the price, it can ensure that the product's sales figures remain unaffected, with the higher price canceling any affect the improved attributes may have on product demand. This is the case for military aircraft, which serve as a test bed for composite materials because the consumer—the U.S. Air Force—is willing to pay a premium for the weight reduction the alternative technology enables [31]. While other market segments such as civilian airlines may also be willing to pay for weight reduction, they may not value it as highly so their respective aircraft are not appropriate test beds—at least from a financial standpoint. The concept of "volume-neutral price" and consumer willingness to pay is also used in Appendix A to predict incremental revenue in its presentation of profit-based materials selection.

Minimizing the cost premium of a test bed is not the only factor that decides whether a firm

Table 4.5: Impact of stylized exercise parameters on a firm's preference for either scenario in Table 4.1.

Volume of Product *A*

		Low	High
Volume of Product *B* (relative to V_{th})	Low	Scenario *1* preferred unless alternative material is cheaper than baseline material on Product *A*.	Scenario *1* preferred because firm does not gain enough experience.
	High	Scenario *2* preferred since savings in later years offset cost of using alternative material on Product *A*.	Scenario *1* preferred since savings are unable to compensate for high cost of test bed.

Learning Scope

Low	High
Scenario *1* preferred since unit cost does not evolve enough to offset added cost of test bed.	Scenario *2* preferred since unit cost evolves enough to compensate for added cost of test bed.

Scenario *1*: Uses alternative material only when necessary
Scenario *2*: Introduces alternative material on a test bed

will take this route: also important are the savings the firm can realize by gaining experience with a test bed and applying that experience to other products. The exercise results show that if the firm cannot gain enough experience, either from the test bed or from simply using the alternative material on its products, to see any evolution in the material's unit cost, it will not spend additional money on developing a test bed because there is no benefit associated with doing so. This leads to a firm's preference for the first scenario at low volumes of Product *B*. Likewise, if the firm does learn, but the cost evolves too slowly or is not anticipated to change enough to offset the test bed's cost, the firm will also prefer the first scenario. Table 4.5 summarizes these observations from the stylized exercise.

Ultimately, a firm's decision whether or not to use a test bed will depend on a number of factors, including input variables, operational and market conditions, government regulations, and so forth. The above exercise suggests what the firm should do in one particular case, and the sensitivity analyses indicate how the firm's decision might change if conditions were different. Either way, the stylized exercise serves its purpose by illustrating that the consideration of shared learning in the assessment of a firm's total manufacturing cost can lead the firm to use a test bed for introducing a new material to its products.

4.4 Materials Selection and the Stylized Exercise

While the stylized exercise shows that shared learning between products can lead to the use of a test bed, it does not explicitly select a material for the test bed, but instead chooses between two pre-defined scenarios that both force the use of the alternative material in later years. A typical firm's decision is, obviously, not as clear-cut. First, the firm has to recognize the need for a new material, whether because of consumer demand, government regulations, or simply because the current material will cease to be an option. It then has to identify which material to adopt for its products in light of design criteria and constraints. Identifying materials, especially when shared learning is involved, in turn requires a means to inform the materials selection decision across multiple products as actions taken for one product can affect decisions for others. Only once the preferred materials are identified can the firm determine whether any of them should be introduced on a test bed—or by another such strategy—and which product would best serve as that test bed.

A multi-product selection framework is therefore needed both to inform the materials selection decision and to analyze strategies a firm can adopt when introducing new materials. The method identifies appropriate materials for use on a firm's many products over a given time horizon and is designed so that the firm has the option to consider shared learning between products in the selection process. The following chapter describes the basic structure and implementation of such a method. Two case studies are then used to demonstrate its application, the first being the stylized exercise from this chapter and the second, the case of an automaker seeking to improve its fleet's fuel economy through use of alternative, lightweight materials. In both instances, the basic multi-product selection framework is tailored to account for additional assumptions and details of each particular case.

Chapter 5

Materials Selection for Multiple Products

Although more complicated, an extended scope of a traditional materials selection method is necessary if a firm wishes to consider property evolution in the selection process—particularly when those properties are affected by shared learning. An analysis at the single product level is unable to comprehend this because it necessarily assumes the product is manufactured independently of the firm's other products; the firm is thus unable to share any experience it has gained from working on that product with any others. This is shown in the preceding chapter with a stylized exercise, which illustrates that the sharing of experience among products is needed to achieve the real benefits of learning. Even though the exercise simply compared the manufacturing cost of two scenarios and did not explicitly consider materials selection, it nonetheless indicates there is a whole class of problems that an analysis limited to a single product is unable to explore—such as the use of a test bed to introduce new materials to a firm's products. This chapter goes beyond the stylized exercise and constructs a framework for incorporating shared learning into a multi-product materials selection method. The framework is then demonstrated with the stylized exercise from Chapter 4, as well as with a case study of an automaker seeking to improve the fuel economy of its fleet through use of alternative lightweight materials.

5.1 Materials Selection Framework

Assessing the impact of the consideration of learning on a firm's preferred materials first requires the development of a selection method that is capable of simultaneously identifying satisfactory materials for use on a number of products or applications. What constitutes "satisfactory" is defined by the design criteria—in this case, the total manufacturing cost of all products given the material selected for each one. Property evolution through learning can then be incorporated into the selection method. If cost is the only property that evolves, the new method can be used in a fashion similar to the traditional method presented in Section 3.1 and applied in Section 3.2: to compare the preferred materials based on short- and long-term (i.e. unlearned and learned) manufacturing costs to those selected by considering cost evolution due to shared learning over the firm's time horizon.

The explicit consideration of evolution in a multi-product manufacturing cost calculation, however, requires that the selection method include multiple periods, much like the stylized exercise. If the method is to be able to consider the use of test beds or other such strategies for introducing materials, it will need the option to alter the preferred materials for any product at various points throughout the firm's time horizon. A test bed, for instance, depends on the firm being able to implement a new material on one of its products, and only later switch to the new material in its other products, once it is comfortable working with that material. More generally, a firm may need to switch materials in response to evolution in material properties, design criteria, or other constraints. In other words, if the firm were limited to selecting one material per application over the entire time frame, it would not be able to adapt to changes in the problem's parameters or context. The selection method proposed in this chapter will therefore not only account for selection across multiple products or applications, but also for the evolution of the manufacturing cost associated with each material and, by extension, the firm's preferred materials over the time horizon. Thus, any proposed solution or *selection decision* made using this method will specify which materials are preferred for each product and when they are used within the problem's time frame.

Figure 5-1 presents a diagram of the multi-product, multi-period selection framework. An integer linear program (ILP) is first used to identify an appropriate material for each product given time-invariant material properties, but time-varying constraints. This optimization method, by

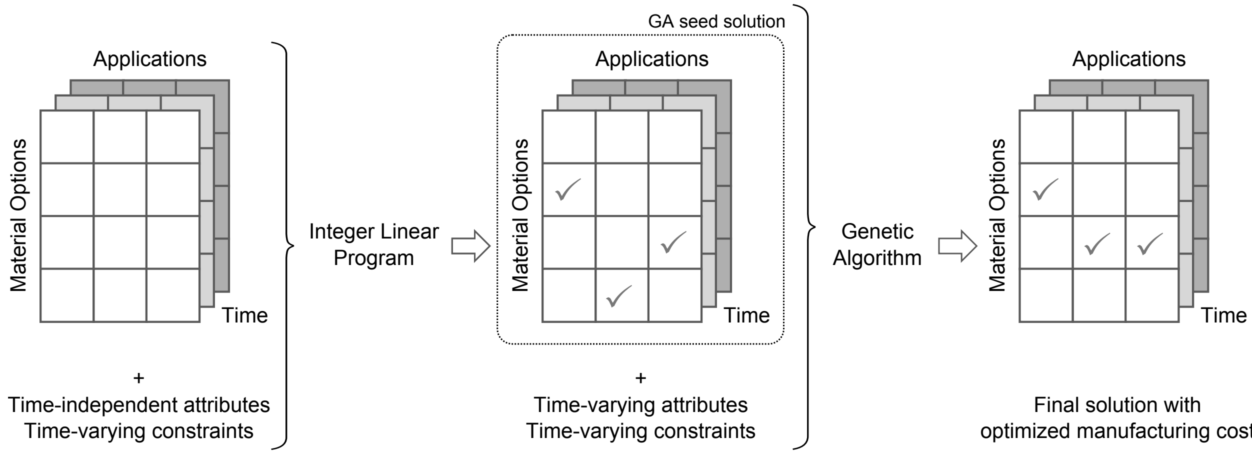

Figure 5-1: Diagram of multi-product materials selection method.

definition, requires that both the objective function and constraints be linear functions of the decision variables (i.e. whether or not a material is selected for a particular product); consequently, it cannot consider learning in the calculation of manufacturing cost because learning is a non-linear operation due to its use of cumulative volume to predict the extent to which costs have evolved. Instead, its resulting combination of materials and products is used to seed a genetic algorithm (GA), which evaluates the same selection problem using the same material, products, and constraint space, but without any linear restrictions. The GA is used to arrive at the final selection decision, one that accounts for learning in the selection process. A seed value, however, is used because a GA is not guaranteed to arrive at the *best* solution—merely an acceptable one; the seed from the ILP guarantees that the final solution will be at least as satisfactory as the one provided by the ILP.

The remainder of this section describes the selection method in greater detail for the general case of materials selection across a firm's products. Later, when the model is demonstrated with case studies, case-specific calculations and constraints are introduced.

5.1.1 Manufacturing Cost

For simplicity, the multi-product selection decision is made according to a single metric: total manufacturing cost of all the products over the firm's time horizon. Product-material combinations that have lower total manufacturing cost are deemed "better" than those incurring higher costs. While, again, it is more economically accurate to select materials based on firm profitability rather than production cost, working with the latter as the decision metric is the easier option as it does not require knowledge of either the price elasticity of demand or the profit margins of the products. Profitability, though, can be substituted for production cost in the model framework and is anticipated to make a difference in the selection decision, particularly when demand elasticity is a strong function of vehicle platform; this is illustrated in Appendix A.

Calculating total manufacturing cost is a relatively straightforward process when unit costs are time-invariant and do not evolve—as is the case for analyses that consider only short- or long-term manufacturing cost. Since unit costs do not change, the total manufacturing cost of any given combination of materials and products is simply the sum of the unit cost of each selected material, multiplied by the annual production volume of the associated product. The calculation, however, is more complicated when learning by doing is explicitly considered. Even if the decision's context (such as operational conditions or market conditions) does not change, manufacturing cost will evolve because the firm gains experience and, in doing so, reduces its expenses. Adding learning curves to the manufacturing cost calculation creates links among products that share manufacturing processes (or more generally, resources) in that each product contributes to a firm's experience working with its selected material. Experience is measured by cumulative production volume and is used to predict the extent of evolution in a firm's cost. The consideration of learning also has inter-temporal consequences because actions taken today can affect future costs and therefore future decisions made by the firm.

The calculation of the total manufacturing cost for an arbitrary selection decision (i.e. any given combination of products and materials) involves dividing the decision's time horizon into periods of equal length. The same period-based approach was used in the stylized exercise and is continued here because it simplifies the calculation process and limits the number of decisions the firm has to make over the time horizon. If cost evolves as the firm learns, the average unit cost of a material is approximated based on the cumulative volume of the manufacturing process used to form that material at the beginning of the current period and at the end of the period; otherwise, short-term or long-term unit costs are used. The resulting unit cost is then multiplied by the annual production volume of its associated product to arrive at the manufacturing cost of that material-product pair in the current period. These calculations are repeated for all the preferred materials in all periods and summed to generate the final value of the objective function. If necessary, material options are separated according to manufacturing process to accommodate the differences in cost structure and in learning rate and scope between the processes.

5.1.2 Combinatorial Optimization

So far, the methodology discussion has revolved around calculating the decision metric—total manufacturing cost given the preferred materials for a number of products over the firm's time horizon. While the methodology thus far can now perform the necessary calculations to evaluate any given selection decision, it still lacks a means to arrive at a *good* decision. For problems of limited scope with only one or two products and few material options for each, such as the stylized exercise, this additional step is unnecessary because it is possible to generate all feasible combinations of materials and products, and identify the one that minimizes total manufacturing cost. An enumeration approach, however, is impractical for problems with larger scopes, particularly when feasible combinations have to be identified for each period during the time horizon: the number of these combinations grows exponentially with each additional period, product, or material. Consequently, an alternative means to handle larger scale problems is necessary.

One possible way to address the issue is through the use of combinatorial optimization. These optimization algorithms are specifically designed to cope with problems which have a discrete, finite set of alternatives, but for which the alternatives may be too numerous for exhaustive search. As with most optimization approaches, they optimize an objective function subject to problem constraints. For the materials selection problem described above, the objective function is represented by manufacturing cost and the decision variables by the material options for each product. Each decision variable thus indicates whether a material is used on its associated product in a particular period.

Constraints

Often, optimization problems have constraints that restrict the values of the decision variables. The constraints for this particular problem place restrictions on which materials can be used to ensure that the resulting combination of products and materials is feasible. While the exact nature of the constraints will depend on the implementation of the method, the more common restrictions and how they relate to the problem at hand can still be discussed in the general sense.

The constraints used by the multi-product materials selection problem concern whether a material cannot, can, or should be used given the implementation of another material. In the simplest case, no constraint exists and the decision to use one material is completely independent of the decision to use another. This would be the situation for material decisions concerning two distinct products or applications: for example, an automaker's selection of aluminum for the body of a midsize car should have no bearing on its choice of material for the floor mats of a large SUV. In another case, the two materials are incompatible or "conflict" with each other so the firm is forced to choose between them because it cannot implement them both. An example of this is the selection of two or more materials for an application designed to only implement one. The outer panel of a car door, for instance, cannot be manufactured from both aluminum and high-strength steel: at most one material can be selected. A conflict constraint can also apply to the selection of materials for separate applications within the same product, particularly when there will be corrosion or other reliability problems if certain materials are selected for those applications. Finally, the two materials may "require" each other for implementation so that if one material is present, they both must be and vice versa. This can happen when two or more manufacturing processes are required to form a single material design (as described in Section 5.1.1), but each process is assigned a separate decision variable. The selection method will therefore need a means to ascertain that all pieces of that design are included in order to account for its full cost.

Optimization Methods

At this point, it is now possible to identify optimization methods that can be applied to solve the above materials selection problem. Potential approaches include

- Exhaustive search
- Linear programming
- Non-linear programming
- Dynamic programming
- Genetic algorithms
- Simulated annealing

These approaches vary in whether they are able to find the global optimum or simply a satisfactory solution. Exhaustive search, linear programming, and dynamic programming fall into the former category, but each is suited to different classes of problems. Exhaustive search can be applied to just about any optimization problem, but is impractical for large-scale problems because it considers every possible combination of decision variable values in its quest to find the optimal one. Linear programming, on the other hand, is applicable to large-scale problems, but is limited to those whose objective functions and constraints are linear functions of the decision variables. Unfortunately, this requirement excludes the material selection problem described above due to its inclusion of learning in the objective function calculation: learning is a non-linear operation because of its use of cumulative volume to predict the extent to which a firm's costs have evolved. The cumulative production volume of a common resource is a function of the decision variables of the applications (or products) that use that particular resource. Since this cumulative volume represents a firm's experience and thereby, how far down the learning curve the firm has progressed, it is required to predict the evolved unit costs of each of those applications that share the common resource. These unit costs are, in turn, multiplied by the production volumes of their respective applications and by decision variables, which indicate whether or not the application is used. If the learning curve is linear with respect to cumulative volume, this leads to an objective function (represented by manufacturing cost) that is a second degree polynomial and clearly not linear.

Dynamic programming, on the other hand, is capable of handling large-scale problems with non-linear objective functions and constraints; however, accounting for shared learning among applications is not feasible with this method. Dynamic programming copes with large-scale problems by breaking them down into independent sub-problems. Creating separable sub-problems in the presence of shared learning is not possible because of the system-wide interaction of materials options through shared manufacturing processes, and the use of cumulative volume to predict cost. For instance, without learning, each period within the time horizon could represent a sub-problem. The method could then identify the preferred material options in each sub-problem independently of all other sub-problems. With shared learning, however, the sub-problems—periods—are no longer independent because decisions made in earlier periods affect the cumulative volumes seen in later periods and thereby, the materials selected in those later periods.

The other methods, non-linear programming, genetic algorithms, and simulated annealing do not guarantee they will find the global optimum, but have fewer restrictions on the classes of problems they are suited for. Non-linear programming is like linear programming, but without the linear requirements and consequently can be more resource-intensive to solve. Genetic algorithms, a subclass of evolutionary algorithms, are a good general-purpose optimization tool. They do not exhaustively search the solution space, however, and therefore may only find a local optimum. Finally, simulated annealing is comparable genetic algorithms, but adopts a different approach to exploring the solution space and covers less ground in contrast to genetic algorithms. The preferred method—genetic algorithms or simulated annealing—will ultimately be problem-dependent as for different problems, simulated annealing can perform better or worse than a genetic algorithm [48, 73, 82].

5.1.3 Implementation

Ultimately, a combination of an integer linear program (ILP) and a genetic algorithm (GA) was chosen for implementing the multi-product, multi-period materials selection framework. The ILP is first used to optimize manufacturing cost in the absence of cost evolution from learning. An *integer* linear program, rather than a linear program, was picked because the problem formulation requires binary decision variables. The resulting selection decision from the ILP is then used as a seed value for the GA to provide the GA with a starting point for its optimization process. Manufacturing cost is then optimized using the GA which, unlike the ILP, considers cost evolution in the calculation of the objective function. The combination of materials and products suggested by the GA represents the final solution to the materials selection problem.

Seeding the GA ensures that its solution will be at least as good (in terms of objective function value) as that from the ILP. Without the seed, the GA randomly selects its starting points and therefore cannot be guaranteed to arrive at a better solution than the ILP. The GA can be run without a seed, however, and in some cases, this may be advantageous, especially if the consideration of learning has a significant impact on the preferred materials suggested by the GA (in contrast to those selected by the ILP). Using both models also facilitates the comparison of selection decisions made with and without considering cost evolution from learning to assess whether its consideration impacts the preferred materials.

Integer Linear Programming

The integer linear program was written with the LINGO software package. Code for the automotive case study (Chapter 6) can be found in Appendix B. Given a linear objective function and linear constraints, the ILP uses a branch-and-bound algorithm to find the decision variable values that optimize the objective function, total manufacturing cost.

Binary decision variables are used for the materials selection problem, with each variable indicating whether or not a material is implemented on its associated product in a particular period. For instance, if a firm is seeking to select materials for three products, each with four material options, over a time horizon divided into three periods, 36 ($= 3 \times 4 \times 3$) decision variables are required for a total of $2^{36} \approx 68 \times 10^9$ possible combinations of materials, products, and periods. Not all these combinations, however, are feasible: assuming each product can only accommodate one material at a time, only $(4^3)^3 \approx 262 \times 10^3$ will be permissible. Infeasible combinations are eliminated with the conflict and requirement constraints described above.

Genetic Algorithms

A materials selection method based solely on an ILP is unable to account for cost evolution in the objective function calculation. In order to overcome this problem, a genetic algorithm also evaluates the selection decision and assesses the impact of considering learning on the preferred materials. The ILP's solution is used to seed the GA if the latter is unable to identify a better solution—that is, one with a lower total manufacturing cost—than the ILP.

The GA was implemented using MATLAB's Global Optimization toolbox; Appendix C contains the code for the automotive case study. A typical genetic algorithm begins with a set or *population* of randomly generated candidate solutions, known as *chromosomes*, to the optimization problem. This population represents the first *generation* of chromosomes. The best or *fittest* solutions, as dictated by the objective or *fitness* function, are identified and others discarded. The remaining solutions become the *parents* and are combined via *crossover* functions and *mutated* to create a new generation of chromosomes. This process repeats until the solutions converge or a predetermined generation limit or time limit has been reached. Since only fitter chromosomes

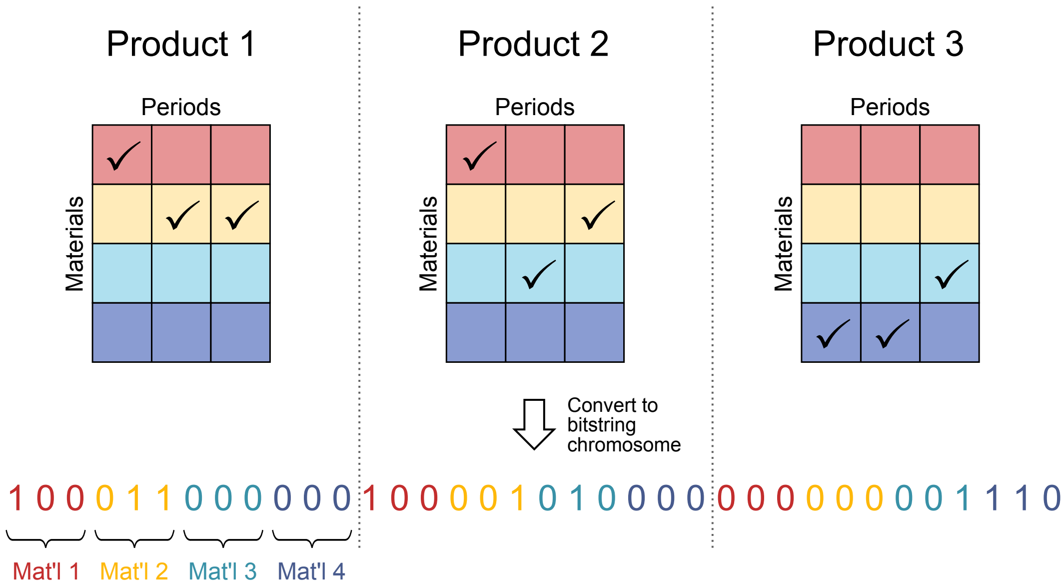

Figure 5-2: Example of a bitstring chromosome for a genetic algorithm.

survive and are used to create each new generation, the algorithm is able to improve the solution.

A genetic algorithm can be designed to either use the same binary decision variables as the ILP, resulting in a bitstring chromosome, or a chromosome that is customized for the problem at hand. Figure 5-2 contains an example of a bitstring chromosome for a firm seeking to choose between four material options for three products, over a three-period time horizon. Each bit or *gene* indicates whether or not a material is implemented on its associated product in a particular period.

Since the firm in the stylized exercise has only two material options—the baseline and an alternative—the corresponding case study with the multi-product selection method uses a bitstring chromosome. In contrast, a different chromosome is employed by the automotive case study to reduce its search space (see Section 6.2.5).

Aside from the seed value, the initial population for a GA is randomly generated; consequently, the initial generation of chromosomes is not guaranteed to satisfy all problem constraints. If the GA were then to eliminate all chromosomes that do not satisfy all constraints, it is highly probable that the resulting population would be an empty set. Therefore, instead of viewing constraints as binary—either they are satisfied or they are not—and eliminating chromosomes that do not meet all the criteria, the GA in this study keeps chromosomes even when they

Table 5.1: Genetic algorithm parameters used in automotive case study.

GA Parameter	Value
Population Size	8 × # of decision variables
Max Generations	100
`StallGenLimit`	50
`TolFun`	1×10^{-6}
Selection Function	Tournament
Crossover Fraction	60%
Crossover Function	Two-point crossover
Mutation Function	Uniform
Mutation Rate	1%

do not satisfy all constraints but takes this into account when calculating each chromosome's fitness. By treating the constraints as additional objective functions to be optimized (e.g. minimize the number of conflicting materials in the chromosome), the GA is able to preserve the genetic diversity of its candidate solutions and avoid eliminating its entire initial population.

Genetic algorithms, however, are best suited to solving single-objective problems, so some modification is required. One option is to assign each objective function a weight and combine all weighted objective function values into a single metric that is then optimized by the algorithm. Another approach, described by Roth et al. [63], is to compare a randomly chosen objective between two candidate solutions. The fitter solution (according to the objective) is kept and the other discarded to be later replaced in the next generation. The proposed selection model uses the former option because, of the two options, it is more adaptable to MATLAB's genetic algorithm framework.

Choosing appropriate values for algorithm parameters is important, as these parameters govern the GA's performance and its ability to optimize the objective function. Unsuitable values may lead to wasted computational effort or prevent the algorithm from finding a reasonable solution. Table 5.1 shows the parameters and their values used in the GA model for the automotive case study. A dynamic population size that changes according to the number of decision variables in the problem is used to accommodate different problem scales. The next three parameters in Table 5.1 set the algorithm's termination conditions: calculations stop when the GA either reaches the maximum number of generations allowed, or has run for a set number of generations with minimal change in the value of the objective function. The selection function defines the selection algorithm used to choose chromosomes that serve as parents to the next generation.

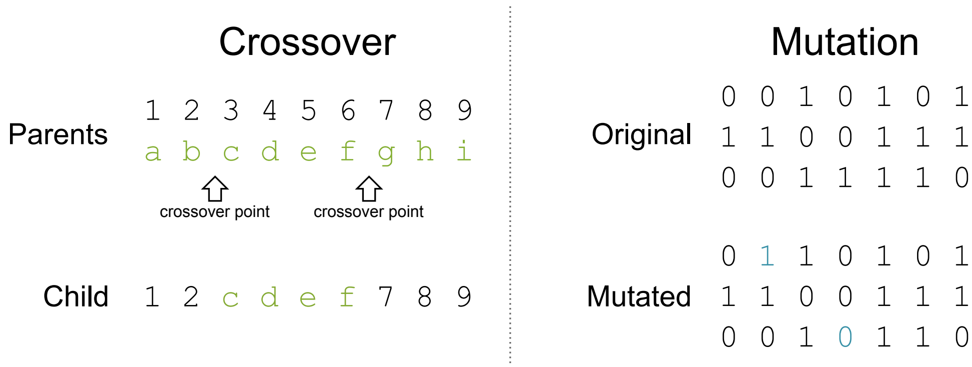

Figure 5-3: Examples of chromosome crossover and mutation.

Crossover fraction determines the fraction of the population that is replaced in each generation, and the related crossover function, the algorithm the GA uses to combine chromosomes and create the next generation. Two-point crossover is chosen for the case studies because it is one of the more common methods [54]. Finally, the mutation function completes the evolution to the next generation by selecting zero or more decision variables in each chromosome for mutation, with the mutation rate determining the probability each variable has of being mutated. Figure 5-3 shows crossover and mutation examples for a sample chromosome.

5.2 Case Study: Stylized Exercise

In the following section, the generalized materials selection framework is adapted to the stylized exercise from Chapter 4 with the goal to test the capability of the proposed computational approach by comparing its results to what is already known from stylized exercise. The same products, materials, time frame, and constraint space are used as before. Unlike the stylized exercise, though, the selection method has the ability to choose which material each product should implement in each year. Therefore, instead of defining scenarios like those from exercise (see Table 4.1), the method selects materials for each product by minimizing total manufacturing cost given the problem's constraints. These constraints are used to ensure that the alternative material is the only permissible option in later years—same assumption as the stylized exercise. Manufacturing cost is then minimized both with and without evolution in the alternative material's unit cost. Finally, the two resulting scenarios or *selection decisions* are compared to assess whether the consideration of learning leads to a different decision and potentially, to the use of a test bed for deliberate learning by the firm. While the difference in the resulting decisions alone should indicate whether a test bed makes financial sense for the firm, the total manufacturing cost for each decision can also be compared to confirm the initial conclusion.

5.2.1 Case Study Specifics

The case study uses the same inputs as in Table 4.2 to calculate the firm's total manufacturing cost. This cost is minimized first assuming unit costs do not evolve and second, assuming cost evolution due to learning takes place. When cost evolution is not a factor, either the short- or long-term unit cost (C_i or C_f) of the alternative material is used in the integer linear program to arrive at a selection decision. Introducing cost evolution into the selection process, however, is not possible with an integer linear program for reasons discussed previously, and therefore requires use of the genetic algorithm. The second selection approach is seeded with the result from the integer linear program to ensure it performs at least as well as the first (although in this case, it does not matter because of the small search space).

To simplify cost calculations, the genetic algorithm approximates the area under the alternative material's learning curve using the unit cost at three points: the cumulative volume at the beginning of the year, midway through the year, and at the end of the year. Shared learning inputs indicate which materials (in this case, the alternative material on Product *A* and on Product *B*) share manufacturing processes and thus, experience. As before, the time frame is divided into six years, each year representing a period in the model with the firm permitted to switch materials between periods. A constraint that places a minimum on average product mass is also introduced during the final few years of the time frame. This constraint is formulated to favor the alternative material, whose design is assumed to have a lower mass than that of the baseline material. Since the stylized exercise involves only two materials—a baseline material and an alternative—there is no need for the conflict or required constraints discussed in Section 5.1.2. Neither material is divided by manufacturing process and the optimization models are designed to default to the baseline option when an alternative material is not selected. If there were two or more alternative materials however, a conflict constraint would be necessary.

Table 5.2: Selection decisions identified by the selection method for the stylized exercise.

	Decision 1		**Decision 2**	
Year	Product A	Product B	Product A	Product B
1	BL	BL	ALT	BL
2	BL	BL	ALT	BL
3	BL	BL	ALT	BL
4	ALT	ALT	ALT	BL
5	ALT	ALT	ALT	BL
6	ALT	ALT	ALT	BL

Table 5.3: Total manufacturing costs of each decision in the stylized case study.

Cost Evaluation Approach	Decision 1	Decision 2
Short-Term	**$687 M**	$690 M
Long-Term	$615 M	**$608 M**
Evolving	$670 M	**$669 M**

5.2.2 Results

Figure 5.2 shows the preferred materials based on minimizing the firm's total manufacturing cost over a six-year time horizon. These selection decisions are the same as the scenarios defined for the stylized exercise in Table 4.1. The decision on the left of Figure 5.2 results from considering only short-term manufacturing costs—without any cost evolution from learning—in the selection process. Under these assumptions, the alternative material is the more expensive option so the firm prefers to use the baseline material in both products for as long as possible, and switches to the alternative material only when it is absolutely necessary in the fourth year. If, instead, the cost is permitted to evolve, the result is the selection decision on the right of Table 5.2, in which the firm uses a test bed to deliberately gain experience with the alternative material. Removing the test bed from the picture increases total manufacturing cost (Table 5.3) despite the presence of learning in both calculations. This, in turn, indicates that using Product A as a test bed for the introduction of the alternative material reduces overall cost.

The right-hand selection decision is also the result of considering only long-term manufacturing costs without any cost evolution in the selection process. In this case, the alternative material is always favored on Product A because its long-term unit cost is lower than the unit cost of the baseline material. Table 5.3 shows results for the total manufacturing costs of both selection decisions in Table 5.2 using all three cost evaluation approaches. The preferred selection deci-

Table 5.4: Test bed start date given year constraint becomes binding.

Year constraint is binding	Year to start test bed
1	N/A
2	1
3	1
4	1
5	never
6	never

sion in each row is highlighted in bold, with the bottom row confirming the stylized exercise's conclusion that under certain conditions, introducing a new material on a test bed can lead to reductions in overall cost.

The selection model can also replicate the results of the sensitivity analysis of the stylized exercise—in particular, the timing of the constraint and whether it affects a firm's decision to use a test bed. The earlier exercise concluded that a firm will not use a test bed if the constraint is sufficiently delayed because the cost of doing so outweighs any savings the firm may realize from the extra learning. The analysis, however, contained the assumption that the firm always started using a test bed in Year 1 in order to avoid evaluating all possible combinations of when to start with the test bed given the timing of the constraint. A similar assumption is unnecessary for the current selection model because it has the freedom to decide whether or not to use the alternative material in the earlier years before the constraint becomes binding. The model can therefore be run to assess whether it ever makes sense to start the test bed after the first year.

Table 5.4 shows the results of this analysis, which, as it turns out, concludes that the firm should use the test bed starting in Year 1 or not at all. This is because when the cumulative volume of the alternative material by the end of the time frame is too low (i.e. cumulative volume $\sim V_{th}$), the firm does not realize enough savings to offset the test bed's cost. On the other hand, when a large volume of products is manufactured with the alternative material (i.e. cumulative volume $\sim V_{hi}$), each year the firm uses Product A as a test bed allows it to shift the units of Product B to the right along the x-axis of the alternative material's learning curve. This in turn increases the number of Product B manufactured at long-term unit cost and decreases the number manufactured at short-term unit cost. Therefore, as long as the savings—approximated by the difference in short- and long-term unit costs of the alternative material on Product B—outweigh the added cost of using a test bed—defined by the difference between the short-term unit cost of the alternative material and the unit cost of the baseline material on Product A—the firm will have an incentive to use as many years of the test bed as possible. This, of course, is subject to discount rates and the problem's time horizon: if, for instance, the constraint is several years into the future, the firm may not want to begin immediate use of a test bed.

5.2.3 Conclusions

The results of the above case study illustrate that the proposed materials selection method is able to account for cost evolution through learning in its evaluation of a firm's total manufacturing cost, and thus conclude that under certain circumstances, it is financially beneficial for the firm to use a test bed for the introduction of a new material. This is the same conclusion reached by the earlier stylized exercise and therefore demonstrates that the formalized framework is capable of performing the same analysis. The materials selection method, though, is more flexible and able to choose among different material options—in contrast to the stylized exercise which simply compares pre-defined scenarios. Consequently, it expands the potential analyses beyond what the stylized exercise is capable of. The next section presents a larger multi-product case study concerning the materials selection in the automotive industry.

Chapter 6

Case Study: Selection in the Automotive Industry

The goal of the preceding chapter was to apply the multi-product selection method to the stylized exercise and demonstrate that the selection method was able to arrive at the same conclusion as when the exercise was first presented in Chapter 4. The stylized exercise case study was necessary to confirm that the proposed selection framework behaves as expected before presenting a larger and more complex case study—one that is not practical to analyze by hand, as was done for the stylized exercise. The new case study, presented in this chapter, is used for further exploration of the consequences of considering cost evolution due to learning on the selection decision and of potential strategies a firm can adopt when introducing new materials. This case study, as with the single-product study, concerns an automaker seeking to improve fuel economy by reducing vehicle weight via the use of alternative, lightweight materials—only this time, the focus is on multiple applications of materials within the automaker's fleet rather than just on the body-in-white of a single car. Multiple vehicle platforms are included in the scope, which also considers shared learning and thereby the ability to evaluate an automaker's use of test beds to deliberately gain experience with new materials.

Before the case study analysis can be performed, however, the proposed selection method has to be tailored to this particular case with relevant assumptions and constraints. Once this is completed, case study inputs, along with selection results and a sensitivity analysis are presented. The sensitivity analysis investigates the conditions necessary for cost evolution to have an impact

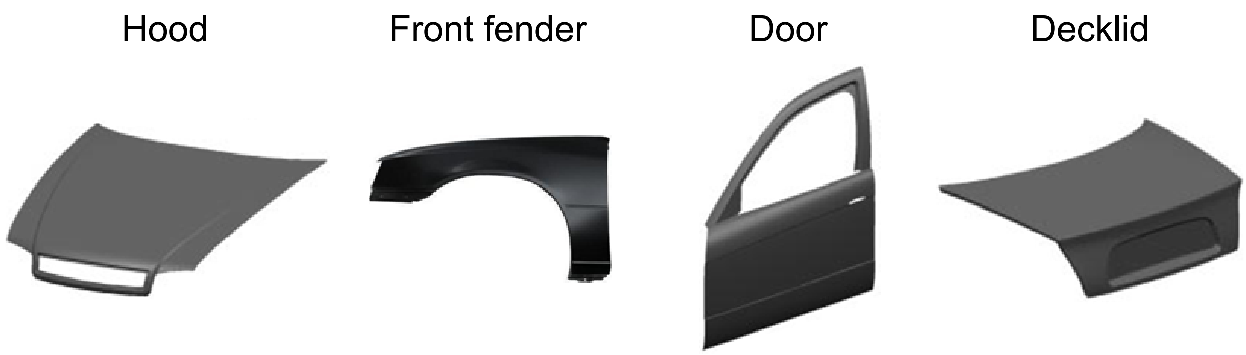

Figure 6-1: Vehicle closures [76].

on the selection decision and when using a test bed makes fiscal sense for the automaker.

6.1 Scope

In this case study, an automaker is seeking to select alternative materials for the body-in-white (Figure 3-2) and the closure (Figure 6-1) subsystems in three vehicle platforms within its fleet over the course of a 16-year time frame. The vehicle platforms include a compact, a midsize, and a large car. Four alternative materials are available for each subsystem, plus a baseline material which the automaker currently uses. In all cases, the baseline is mild steel. The alternative materials include high-strength steel, aluminum, glass fiber composite, and carbon fiber composite for the body—same options as in the earlier case study (Section 3.2)—and high-strength steel, aluminum, magnesium-aluminum, and sheet molding compound (SMC) for the closures. Figure 6-2 presents the material choices facing the automaker for a single period within the time frame.

6.2 Case Study Specifics

The multi-product materials selection method presented in Chapter 5 is for a generalized problem. The following section details some of the modifications and assumptions made to tailor the generic method to one for materials selection in the automotive industry.

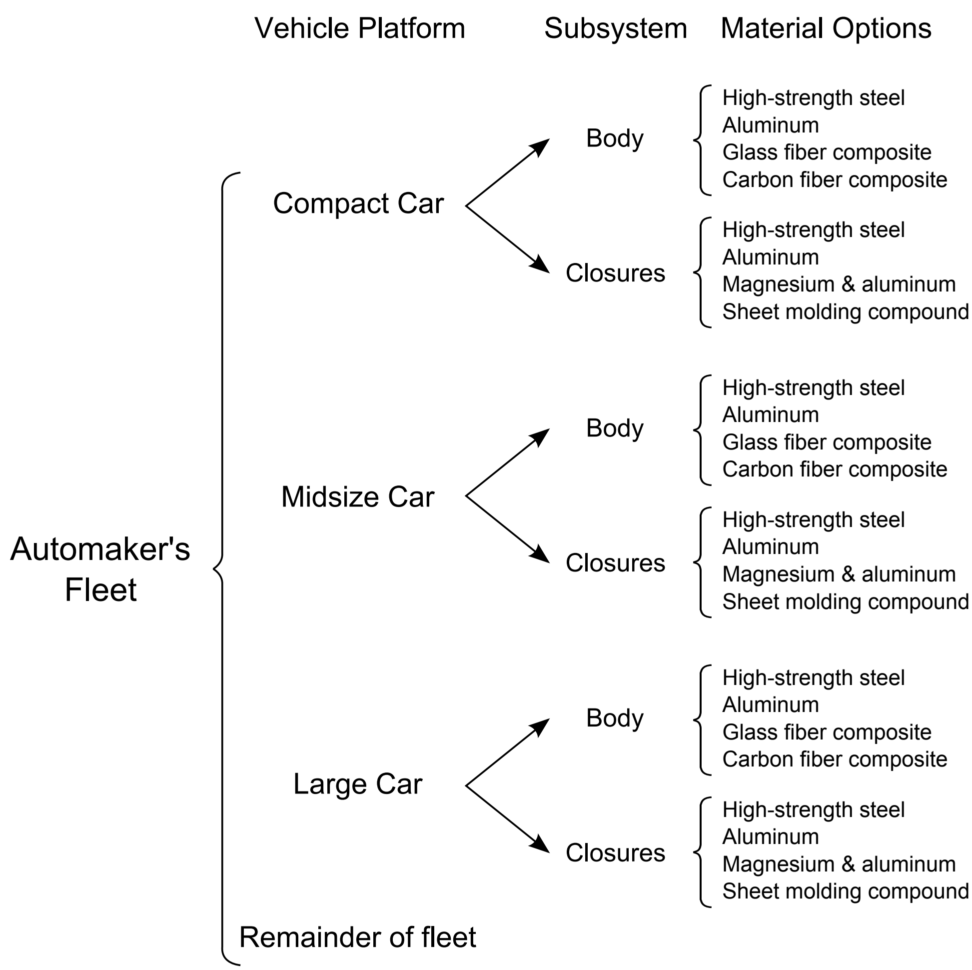

Figure 6-2: Vehicle platforms, subsystems, and material options included in case study.

6.2.1 CAFE Constraint

The first of these modifications is a constraint on the decision variables. This new constraint represents CAFE or the Corporate Average Fuel Economy regulation, which sets a minimum target for the average fuel economy of all new vehicles sold by a manufacturer within a model year. When a lighter weight material replaces the current material of a vehicle's subsystem, it changes that vehicle's mass and thereby its fuel economy as well as the automaker's CAFE number. This CAFE number is calculated using a sales-weighted harmonic average following the equation

$$\text{CAFE Target}_y \leq \frac{\sum_v vol_{v,y}}{\sum_v vol_{v,y}/FE_{v,y}} \tag{6.1}$$

where $vol_{v,y}$ is the sales volume of vehicle v in model year y and $FE_{v,y}$ is the fuel economy of that vehicle. In this case study, the automaker has to select materials that enable it to improve the fuel economy of its vehicles and satisfy the CAFE target. Automakers that fail to meet CAFE are penalized. The regulation, however, does permit some flexibility in the calculation of CAFE numbers, primarily by granting surplus "credits" to automakers that exceed the CAFE target in a given model year. The automakers can then apply these credits to any of the three years preceding or following that year of surplus to make up for any past or future deficits. Alternative fuel vehicles are also accounted for in CAFE calculations, but there is a limit to how much they can contribute to an automaker's CAFE number [8].

Imposing a CAFE constraint not only makes the case study more representative of the conditions an automaker faces, but also forces the materials selection model to consider more than just the lowest-cost options. Over the next decade, the CAFE target is scheduled to increase, as required by the Energy Independence and Securities Act of 2007 [70], and affirmed by President Barack Obama in his 2009 announcement for a national fuel economy and greenhouse gas standard [39]. The standard has also been reformed to account for the size distribution within an automaker's fleet. NHTSA has set the CAFE targets through the year 2016 [9] (Figure 6-3); the rapid change of the standard, after nearly two decades of holding steady, will force traditionally conservative automakers to evaluate and adopt alternative materials, as well as other technologies (e.g. hybrid powertrains or continuously variable transmissions), to increase their vehicles' fuel economies and avoid penalties. Automakers seeking to improve their fleet's fuel economy

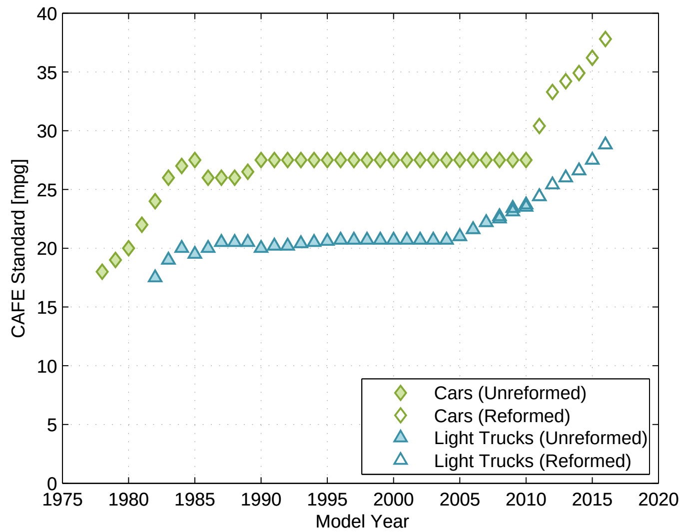

Figure 6-3: Car and light-duty truck CAFE targets over the years [3, 6, 9].

will likely need a systematic means to identify the best technology and material options.

Calculating the Automaker's CAFE Number

A few assumptions are made to simplify the CAFE calculation for this case study. First, the automaker is required satisfy the target in each model year: surplus or deficit credits are not factored into the calculations. As a result, penalties are not imposed because the constraint implies that the automaker *will* meet CAFE. Calculations also assume the unreformed version of CAFE—so without the size-based metric. Second, alternative fuel vehicles are not included in the problem. All the vehicles in this case study use a standard internal combustion engine. Finally, the automaker is assumed to sell all the vehicles it manufacturers in a model year so that the sales volume of each vehicle is equal to its production volume. This is a realistic assumption for an automaker because its production plans are often set years in advance and are difficult, as well as costly, to alter. Consequently, the automaker will often prefer to follow its plans and do what it can to sell its vehicles.

Even with these simplifications to CAFE, Equation (6.1) cannot be used as presented to calculate the automaker's CAFE number in a linear program because it is non-linear with respect to vehicle fuel economy and thereby the decision variables. Instead, the constraint equation has to be rewritten in terms of fuel consumption—the reciprocal of fuel economy—which describes the amount of fuel a vehicle requires to travel a fixed distance. Fuel consumption is often expressed in gallons per mile or in liters per 100 kilometers; improvements correspond to a *reduction* in fuel consumption so that less fuel is used per unit distance. In fuel consumption space, Equation (6.1) becomes

$$
\begin{aligned}
\text{CAFC Target}_y &= \left(\text{CAFE Target}_y\right)^{-1} \\
&\geq \frac{\sum_v vol_{v,y} \cdot FC_{v,y}}{\sum_v vol_{v,y}}
\end{aligned}
\tag{6.2}
$$

where $FC_{v,y}$ is the fuel consumption of each vehicle and equal to $FE_{v,y}^{-1}$.

The initial fuel consumption of each vehicle is calculated from the inverse of its fuel economy. Improvements in fuel consumption, like those in fuel economy, depend on both the weight savings associated with using an alternative, lightweight material, as well as how a vehicle's fuel consumption responds to changes in vehicle weight. This latter parameter is often expressed as a linear relationship between the percent improvement in fuel consumption and the percent reduction in vehicle weight. Typically, the improvement in fuel consumption is normalized to its change per 10% reduction in vehicle weight—for example, a vehicle may experience a 5% improvement in fuel consumption for every 10% reduction in weight. The exact numbers will be sensitive to the specifics of the vehicle's powertrain and its driving cycle [60, 78]. The weight savings associated with an alternative material is also design-specific and depends on the new material as well as the design of the subsystem using it. A vehicle's final fuel economy after material substitution is therefore

$$
\begin{aligned}
FC_v &= FC_i + \Delta FC \\
&= FC_i + \frac{\eta}{0.10}\frac{\Delta M}{M_i} FC_i
\end{aligned}
\tag{6.3}
$$

FC_i and M_i are, respectively, the initial fuel consumption and mass of the vehicle, and η, the percent improvement in fuel consumption per 10% reduction in vehicle weight. ΔM, the change in vehicle mass, is the only variable affected by the different alternative materials and represents the sum of the weight savings from each subsystem or application that implements a new material:

$$\Delta M = \sum_{s} \Delta m_s$$

where Δm_s is the weight savings attributed to the use of a new material on subsystem or application s. Once the total weight savings for a vehicle is known, the improvement in that vehicle's fuel consumption can be calculated. This calculation, in turn, is performed for each vehicle and the results used to calculate the automaker's new CAFE—or "CAFC"—number.

6.2.2 Time Frame Periods

In addition to the CAFE constraint, the automotive case study makes assumptions regarding the selection model's time horizon. As was previously noted, a defined time frame is necessary in order to assess the impact of learning. This time frame is divided into periods to permit the firm to switch materials and potentially expand use of a new material beyond a test bed. Accordingly, the case study's 16-year time frame is divided into four periods of four years each—meaning that the automaker effectively chooses materials for its vehicle subsystems four times. A four-year design cycle was chosen not only because it reduces the problem scale from having to select a new material every year, but also because it represents the approximate frequency at which automakers redesign their vehicles and implement substantial changes such as new materials for the body or closures [5][1]. Four years, however, is a long time between each recalculation of learned manufacturing unit cost. Therefore, this cost is updated yearly rather than every design cycle (a.k.a. period).

To simplify the implementation of the selection model, the design cycles of the three vehicle platforms take place simultaneously so all three platforms are up for redesign during the same model year. The CAFE target, likewise, changes only when the automaker modifies its vehicles—that is, at the beginning of each design cycle. While this scenario is not wholly realistic given an automaker's limited vehicle development resources and the rapidly changing CAFE target in the coming years, it is nonetheless sufficient for assessing the impact of considering learning on an automaker's material preferences.

[1]A four-design cycle falls somewhere between the 2 to 3-year average refresh cycle and the 5-year average design cycle used in industry.

6.2.3 Cost Minimization Versus Profit Maximization

The general selection method presented in Section 5.1 defines a metric, total manufacturing cost, for comparing different combinations of materials and products and for determining which combinations are better. Although this case study relies on the same metric, it should be noted that, in multi-product problems, the cost-minimizing solution is not necessarily the same as the profit-maximizing solution—and in general, one would expect a firm to maximize profit rather than minimize cost. The difference between these two objectives arises from a firm's ability to alter the profit margins as well as production volumes of its products in order to increase its revenues. Introducing these levers, however, makes profit maximization a more difficult problem because additional factors such as the products' price elasticity of demand have to be taken into account.

Rather than identifying these price elasticities and profit margins necessary to maximize the total profit of all vehicle platforms in an automaker's fleet, this case study instead simply minimizes the total cost when selecting materials for each platform in the fleet. The two approaches, however, are expected to lead to the same selection decision—at least in this particular case—because of the CAFE constraint which, as it becomes more binding, necessitates that certain materials be used to improve fuel economy, regardless of cost or profit. The production volume lever can be removed to further simplify the problem because changes in volume are not only difficult and costly for the manufacturer, but will also affect the automaker's CAFE number and lead to an even more complex optimization problem. Additionally, the automaker can set the price of its products to ensure that consumers purchase the same number of vehicles with the new material as with the original material. This will alter profit margins, but likely not enough to affect the materials selection decision under a CAFE constraint (see Appendix A).

An approach to calculating total manufacturing cost is presented in the next section. This metric consists of the cost of using the selected material on each product in each period within the model's time frame. The calculation itself is somewhat more complex than the one used by the stylized exercise because it tries to accurately represent an automaker's cash flow.

6.2.4 Manufacturing Cost Decomposition

Manufacturing costs are limited to those of the subsystems that are the focus of the automaker's materials selection decision. The remainder of the vehicle is assumed to be the same regardless of the materials chosen for the subsystems; therefore, its cost is not included in the calculations. The manufacturing costs of the material options available in this case study are decomposed by manufacturing process and again by cost category, as shown in Figure 6-4 for the aluminum and the glass fiber composite bodies. At the highest level, cost is broken down according to the different manufacturing processes required to produce the body (or closures) from the selected material. This initial division is adopted so the selection model can more accurately account for the differences in cost structure, as well as learning rate and scope, between the different processes. For instance, manufacturing aluminum bodies requires four distinct processes: three forming processes—die casting, extrusion, and stamping—plus assembly, during which the individual parts are joined.

The cost of each manufacturing process is, in turn, broken down into dedicated investment cost, non-dedicated investment cost, and variable cost, this time so that the model calculations can more accurately represent a manufacturer's cash flow. For the forming processes, dedicated investment cost represents tooling—the stamping, extrusion, or other dies required to form a specific part. These costs are typically paid at the beginning of a design cycle, when the automaker has to invest in new dies to manufacture its redesigned vehicle, and on an as-needed basis throughout the cycle depending on how quickly the dies wear out. For the most part, though, this cost is a fixed cost and independent of production volume. Variable costs, on the other hand, scale with production volume. Materials, energy, and labor costs all fall under this category. The final category, non-dedicated investment cost, represents investments, for example, a stamping press, that can be used to manufacture a wide variety of parts. Equipment, building, overhead[2], maintenance, and working capital[3] costs are all considered to be non-dedicated investment costs. Since these investments are assumed to be completely non-dedicated and operate at full capacity, each product "pays" only for the time it uses a particular investment. Consequently,

[2]Overhead costs in the PBCMs represent the cost of indirect labor—so workers assigned to oversee other workers or the process line.

[3]Working capital cost represents the revenue a firm could have realized had it taken the money used to manufacture its products and invested it elsewhere for capital gains. The cost is incurred because the firm presumably has to pay upfront for materials, energy, etc. and only later sees revenue from product sales.

Body Design	Material	Mfg Process	Cost Category	Type
Glass Fiber Composite	Glass fiber composite	SRIM	Dedicated Non-dedicated Variable	Tooling Equipment, building, overhead... Material, energy, labor
	Mild steel	Stamping	Dedicated Non-dedicated Variable	Tooling Equipment, building, overhead... Material, energy, labor
		Assembly	Dedicated Non-dedicated Variable Equipment	Tooling Building, overhead... Material, energy, labor
Aluminum Spaceframe	Aluminum	Die Casting	Dedicated Non-dedicated Variable	Tooling Equipment, building, overhead... Material, energy, labor
	Aluminum	Extrusion	Dedicated Non-dedicated Variable	Tooling Equipment, building, overhead... Material, energy, labor
	Aluminum	Stamping	Dedicated Non-dedicated Variable	Tooling Equipment, building, overhead... Material, energy, labor
		Assembly	Dedicated Non-dedicated Variable Equipment	Tooling Building, overhead... Material, energy, labor

Figure 6-4: Manufacturing cost breakdown of aluminum and glass fiber composite bodies.

the cost attributed to any given product scales with that product's production volume, much like variable cost. The same amortization assumptions from the earlier case study—a 13-year equipment lifetime and a 40-year building lifetime—are used to calculate the per-unit equipment or other non-dedicated investment costs.

Assembling vehicles is different from the processes used to form their individual parts (see Section 3.2.1 for a description). The primary consequence of this difference is the need for a fourth cost category, equipment cost, in addition to the three listed above. In assembly, the equipment is neither completely dedicated nor completely non-dedicated. Rather, it is dedicated to a vehicle design for the time that particular design is in production—in this case study, for a design cycle or four years out of the equipment's 13-year lifetime. Once the cycle is over, the assembly line is dismantled and the equipment reused in a new assembly line for the next design. The cost calculations represent this by assuming the automaker pays the full cost of the assembly equipment up front at the beginning of the cycle, and then sells the equipment for its residual value at the end of the cycle. A database is required by the selection method to estimate assembly equipment, as well as tooling and building costs at different annual production volumes. Because assembly lines scale serially with volume, there is no single number that represents the process' fixed costs at all volumes. However, once the assembly costs have been obtained, they are treated in the same manner as forming costs.

In addition to accommodating the different cost structures and learning rates and scopes of each manufacturing process, breaking down manufacturing cost has other benefits. First, it permits a more accurate representation of an automaker's cash flow. Investment costs are paid at the beginning of each design cycle, while variable costs and non-dedicated investment costs are made on a yearly basis (Figure 6-5). Future costs are then discounted to account for the time value of money. Second, the cost breakdown enables more accurate accounting of any experience the automaker gains during the model's time frame. In some cases, subsystem designs include small parts (often reinforcements) that are manufactured using different materials or different processes from the subsystem's primary material (see Figure 6-4). The automaker can still learn from producing these parts and share that knowledge with other designs or subsystems that use the same material and process. For example, instead of experience gained with aluminum closures being applied only to aluminum closures in other vehicles, it can also be applied to other

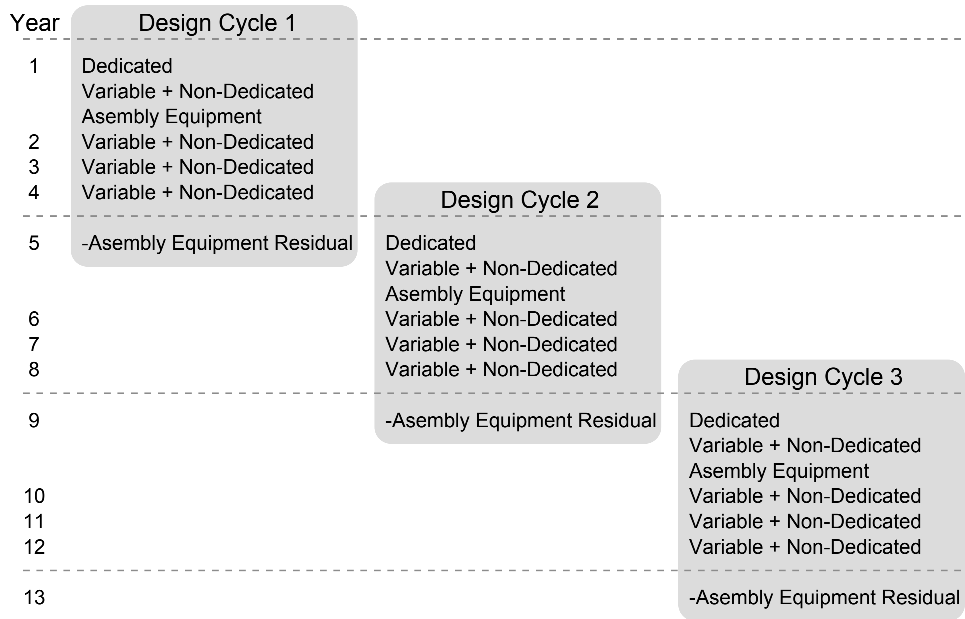

Figure 6-5: Manufacturing cost cash flow.

designs or subsystems that use aluminum, such as Mg-Al closures or aluminum bodies. Third, the separation of fixed costs from variable costs facilitates the calculation of manufacturing cost at various production volumes. Once the inputs are generated, the selection model can then be run independently of the process-based cost models that are required to predict manufacturing cost (see Section 3.2.1). And finally, dividing manufacturing cost means that not only can the learning rate and scope be chosen for each manufacturing process, but so can the components of the cost that evolve as the automaker gains experience. In the case study, only variable costs and non-dedicated investment costs are assumed to evolve since learning takes place by doing—that is, by the automaker physically producing vehicles. Dedicated investment costs do not change because there is presumably minimal opportunity to learn and improve the tooling, especially if purchases take place only once per design cycle.

6.2.5 Chromosome Modifications

While there are many benefits to decomposing the cost of each material option into its component manufacturing processes, there are also consequences. Instead of one binary decision variable assigned to a material option, each option can now be associated with several decision variables, each variable representing a manufacturing process used by that option in each de-

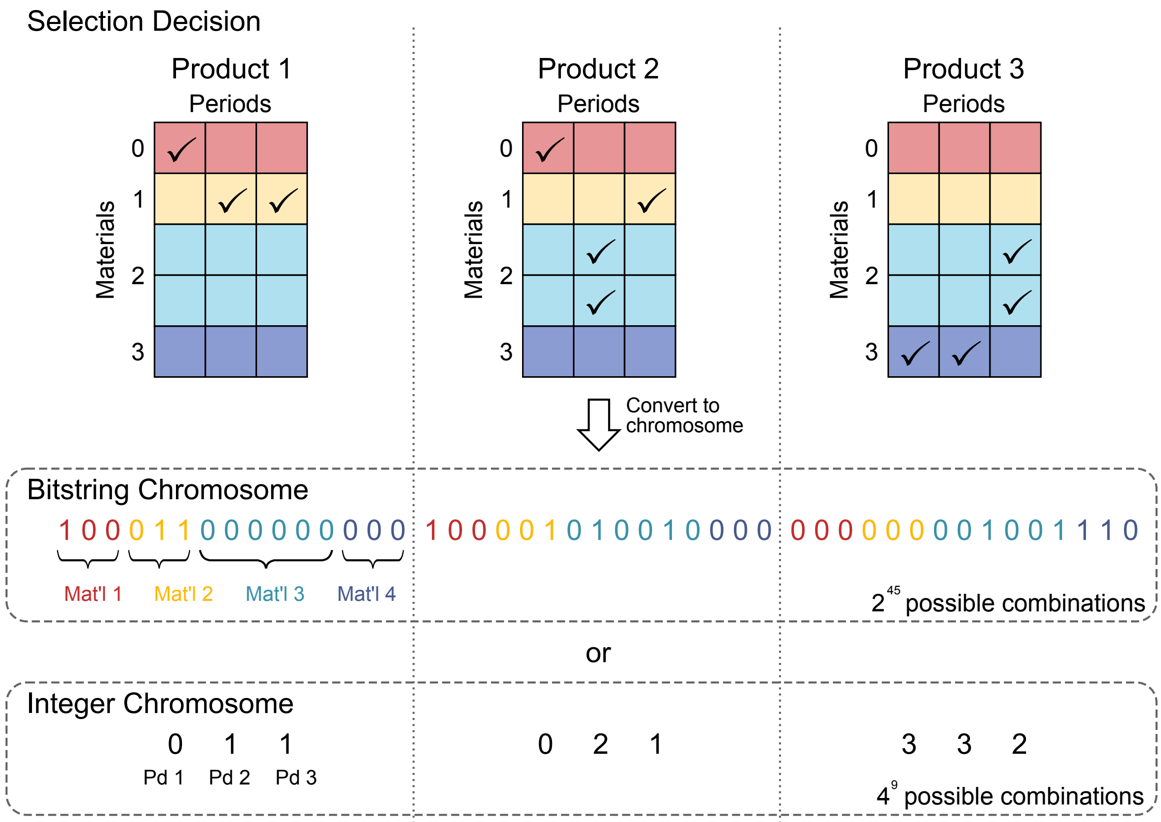

Figure 6-6: Example of a bitstring and integer chromosome for the genetic algorithm in the automotive case study.

sign cycle (bitstring chromosome in Figure 6-6). This not only increases the number of decision variables—and therefore problem scale—but also necessitates use of the aforementioned conflict and requirement constraints (see Section 5.1.2) to ensure that, at most, one material is selected for each subsystem and that all components of that material are included in the cost and fuel economy calculations. While an integer linear program can handle a problem implemented in this manner, the scale of the search space and its fragmentation by the constraints can be problematic for a genetic algorithm.

Therefore, the integer linear program uses binary decision variables as in the stylized exercise, but the genetic algorithm relies on a modified chromosome with integer decision variables (bottom chromosome in Figure 6-6). In this alternative implementation, each gene or decision variable represents the material choice for a subsystem in a single design cycle. Decision variable values are restricted to a range of integers, each number corresponding to a material option for its respective subsystem—for example, zero represents the baseline material; one, high-strength steel; two, aluminum, and so forth. The modified chromosome thus reduces the number of decision variables in exchange for an increase in the number of states or values each variable can adopt. Overall, this still reduces the search space—from 68×10^9 to 262×10^3 potential combinations for a three-product, four-material, three-period problem (see Section 5.1.3). The fitness function for this chromosome then recreates the bitstring chromosome to calculate total manufacturing cost so the genetic algorithm can use the same inputs, complete with materials costs decomposed into component manufacturing processes, as the integer linear program.

Adopting this alternative implementation of the genetic algorithm's chromosome not only reduces the size of the problem's search space, but also eliminates the need for conflict and required constraints. The smaller search space of the integer chromosome can have a drastic effect on the algorithm's runtime, particularly for larger-scale problems, and facilitates the identification of feasible solutions, even when the genetic algorithm is run without a seed.

6.3 Inputs

The above details can now be used to build integer linear program and genetic algorithm selection models that will identify preferred materials for each subsystem and design cycle. These preferred materials minimize total manufacturing cost, but still enable the automaker to satisfy CAFE and other constraints. To calculate the manufacturing cost and fuel economy improvement of the material options, however, the selection models require both vehicle and material attributes, which are detailed below.

6.3.1 Vehicles

Vehicle attributes are shown in Table 6.1 for the three vehicle platforms that are available for modification, plus a fourth "platform" that represents the remainder of the automaker's fleet and rounds out the fleet's CAFE number. Sales fractions from 2009 EPA data [11] are used to estimate the annual production volume of each platform, assuming a high-volume fleet. Fuel economy is also taken from the same source, although the numbers used in this model are increased by 5% from the EPA's numbers. The change in fuel consumption with respect to

Table 6.1: Case study vehicle attributes.

Platform	Curb Weight [kg]	Initial FE [mpg]	% ΔFC per 10% Δwt	Ann. Production Volume [#/yr]
Compact	1,357	26.99	6%	200,000
Midsize	1,549	26.36	6%	170,000
Large	1,613	23.31	6%	120,000
Remainder	1,814	30.00	0%	400,000

Table 6.2: Abbreviations for body and closure material options.

	Body		**Closures**
ID	Primary Material	ID	Primary Material
MS	Mild Steel	MS	Mild Steel
HSS	High-Strength Steel	HSS	High-Strength Steel
AL	Aluminum	AL	Aluminum
GF	Glass Fiber Composite	MG	Magnesium & Aluminum
CF	Carbon Fiber Composite	SMC	Sheet Molding Compound

vehicle weight reduction is based on work by Wohlecker et al. [78] and assumed to be constant over all platforms.

6.3.2 Materials

All available material options for the bodies and closures are listed in Table 6.2; specific inputs for a few midsize car bodies are shown in Table 6.3, with the remainder of the inputs for the other materials and vehicle platforms detailed in Appendix D. The body designs are the same as those used in the single-product case study. Since actual designs for three different car sizes were not available, the dimensions and masses of the subsystems for the compact and large cars were scaled from those of the midsize car. These designs served as the basis for generating inputs used in the materials selection models.

The change in fuel consumption associated with each material option is related to the weight of the alternative design relative to that of the current design (i.e. the mild steel design). Lighter weight subsystems have more impact on fuel consumption, but are often associated with higher costs. Calculating total manufacturing cost is more complicated than calculating vehicle fuel consumption as it needs the individual costs of each design. As with the preceding automotive case study, determining these costs requires more than just knowledge of the material and subsystem because of their context-dependent nature. Process-based cost models are again employed to

Table 6.3: Material options for three midsize car body designs. Long-term non-dedicated and variable costs are shown.

ID	Mild steel (MS)	Aluminm (AL)			Glass Fiber (GF)	
Mass	322 kg	193 kg			219 kg	
Material	Mild steel	Al	Al	Al	GF composite	Mild steel
Process	Stamping	Stamping	Die casting	Extrusion	SRIM	Stamping
Learning scope	Low	Medium	Medium	Medium	Medium	Low
Learning rate	Slow	Fast	Slow	Slow	Slow	Slow
Forming Costs						
Variable	$477	$223	$151	$267	$900	$42
Non-ded. investment	$144	$67	$139	$226	$191	$11
Dedicated investment	$70.3 M	$33.4 M	$5.0 M	$420,000	$13.2 M	$1.6 M
Allocation fraction	3.2%	3.8%	81%	29%	40%	0%
Investment life	500 M	500 M	120,000	34,800	165,000	500 M
Assembly Costs @ 170k						
Variable	$135	$166			$58	
Equipment	$41.2 M	$61.1 M			$17.1 M	
Tooling	$15.8 M	$22.1 M			$8.92 M	
Building	$14.2 M	$19.8 M			$6.62 M	

predict dedicated investment (and assembly equipment) cost, as well as long-term variable and non-dedicated investment cost. The same values for the exogenous variables and material prices (Tables 3.2 and 3.3) are used. For this case study, however, two additional cost model outputs are required: allocation fraction and investment lifetime, both of which are needed to more accurately predict dedicated investment cost over a design cycle. Allocation fraction is primarily used to determine the number of parallel lines necessary to meet a desired annual production volume, while investment lifetime indicates how often tools wear out and whether they will have to be replaced mid-cycle.

Once long-term variable and non-dedicated investment costs have been obtained from the cost models, short-term (or unlearned) cost is then backed out based on the learning scope of their respective manufacturing processes as in Section 3.2.1. The learning curves for the materials use the same functional form as those in the single-product case study (see Equation (3.2) and Figure 3-5). To simplify matters, manufacturing processes are assigned a low, medium, or high learning scope and a slow or fast learning rate. These scopes and rates translate to the values of S-curve parameters shown in Table 6.4. As before, these numbers are loosely based on the parameters used by NHTSA in its 2011 CAFE target analysis [4], although they do not exactly match those used for the midsize body-in-white in the first case study (Table 3.4). Finally, an input table keeps track of which manufacturing processes are shared between subsystems and

Table 6.4: S-curve parameter values for various learning scopes and rates.

Rate	Scope	Scope (σ)	Threshold Volume (V_{th})	Max Volume (V_{hi})
slow	low	0.0	300,000	1,200,000
slow	medium	0.4	300,000	1,200,000
slow	high	0.6	300,000	1,200,000
fast	low	0.0	175,000	700,000
fast	medium	0.4	175,000	700,000
fast	high	0.6	175,000	700,000

Table 6.5: CAFE targets for each design cycle in the selection model case study.

Design Cycle	CAFE Target [mpg]	CAFC Target [L/100km]
1	27.5	8.56
2	27.5	8.56
3	28.5	8.25
4	28.5	8.25

vehicle platforms.

6.3.3 Other Inputs

Other inputs to the selection model include a discount rate and CAFE targets for each of the four design cycles. An 8% discount rate was used for this study. CAFE targets are shown in Table 6.5 and are selected so that the automaker will see a step increase in the constraint (much like in the stylized exercise). Although the target values are well below the CAFE standards of the coming years (Figure 6-3), they were chosen to avoid over-constraining the problem, since only materials solutions are available to the automaker in this case study.

6.4 Additional Assumptions

Before the selection results are presented, there are a few additional assumptions the integer linear program and genetic algorithm selection models make in order to simplify their manufacturing cost and fuel consumption calculations:

- The material options have the same risk and are functionally equivalent in all respects other than weight and cost; the automaker will therefore have no preference for any design other than for its contribution to improving fuel consumption and its impact on manufacturing cost.

- A vehicle's weight savings from its two subsystems are additive and independent of the subsystem's specific location in the vehicle. Fuel consumption improvement is also assumed to behave linearly with the vehicles weight (see Equation (6.3)).
- For shared learning, each subsystem "unit" adds one unit to the cumulative volume of a manufacturing process, despite the fact that it is likely composed of numerous parts that are manufactured with that process.

6.5 Selection Results

The purpose of this case study is not only to demonstrate the capabilities of the proposed selection model with a larger, more realistic example, but also to assess the impact of expanding the selection problem's scope to encompass several products and the shared learning among those products. An expansion in problem scope also enables the evaluation of strategies a firm can adopt to introduce new materials; in particular, the use of a test bed and conditions that lead a firm to opt for this approach. As with the preceding case studies, materials are selected using the three representations of manufacturing cost: short-term cost, long-term cost, and evolving cost, in which unit costs evolve due to learning by the automaker. The total manufacturing cost over the case study's time frame is calculated and the resulting selection decisions compared to assess whether the consideration of learning impacts the decision and if so, in what ways. The analysis is first run using the inputs as described in Tables 6.4 and 6.5. Additional analyses are later performed to assess whether changes in conditions can impact the automaker's preferred materials and to identify potential strategies for introducing new materials.

Selection decision results given the above inputs are presented in Figure 6-7 for each of the three approaches to calculating manufacturing cost (labeled on the left); preferred materials that differ between approaches are highlighted in the figure. Table 6.6 contains corresponding manufacturing costs for each of the three decisions. In order to directly compare decisions, the cost of each is evaluated using all three calculation methods. That is, the total manufacturing cost of the materials preferred when the automaker pays short-term costs is evaluated not only using the short-term costs of the material options, but also as if the automaker had, instead, paid the

		Compact Car		Midsize Car		Large Car	
	Cycle	Body	Closures	Body	Closures	Body	Closures
Short-Term	1	MS	MS	MS	MS	MS	MS
	2	MS	MS	MS	MS	MS	MS
	3	AL	AL	AL	MG	HSS	MG
	4	AL	AL	AL	MG	HSS	MG
Long-Term	1	MS	MS	MS	MS	MS	MS
	2	MS	MS	MS	MS	MS	MS
	3	AL	MG	AL	AL	HSS	MG
	4	AL	MG	AL	AL	HSS	MG
Evolving	1	MS	MS	MS	MS	MS	MS
	2	MS	AL	MS	MS	MS	MS
	3	AL	MG	AL	AL	HSS	MG
	4	AL	MG	AL	AL	HSS	MG

Figure 6-7: Preferred materials in the automotive case study for each manufacturing cost calculation approach. CAFE targets in Table 6.5.

Table 6.6: Total manufacturing cost of selection decisions presented in Figure 6-7.

	Cost Evaluation Approach		
Select according to...	Short-Term	Long-Term	Evolving
Short-Term Cost	**$7.503 B**	$6.0951 B	$6.315 B
Long-Term Cost	$7.508 B	**$6.0950 B**	$6.316 B
Evolving Cost	$7.660 B	$6.1590 B	**$6.309 B**

options' long-term costs or experienced cost evolution. Similar calculations are performed for the other two decisions and show that the integer linear program and the genetic algorithm solutions are indeed optimal given their respective search spaces.

Since subsystem costs do not evolve for the first two selection decisions presented in Figure 6-7, their analysis does not require use of the genetic algorithm. These two decisions are almost identical, with the exception of the closures for the compact and midsize cars, which swap between the aluminum and magnesium[4] material options. While both combinations clearly enable the automaker to satisfy its CAFE constraint, magnesium closures in the midsize car are pre-

[4]Magnesium closures are not a pure magnesium design, but rather a combination of magnesium and aluminum, because sheet magnesium cannot yet meet the surface finish requirements of a vehicle's exterior [34].

ferred over aluminum closures when short-term cost is considered because of the high cost of magnesium in comparison to the cost of aluminum and therefore, the automakers desire to limit the number of vehicles that use the more expensive material. In contrast, when only long-term cost is considered, the cost of magnesium is only slightly greater than that of aluminum in a given vehicle platform; consequently, the cost the automaker pays to use magnesium on the compact car, which has the higher annual production volume, is more than offset by the savings it realizes from switching to aluminum on the midsize car.

The third selection decision presented in Figure 6-7 considers cost evolution due to learning in the decision-making process and thus requires use of the genetic algorithm to analyze the problem. Since one of the trade-offs of genetic algorithms is that they cannot guarantee the final result will be the global optimum, the decisions based on short- and long-term costs from above are used to seed the genetic algorithm and help it identify a reasonable combination of materials. The genetic algorithm, however, is also run without the seed chromosomes, should the set of preferred materials from considering cost evolution be very different from those proposed by the integer linear program.

As it turns out for the set of inputs, the selection decision made by the genetic algorithm is nearly identical to those from the integer linear program, the one difference being the use of the compact car's closures as a test bed for the introduction of aluminum in the second design cycle—one design cycle before the CAFE constraint jumps and necessitates the use of any alternative materials. All other materials for the bodies and closures are introduced on an as-needed basis, as determined by the CAFE constraint. The presence of a test bed in the third selection decision shows that deliberately introducing aluminum in anticipation of using it later enables the automaker to move down the learning curve and realize savings during the following design cycles. A comparison of manufacturing costs in the last column of Table 6.6 confirms this conclusion and indicates that the use of a test bed lowers the automaker's total cost.

Given that the automaker pays a cost premium to introduce a material on a test bed, but in exchange realizes cost savings when that material is used again in the future, the automaker's return on investment can be calculated for this action. The investment is equivalent to the cost premium the firm pays to use a test bed in the second design cycle, and likewise, the return is represented by the savings in the third and fourth design cycles. These quantities are calculated

Table 6.7: Return on investment for follwing the third selection decisions in Figure 6-7.

	Calculated relative to	
	Short-Term Decision	Long-Term Decision
Added Cost (Cycles 1–2)	\$107.8 M	\$107.8 M
Savings (Cycles 3–4)	\$114.4 M	\$115.5 M
Return on Investment	6.1%	7.1%

for the selection decision made assuming evolving cost, relative to the decisions made using short- and long-term costs. Numbers are shown in Table 6.7 and represent the cost premium of the test bed and the amount the firm saved by following the third selection decision in Figure 6-7 instead of the first or second one.

When deciding whether to first implement a particular material on a test bed, an automaker has to confront the trade-off between minimizing the upfront cost of the test bed versus gaining sufficient experience to move down the learning curve so it can realize future savings and offset the test beds cost. On one hand, using products that are manufactured at lower production volumes leads to lower test bed costs simply because the automaker has to produce fewer units with a new material. These products, however, are also associated with lower future savings because the firm gains less experience (as measured by cumulative production volume) than if it had chosen a higher-volume product. This trade-off is illustrated by the stylized exercise (see Figure 4-7), but is relevant to the automotive case study.

In the stylized exercise, the lower-volume product is also preferred for a test bed because of its lower unit cost premium for the alternative material. In contrast, the firm in the automotive case study does not have the same low-volume versus high-volume choice: all three cars are manufactured at production volumes of 120,000 units per year or higher. However, the firm does have a choice of bodies or closures for any of the three cars—and chooses the closures of the compact car as the test bed for aluminum. Although they are not the cheapest option in terms of added cost (Figure 6-8), the compact car's closures have the lowest short-term *unit* cost premium and enable the automaker to learn the most about each manufacturing process due to the vehicle platform's high annual production volume. The firm's preference for closures over the body is also motivated by the assumption that the subsystems are weighted equally in their contributions to cumulative volume (see Section 6.4); consequently, the automaker gains as much experience producing aluminum closures as it would producing an aluminum body—but at lower cost. The

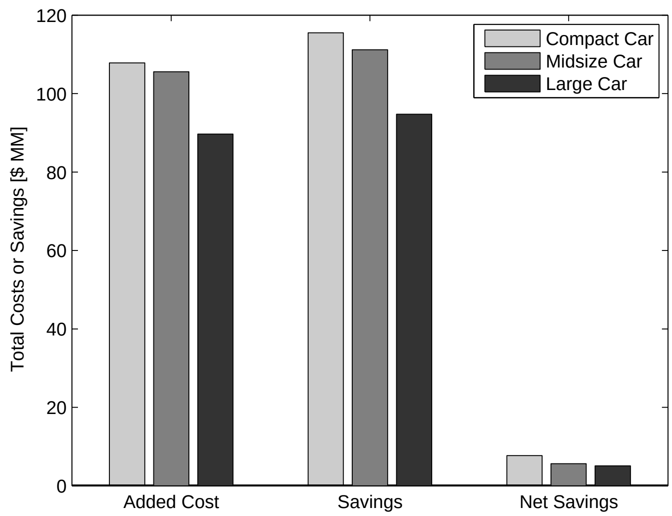

Figure 6-8: Costs and savings of using the closures of alternative vehicle platforms as test beds for aluminum.

preference for the compact car's closures is confirmed in Figure 6-8, which compares the added cost of a test bed in the second design cycle to savings from the third and fourth cycles. The same calculation is performed assuming the automaker chose to introduce aluminum on the midsize car's closures or the large car's closures, in place of the compact car's closures. Vehicle bodies are not considered because they only increase test bed cost without providing any additional benefits.

6.6 Sensitivity Analysis

6.6.1 Test Bed Annual Production Volume

The results from the automotive case study's analysis indicate that an automaker should use the highest-volume car as its test bed and should use that test bed for the minimum amount of time—in this case, one design cycle. These results, however, appear at odds with the conclusions from the stylized exercise in Chapter 4, which suggest that a manufacturing firm should use the lower-volume product as its test bed, and do so for as many years as the time frame will allow. In order to reconcile these observations, a sensitivity analysis that evaluates the impact of changes to the annual production volume of the product that is selected as the test bed is performed.

Although a similar analysis was already done for the stylized exercise, the one in this case study is used to better understand the consequences of changes to the test bed's volume under more realistic conditions. Additional goals of the current sensitivity analysis include assessing the impact of the changes on manufacturing cost and on an automaker's decision to use a test bed, as well as evaluating the robustness of the material and application choice for a test bed.

The sensitivity analysis starts by comparing two scenarios or selection decisions—much like the original approach of the stylized exercise—to gain a clearer understanding of how cost behaves with changes in the test bed's annual production volume. Later, the genetic algorithm is used to select the material and application for the test bed; for now, the costs of the long-term and evolving cost selection decisions from Figure 6-7 are evaluated at different volumes of the compact car. The annual production volume of the compact car is changed only in the first and second design cycles to avoid affecting the automaker's CAFE number and therefore material preferences in the third and fourth cycles. Lower volumes of the compact car will not affect the automaker's ability to satisfy CAFE in the first two design cycles: the fuel economy of the compact car is below the CAFE target in those years so decreasing that car's volume will increase the automaker's CAFE number. As in previous analyses, the added cost of using a test bed from the first two design cycles, future savings from the final two design cycles, and net savings, which represent the difference between future savings and added cost—or more precisely, the difference in manufacturing cost between the two selection decisions—are all calculated and plotted.

Analysis results are shown in Figure 6-9. As annual production volume decreases, the additional cost the automaker pays to learn decreases as well because the new material is applied to fewer vehicles. Savings also decrease because the lower volumes imply less experience for the automaker and therefore less cost evolution. These savings depend on the functional form of the learning curve, as well as the number of vehicles that implement alternative materials whose costs are affected by the use of a test bed. For instance, the jump in savings at around a test bed production volume of 50,000 units per year is the result of the curve switching from constant to log-linear at V_{th} (see Figure 3-5).

The net savings illustrates the competing effects between the added cost of a test bed and the savings the automaker realizes later on. As shown in Figure 6-9, net savings is negative at low volumes, indicating that given a choice, the automaker would not introduce aluminum on

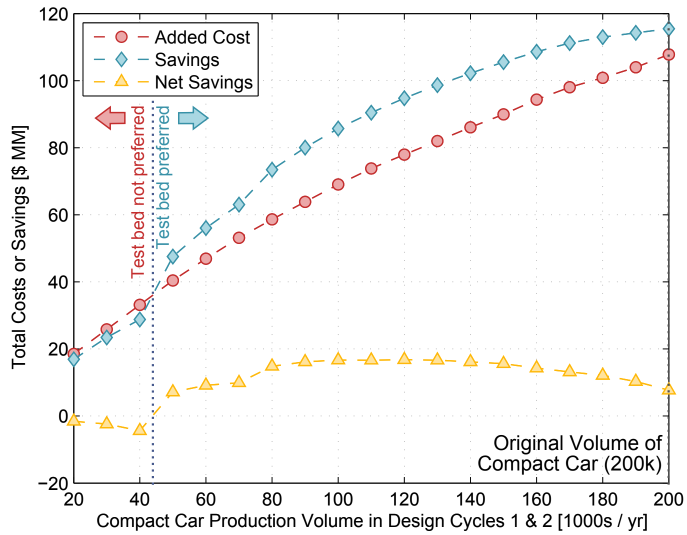

Figure 6-9: Costs and savings as functions of the annual production volume of compact car during first and second design cycles.

the compact car's closures in the second design cycle. As production volume increases, so do net savings because the automaker is able to move further down the learning curve and increase savings. There is a limit, however, to the amount of savings an automaker can realize because the learning curve saturates above V_{hi}. Once this saturation point is reached, increasing the volume of the test bed to deliberately gain more experience and move further down the learning curve will not have a corresponding increase in benefits—only cost—so net savings starts to decrease with increasing production volume. This point is sensitive to the discount rate, but starts happening around 120,000 units per year in Figure 6-9.

The results presented in Figure 6-9, when compared with the stylized exercise's results in Figure 4-7 appear to validate the initial observation made at the beginning of this section: that the stylized exercise and the automotive case study are contrary in their suggestions of when to use a test bed. These results, as it turns out, can be reconciled by taking a closer look at the unit cost premium (assuming short-term unit cost of the new material) of the test beds in both cases, shown in Figure 6-10. For the stylized exercise, this unit cost premium decreases with decreasing annual production volume so the firm will prefer a low-volume product for its test bed. In contrast, the unit cost premium for aluminum in the automotive case study increases with decreasing production volume, so at low volumes (below 50,000 units per year), the savings the

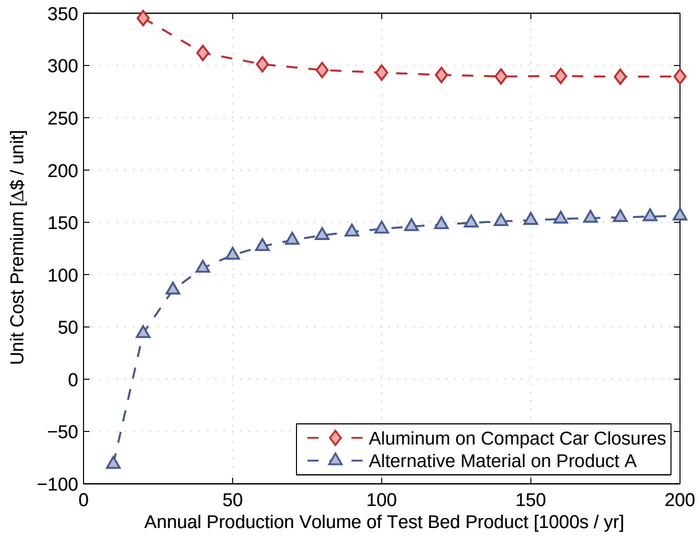

Figure 6-10: Unit cost premium of test beds in the stylized exercise and the automotive case study.

Table 6.8: Compact closure material choice for test bed given production volume.

Compact Car Volume [units / yr]	Test Bed Subsystem	Test Bed Material	No. of Design Cycles
20,000	Compact car closures	SMC	2
30,000	Compact car closures	SMC	1
40,000	Compact car closures	SMC	1
50,000+	Compact car closures	AL	1

automaker realizes are not enough to offset the added cost of using the test bed. Consequently, the automaker will avoid using the compact car's closures as a test bed for aluminum at low production volumes of the compact car.

While the above sensitivity analysis explains the difference in preference for a low or high volume test bed between the stylized exercise and the current case study, it does not provide any insight to the number of years a test bed should be used. This question is answered by repeating the sensitivity analysis, but instead of comparing two selection decisions, the material and application used for the test bed are chosen by the genetic algorithm given different production volumes of the compact car in the first and second design cycles. The results of this analysis are summarized in Table 6.8.

The compact car's closures are identified by the genetic algorithm as the preferred subsystem for the test bed, regardless of the car's production volume. The preferred material, however, changes so that at 50,000 units per year and above, aluminum is chosen for the test bed, and below 50,000 units per year, SMC is preferred. This is consistent with the materials' manufacturing costs: SMC has lower fixed (i.e. tooling) costs and is therefore the cheaper option at low volumes. As for why SMC, its design contains extruded aluminum parts, from which the automaker can still gain experience and have that experience count as much as if it were gained from manufacturing aluminum closures. This result is partly due to the assumption that all subsystems, regardless of their size or number of parts, are weighted equally when it comes to learning.

At very low production volumes of the compact car (20,000 units per year in the above analysis), a test bed is employed in both the first and second design cycles. The extended use of a test bed enables the automaker to move further down the learning curve than if it had implemented a test bed for only one cycle, and indicates that at very low volumes (at least relative to V_{th}), future savings can outweigh the costs of gaining additional experience. This result is consistent with the observation from the stylized exercise that a firm will prefer to use a test bed as long as its time frame allows. (It should be noted, though, that the stylized exercise marks time in years, whereas the automotive case study does so in design cycles, each of which equals four years.) At higher production volumes of the compact car, the use of a test bed is limited to one design cycle because cumulative volume by the end of that cycle is high enough that the cost of a second cycle of using the test bed will be greater than the savings the automaker can realize later on. The lack of additional benefits is partly because the learning curve saturates; this saturation is reflected in the slowing rise of savings in Figure 6-9.

6.6.2 CAFE Target

The second sensitivity analysis investigates the consequences of changing the CAFE constraint on the preferred materials and the total manufacturing cost. Rather than changing the constraint timing, as was done in the sensitivity analysis of the stylized exercise, the automaker's constraint space is gradually tightened over time. Gradually tightening the CAFE standard presents a more interesting problem than simply increasing the constraint from 27.5 mpg to 28.5 mpg because the automaker is forced to continually readjust in order to accommodate the changing standard. The larger scope of the current case study also means the automaker will have more flexibility

		Compact Car		Midsize Car		Large Car	
	Cycle	Body	Closures	Body	Closures	Body	Closures
Short-Term	1	MS	MS	MS	MS	MS	MS
	2	HSS	MS	HSS	MS	HSS	MS
	3	AL	AL	AL	MG	HSS	MG
	4	CF	AL	CF	AL	CF	AL
Long-Term	1	MS	MS	MS	MS	MS	MS
	2	HSS	MS	HSS	MS	HSS	MS
	3	AL	MG	AL	AL	HSS	MG
	4	CF	MG	CF	AL	CF	SMC
Evolving	1	MS	MS	MS	MS	MS	MS
	2	HSS	AL	HSS	MS	HSS	MS
	3	AL	HSS	AL	AL	AL	AL
	4	CF	AL	CF	AL	CF	AL

Figure 6-11: Preferred materials assuming a linearly increasing CAFE target.

Table 6.9: Total manufacturing cost of selection decisions presented in Figure 6-11

	Cost Evaluation Approach		
Select according to...	Short-Term	Long-Term	Evolving
Short-Term Cost	**$9.842 B**	$6.831 B	$7.757 B
Long-Term Cost	$9.858 B	**$6.823 B**	$7.771 B
Evolving Cost	$10.008 B	$6.917 B	**$7.737 B**

in choosing how to satisfy the constraint; in contrast, the firm in the stylized case study was limited to two material options for two products. The goal of this analysis is to assess whether the consideration of cost evolution will continue to affect the automaker's material preferences, as well as lead to its decision to introduce any of the alternative materials on a test bed.

The new constraint space assumes the CAFE target increases linearly from 27.5 to 29.0 mpg in increments of 0.5 mpg over the four design cycles. The selection models are then run using the new CAFE targets and the three manufacturing cost calculation approaches; resulting selection decisions are shown in Figure 6-11, with costs in Table 6.9 and return on investment in Table 6.10. As expected, the integer linear program results in the first and third design cycles match those presented in Figure 6-7 since both analyses have the same CAFE targets for their respective

Table 6.10: Return on investment for following the third selection decisions in Figure 6-11.

	Calculated relative to	
	Short-Term Decision	Long-Term Decision
Added Cost (Cycles 1–2)	$107.8 M	$107.8 M
Savings (Cycles 3–4)	$127.9 M	$141.5 M
Return on Investment	19%	31%

periods. The manufacturing costs in Table 6.9 confirm that the integer linear program and the genetic algorithm results are optimal given their respective search spaces; these costs, though, are higher than those in Table 6.6 because of the tighter constraints.

A comparison of the integer linear program and genetic algorithm selection decisions in Figure 6-11 shows that accounting for cost evolution in the manufacturing cost calculations has an impact on the preferred materials. First, the genetic algorithm's results indicate that it is financially beneficial for the automaker to use the compact car's closures as a test bed for the introduction of aluminum in the second design cycle—same as in the initial analysis (Figure 6-7). The presence of learning also causes the genetic algorithm to exhibit a strong preference for aluminum in the bodies and closures of all vehicle platforms during the third design cycle. This emphasis on one material allows the firm to capitalize on its experience with that material and thereby reduce overall cost. This is reflected in the firm's return on investment (Table 6.10), which is notably higher than in the previous analysis (Table 6.7), because the firm is able to use aluminum in nearly every subsystem and pay close to long-term cost for each—thereby realizing a sizeable return for the same investment. In contrast, switching materials between design cycles, or using different materials within one cycle, leads the firm to pay short-term cost for more of the subsystems, which in turn increases its total manufacturing cost. Switching materials was less of a concern in the initial analysis when the CAFE target changed only between the second and third design cycles and the automaker was not forced to continually readjust its material choices.

6.6.3 Planning for Uncertainty in CAFE

While the above analyses assume the CAFE target is a known quantity for the next 16 years, the reality is that the target is set for no more than a few years in advance. For instance, NHTSA only finalized the light-duty vehicle standards for 2012 through 2016 in the beginning of 2010, while the light truck standards for 2008 through 2011 were set in 2006 [3, 9]. And even then, there is the possibility that the standards may change: in 1985, Ford and GM lobbied for—and were granted—a reduction in the 1986 CAFE target to 26 mpg down from 27.5 mpg [28]. The 2008 to 2011 light truck standards were also found to be "arbitrary and capricious" by the 9th Circuit Court of Appeals, and were returned to NHTSA for further review [1]. Likewise, the Center for Biological Diversity also challenged the 2011 passenger car and light truck CAFE standards [4] for not being set at the "maximum feasible level" [30]. Uncertainty in the CAFE target creates difficulties for an automaker because it does not know which target to design for. This can lead the automaker to paying more (or less) than originally anticipated, particularly if it has limited time to accommodate the new constraint. This section investigates whether the consideration of cost evolution through learning can help an automaker cope with uncertainty in its constraint space.

In the analysis, the costs of "what-if" scenarios, in which the automaker plans for a specific CAFE target in the fourth design cycle, but partway through the third cycle finds out that the target for the final cycle has been either increased or decreased, are evaluated. Because the automaker is only alerted to the change in the third design cycle, only the preferred materials in the final cycle are affected. The scenarios' costs are then compared to the cost the automaker would have paid if it had planned for the revised CAFE target in the first place.

The first scenario asks "what if" the CAFE constraint is relaxed after the automaker has planned to meet an original, tighter constraint. For the original constraint, the CAFE target increases linearly from 27.5 to 29.0 mpg by the end of the time frame. The automaker therefore plans to use the materials from the "evolving cost" selection decision in Figure 6-11. At some point, however, the automaker learns the constraint for the final four years of its time frame has been relaxed by 0.5 mpg so the targets over the time frame progress as follows: 27.5, 28.0, 28.5, and 28.5 mpg. Consequently, the automaker has the option to lower its costs by choosing a new set of materials for the fourth design cycle, given the new target. Optimizing for the fourth design cycle assuming the automaker has already chosen and implemented materials during the first three cycles, yields the selection decision in Figure 6-12, which has a manufacturing cost of $6.453 billion (calculated assuming costs evolve).

The next step is to compare the above cost to the cost the automaker would have paid had it planned ahead for the 28.5-mpg target in the fourth design cycle. Figure 6-13 shows the

	Compact Car		Midsize Car		Large Car		
Cycle	Body	Closures	Body	Closures	Body	Closures	CAFE
1	MS	MS	MS	MS	MS	MS	27.5 mpg
2	HSS	AL	HSS	MS	HSS	MS	28.0 mpg
3	AL	HSS	AL	AL	AL	AL	28.5 mpg
4	AL	HSS	AL	AL	AL	AL	28.5 mpg

Figure 6-12: Preferred materials under relaxed constraint.

		Compact Car		Midsize Car		Large Car		
	Cycle	Body	Closures	Body	Closures	Body	Closures	Evolving Cost
Short-Term	1	MS	MS	MS	MS	MS	MS	$6.456 B
	2	HSS	MS	HSS	MS	HSS	MS	
	3	AL	AL	AL	MG	HSS	MG	
	4	AL	AL	AL	MG	HSS	MG	
Long-Term	1	MS	MS	MS	MS	MS	MS	$6.457 B
	2	HSS	MS	HSS	MS	HSS	MS	
	3	AL	MG	AL	AL	HSS	MG	
	4	AL	MG	AL	AL	HSS	MG	
Evolving	1	MS	MS	MS	MS	MS	MS	$6.450 B
	2	HSS	AL	HSS	MS	HSS	MS	
	3	AL	MG	AL	AL	HSS	MG	
	4	AL	MG	AL	AL	HSS	MG	

Figure 6-13: Preferred materials and total evolving manufacturing cost assuming assuming automaker plans ahead for relaxed constraint.

Cycle	Compact Car Body	Compact Car Closures	Midsize Car Body	Midsize Car Closures	Large Car Body	Large Car Closures	CAFE
1	MS	MS	MS	MS	MS	MS	27.5 mpg
2	HSS	AL	HSS	MS	HSS	MS	28.0 mpg
3	AL	MG	AL	AL	HSS	MG	28.5 mpg
4	CF	AL	CF	AL	CF	AL	29.0 mpg

Figure 6-14: Preferred materials under tightened constraint.

preferred materials under the modified CAFE targets given all three approaches to calculating manufacturing cost; the cost of each selection decision as evaluated with cost evolution is shown beside each selection decision. The costs of the integer linear program decisions represent the cost to the automaker if it had not considered learning when selecting materials, but costs evolved regardless. A comparison of the numbers in Figure 6-13 and the cost of the decision in Figure 6-12 reveals that, given the inputs for this case study, planning for a 29.0-mpg target but making a last-minute switch to 28.5-mpg has a minimal impact on the automaker's cost—$6.453 billion versus $6.450 billion. In fact, the decision in Figure 6-12 has a lower cost than the integer linear program results in Figure 6-13 because it emphasizes materials the automaker has already worked with.

The reverse scenario, in which the CAFE constraint is tightened, can be evaluated using a similar approach. In this case, the automaker plans for a 28.5-mpg target in the final two design cycles, but is instead forced to accommodate a 29.0-mpg target in the final four years. The resulting selection decision is shown in Figure 6-14 and has a manufacturing cost of $7.751 billion (again, assuming costs evolve). The first three design cycles of this decision are the same as those in the genetic algorithm's solution in Figure 6-13; the materials selected for the fourth cycle, however, have been altered in response to the higher CAFE target. If the automaker were, instead, to have originally planned for the increased target in the fourth design cycle, it would have arrived at the results in Figure 6-11 and Table 6.9. Comparing these numbers indicates that *by considering learning when selecting materials, the automaker is able to reduce its costs*. It still pays $14 million ($7.751 billion versus $7.737 billion) more than if it had planned for the higher target in the first place, but the total amount is still less than if it did not consider learning in the selection process ($7.751 billion versus $7.757 billion or $7.771 billion, according to Table 6.9).

The above analysis concerning uncertainty in the CAFE constraint suggests that the firm can lower its costs by accounting for learning in the materials selection process. The firm is also better off planning for a tighter constraint when considering cost evolution: even if the constraint changes to become less binding, the automaker can still choose materials it has experience in manufacturing and thereby keep the additional cost of modifying its production plans to a minimum (in this case, $3 million). On the other hand, if the automaker does not plan for the tighter constraint, there is the possibility its modified plans will include the use of an unfamiliar material and limit its ability to capitalize on its experience from the previous design cycles. In the above analysis, tightening the constraint costs the automaker an additional $14 million. Not considering cost evolution will also lead to higher costs for similar reasons: new or unfamiliar materials are selected, for which the automaker will have to pay higher costs without receiving any benefit. Therefore, considering cost evolution, as well as planning for tighter constraints, can help an automaker better cope with uncertainty in its constraint space.

6.7 Summary

The purpose of the automotive case study was, first, to illustrate the application of the proposed multi-product selection framework to a larger-scale problem; second, to further explore the consequences of considering cost evolution on a firm's preferred materials; and finally, to identify potential strategies the firm could adopt when introducing new materials. To perform the desired analyses, the proposed framework was tailored to analyze an automaker's selection problem and to calculate the total manufacturing cost of its vehicle fleet. Cost was used in place of profit, not only because calculating net revenue requires additional information, but also because of the presence of the CAFE constraint.

The case study's results show that the consideration of cost evolution can affect an automaker's preferred materials and point towards strategies the automaker can adopt when introducing the new materials to its fleet. One of these strategies is the use of a test bed, which the automaker can employ to gain experience with unfamiliar materials before a widespread application of those materials throughout its fleet. Although the stylized exercise already indicated that, under certain conditions, using a test bed is financially beneficial to a firm, the exercise's observations were incomplete due to its limited scope. This limited scope is addressed with the automotive case study, which was designed to provide a more realistic example for the application of the multi-product selection framework. Despite its larger scale, the automotive case study's results suggest that the use of a test bed continues to be favored, even at larger annual production volumes. As long as the automaker has a means to reduce the upfront cost of a test bed or increase any future savings it will realize, it will choose to use a test bed.

The consideration of cost evolution from learning also encourages a firm to become comfortable working with one material and use that material as often as the constraint space will allow. This is best illustrated with analysis concerning the linearly increasing CAFE target: both integer linear program solutions, which do not account for learning, select a mix of materials, whereas the genetic algorithm solution emphasizes the use of aluminum in both bodies and closures. Finally, the results of the sensitivity analysis suggest that, when selecting materials, an automaker should plan for tighter constraints as a means of coping with uncertainty in the constraint space. By aiming high, the automaker avoids being "side-tracked" by materials it may not be able to use for more than one design cycle and is easily able to adjust should the constraint be relaxed. On the other hand, planning for a lower constraint can have a larger impact on the automaker's manufacturing cost, especially if the constraint tightens and the automaker is forced to manufacture its vehicles with unfamiliar materials in order to satisfy the tighter constraint.

It should be noted that the above conclusions are drawn despite the knowledge that a genetic algorithm does not necessarily identify *the* best selection decision—merely a satisfactory one. However, that satisfactory decision is clearly still good enough in that its total manufacturing cost is lower than those of decisions resulting from the integer linear program, assuming cost evolution due to learning takes place.

Chapter 7

Conclusions

7.1 Summary

Selecting materials for use in a product is an important component of the development process because poor choices can negatively impact the product's market share or profitability. Identifying satisfactory materials, though, is a complex process: often, a firm or product designer has many options to choose among and has to make selections among competing criteria. In response to this problem, a large number of methods have been developed to inform a firm's selection decision. These methods, however, contain assumptions in order to make the selection process more manageable. Such assumptions include

- Material properties are invariant over the decision's time horizon
- Materials for a product or application are selected independently of other products or applications
- The implementation strategy a firm uses to introduce a new material to its products is independent of that material and its properties.

While these assumptions are true in most cases, there are occasionally conditions under which they do not hold.

One such instance is when emergent properties are key selection criteria. These properties are, by definition, context-dependent and likely are *not* time-invariant over the decision's time horizon: as the context evolves, whether driven by changes in operational conditions, market conditions, consumer preferences, or for other reasons, so too do these properties. A product's manufacturing cost is a prime example of an emergent property—and one that factors into almost every materials selection decision. Evolution in manufacturing cost can be driven by external factors such as changes in raw material prices, interest rates, the job market, and so forth, or by internal factors in which a firm continually improves its manufacturing processes and product designs. This study focuses on the internal causes of cost evolution, specifically on "learning by doing" as the driver of changes in cost. According to the theory of learning, firms gain experience through repetition and are able to improve their manufacturing processes or product designs and thereby reduce their costs. This is particularly true for new or unfamiliar materials, which firms have minimal experience working with, but are often forced to consider in the materials selection process. Since learning happens internally, a firm can control its own learning process and can choose to deliberately gain experience and apply that knowledge to improving its product's design or its manufacturing processes.

Through learning, the manufacturing cost of a product can evolve over the firm's time horizon. For firms that manufacture several similar products, however, it is likely that experience transfer can take place among products that share a common resource—such as a manufacturing process line for a specific material. As the firm gains experience and improves its facility with that resource, all products that rely on it benefit and contribute to its cumulative production volume. In this way, *shared learning* takes place among products and selection decisions made for one product can affect future decisions for others. Considering learning can also influence a firm's decision to adopt a new material and on which product. For instance, the firm may choose to introduce a material first on a test bed before adopting it on the remainder of its products; this decision, though, will depend on that material's manufacturing cost and whether it is expected to evolve and to what extent. This research, therefore, incorporates manufacturing cost evolution through learning into the materials selection process in order to evaluate its impact on a firm's preferred materials and to identify strategies the firm can adopt to introduce those materials to its products.

The first part of this investigation focused on incorporating learning into a traditional, single-product selection method. A case study for the body-in-white of a midsize car was used to illustrate the application of the modified selection method. The study's results indicated that, while cost evolution affected the ordering of material options, it only rarely suggested alternatives that a purely static analysis would not have also chosen. It was recognized, however, that restricting a selection method's scope to a single product—as is the case of traditional methods and the body-in-white case study—limited the method's usefulness in accurately accounting for the impact of learning on a firm's selection decision, particularly when learning can be shared among multiple products.

Therefore, the selection problem scope was expanded in the second part of this research to include multiple products and account for a firm's materials choices over its time horizon. A stylized exercise was first used to motivate the need for multiple products and demonstrate how shared learning can lead to lower manufacturing costs through the use of a test bed. In this exercise, the total manufacturing cost was calculated for two scenarios, but only one of which uses a test bed to introduce the alternative material. Learning was shared between both products by assuming that the firm gained experience working with the alternative material regardless of which product implemented that material. The exercise's results suggested that shared learning can lead to the use of a test bed—and indeed, is a necessary assumption if the test bed is to lower the firm's total manufacturing cost. The results also showed that the firm's decision to introduce the alternative material on a test bed depended on the production volume of the product used as a test bed, the alternative material's cost relative to that of the baseline material, and the total volume of products that ultimately use the alternative material, among other factors (e.g. discount rate).

Once it was shown that the consideration of shared learning could reduce manufacturing cost through the use of a test bed for introducing new materials, a more formal multi-product, multi-period materials selection framework was developed. This revised framework expanded the scope of the selection problem with the objective of finding the combination of materials and products that minimizes a firm's total manufacturing cost over a specified time horizon; shared learning is accounted for in the manufacturing cost calculation. The selection framework relied on a combination of an integer linear program and a genetic algorithm to optimize manufacturing cost. The former found a feasible solution in the absence of learning (which involves non-linear calculations, owing to its use of cumulative production volume to predict evolving cost), and the latter optimized that solution assuming costs evolved.

The selection framework was then applied to two case studies, the first of which revisited the stylized exercise from before in order to validate the expanded selection approach. The second case study posed a more realistic selection problem in the context of an automaker seeking to improve the fuel economy of its vehicles through use of alternative lightweight materials. Results from the automotive case study showed that, while the consideration of cost evolution did not always alter an automaker's material preferences, it could still impact *when* the automaker decided to introduce a new material to its products. Specifically, the results indicated that accounting for shared learning led the automaker to introduce aluminum on the closures of its compact car a full design cycle before constraints necessitated the use of any alternative material in the fleet (see Figure 6-7). This use of the compact car's closures as a test bed, in turn, allowed the automaker to deliberately gain experience and reduce its overall manufacturing cost. A test bed, though, was not the only strategy for introducing new materials suggested by the case study's results: in the sensitivity analysis, it was shown that the automaker could emphasize the use of a particular material across different applications to reduce its costs, or plan for tighter design constraints as a means to cope with uncertainty in the constraint space.

7.2 Future Work

7.2.1 Risk and Uncertainty

There are several factors which the analyses in this study do not consider but may matter to the decision-maker. Risk and uncertainty are two such factors. Often, a firm is reluctant to change a product's materials because with the new materials come the risks of development delays, supply chain or manufacturing problems, consumer rejection, and so forth. Therefore, unless there is a compelling reason to switch materials (e.g. the new material lowers costs, meets tighter performance requirements, enables the firm to satisfy government regulations, matches consumer preferences, etc.), the firm will prefer to continue using its current materials because it knows they work. The presence of risk is also a reason why firms prefer lower-volume products as a test beds: it enables them to minimize the downside should there be manufacturing or other problems, or should consumers not accept the modified product.

Even when the firm decides it needs to consider alternative materials for its products and thus gather information to inform its selection decision, it still faces uncertainty in material properties and constraints used in the selection process. Some of these uncertainties can be removed (for example, contracts with suppliers to lock-in raw material prices), but others, such as the rate or scope of learning, manufacturing process cycle time, and so forth, will remain. If these uncertainties—and the risk a firm takes when using a new material—are to be considered in the materials selection process, selection framework will have to be modified in order to account for these factors, and additional sensitivity analyses run to assess when variations in input parameters will impact the preferred materials.

7.2.2 Net Revenue

The analyses conducted in this study optimized for cost rather than net revenue. For a single-product analysis in which the firm is a price taker, the two metrics are equivalent. This was also the case for the multi-product analysis because of CAFE regulations combined with other assumptions. In the more general case, however, there are situations in which optimizing for net revenue will lead to different results compared with optimizing for cost—for instance, when several products are involved or when the firm can set product prices. As a case in point, improving a product by implementing a new material will affect that product's revenue according to how many consumers are willing to purchase that improved product and the extent to which the firm changes its profit margin and production volume. The revenue of other products can also be affected if consumers who would have purchased those products instead opt to buy the improved product. Therefore, calculating revenue change will require product profit margins, price elasticities of demand, and consumer willingness to pay for various improvements.

Including these parameters in the analysis can also affect a firm's choice of a product for a test bed (as well as its preferred materials). For example, using a product with low price elasticity for a test bed can help a firm recover the cost premium it pays to implement a test bed in the first place because the firm would be able to increase that product's price with minimal losses in demand; this is discussed in more detail in Section 4.3.3. Accounting for revenue through the use of price elasticity and other parameters in the multi-product materials selection framework, however, results in a more complex analysis. As is, the search space is already quite large, even when limited to the selection of materials for a set of products over a given time horizon;

permitting the annual production volume and the prices or profit margins of those products to vary on top of this will only widen the search space and make optimization that much more difficult.

There is a possible way, though, to still account for revenue, but simplify the analysis so it fits within the proposed selection framework: through the use of volume-neutral price. Volume-neutral price represents how much a firm can increase the price of an improved product, and still sell the same volume of that product as it did before the product was improved. Its calculation reflects how much a consumer is willing to pay for that particular improvement in the product. This willingness to pay, in turn, is calculated based on the product's price elasticity of demand—so how much its demand will decrease in response to an increase in its price—and its elasticity of demand with respect to the improved attribute—so how much its demand will increase in response to the improvement in that attribute. These attributes can include product weight or performance, such as the acceleration time or fuel economy of a vehicle. By relying on volume-neutral price instead of price elasticities or profit margins to calculate revenue, the selection framework can make the assumption that the annual production volumes of products in the analysis do not change—by definition, this is the case for volume-neutral price—and thus simplify the optimization process. Appendix A demonstrates a possible approach for incorporating volume-neutral price—or more precisely, consumer willingness to pay—into the materials selection framework and accounting for incremental revenue, as well as cost, in the selection process.

7.2.3 Automotive Case Study Limitations

Much can also be done to improve the multi-product automotive case study. In particular, this case study makes several assumptions to simplify the problem. These assumptions, though, are not necessarily representative of manufacturing practices within the automotive industry and at the very least should be investigated to determine whether changing them will have any bearing on an automaker's material preferences. For one, the case study assumes that the design cycles of all vehicle platforms begin and end at the same time. The reality is that an automaker does not have the resources to redesign all vehicle platforms simultaneously, although one can argue that the other vehicle platforms encompassed within the remainder of the fleet are on different cycles. Staggering platform redesigns will affect not only the cost of each vehicle (through discounting), but also the experience an automaker has gained by the manufacturing start date. Additionally, the model ignores design refreshes, which occur between redesigns and in which automakers make small modifications to a vehicle. This assumption is valid when the automaker's options are limited to alternative materials, which usually require significant changes to the vehicle's design and are therefore only implemented during redesigns. Other technology options, however, may be simple enough to be implemented during design refreshes and thus necessitate the consideration of more frequent modifications to the vehicle's design—especially if the automaker is facing a rapidly increasing CAFE target.

Also in question is the freedom the model gives the automaker to switch materials without any adverse consequences other than the potential lack of experience with new materials. As a result, the proposed selection decisions freely jump between materials, even within a single application—for example, the bodies of the compact and midsize cars in Figure 6-11 use a different material in each design cycle. One possible fix to the selection method is to add the assumption that once an alternative material is implemented, the automaker cannot switch it out and has to use it for the remainder of the time frame. Or alternatively, the model could limit the number of times the automaker can switch the material of each subsystem. This approach may work best if more subsystems or other applications are included in the problem; too few options, such as in the above case study, would limit the manufacturer's flexibility and force it to design to the highest CAFE standard. Another option is to include the investment the automaker makes in designing and engineering vehicles with new materials in the selection method's objective function so that each time the automaker has to redesign a subsystem from a new material, it pays an additional price. Learning can also take place in this investment cost so that the more times an automaker designs, for example, an aluminum body for any of the vehicles, the lower that body's engineering cost. This consideration of learning in engineering cost would also discourage the selection model from switching materials every design cycle.

Another assumption that may have to be altered is the genetic algorithm's approach to calculating a material's cumulative production volume. In learning theory, cumulative volume is used as a proxy for a firm's experience. The model implemented for the preceding analysis assumes that each vehicle body or closure set the automaker manufactures adds one unit to the cumulative volume of each of the materials that compose the subsystem. Consequently, an automaker gains the same amount of experience from a closure set with three aluminum reinforcements as from an entire aluminum body. This leads the model not only to favor closures for the test bed, but also to choose designs like SMC just because they contain those few aluminum parts (see Table 6.8). One way to correct for this would be to weight the learning function according to the number of individual parts manufactured from a specific material (and manufacturing process) within each subsystem. In doing so, the model would be able to more accurately represent the experience gained by producing an aluminum body versus aluminum closures.

7.2.4 Potential Analyses with the Automotive Case Study

The analyses in this study focused on the consequences of considering cost evolution in materials selection within the automotive industry and identified strategies an automaker could adopt when introducing new materials to its vehicles. The proposed selection framework, however, can be used to conduct other automotive-related analyses concerning learning and materials—or more generally, technology—selection. All the case studies presented in this thesis assumed the impact of learning is limited to a product's manufacturing cost. In reality, there are likely other aspects of the product development process in which a firm can learn—for instance, learning in product engineering. Learning is also not limited to cost: a product's performance or energy efficiency can also improve as the firm gains experience. And in some cases, performance improvements can translate to higher revenue for the firm if consumers are willing to pay for those improvements; if the firm's selection metric is based on net revenue, these improvements can potentially affect the preferred materials.

Future analyses can also compare other technologies that improve a vehicle's fuel economy to the alternative lightweight materials that were investigated in this study. Often, lightweight materials are perceived to be an enabling technology in that they are not used to directly improve fuel economy, but rather to compensate for additional mass from other technologies (e.g. hybrid powertrains) that are used to improve fuel economy. The multi-product selection framework can be tailored to evaluate this use of lightweight materials in vehicles.

In the end, the selection methods proposed in this study are not designed to select the best option for the firm, but rather to inform the selection decision and aid the firm by providing a means to systematically evaluate its material options. Since the results are context-dependent, the firm is responsible for defining the input parameters and ultimately making the final decision, which may or may not coincide with what the methods suggest to be the preferred materials.

Appendix A

Profit-Based Materials Selection

The selection methods presented in the main body of this thesis are designed to identify materials that enable a firm to satisfy design criteria at minimal manufacturing cost. While manufacturing cost is a suitable metric for materials selection, it is not necessarily the economically accurate one because firms are more likely to base their decisions on profitability rather than on cost alone. Optimizing a product's profit, however, requires more complex calculations than optimizing its manufacturing cost because profit is a function of both manufacturing cost as well as revenue the firm earns from selling that product.

Both cost and revenue, in turn, depend respectively on the production and sales volume of the product. Although unit manufacturing cost typically decreases with production volume (see Figures 3-4 and 4-4) because fixed costs are distributed over more units, total manufacturing cost will always increase with production volume because of the added variable cost required to produce the marginal unit. Revenue, on the other hand, is not guaranteed to always increase with sales volume. If a firm wishes to raise a product's price, it will likely have to expect that it will sell fewer units, since demand typically falls with a rise in price. This is captured by the equation:

$$E_P \equiv \frac{\Delta Q/Q}{\Delta P/P} \quad \Rightarrow \quad \Delta P = \frac{1}{E_P} \cdot P\frac{\Delta Q}{Q}$$

where E_P is the price elasticity of demand, and P and Q represent price and demand, respectively. With very few exceptions, $E_P < 0$, so if the firm changes its demand by ΔQ, its change in revenue, ΔP, will have the opposite sign. Whether revenue increases or decreases, though, will depend on how responsive consumer demand is to changes in that product's price, which is captured by the price elasticity of demand for that product. If demand is relatively inelastic ($|E_P| < 1$), the firm will likely be able to realize an increase in revenue because it can manufacture slightly fewer units, but sell each one for a higher price. In contrast, if demand is relatively elastic ($|E_P| > 1$), increasing product price will have a large impact on the number of units the firm can sell and may reduce overall revenue (and profit).

Consequently, identifying materials that optimize a firm's profit while still satisfying design constraints is not as simple as identifying ones that optimize the firm's manufacturing cost. The calculations become even more complex when the scope of the materials selection problem is expanded to include multiple products, particularly if the cross-price elasticities of demand are non-zero for the different products. In such cases, changes in the price of one product will affect the demand of other products.

A.1 Volume-Neutral Price

Nonetheless, profit-based analyses are still feasible within the materials selection frameworks presented in this thesis, albeit with some simplifying assumptions. This appendix illustrates one approach to conducting a profit-based analysis, specifically through the use of volume-neutral price to calculate a firm's revenue. The volume-neutral price of a product represents consumer willingness to pay for improvements to one or more of that product's attributes and thus, the extent to which a firm can alter the improved product's price without affecting that product's sales volume. Changes to both a product's attributes and its price are necessary to maintain a constant sales volume: unless consumers are completely indifferent to these changes (e.g. demand is perfectly inelastic), either alone will lead to higher or lower sales for the firm. In this study, product attributes are improved by the use of alternative materials. Although the exact nature of the improvement will be case-specific, the alternative materials, because of their different properties, will enable product performance the current material cannot. Ideally this performance will be valued by consumers, who are willing to pay for it; the firm captures this consumer value by raising product price and thus increasing its revenue—but not its sales volume.

A product's volume-neutral price depends on the changes the firm makes to one or more of that product's attributes, the elasticity of demand for those attributes, and the price elasticity of demand for that product. To begin with, the change in demand, ΔQ_A, resulting from a change in an attribute, ΔA, can be calculated from the definition of that attribute's elasticity of demand, E_A:

$$E_A \equiv \frac{\Delta Q_A / Q_o}{\Delta A / A_o} \quad \Rightarrow \quad \Delta Q_A = E_A \cdot Q_o \frac{\Delta A}{A_o} \tag{A.1}$$

Since ultimately, the firm does not want the product's volume to change, it will want to change the product's price so that the demand will change by $-\Delta Q_A$. Again, this change in price, ΔP, can be calculated using the product's price elasticity of demand, E_P:

$$\Delta P = \frac{1}{E_P} \cdot P_o \frac{-\Delta Q_A}{Q_o} \tag{A.2}$$

Volume-neutral price is therefore the initial price, P_o, plus the change in price from Equation (A.2). Substituting both Equation (A.1) and (A.2) to calculate ΔP yields the volume-neutral price in terms of the elasticities of demand, the change in attribute A, and the initial values of the product's attribute and price:

$$\begin{aligned} VNP &= P_o + \Delta P \\ &= P_o + \frac{1}{E_P} \cdot P_o \frac{-\Delta Q_A}{Q_o} \\ &= P_o - \left(\frac{1}{E_P} \cdot P_o \frac{1}{Q_o} \right) \cdot \left(E_A \cdot Q_o \frac{\Delta A}{A_o} \right) \\ &= P_o - \frac{E_A}{E_P} \cdot \frac{\Delta A}{A_o} P_o \end{aligned} \tag{A.3}$$

Consumer willingness to pay can also be calculated from the quantities above, either per unit improvement of the attribute, $\Delta P / \Delta A$, or in absolute terms, ΔP. The initial price of the product and value of the attribute, as well as the change in the attribute are all set by the firm; elasticities of demand can be estimated through market research, either using a market model or by other means. Equation (A.3) is then used to predict the change in revenue associated with using an alternative material in a product. This change in revenue is then incorporated into the selection model, along with manufacturing cost, to calculate profit. For simplicity, sales volume (revenue) is assumed to be equal to production volume (manufacturing cost) so that all products the manufacturers are eventually sold.

The following sections demonstrate the incorporation of volume-neutral price into both the single-product and the multi-product materials selection frameworks that were originally designed for optimizing manufacturing cost. For the single product framework from Chapter 3, the metric is changed to account for profit, represented by the incremental revenue a firm can realize by increasing a product's price to its volume-neutral price less that product's manufacturing cost; the material options are then ranked according to this new metric. Similar actions are taken with the multi-product selection framework so that its metric accounts for the incremental revenue of each product included in the selection problem's scope. In both cases, cost evolution through learning is still considered in the manufacturing cost calculations. Both selection methods are illustrated with their respective automotive case studies. Volume-neutral pricing is applicable to vehicles because each alternative material option reduces a vehicle's weight and thereby improves its fuel economy—something consumers are willing to pay for. Likewise, volume-neutral pricing is particularly useful to automakers because it allows the firms to modify their products and reap benefits through changes in price, but avoid altering production volumes or product-plant allocation, time-consuming and costly tasks.

A.2 Profit in Single-Product Selection

The single product case study from Chapter 3 is revisited to demonstrate the selection framework's adaptability to optimizing for profit in place of manufacturing cost. In this particular instance, the change to the methodology is straightforward: materials for the midsize body-in-white are ranked to maximize the automaker's profit rather than to minimize its manufacturing cost. Profit, in turn, is represented by the incremental revenue (ΔP) associated with each material option less that option's manufacturing cost; P_o from Equation (A.3) is not used because it is the same regardless of material and therefore will not influence the preferred option.

The same manufacturing costs as those calculated for the original case study are used in this analysis; this leaves incremental revenue, which is based on the volume-neutral price of each material option—or more precisely, consumer willingness to pay ($\Delta P/\Delta A$) for an improvement in a vehicle attribute because P_o is not included in the calculations. Since the use of any alternative material reduces the weight of the midsize car's body-in-white relative to its weight when manufactured from mild steel, the material's volume-neutral price is calculated for fuel economy improvement and thus represents the amount consumers are willing to pay for higher fuel

Table A.1: Weight savings, fuel economy improvements, and incremental revenue for each material option for the midsize body.

Strategy ID	Weight [kg]	Weight Savings [Δkg]	FC Improvement [ΔL/100km]	FE Improvement [Δmpg]	Incr. Revenue [Δ$]
MS	322	0	0.0	0.0	$0
HSS	243	79	−0.274	0.84	$79
AL	193	129	−0.445	1.38	$132
GF	219	103	−0.356	1.10	$104
CF	138	184	−0.637	2.03	$193

economy.

A.2.1 Calculating Profit

Calculating the incremental revenue of each material option based on the volume-neutral price first requires a knowledge of the extent to which each option impacts fuel economy, as well as of the demand elasticities of the product—in this case, a typical midsize car. The former can be estimated from the difference in weight between an alternative body design and a body formed from mild steel, and then translated to fuel economy improvement, either through use of a powertrain model or by relying on the 5%–10% rule of thumb. This example assumes a 6% improvement in fuel consumption for every 10% reduction in weight—same as for the vehicles in the multi-product case study (see Table 6.1). Table A.1 lists weight savings and changes in fuel consumption and in fuel economy associated with each alternative material for the midsize car's body; calculations assume an initial curb weight of 1,549 kg and an initial fuel economy of 26.36 mpg (or 8.92 L/100 km).

The elasticities of demand for a representative midsize car are estimated with a proprietary market model. This model predicts the change in the fractional market share of a particular vehicle nameplate (e.g. a Volkswagen Beetle), given an absolute or percentage change in that vehicle's attributes. It is therefore used to obtain changes in market share as a function of the vehicle's fuel economy and price; an example of the results is shown in Figure A-1. This relationship, in turn, is used to calculate consumer willingness to pay and—if P_o were known—volume-neutral price. Assuming that fuel economy and price are independent variables, the change in market

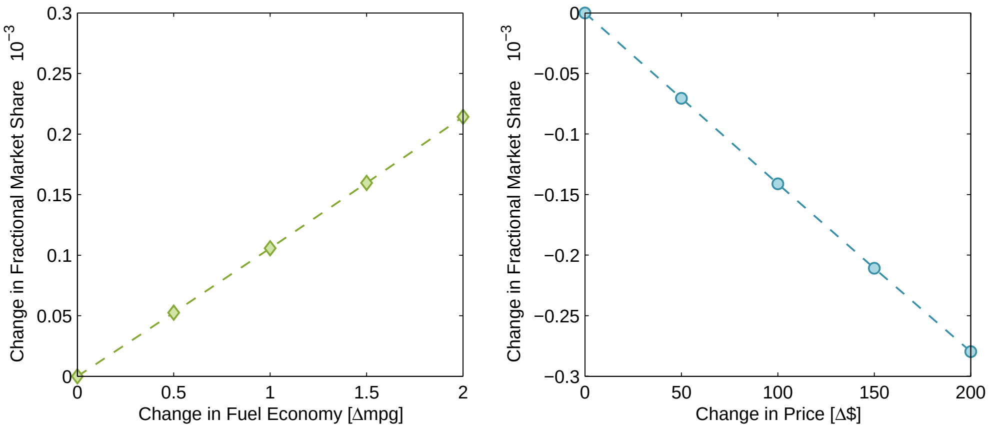

Figure A-1: Example market model results.

share as a linear function of changes in these variables can be written as

$$\Delta MktSh = m_{FE}\Delta FE + m_P \Delta P$$

where m_{FE} and m_P represent the slopes in Figure A-1 or change in market share ($\Delta MktSh$) per change in fuel economy (ΔFE) or change in price (ΔP). These slopes are related to their respective elasticities of demand, but are not the same because they account for absolute changes in fuel economy and price, and not percent changes (as is the case for elasticity). Consumer willingness to pay is therefore the change in price at which, given a change in fuel economy, the change in market share is zero:

$$\Delta MktSh = 0 \quad \Rightarrow \quad \Delta P_{WTP} = -\frac{m_{FE}}{m_P}\Delta FE$$

Calculating this for a typical midsize car yields a consumer willingness to pay of around $95 for a one-mpg increase in fuel economy[1]. This translates to an incremental revenue for the automaker between $0 and $193 for each of the body designs (Table A.1). These results are then factored into the profit metric for the single-product selection analysis.

A.2.2 Ranking Results

Table A.2 presents the results of the incremental revenue minus manufacturing cost associated with each material option. Metric values are based on the manufacturing costs from Table 3.5;

[1]Based on data collected during the spring of 2009

Table A.2: High volume unit profit (incremental revenue less unit manufacturing cost) of the midsize body-in-white.

	Mfg Cost Calculation Approach		
Strategy ID	Short-Term	Long-Term	Average
MS	−\$868	−\$868	−\$868
HSS	−\$1,061	−\$947	−\$988
AL	−\$1,949	−\$1,325	−\$1,673
GF	−\$1,786	−\$1,219	−\$1,585
CF	−\$4,196	−\$2,002	−\$3,664

Table A.3: Ranking of material options according to high volume unit profit.

Rank	Short-Term	Long-Term	Average
1	MS	MS	MS
2	HSS	HSS	HSS
3	GF	GF	GF
4	AL	AL	AL
5	CF	CF	CF

ordering of preferred materials according to profit is shown in Table A.3. From the negative numbers in Table A.2 and the resulting material order, it is clear that the incremental revenue alone is not enough to offset the added cost of using any of the alternative materials: mild steel is consistently favored regardless of cost calculation method. The results therefore indicate that, for this particular vehicle and context, consumer willingness to pay is not sufficient to affect the ordering of the material options. Nonetheless, the modified case study serves to illustrate one possible approach—the use of volume-neutral price and consumer willingness to pay—for selecting materials according to firm profitability rather than according to manufacturing cost.

A.3 Profit in Multi-Product Selection

The use of consumer willingness to pay in calculating an automaker's incremental revenue can also be incorporated into the multi-product, multi-period case study from Chapter 6. As with the midsize body-in-white case study, the decision metric is revised to account for profit—incremental revenue minus manufacturing cost—over the automaker's time horizon; only this time, incremental revenue is calculated for each of the subsystems within each vehicle platform. Since the platforms have different attributes and are targeted at different market segments, con-

Table A.4: Consumer willingness to pay (WTP) and incremental revenue for case study vehicle platforms and respective subsystems.

Platform	Compact Car	Midsize Car	Large Car
Initial FE [mpg]	26.99	26.36	23.31
WTP [Δ$ / mpg]	$77	$95	$233
Incremental Revenue: Body-in-White Material Options			
MS	$0	$0	$0
HSS	$72	$80	$171
AL	$118	$132	$279
GF	$93	$104	$221
CF	$173	$193	$408
Incremental Revenue: Closures Material Options			
MS	$0	$0	$0
HSS	$5	$6	$13
AL	$25	$29	$65
MG	$28	$33	$74
SMC	$12	$14	$31

sumers within each of these segments will respond differently to changes in fuel economy and price, and therefore be willing to pay different amounts for an increase in fuel economy. The market model is consequently needed to evaluate consumer willingness to pay for each vehicle platform within this case study.

Table A.4 shows consumer willingness to pay, as estimated by the market model, for a one-mpg fuel economy improvement in representative vehicle platforms. The calculations used to generate these revenue estimates assume negligible cross-elasticities of products. The estimates illustrate how product-specific the results can be. Incremental revenue associated with each subsystem is also shown in Table A.4; these numbers are used by the multi-product selection framework to calculate the profit.

In order to estimate the total profit of the fleet, both the integer linear program and the genetic algorithm are modified to maximize the automaker's profit rather than minimize its total manufacturing cost. For simplicity, sales volume is assumed to equal production volume so the automaker sells every car it produces. The same CAFE targets as in Table 6.5 and other inputs from Section 6.3 are used in this analysis.

Preferred materials are shown in Figure A-2, with respective profits in Table A.5. Profits are calculated using the same three cost calculation approaches used in the original presentation of this case study; highlighted material options indicate the difference between the three calculation

	Cycle	Compact Car Body	Compact Car Closures	Midsize Car Body	Midsize Car Closures	Large Car Body	Large Car Closures
Short-Term	1	MS	HSS	MS	HSS	HSS	HSS
	2	MS	HSS	MS	HSS	HSS	HSS
	3	AL	AL	AL	MG	HSS	MG
	4	AL	AL	AL	MG	HSS	MG
Long-Term	1	MS	HSS	MS	HSS	HSS	HSS
	2	MS	HSS	MS	HSS	HSS	HSS
	3	GF	AL	AL	AL	AL	MG
	4	GF	AL	AL	AL	AL	MG
Evolving	1	MS	HSS	MS	HSS	HSS	HSS
	2	MS	HSS	MS	HSS	HSS	AL
	3	AL	HSS	AL	AL	AL	AL
	4	AL	HSS	AL	AL	AL	AL

Figure A-2: Preferred materials according to profit given each approach to calculating manufacturing cost.

Table A.5: Profit of selection decisions presented in Figure A-2.

	Cost Evaluation Approach		
Select according to...	Short-Term	Long-Term	Evolving
Short-Term Cost	**-$7.142 B**	-$5.734 B	-$5.954 B
Long-Term Cost	-$7.343 B	**-$5.732 B**	-$6.135 B
Evolving Cost	-$7.351 B	-$5.777 B	**-$5.913 B**

approaches. A comparison of Figures 6-7 and A-2 indicates that optimizing total profit can lead to different preferred materials compared to optimizing total manufacturing cost. One of the key changes is the use of high-strength steel in place of mild steel, even when the former is not necessitated by CAFE. This is because the incremental revenue due to consumer willingness to pay more than compensates for the additional cost of using high-strength steel; consequently, high-strength steel is preferred over mild steel for the body of the large car and the closures of all three vehicle platforms. Mild steel, though, is still preferred for the compact and midsize bodies because consumers who purchase their respective cars are not as willing to pay for improvements in fuel economy as consumers who purchase large cars.

Beyond the use of high-strength steel in the first and second design cycles, the short-term cost (i.e. top) selection decision in Figure A-2 does not differ from the one in Figure 6-7. The long-term decision, however, indicates a preference for aluminum over high-strength steel for the large car's body. The cost of this material—in comparison to that of high-strength steel—is offset, not only by the high consumer willingness to pay for the large car, but also by its higher fuel economy improvement which enables the automaker to choose less expensive materials for the compact car. Specifically, glass fiber composite and aluminum are, respectively, selected for the compact car's body and closures when profit is considered (Figure A-2), whereas aluminum and magnesium—both more expensive, but associated with greater impact on the vehicle's fuel economy—are selected for those same parts when only cost matters (Figure 6-7).

The original formulation of the case study in Chapter 6 showed that the consideration of cost evolution in the selection process can lead to the use of a car's closures as a test bed for aluminum. This still holds true when profit is substituted for the metric—only this time, the closures for the large car are chosen in place of those for the compact car. The choice of the large car as the test bed is, again, due to the higher willingness to pay for fuel economy of consumers who purchase that car: this willingness to pay enables the firm to offset some of the added cost of using the closures a test bed. Consumer willingness to pay, combined with the lowered cost of aluminum due to the use of a test bed, also leads the automaker to select aluminum for the large car's body in the third and fourth design cycles. This use of aluminum in turn enables the automaker to "downgrade" from magnesium closures in the compact and large cars (Figure 6-7) to the less expensive high-strength steel and aluminum closures, respectively.

In conclusion, the above analysis shows that accounting for revenue—specifically, incremental revenue due to consumer willingness to pay for the fuel economy improvements associated with the use of alternative, lightweight materials in vehicle design—is feasible within the multi-product, multi-period materials selection framework. The results from Figure A-2 indicate that, when optimizing for profit, the automaker continues to use a test bed to deliberately learn and lower manufacturing costs, but chooses a different subsystem—closures for the large car—because the incremental revenue enables it to partially offset the added cost of employing a test bed. This incremental revenue depends on consumer willingness to pay and is a function of the product's price and fuel economy elasticities of demand (Equation (A.3)), which represent consumer response to changes in both these vehicle attributes. Thus, the results confirm the statement made in Section 4.3.3 that a firm will select the product whose consumers are most willing to pay for improvements and try to capture that consumer value to offset the test bed's cost. The selection results also support the earlier observation that when learning is considered, the automaker will prefer to emphasize the use of a single material rather than introducing different materials and be forced to start from scratch with each one. Regardless, the use of consumer willingness to pay and volume-neutral price is just one approach to incorporating profit into the materials selection process, but one that is clearly feasible and can provide additional information to a firm faced with having to select new materials for its products.

Appendix B

Integer Linear Programming Code

The following code was written for LINGO 11.0 linear programming software. In order to improve the model's performance, preliminary calculations are run in an Excel workbook and their results read by the model.

```
MODEL:
 TITLE MATLS_SELN;

!*** No learning                  ***;
!*** New vehicle designed every period   ***;
!***  (with corresponding tooling investments)  ***;


 SETS:
  TECHS: PD_AVAIL, D_CAFC, NPV_PREM, NUM_CONFL, NUM_REQD;
  PDS: CAFC_TARGET, DISCOUNTED, DEFICIT;

  TECH_APPL (TECHS, PDS): X;

  TECH_RELNS (TECHS, TECHS): CONFLICTS, REQUIRED;
  EXCHANGE (PDS, PDS): CREDITS;
 ENDSETS


 DATA:
  TECHS, PD_AVAIL, D_CAFC = @OLE('matl_data.xlsx', 'TECHS', 'PD_AVAIL', 'D_CAFC');
  NPV_PREM = @OLE('matl_data.xlsx', 'NPV_PREM');
```

```
  PDS, CAFC_TARGET, DISCOUNTED = @OLE('matl_data.xlsx', 'PDS', 'CAFC_TARGET', 'DISCOUNTED');
  CONFLICTS, REQUIRED = @OLE('matl_data.xlsx', 'CONFLICTS', 'REQUIRED');

  VEH_CAFC = @OLE('matl_data.xlsx', 'VEH_CAFC');   ! initial fleet CAFC number ;
  FLEET_VOL = @OLE('matl_data.xlsx', 'FLEET_VOL');  ! total fleet volume ;

  A_NUMBER = 500;
 ENDDATA


 CALC:
  !*** Count number of conflicting or required technologies ***;
  @FOR(TECHS(T):
     NUM_CONFL(T)  = @SUM(TECHS(U): CONFLICTS(T,U));
     NUM_REQD(T)  = @SUM(TECHS(U): REQUIRED(T,U));
  );
 ENDCALC


!*** Objective function ***;
MAX = @SUM(PDS(P):
      DISCOUNTED(P) * @SUM(TECHS(T): X(T,P) * NPV_PREM(T)));


!*** Decision variable constraints ***;
@FOR(TECH_APPL(T,P):
   @BIN(X(T,P));
);


!*** CAFE target constraint ***;
@FOR(PDS(P):
   VEH_CAFC + @SUM(TECHS(T): X(T,P) * D_CAFC(T) / FLEET_VOL) ≤CAFC_TARGET(P)
);


!*** Technology implementation constraints ***;
@FOR(TECH_APPL(T,P):
   NUM_CONFL(T) * X(T,P) + @SUM(TECHS(U): CONFLICTS(T,U) * X(U,P)) ≤NUM_CONFL(T); ! tech conflicts ;
   NUM_REQD(T) * X(T,P) - @SUM(TECHS(U): REQUIRED(T,U) * X(U,P)) ≤0;   ! tech symm requirements ;
);

```

```
!*** Period availability constraints ***;
@FOR(TECHS(T):
   @FOR(PDS(P) | P #LT# PD_AVAIL(T):
      X(T,P) = 0;
   );
);


 DATA:
   @OLE('matl_data.xlsx', 'X_OUT') = X;
 ENDDATA
```

Listing B.1: Integer linear program code for automotive case study.

Appendix C

Genetic Algorithm Code

The following set of code is used to run a genetic algorithm in MATLAB. It consists of 3 files: `readFiles.m` gathers and processes the data from an Excel file, `fitness.m` evaluates the fitness of each candidate solution or *chromosome*, and `matlSeln.m` initializes and runs the genetic algorithm. Only the code used to set up and run the optimization process, and the fitness function are shown below. The GA code itself is part of MATLAB's Global Optimization toolbox.

```
% Single-objective GA with all objectives combined into one value
% Read inputs from Excel file
clear all
readFiles16;                              % read data from Excel file and generate inputs

% ===========================================================================
% Linear program seed values

ilpST = [0 0 0 0 0 0 2 2 0 0 2 2 0 0 1 1 0 0 2 2 0 0 3 3 0 0 3 3];
ilpLT = [0 0 0 0 0 0 2 2 0 0 2 2 0 0 1 1 0 0 3 3 0 0 2 2 0 0 3 3];
gaOut = [0 0 0 0 0 0 2 2 0 0 2 2 0 0 1 1 0 2 3 3 0 0 2 2 0 0 3 3];

% ===========================================================================
% Initialize GA variables and options

nvars = apps * pds;                       % chromosome (genome) length
popSize = 8 * nvars;                      % number of chromosomes in the population
last = 100;                               % number of generations
replace = 0.6;                            % fraction of chromosomes replaced
```

```
limits = [zeros(1, nvars); maxGeneVal];   % boundaries on each gene

options = gaoptimset();
options = gaoptimset(options, 'PopulationType', 'custom');
options = gaoptimset(options, 'CreationFcn', @chrCreation);
options = gaoptimset(options, 'MutationFcn', @chrMutation);
options = gaoptimset(options, 'CrossoverFcn', @crossovertwopoint);
options = gaoptimset(options, 'SelectionFcn', @selectiontournament);
options = gaoptimset(options, 'PopulationSize', popSize);
options = gaoptimset(options, 'CrossoverFraction', replace);
options = gaoptimset(options, 'Generations', last);
options = gaoptimset(options, 'PopInitRange', limits);
options = gaoptimset(options, 'InitialPopulation', [ilpST; ilpLT]);

% ============================================================================
% Running the GA

fprintf('Running GA: ');
tic
for j = 1:5
  [x fval] = ga(@fitness16, nvars, [], [], [], [], limits(1,:), limits(2,:), [], options);
  optChoice(j,:) = +x;
  netRev(j) = -fval;
% options = gaoptimset(options, 'InitialPopulation', optChoice(j,:));
end
toc

% ============================================================================
% Convert optChoice into something more useful

[maxVal maxInd] = max(netRev);
maxCnt = sum(netRev == maxVal);
maxX = optChoice(maxInd,:);

for a = 1:apps
  maxChr(a,:) = maxX( ((a-1) * pds + 1):(pds*a) );
end
```

Listing C.1: `matlsSeln.m` – Code to read in materials data and initialize and run genetic algorithm for automotive case study.

```
function tf = fitness16(chr)

  % ********************************************************************************
  % Global variables

  global pds yrs apps yrsPerPd btAll discount
  global btAppGrP btAppSubP learnData
  global initCAFCVol cafcTargetVol btDelCAFC cafcUnitPenalty availUnitPenalty
  global btVol btPdAvail btToolsPerYr btLearn btDedInvLife
  global btAsmInv btDedInvCost btDelRev btEffVarUnit btAsmBldg


  % ********************************************************************************
  % convert chromosome into a more useful form

  % initialize gene
  btGene = zeros(btAll, pds);

  % breakdown chromosome into individual materials
  for a = 1:apps
    appUse = chr( ((a-1) * pds + 1):(pds*a) ); % (row) break down chr into [apps x 1 pd] array
    btGene = btGene | ((btAppGrP == a) & (btAppSubP == repmat(appUse, btAll, 1)));
  end


  % ********************************************************************************
  % initialize variables
  volUsed = btVol .* btGene;                   % (mtx) volume at (t,p) is non-zero if tech used
  cumVol = zeros(btAll, 1);                    % (col) cummulative volume for variable cost learning
  asmCost = zeros(1, yrs);                     % (row) assembly equipment cost

  % check if CAFE standard met; need cafcMet(p) ≤0
  cafcChange = sum(btDelCAFC .* btGene);       % (mtx) effect on CAFC at (t,p) if tech used
  cafcDiffPd = initCAFCVol + cafcChange - cafcTargetVol; % (row) CAFC shortfall (or surplus) for each period


  % ********************************************************************************
  % calculate objectives for each period
  for p = 1:pds

    % check for early implementation
    implEarly(:,p) = (btPdAvail > p);

```

```
% years encompassed by period
yMin = (p - 1) * yrsPerPd + 1;
yMax = p * yrsPerPd;

% assembly cost at the beginning of each period
asmCost(yMin) = btAsmInv(p,:) * btGene(:,p);

% dedicated (tooling) investment costs
pdTooling = btGene(:,p) .* btToolsPerYr(:,p); % (col) purchase new tools at beginning of p
toolLifeVol = zeros(btAll, 1);

for y = yMin:yMax

  % CAFC difference for each year
  cafcDiff(y) = cafcDiffPd(p);

  % variable and non-dedicated investment cost of baseline and implemented technologies
  nextCumVol = cumVol + learnData * volUsed(:,p);

  % calculate normalized learning curve cost at current, "midway," and future cumulative volumes
  cLo = sqrt(1 - btLearn(:,1)) .^ max(0, log2( min( cumVol, btLearn(:,3)) ./ btLearn(:,2) ));
  cMd = sqrt(1 - btLearn(:,1)) .^ max(0, log2( min((cumVol + nextCumVol) / 2, ...
                                             btLearn(:,3)) ./ btLearn(:,2) ));
  cHi = sqrt(1 - btLearn(:,1)) .^ max(0, log2( min(nextCumVol, btLearn(:,3)) ./ btLearn(:,2) ));

  % estimate average cost based on normalized learning curve results
  avg2 = (cMd ≠1) .* (cMd + cHi) / 2 + (cMd == 1) .* (cMd + (cMd + cHi) / 2) / 2;
  avg1 = (cMd ≠(1 - btLearn(:,1))) ...
            .* (cLo + cMd) / 2 + (cMd == (1 - btLearn(:,1))) .* ((cLo + cMd) / 2 + cMd) / 2;

  avg1 = (cLo ≠1) .* avg1 + (cLo == 1) .* (cLo + (cLo + cMd) / 2) / 2;
  avg2 = (cHi ≠(1 - btLearn(:,1))) ...
            .* avg2 + (cHi == (1 - btLearn(:,1))) .* ((cMd + cHi) / 2 + cHi) / 2;

  % scale normalized curve
  learnedUnit(:,y) = btEffVarUnit .* (avg1 + avg2) / 2;

  effVarCost(y) = (learnedUnit(:,y)') * volUsed(:,p) + btAsmBldg(p,:) * btGene(:,p);
  cumVol = nextCumVol;

  % replace tools that wear out
  replTooling = (toolLifeVol > (btDedInvLife .* pdTooling)) .* pdTooling;
  dedInv(y) = (pdTooling * (y == yMin) + replTooling)' * btDedInvCost;
```

```
    toolLifeVol = toolLifeVol .* (replTooling ≤0) + volUsed(:,p);

    % total cost and revenue for the year
    totalRev(y) = btDelRev * volUsed(:,p);
   end
  end


  % ********************************************************************************
  % net revenue for technologies
  netRev = (totalRev - (effVarCost + asmCost + dedInv)) * discount;

  % penalties
  availPenalty = sum(sum(implEarly .* volUsed )) * availUnitPenalty;
  cafcPenalty = max(cafcDiff, 0) * discount * cafcUnitPenalty;

  tf = -(netRev - cafcPenalty - availPenalty);

end
```

Listing C.2: `fitness.m` – Fitness function for genetic algorithm.

Appendix D

Automotive Case Study Inputs

This appendix contains body-in-white and closure material inputs for the compact, midsize, and large cars in the multi-product selection case study. Inputs from Table 6.3 are included for consistency. All cost numbers were generated with the process-based cost models (discussed in Section 3.2.1) and assume long-term cost; corresponding numbers for learning scope and rate can be found in Table 6.4.

Table D.1: Material options for compact car body designs.

ID	Mild steel (MS)	HS steel (HSS)	Aluminum (AL)			Glass Fiber (GF)		Carbon Fiber (CF)	
Material	Mild steel	HS steel	Al	Al	Al	GF composite	Mild steel	CF composite	Mild steel†
Mass	304 kg	228 kg	58 kg	45 kg	80 kg	179 kg	28 kg	102 kg	20 kg
Process	Stamping	Stamping	Stamping	Die casting	Extrusion	SRIM	Stamping	SRIM	Stamping
Learning scope	Low	Low	Medium	Medium	Medium	Medium	Low	High	Low
Learning rate	Slow	Slow	Fast	Slow	Slow	Slow	Slow	Slow	Slow
Forming Costs									
Variable	$452	$687	$212	$144	$243	$790	$40	$1,720	$40
Non-ded. investment	$142	$99	$67	$137	$205	$172	$11	$176	$11
Dedicated investment	$68.8 M	$48.1 M	$32.8 M	$5.0 M	$420,000	$12.8 M	$1.6 M	$12.8 M	$1.6 M
Allocation fraction	3.2%	8.3%	3.8%	79.6%	26.4%	40.00%	0.1%	40.0%	0.1%
Investment life	500 M	500 M	500 M	120,000	36,400	165,000	500 M	165,000	500 M
Assembly Costs @ 200k									
Variable	$132	$97	$163			$58		$58	
Equipment	$49.8 M	$50.2 M	$75.0 M			$24.3 M		$24.3 M	
Tooling	$18.8 M	$17.9 M	$26.9 M			$13.1 M		$13.1 M	
Building	$16.8 M	$15.3 M	$24.2 M			$9.7 M		$9.7 M	

† Mild steel costs were generated assuming a weight of 28 kg for the compact car.

Table D.2: Material options for compact car closure designs.

ID	Mild steel (MS)	HS steel (HSS)	Aluminum (AL)		Magnesium - Aluminum (MG)			Sheet molding compound (SMC)		
Material	Mild steel	HS steel	Al	Al	Al	Mg	Al, Mg	Mild steel	Al	SMC
Mass	66 kg	61 kg	36 kg	3 kg	19 kg	11 kg	6 kg	15 kg	3 kg	35 kg
Process	Stamping	Stamping	Stamping	Extrusion	Stamping	Warm forming	Extrusion	Stamping	Extrusion	SMC
Learning scope	Low	Low	Medium	Medium	Medium	High	Medium	Low	Medium	Medium
Learning rate	Slow	Slow	Fast	Slow	Fast	Fast	Slow	Slow	Slow	Slow
Forming Costs										
Variable	$97	$97	$180	$9	$98	$84	$21	$23	$10	$111
Non-ded. investment	$28	$30	$34	$5	$17	$18	$9	$4	$6	$24
Dedicated investment	$14.4 M	$15.9 M	$19.3 M	$80,000	$9.6 M	$7.7 M	$340,000	$520,000	$80,000	$2.3 M
Allocation fraction	2.8%	2.8%	4.2%	7.5%	5.2%	18.9%	2.5%	2.2%	7.9%	47.2%
Investment life	500 M	500 M	500 M	260,000	500 M	500 M	280,000	500 M	260,000	1.2 M
Assembly Costs @ 200k										
Variable	$17	$17	$24		$39			$26		
Equipment	$4.8 M	$4.8 M	$8.5 M		$8.6 M			$5.4 M		
Tooling	$3.7 M	$3.7 M	$4.4 M		$4.4 M			$9.1 M		
Building	$4.2 M	$4.2 M	$7.1 M		$4.8 M			$4.7 M		

Table D.3: Material options for midsize car body designs.

ID	Mild steel (MS)	HS steel (HSS)	Aluminum (AL)			Glass Fiber (GF)		Carbon Fiber (CF)	
Material	Mild steel	HS steel	Al	Al	Al	GF composite	Mild steel	CF composite	Mild steel†
Mass	322 kg	242 kg	61 kg	47 kg	85 kg	189 kg	30 kg	108 kg	21 kg
Process	Stamping	Stamping	Stamping	Die casting	Extrusion	SRIM	Stamping	SRIM	Stamping
Learning scope	Low	Low	Medium	Medium	Medium	Medium	Low	High	Low
Learning rate	Slow	Slow	Fast	Slow	Slow	Slow	Slow	Slow	Slow
Forming Costs									
Variable	$477	$700	$223	$151	$267	$900	$42	$1,181	$42
Non-ded. investment	$144	$100	$67	$139	$226	$191	$11	$182	$11
Dedicated investment	$70.3 M	$49.0 M	$33.4 M	$5.0 M	$420,000	$13.2 M	$1.6 M	$13.2 M	$1.6 M
Allocation fraction	3.2%	8.3%	3.8%	81%	29%	40%	0%	40%	0%
Investment life	500 M	500 M	500 M	120,000	34,800	165,000	500 M	165,000	500 M
Assembly Costs @ 170k									
Variable	$135	$97		$166		$58		$58	
Equipment	$41.2 M	$41.5 M		$61.1 M		$17.1 M		$17.1 M	
Tooling	$15.8 M	$14.9 M		$22.1 M		$8.92 M		$8.92 M	
Building	$14.2 M	$12.8 M		$19.8 M		$6.62 M		$6.62 M	

† Mild steel costs were generated assuming a weight of 30 kg for the midsize car.

Table D.4: Material options for midsize car closure designs.

ID	Mild steel (MS)	HS steel (HSS)	Aluminum (AL)		Magnesium - Aluminum (MG)			Sheet molding compound (SMC)		
Material	Mild steel	HS steel	Al	Al	Al	Mg	Al, Mg	Mild steel	Al	SMC
Mass	72 kg	67 kg	40 kg	3 kg	20 kg	12 kg	7 kg	17 kg	4 kg	38 kg
Process	Stamping	Stamping	Stamping	Extrusion	Stamping	Warm forming	Extrusion	Stamping	Extrusion	SMC
Learning scope	Low	Low	Medium	Medium	Medium	High	Medium	Low	Medium	Medium
Learning rate	Slow	Slow	Fast	Slow	Fast	Fast	Slow	Slow	Slow	Slow
Forming Costs										
Variable	$106	$106	$197	$10	$107	$92	$24	$25	$11	$121
Non-ded. investment	$30	$32	$35	$6	$18	$21	$9	$4	$6	$27
Dedicated investment	$14.9 M	$16.4 M	$19.9 M	$80,000	$9.9 M	$7.9 M	$340,000	$540,000	$80,000	$2.4 M
Allocation fraction	2.8%	2.8%	4.2%	8.2%	5.2%	18.9%	2.8%	2.2%	8.7%	47.2%
Investment life	500 M	500 M	500 M	260,000	500 M	500 M	280,000	500 M	260,000	1.2 M
Assembly Costs @ 170k										
Variable	$17	$17	$25			$40			$26	
Equipment	$4.7 M	$4.7 M	$7.0 M			$6.4 M			$5.0 M	
Tooling	$3.8 M	$3.8 M	$3.6 M			$3.2 M			$8.0 M	
Building	$4.5 M	$4.5 M	$5.5 M			$4.4 M			$4.5 M	

Table D.5: Material options for large car body designs.

ID	Mild steel (MS)	HS steel (HSS)	Aluminum (AL)			Glass Fiber (GF)		Carbon Fiber (CF)	
Material	Mild steel	HS steel	Al	Al	Al	GF composite	Mild steel	CF composite	Mild steel†
Mass	328 kg	246 kg	63 kg	48 kg	86 kg	193 kg	30 kg	110 kg	21 kg
Process	Stamping	Stamping	Stamping	Die casting	Extrusion	SRIM	Stamping	SRIM	Stamping
Learning scope	Low	Low	Medium	Medium	Medium	Medium	Low	High	Low
Learning rate	Slow	Slow	Fast	Slow	Slow	Slow	Slow	Slow	Slow
Forming Costs									
Variable	$485	$704	$227	$153	$274	$912	$42	$1,841	$42
Non-ded. investment	$144	$101	$67	$140	$232	$193	$11	$184	$11
Dedicated investment	$70.7 M	$49.3 M	$33.6 M	$5.0 M	$420,000	$13.4 M	$1.6 M	$13.4 M	$1.6 M
Allocation fraction	3.2%	8.3%	3.8%	82.0%	30.0%	40.0%	0.1%	40.0%	0.1%
Investment life	500 M	500 M	500 M	120,000	34,800	165,000	500 M	165,000	500 M
Assembly Costs @ 120k									
Variable	$135	$97		$166		$58		$58	
Equipment	$32.8 M	$33.8 M		$46.2 M		$16.1 M		$16.1 M	
Tooling	$12.9 M	$12.0 M		$16.8 M		$8.9 M		$8.9 M	
Building	$11.6 M	$10.4 M		$15.1 M		$6.6 M		$6.6 M	

† Mild steel costs were generated assuming a weight of 30 kg for the large car.

Table D.6: Material options for large car closure designs.

ID	Mild steel (MS)	HS steel (HSS)	Aluminum (AL)		Magnesium - Aluminum (MG)			Sheet molding compound (SMC)		
Material	Mild steel	HS steel	Al	Al	Al	Mg	Al, Mg	Mild steel	Al	SMC
Mass	78 kg	72 kg	43 kg	4 kg	22 kg	13 kg	7 kg	18 kg	4 kg	41 kg
Process	Stamping	Stamping	Stamping	Extrusion	Stamping	Warm forming	Extrusion	Stamping	Extrusion	SMC
Learning scope	Low	Low	Medium	Medium	Medium	High	Medium	Low	Medium	Medium
Learning rate	Slow	Slow	Fast	Slow	Fast	Fast	Slow	Slow	Slow	Slow
Forming Costs										
Variable	$114	$113	$211	$11	$115	$99	$25	$27	$11	$129
Non-ded. investment	$30	$32	$35	$7	$18	$21	$10	$4	$7	$26
Dedicated investment	$15.3 M	$16.8 M	$20.4 M	$80,000	$10.2 M	$8.1 M	$340,000	$550,000	$80,000	$2.5 M
Allocation fraction	2.8%	2.8%	4.2%	9.0%	5.2%	18.9%	3.0%	2.2%	9.5%	47.2%
Investment life	500 M	500 M	500 M	260,000	500 M	500 M	280,000	500 M	260,000	1.2 M
Assembly Costs @ 120k										
Variable	$19	$19	$26			$45			$28	
Equipment	$4.2 M	$4.2 M	$6.0 M			$5.4 M			$4.8 M	
Tooling	$3.4 M	$3.4 M	$3.3 M			$2.8 M			$8.0 M	
Building	$3.9 M	$3.9 M	$5.0 M			$3.8 M			$4.5 M	

Bibliography

[1] *Center for Biological Diversity v. NHTSA*. 06-71891 (9th Cir. 2007).

[2] MEPS – World Carbon Steel Prices. Available from: http://www.meps.co.uk/WorldCarbonPrice.htm; Accessed: Sep. 2009.

[3] Average Fuel Economy Standards for Light Trucks Model Years 2008–2011; Final Rule. *Federal Register*, 71(6):17566–17679, 6 April 2009. Section IV.F.

[4] Average Fuel Economy Standards Passenger Cars and Light Trucks Model Year 2011; Final Rule. *Federal Register*, 74(59):14196–14456, 30 March 2009. Section IV.C.1.

[5] Corporate average fuel economy for MY2012–MY2016 passenger cars and light trucks (preliminary impact regulatory analysis). Technical report, National Highway Traffic Safety Administration, August 2009. Available from: http://www.nhtsa.dot.gov/portal/fueleconomy.jsp; Accessed: Nov. 2009.

[6] Summary of fuel economy performance, 9 December 2009. Available from: http://www.nhtsa.dot.gov/; Accessed: Apr. 2010.

[7] Weekly prices. *Platt's Metals Week*, 80(19):22, 11 May 2009.

[8] CAFE overview, 2010. Available from: http://www.nhtsa.dot.gov/; Accessed: Apr. 2010.

[9] Light-Duty Vehicle Greenhouse Gas Emission Standards and Corporate Average Fuel Economy Standards; Final Rule. *Federal Register*, 75(88):25324–25728, 7 May 2010.

[10] Paul S. Adler. Shared learning. *Management Science*, 36(8):938–957, August 1990.

[11] Environmental Protection Agency. Light-duty automotive technology, carbon dioxide emissions, and fuel economy trends: 1975 through 2009. Technical Report EPA420-R-09-014, November 2008. Available from: http://www.epa.gov/otaq/fetrends.htm; Accessed: Apr. 2010.

[12] R. Amen and P. Vomacka. Case-based reasoning as a tool for materials selection. *Materials & Design*, 22(5):353–358, 2001.

[13] Linda Argote and Dennis Epple. Learning curves in manufacturing. *Science*, 247(4945):920–924, 23 February 1990.

[14] M. F. Ashby, Y. J. M. Brechet, D. Cebon, and L. Salvo. Selection strategies for materials and processes. *Materials & Design*, 25(1):51–67, 2004.

[15] Mike Ashby. The CES EduPack; Unit 4. Ranking: material indices, 2010. Available from: `http://www.grantadesign.com/userarea/teachingresource/`.

[16] H. Assler and J. Telgkamp. Design of aircraft structures under special consideration of NDT. In *European Conference on Non-Destructive Testing*, Berlin, September 2006.

[17] Dale Brosius. Carbon fiber raises consumer performance index. *High Performance Composites*, 13(1):20–25, January 2005. `http://www.compositesworld.com/articles/carbon-fiber-raises-consumer-performance-index.aspx`.

[18] John G. Carlson. How management can use the improvement phenomenon. *California Management Review*, 3(2):83–94, 1961.

[19] Joseph W.K. Chan and Thomas K.L. Tong. Multi-criteria material selections and end-of-life product strategy: Grey relational analysis approach. *Materials & Design*, 28(5):1539–1546, 2007.

[20] Shyi-Ming Chen. A new method for tool steel materials selection under fuzzy environment. *Fuzzy Sets and Systems*, 92(3):265–274, 1997.

[21] National Research Council, editor. *Materials Science and Engineering for the 1990s: Maintaining Competitiveness in the Age of Materials*. National Academy Press, Washington, D.C., 1989.

[22] George S. Day and David B. Montgomery. Diagnosing the experience curve. *The Journal of Marketing*, 47(2):44–58, 1983.

[23] John DeCicco, Feng An, and Marc Ross. Technical options for improving the fuel economy of U.S. cars and light trucks by 2010-2015. Technical report, American Council for an Energy-Efficient Economy, 2001.

[24] John Decicco and Marc Ross. Recent advances in automotive technology and the cost-effectiveness of fuel economy improvement. *Transportation Research Part D: Transport and Environment*, 1(2):79–96, 1996.

[25] L. Dudley. Learning and productivity change in metal products. *The American Economic Review*, 62(4):662–669, September 1972.

[26] Goutam Dutta, Robert Fourer, Akhilesh Majumdar, and Debabrata Dutta. An optimization-based decision support system for strategic planning in a process industry: The case of a pharmaceutical company in india. *International Journal of Production Economics*, 106(1):92–103, March 2007. Special section on contextualisation of supply chain networks, 9th International Symposium in Logistics.

[27] Dennis Epple, Linda Argote, and Rukmini Devadas. Organizational learning curves: A method for investigating intra-plant transfer of knowledge acquired through learning by doing. *Organization Science*, 2(1):58–70, March 1991.

[28] Daniel J. Evans and John Heinz. Reject the plea from Ford, G.M. *New York Times*, page 1.23, 20 April 1985.

[29] Frank Field, Randolph Kirchain, and Richard Roth. Process cost modeling: Strategic engineering and economic evaluation of materials technologies. *JOM Journal of the Minerals, Metals and Materials Society*, 59(10):21–32, October 2007.

[30] Center for Biological Diversity. Lawsuit challenges Obama fuel economy standards. Press Release, 2 April 2009. Available from: `http://www.biologicaldiversity.org/news/press_releases/2009/cafe-standards-04-02-2009.html`; Accessed: Apr. 2010.

[31] E. Forster, S. Clay, R. Holzwarth, D. Partt, and D. Paul. Flight vehicle composite structures. In *26th Congress of International Council of the Aeronautical Sciences, Anchorage, AK*, number AIAA 2008-8976. American Institute of Aeronautics and Astronautics, 14–19 September 2008.

[32] Gordon R. Foxall and Brian R. Johnston. Innovation in grand prix motor racing: the evolution of technology, organization and strategy. *Technovation*, 11(7):387–402, November 1991.

[33] Alexandra Frangi. Lightweight body designs as enablers for alternative powertrain technologies: Understanding cost and environmental performance tradeoffs. Master's thesis, Massachusetts Institute of Technology, June 2001.

[34] H. Friedrich and S. Schumann. Research for a "new age of magnesium" in the automotive industry. *Journal of Materials Processing Technology*, 117(3):276–281, 23 November 2001.

[35] Erica R. H. Fuchs. The significance of production cost inputs in regional technology choice: Composite automotive body-in-whites in the us versus china. Master's thesis, Massachusetts Institute of Technology, June 2003.

[36] Anshuman Gupta and Costas D. Maranas. Market-based pollution abatement strategies: Risk management using emission option contracts. *Industrial & Engineering Chemistry Research*, 42(4):802–810, 2003.

[37] Nile W. Hatch and David C. Mowery. Process innovation and learning by doing in semiconductor manufacturing. *Management Science*, 44(11):1461–1477, November 1998.

[38] Darrell R. Herling. Cost assessment of emerging magnesium sheet production methods. In Edward J. Wall, Rogelio Sullivan, and Joseph Carpenter, editors, *Progress Report for Automotive Lightweighting Materials*, section 2J, pages 85–94. August 2006. Available from: `http://www1.eere.energy.gov/vehiclesandfuels/resources/fcvt_alm_fy05.html`.

[39] U.S. White House. President Obama announces national fuel efficiency policy. Press Release, 19 May 2009. Available from: `http://www.whitehouse.gov/the_press_office/President-Obama-Announces-National-Fuel-Efficiency-Policy/`; Accessed Apr. 2010.

[40] Frank R. Field III. *Application of Multi-Attribute Utility Analysis to Problems in Materials Selection*. PhD thesis, Massachusetts Institute of Technology, 1985.

[41] Douglas A. Irwin and Peter J. Klenow. Learning-by-doing spillovers in the semiconductor industry. *The Journal of Political Economy*, 102(6):1200–1227, December 1994.

[42] Dong-Hyun Jee and Ki-Ju Kang. A method for optimal material selection aided with decision making theory. *Materials & Design*, 21(3):199–206, 2000.

[43] Ashish M. Kar. A cost modeling approach using learning curves to study the evolution of technology. Master's thesis, Massachusetts Institute of Technology, June 2007.

[44] Ashish S. Kelkar. Analysis of aluminum in auto body designs and its strategic implications for the aluminum industry. Master's thesis, Massachusetts Institute of Technology, February 2001.

[45] Aida Khajavirad, Jeremy J. Michalek, and Timothy W. Simpson. An efficient decomposed multiobjective genetic algorithm for solving the joint product platform selection and product family design problem with generalized commonality. *Structural and Multidisciplinary Optimization*, 39(2):187–201, August 2008.

[46] Michelle Rene Kirby. *A Methodology for Technology Identification, Evaluation, and Selection in Conceptual and Preliminary Aircraft Design*. PhD thesis, Georgia Institute of Technology, March 2001.

[47] Randolph E. Kirchain and Frank R. Field III. *Cost modeling of Materials and Manufacturing Processes*, volume 2 of *Encyclopedia of Materials: Science and Technology*, pages 1718–1727. Elsevier, New York, 2001.

[48] Jussi Lahtinen, Petri Myllyäki, Tomi Silander, and Henry Tirri. Empirical comparison of stochastic algorithms. In Jarmo T. Alander, editor, *Proceedings of the Second Nordic Workshop on Genetic Algorithms and their Applications*, pages 45–59, Vaasa, Finland, August 1996. Available from: `http://cosco.hiit.fi/publications.html`.

[49] Shanling Li and Devanath Tirupati. Dynamic capacity expansion problem with multiple products: Technology selection and timing of capacity additions. *Operations Research*, 42(5):958–976, 1994.

[50] Marvin B. Lieberman. The learning curve and pricing in the chemical processing industries. *The Rand Journal of Economics*, 5(2):213–228, 1984.

[51] Franco Malerba. Learning by firms and incremental technical change. *The Economic Journal*, 102(413):845–859, July 1992.

[52] A. McDonald and L. Schrattenholzer. Learning curves and technology assessment. *International Journal of Technology Management*, 23(7):718–745, 2002.

[53] W. S. Miller, L. Zhuang, J. Bottema, A. J. Wittebrood, P. De Smet, A. Haszler, and A. Vieregge. Recent development in aluminium alloys for the automotive industry. *Materials Science & Engineering A*, 280(1):37–49, 15 March 2000.

[54] Melanie Mitchell. *An Introduction to Genetic Algorithms*. The MIT Press, Cambridge, 1998.

[55] Edward A. Muccio. *Plastic Part Technology*. ASM International, Materials Park, OH, 1997.

[56] Marie-Claude Nadeau. Evaluating manufacturing flexibility driven by learning. Master's thesis, Massachusetts Institute of Technology, June 2009.

[57] Gregory F. Nemet. Beyond the learning curve: factors influencing cost reductions in photovoltaics. *Energy Policy*, 34(17):3218–3232, November 2006.

[58] Committee on the Survey of Materials Science and Engineering, editors. *Materials and Man's Needs*. National Academy of Sciences, Washington, D.C., 1974.

[59] J. H. Owen and S. P. Marin. A multi-period optimization model for CAFE technology selection and credit allocation, 2004. General Motors Internal Report.

[60] S. Pagerit, P. Sharer, and A. Rousseau. Fuel economy sensitivity to vehicle mass for advanced vehicle powertrains. In *SAE World Congress*, number 2006-01-0665, Detroit, MI, April 2006. Society of Automotive Engineers.

[61] John Remy, David Hamilton, Edward D. Moss, Boyd, Gary Zu, and Roger Johnson. 2006 Corvette Z06 carbon fiber structural composite panels – design, manufacturing and material development considerations. In *SAE World Congress*, number 2005-01-0469, Detroit, MI, April 2005. Society of Automotive Engineers.

[62] Billy Roeseler, Branko Sarh, and Max Kismarton. Composite structures – the first 100 years. In *16th International Conference on Composite Materials*, Kyoto, July 2007.

[63] Bryce Roth, Brian German, Dimitri Mavris, and Noel Macsotai. Adaptive selection of engine technology solution sets from a large combinatorial space. In *37th AIAA/ASME/SAE/ASEE Joint Propulsion Conference and Exhibit*, number AIAA 2001-3208, Salt Lake City, UT, July 2001. American Institute of Aeronautics and Astronautics.

[64] Richard Roth, Frank R. Field, and Joel P. Clark. Materials selection and multi-attribute utility analysis. *Journal of Computer-Aided Materials Design*, 1(3):325–342, 1994.

[65] B. Roy, M. Présent, and D. Silhol. A programming method for determining which Paris metro stations should be renovated. *European Journal of Operational Research*, 24(2):318–334, February 1986.

[66] Bernard Roy. The outranking approach and the foundations of ELECTRE methods. *Theory and Decision*, 31(1):49–73, 1991.

[67] S. M. Sapuan. A knowledge-based system for materials selection in mechanical engineering design. *Materials & Design*, 22(8):687–695, 2001.

[68] John Van Schalkwyk, Kevin Green, Don Pickrell, Ryan Harrington, and Mark Shaulov. *2012–2016 Corporate Average Fuel Economy Compliance and Effects Modeling System Documentation*. U.S. Department of Transportation, National Highway Traffic Safety Administration, March 2010.

[69] A. Shanian and O. Savadogo. A material selection model based on the concept of multiple attribute decision making. *Materials & Design*, 27(4):329–337, 2006.

[70] Fred Sissine. Energy independence and security act of 2007: A summary of major provisions. Technical Report RL34294, Congressional Research Service, Library of Congress, 21 December 2007. Available from: National Council for Science and the Environment; Accessed: Feb. 2008.

[71] Julie Ann Stuart, Laura J. Turbini, and Jane Chumley Ammons. Investigation of electronic assembly design alternatives through production modeling of life cycle impacts, costs, and yield. *IEEE Transactions on Components, Packaging, and Manufacturing Technology, Part C*, 20(4):317–326, October 1997.

[72] Jürg M. Thölke, Erik Jan Hultink, and Henry S. J. Robben. Launching new product features: A multiple case examination. *Journal of Product Innovation Management*, 18(1):3–14, January 2001.

[73] Dale R. Thompson and Griff L. Bilbro. Comparison of a genetic algorithm with a simulated annealing algorithm for the design of an ATM network. *IEEE Communications Letters*, 4(8):267–269, August 2000.

[74] ULSAB-AVC Consortium. ULSAB-AVC Model 2. Available from: `http://www.worldautosteel.org/`; Accessed: Nov. 2009.

[75] ULSAB-AVC Consortium. ULSAB-AVC engineering report, October 2001. Chapter 6, Body-in-White Concepts; Available from: `http://www.worldautosteel.org/`.

[76] ULSAC Consortium. ULSAC, concept engineering report, May 2001. Available from: `http://www.worldautosteel.org/`; Accessed: Nov. 2009.

[77] Mao-Jiun J. Wang and Tien-Chien Chang. Tool steel materials selection under fuzzy environment. *Fuzzy Sets and Systems*, 72(3):263–270, June 1995.

[78] Roland Wohlecker, Martin Johannaber, and Markus Espig. Determination of weight elasticity of fuel economy for ICE, hybrid, and fuel cell vehicles. In *SAE World Congress*, number 2007-01-0343, Detroit, MI, April 2007. Society of Automotive Engineers.

[79] T.P. Wright. Factors affecting the cost of airplanes. *Journal of the Aeronautical Sciences*, 3:122–128, February 1936.

[80] Yingxia Yang. *A Screening Model to Explore Planning Decisions in Automotive Manufacturing Systems Under Demand Uncertainty*. PhD thesis, Massachusetts Institute of Technology, 2009.

[81] Louis E. Yelle. The learning curve: Historical review and comprehensive survey. *Decision Sciences*, 10(2):302–328, 1979.

[82] Habib Youssef, Sadiq M. Sait, and Hakim Adiche. Evolutionary algorithms, simulated annealing and tabu search: a comparative study. *Engineering Applications of Artificial Intelligence*, 14(2):167–181, April 2001.