Integrating Conjoint Analysis and Engineering Design

Rajeev Kohli
Hitendra Wadhwa
Jim Christian [1]

August 3, 2006

[1]Rajeev Kohli is a Professor, and Hitendra Wadhwa is an Assistant Professor, at the Graduate School of Business, Columbia University. Jim Christian is Technical Director of Global Product Research at General Motors Corporation. The authors acknowledge the contributions of Dr. Richard Smallwood and Dr. Jim Weisinger, who were both partners at Applied Decision Analysis, Inc. in the early 1990s and played a central role, in collaboration with others at General Motors, in the design and implementation of the methodology described in this paper.

Integrating Conjoint Analysis and Engineering Design

Abstract

The authors describe an approach to integrating conjoint analysis and engineering design. Developed in collaboration with General Motors, the proposed approach can be used for designing products comprising several interconnected subsystems. The proposed approach has four key steps: (i) representation, (ii) linking, (iii) constraining, and (iv) costing. Representation concerns the depiction of a product in ways that are meaningful to consumers and to engineers. Linking refers to the mapping of engineering attributes onto the conjoint attributes. Constraining concerns the restriction of designs to those that are feasible from an engineering standpoint; these constraints are propagated, using the linking procedures, to the conjoint designs. Costing refers to the methods used to estimate fixed and variables costs associated with module designs. The linking procedure then maps these costs onto the product concepts evaluated by consumers, and allows the implementation of a profit-maximizing objective in the selection of a product or product line.

The proposed procedure allows an assessment of consumer and market response to engineering choices. It enables product designers to optimize the trade-off between increased customization and increased standardization via the use of common platforms and modules. It allows for the evaluation of the economics of product variety, and the cannibalization in sales due the use of common platforms in product-line designs. It makes possible the imposition of technological and cost constraints both in conjoint simulations and in algorithms for identifying optimal products and product lines. Finally, it permits feasibility assessments of alternative technologies, helps prioritize R&D projects, and facilitates the evaluation of cost-cutting efforts on the market performance of products. We illustrate aspects of the methodology using as a running example the design of a line of mid-sized cars by General Motors. We also summarize the key learnings that have accumulated in over a decade of use of these methods by the company.

Keywords: New-product design; engineering design; marketing/manufacturing interface; conjoint analysis.

1 Introduction

New product designers need to simultaneously consider financial, technical and market factors in evaluating alternative design choices and selecting the optimal design to launch in the market. This problem is simple in conceptual terms — designers need to establish which product designs are feasible, estimate the cost and demand for each feasible design, and then use business objectives to guide the selection of one or more feasible designs. In practice, however, product designers face a number of challenges. The full range of feasible designs may be too large to enumerate and evaluate individually. Product engineers, who are experts at making technical tradeoffs and judgments, may lack a clear understanding of the market or financial impact of their design decisions. And marketers charged with studying customer needs and making product design recommendations may lack an understanding of cost and technological factors that, in addition to customer factors, will ultimately drive the financial performance of the product. These challenges grow exponentially as we move from a single product design to a product portfolio design context because of cross-product effects such as cannibalization and economies of scope.

Existing methodologies for product design are at best only partially effective in addressing these challenges. Conjoint analysis is widely recognized as a powerful methodology for quantifying customer needs and estimating market response. But researchers have faced significant hurdles in extending this approach to accommodate costs and technological constraints. As Srinivasan, Lovejoy and Beach (1997) observe, the cost of a product described by conjoint attributes can vary substantially depending on the product architecture, the choice of materials and the processes used in manufacturing. Similarly, technological constraints arise naturally in the engineering domain, and are not trivial to map onto conjoint representations of product concepts. Quality Function Deployment (QFD), which links the voice of the customer to the engineering dimension of product development (Hauser and Clausing 1988, Griffin 1992, Griffin and Hauser 1993), is useful for detailed design, but is less relevant for relating engineering decisions to such strategic issues as product-line positioning, market segmentation and cannibalization. As it uses aggregate customer data, QFD does not provide a direct link to the market simulations and optimal product design algorithms used with conjoint data (e.g., Kohli and Krishnamurti 1987, Nair, Thakur and Wen 1995, Dobson and Kalish 1993), and it does not directly incorporate costs or provide the kind of quantitative assessments of market response that are needed to assess the financial performance of alternative product designs.

The purpose of this paper is to describe methods for linking engineering and marketing design of new products. Initiated as a collaboration between Applied Decision Analysis, Inc. (ADA), a quantitative management consulting firm, and General Motors Corporation (GM) in the early 1990s, these methods have been tested and improved for over a decade. We illustrate the key aspects of the methods in the context of a GM initiative focusing on the redesign of Buick, Chevrolet, Oldsmobile and Pontiac lines of mid-size cars. Teams of engineers, marketers and designers across several GM divisions collaborated in, and contributed to, the development, implementation and testing of the method. Since then, the method has been used for several different lines of GM automobiles.

The methodology has broad applicability, beyond the design of automobiles. It integrates market data based on conjoint analysis with engineering and cost data. It can be used as a comprehensive tool to evaluate alternative designs for products comprising several interconnected components and design features. It can be used to assess the impact of choices in engineering design on consumer responses, and on the market performance of products and product lines. It can also be used to assess the impact of technological changes, alternative R&D investments, and major or minor changes in designs on the performance of new products.

The proposed method has four key steps: (i) representation, (ii) linking, (iii) constraining, and (iv) costing. Representation refers to the description of a product in two separate ways: conjoint attributes and levels, based on a customer viewpoint; and engineering attributes and levels, based on an engineering viewpoint. Linking concerns the construction of functional relationships between the engineering attributes and the conjoint attributes. Constraining refers to the specification of constraints that limit the evaluations of alternatives to feasible engineering designs. Costing concerns the estimation of variable and fixed costs for a new product. These costs are often easier to decompose and estimate for engineering designs than they are for conjoint profiles.

Figure 1 gives a schematic overview of the four steps and their relation with an optimizer that selects an engineering design to maximize market share, profits or ROI, for a single product or for a portfolio of products built around a common engineering platform. The optimizer also allows for simulations that use the integration of market, engineering and cost information. It can be used to assess the value of alternative technology choices; to assess the economics of product variety; and to prioritize R&D investments based on an integrated analysis of investment requirements, technology performance, customer needs and manufacturing costs. The details of the optimization method are not presented in this paper. The reason is that we can use existing algorithms described in the marketing literature on optimal product and product line design (Green and Krieger 1985, 1989, Kohli and Krishnamurti 1987, McBride and Zufryden 1988, Dobson and Kalish 1988, Balakrishnan and Jacobs 1996, Nair, Thakur and Wen 1995, Chen and Hausman 2000). The only difference in the use of these methods is that our search space is a set of engineering designs. We link these designs to conjoint profiles, and thus to market response functions maximizing market share, revenue and profits.

To the best of our knowledge, the proposed approach is the first to link engineering design to conjoint-based models. We propose a framework for representing products as a collection of modules, and modules as a set of engineering attributes. We offer a structured approach to mapping engineering modules and attributes onto the conjoint attributes relevant for estimating consumer preference functions. We describe a computationally efficient method for eliminating infeasible choices in the selection of optimal products and product profiles. We describe an approach to estimating cost functions, which are more naturally identified with engineering aspects of a product — the product architecture, the choice of materials, the processes used in manufacturing — than with the attributes used in conjoint analysis (Green and Krieger 1991, Srinivasan, Lovejoy and Beach 1997). [1]

Section 2 provides a brief background on the product design initiative at General Motors that led to the development of the present approach. Section 3 describes the four steps of this approach, using the design of a line of mid-size car at GM as an illustrative example. Section 4 discusses how the proposed methodology can be used to answer questions concerning investments in new technologies by a firm, and to project the impact of cost reductions in one or more design choices on the market performance of products. We couch this discussion in the context of how managers at General Motors have used the methodology to examine such issues over the last ten years. Section 5 concludes with a discussion of the areas in which improvements are still being made, the lessons learnt by users of the methodology and the general impact the method has had on the process of product development at General Motors.

[1] Such techniques as Taguchi methods (Taguchi, Elsayed and Hsiang 1989), and the Boothroyd-Dewhurst approach to cost reduction (Boothroyd and Dewhurst 1983), have gained much use for reducing costs and improving quality. We do not explicitly examine these methods, but assume that these are available to engineers for estimating the baseline and marginal costs we discuss later in the paper.

2 The GM Midsize Car Portfolio Project

The Midsize Car Division (MCD) and the Strategic Decision Center (SDC) at General Motors undertook a project to redesign GM's mid-size car portfolio in 1993-94. We will refer to the project, and the study described below, as the MCD study.

The task faced by the joint team working on the MCD study was to design a portfolio of six new sedans to replace the existing line of mid-size cars produced by GM. At that time, conjoint analysis was already an established market research methodology within General Motors. Customer preferences along different product attributes would be quantified using this technique, and a market simulation model would be developed to allow product planners to explore the market impact of a range of product scenarios. Occasionally, the process of searching for the optimal (revenue or market-share maximizing) product design was automated. To aid the search process in yielding practical results, simple constraints were developed in the conjoint attribute space — disallowing, for example, a certain body style with a certain level of interior room. This process was iterative — the model would recommend a product or product line, product design staff would evaluate these designs, and if these were found to be too costly or technically infeasible, new constraints would be imposed on the optimization model to disallow such designs.

The MCD portfolio project entailed a highly complex set of design choices for General Motors. The products were described in terms of 13 conjoint attributes, such as turncircle (turn radius), fuel economy and bodystyle. Recall that there were 6 mid-size vehicles to be designed for GM. The market simulation incorporated 16 additional competitive vehicles. The team wanted to explore both market share and financial objectives such as revenues, investment requirements, and profits. Given the scale of the problem, it was concluded that the traditional process of imposing a few inter-attribute constraints on the conjoint model and refining these based on early recommendations would break down — there were simply too many technical constraints for this process to be able to conclude with a satisfactory portfolio design within a few iterations. Moreover, the simulations would not help evaluate profits because fixed and variable costs were difficult to associate with the conjoint attributes. And designers wanted to extend the methodology so that the market simulations could be linked to the engineering choices, which involved such things as engines, transmissions and floorpans.

Spurred by the challenging nature of the assignment, the team, in partnership with ADA, developed a modeling approach that integrated a model of customer preferences and purchase behavior — via the traditional conjoint machinery — with a model of engineering design choices. The approach was developed, refined and utilized in an ongoing manner within the team to explore a broad array of strategic design choices, investment requirements and future technological and competitor scenarios. A number of critical insights were identified that went substantially beyond the capabilities of traditional tools to deliver. The approach broadened the perspectives of the designers and allowed them to consider a number of new possibilities. Specific learnings from the project were instrumental in shaping a number of the ultimate design decisions taken by the organization.

Since the project's conclusion in 1994, GM has continued to deploy this integrated engineering-marketing approach in other design contexts. Elements of the approach have been standardized to permit swift implementations on a regular basis. Many of the business insights generated by the methodology have since been incorporated into mainstream portfolio thinking at GM.

3 Integrating Marketing and Engineering Design

There are four key steps in the proposed approach to linking consumer preferences and engineering design: (i) representation, (ii) linking, (iii) constraining, and (iv) costing. Representation concerns the depiction of a product in ways that are meaningful to consumers and to engineers. Linking refers to the mapping of engineering attributes onto the conjoint attributes. Constraining concerns the generation of constraints in the engineering domain; these can be mapped, using the linking procedures, onto the product design optimizer in Figure 1. Costing refers to the methods used to estimate fixed and variables costs associated with module designs and, via the linking procedure, to the conjoint descriptions of product concepts. This approach addresses a problem that can occur when it is not possible to associate costs directly with conjoint attributes and product profiles. We now discuss each of these four aspects of the proposed approach.

Representation

The first step comprises the representation of a product in ways that are meaningful to consumers, and to engineers. We use conjoint analysis for the former representation of a product into attributes that are meaningful to consumers. Table 1 shows the set of attributes used in the MCD study; we refer to these as conjoint attributes, to distinguish them from the engineering attributes, which we discuss below. The conjoint attributes are used in a standard manner to obtain consumer judgments about alternative product concepts, and can be used to simulate the market performance of new product alternatives. A prolific literature exists in this area and so we do not elaborate further on this subject here.

The representation of a product by engineering attributes is accomplished in two steps. First, we represent a product by a small collection of modules. Second, we represent each module in terms of engineering features. Thus, we use a hierarchical representation in which a product is conceived as a collection of modules. The levels of each module are the alternative module designs. And each module design is described by a collection of engineering attributes, each of which has a number of possible levels.

Modules. Each module is a physical assembly, typically implementing one major function. For example, the ten modules shown in Table 2 were used in the MCD study at GM. The process of generating possible module designs typically begins by examining existing products, and then extends the search to newer, more novel alternatives that are, or can be made, available. For example, the design team for the MCD study listed all car engines that were manufactured by GM, by its suppliers, and by suppliers to other automobile manufacturers who could potentially be used to source engines. It then added to this list newer engine designs that, while not available on the market, could be produced, albeit with investments in plant and equipment. Some of these engines were being designed elsewhere for use in hybrid (fuel and electric) cars, and for use in cars using natural gas.

Products can, but do not have to, be built using a modular design. The alternative is to build integrated products. We refer the reader to Ulrich (1992) for a discussion of the key differences between modular and integrated product architectures. Here, we only note that modular design has become an important way of making a variety of products — automobiles, computers, watches, and personal stereo systems. There are three main reasons for this. First, if a product is complex, with many parts and components, it can be virtually impossible to design it as a single, integrated unit. Second, if the design of a product can be related to the design of its component modules, then it is possible to combine variants of a module's designs in a combinatorial manner to develop wide product lines at low cost. For example, every Swatch watch is made by mixing and matching a set of modules, each of which is available to designers in a pre-selected number of options. If there are m modules, each with n possible options, one can, in principle, design n^m possible products by selecting one or another option for each module. Incompatibilities exist among the modules, reducing the number of feasible designs. But the possible number of feasible designs can still be very large. Third, there are economies of scope that a firm can exploit by sharing modules — a common engine, or a common door assembly — across a product line. Of course, not all products in a line can use all identical modules. The idea is to select a small number of designs for each module, and then make each product in a line by combining different module designs.

Table 2 shows the ten modules used in the MCD study. The number of designs in the table refers to the different designs that were available for each module. The GM design team needed to select between 1 to 3 designs for each of the modules numbered 1 to 9 in Table 2, and then decide which combinations of the selected modules to use in specific cars in a product line. Conjoint experts might observe how several of these modules, such as Cowl and Rear Compartment Plan, are unsuitable for capturing customer preference feedback. But this view of the product is very natural to product design engineers. And that is exactly the point of this exercise — to decouple, and then link, the engineering-oriented view of the product and the customer-oriented view.

Engineering attributes. We represent each module by a collection of attributes that describe certain physical or technological characteristics of the module. We call these engineering attributes, or module attributes. Each such attribute is described over a continuous range of values, or over a set of discrete values or levels. Figure 2 illustrates the hierarchical relation among modules, engineering attributes and attribute levels. As an illustrative example, all designs of the motor compartment module were described in the MCD study using four continuous attributes: rail width, rail length, front overhang and dash-to-bumper length. And all designs of the front suspension module were described using two discrete attributes: Track and Type (SLA or Strut).[2]

It is often possible to describe a module over a large number of attributes. Which attributes are used in a study depends on their relevance for two purposes. First, some attributes are necessary to specify the dependencies, and thus the technological constraints, within and across modules. Second, there are attributes that are useful for constructing relationships between the consumer attributes, used in the conjoint analysis, and the engineering modules. An attribute should be selected if it facilitates one or both of these purposes. Otherwise, it can be excluded.

Linking

The second step of the proposed approach is the construction of functional relationships between engineering attributes and conjoint attributes. Design engineers develop a model that translates every feasible combination of module designs into an associated set of conjoint attribute levels. This can be easier for some modules than for others. For example, the 2005 Honda Hybrid Accord has a combination electric and fuel engine, and this is not simply an engineering design feature, but also a feature of considerable interest to consumers. But other aspects of engineering design can be more difficult to map onto consumer attributes. In some cases, there may be no apparent connection between consumer attributes and engineering attributes/modules; in other cases, several different modules might together determine the mapping onto a conjoint attribute.

[2]SLA (short-long arm) refers to a type of suspension that uses upper and lower control arms of unequal length. The upper arm is usually shorter than the lower arm to control camber changes during jounce and rebound. An automobile strut eliminates the need for an upper suspension arm and provides both shock absorption and support for sideways loads that are not along its axis of compression.

The development of a method for linking conjoint attributes and engineering modules is perhaps the most important contribution of the MCD study. Without it, there would have been no integration of engineering and marketing design. We describe the main aspects of the method that eventually emerged after several months of work. It consists of the following three steps: (i) decomposing, which refers to the process by which a conjoint attribute is represented as a function of other variables, some of which might not have any direct relationship with the modules or their attributes; (ii) defining intermediate variables, which bridge relationships between modules and conjoint attributes; and (iii) constructing causal or empirical relationships between intermediate variables, conjoint attributes and engineering modules. For concreteness, we describe the three steps in the context of the method GM designers used to link engineering attribute in the MCD study to fuel efficiency, a conjoint attribute.

Linking fuel efficiency with engineering attributes. Figure 3 shows the two major factors that affect fuel consumption: road load and system efficiency. This is the first decomposition. Our objective is to construct a relationship between fuel efficiency, which is a conjoint attribute listed in Table 1, and the engineering modules shown in Table 2. To this end, we first observe that the load on a car is determined by its weight, and on weather and road conditions that affect the aerodynamic drag and the rolling resistance of a car.[3] We consider the weight to be our first intermediate variable; it is a factor engineers can control when making a car.

Figure 4 shows the empirical relationship between auto weight and fuel efficiency for US automobiles in 2002. A linear relationship explains a substantial

[3]Weather and road conditions have a substantially lesser effect on the published/advertised gas milage of cars, because regulatory agencies specify standard test conditions that eliminate their influence on fuel efficiency. For example, the US Environmental Protection Agency (EPA) requires that a test car be parked overnight at a temperature of about 72° F (22° C). It is then driven on a dynamometer, in a controlled laboratory, for a specific period of time.

part of the variance. Much of the additional variance around the line is explained by specifying the system efficiency. Thus, we redefine our problem as constructing a relationship between a car's fuel efficiency, and its weight and system efficiency. Next, we observe that the system efficiency is determined by three modules — engine, transmission and the final drive.[4] Thus, we obtain the relation

$$\text{Gallons per 100 mile} = a + b\,\text{Curb weight} + \sum_i c_i x_i;$$

where the x_i denote (a suitable number of) dummy variable representing the alternative engine, transmission and the final drive modules. This empirical relationship explains over 99% of the variance in the fuel efficiency of a car.

It still remains for us to specify how the weight of a car is related to the modules. To do so, we first represent vehicle weight as a product of the floor area (overall length $\times$ overall width) and the average weight density of a car:

$$\text{Car weight} = \text{Average weight density} \cdot (\text{Overall length} \times \text{Overall width}).$$

The average weight density, overall length and overall width are new intermediate variables. The average weight density is fairly constant for cars within a particular class. In the MCD study, designers started by using the mean value across the existing mid-size cars; later, as choices of materials and components became known, the value of the constant was refined. The overall length and overall width of a car is next decomposed into the sums of smaller lengths and widths, in the manner shown in Table 3. We continue the decomposition until each length and width can be identified as an attribute of a module — in the present instance, the modules comprising the Motor Compartment, Floorpan and Rear Compartment. We have thus completed our mapping of a conjoint attribute — fuel efficiency —

[4]At this point in the design, we ignore the effect of accessories line air conditioning on the system efficiency.

to the modules shown in Table 2 for the MCD study. A summary of the steps appears in Table 3.

The above example illustrates how constructing a relationship between modules and conjoint attributes can be a non-trivial exercise. It can require knowledge of regulations and testing procedures; the use of empirically established relationships; the recursive decomposition of intermediate variables via identities; the approximation of constants, and the refinement of their values as the design team makes finer, lower-level decisions about the materials used to make a product. The three steps — decomposing, introducing intermediate variables, and estimating relationships — can occur several times in the process of linking conjoint attributes and engineering modules. The process is typically time-consuming the first time it is done. Thereafter, it becomes easier, can be standardized, and further refined over time. Many of the relationships are known to design engineers. And some of the attributes in terms of which the modules need to be defined become known only after the linking process is begun. Sometimes, it is easier to begin the linking process from the conjoint side and moved into the engineering module side. At other times, the reverse flow is more convenient.

Constraining

The third step of our procedure concerns the development of methods that prevent the selection of technologically infeasible product concepts. The existing literature on conjoint-based methods for new product design typically assumes that designers know which combinations of consumer attribute levels are infeasible, and that iterative algorithms, such as those described by Kohli and Krishnamurti (1987), can be used to ensure that an optimal product design is feasible. On the other hand, constraints arise naturally in the engineering domain: a battery is too wide to be used in a laptop computer, or an engine is too small for a car. These constraints in turn determine which conjoint designs can be offered to consumers.

One way of generating constraints is to list them out, one by one. This enumerative method is impractical because the number of module combinations can be enormously large. For example, there are $16 \times 12 \times \ldots \times 6 = 46,558,955,520$ possible ways of selecting one each of the ten engineering modules shown in Table 2 — far too many for exhaustive enumeration and evaluation by engineers. Thus, we need a procedure for efficiently identifying constraints in the engineering domain and then propagates these constraints to the domain of conjoint profiles using the proposed linking procedure. We describe such a procedure next.

The first step is for design engineers to describe the compatibility constraints in terms of engineering attributes, not the module designs. For example, each design of the rear suspension module is compatible with only certain designs of the Transmission and Front Suspension modules. A module-level approach would consider each Rear Suspension design separately and list the feasible subsets of the eleven Transmission designs and the nine Front Suspension designs with which it can be used. A more parsimonious description of the constraints is obtained by tracing the source of the incompatibilities to the engineering attributes. In the MCD study, the Rear Suspension was constrained by the Drive Configuration (an attribute of the Transmission module) and the Track (an attribute of the Front Suspension module). Since there are only 4 Drive Configurations and 3 Track levels, this reduces the constraint specification work considerably.

The second step is to separate the modules into three groups: master modules, secondary modules, and independent modules. Master modules are those that have a major influence on other modules, in terms of compatibility. Secondary modules are those which largely do not influence each other, in terms of compatibility, but which are influenced by one or more master modules. Independent modules are those that are not associated with any compatibility constraints; i.e., for an independent module, each possible design is compatible with any feasible combination of designs for the other modules. Observe that independent modules are readily identified, because they have no dependencies on other modules. The partitioning of the other attributes into primary and secondary classes is obtained by solving a maxcut problem. We describe the procedure in the appendix. Here, we observe that this separation of modules substantially reduces the complexity of the constraint specification process, as the following numerical example will show.

Suppose we describe a product using six modules, each module having five possible designs. The total number of design combinations whose technical feasibility needs to be specified is then $5^6 = 15,625$. Suppose that we can identify one of the modules to be independent. Then the total number of design combinations whose technical feasibility needs to be specified is $5^5 = 3,125$, since the independent module does not need to be considered. Further, suppose we can separate these remaining 5 modules into two master modules and three secondary modules. There are then $5 \times 5 = 25$ possible design combinations of the 2 master modules. For each of these 25 combinations, we need to specify technical feasibility vis-a-vis the 5 designs for each of the three secondary modules. This leads to a total of $25 \times (5 + 5 + 5) = 375$ design combinations — a substantial reduction from the original 15,625.

Figure 5 shows a graph in which the vertices, numbered 1–9, correspond to modules 1–9 in Table 2. An edge connects a pair of vertices if the choice of one module constrains the choice of another module. We call this graph a dependency graph. Observe that module 9 (sideframe) has no dependencies; it is the only independent module. The remaining modules constrain the choices of one to five other modules. Figure 6 shows a (maxcut) partition: modules 1, 3 and 5 are in one partition (cut) and the remaining modules are in the other partition. If we associate equal weights with the edges of the dependency graph, the sum of edge weights (which is then equal to the number of vertices across the partitions) is 10. In the MCD study, we labeled modules 1, 3 and 5 as the primary modules and the modules in the other partition as the secondary modules. Note that there are no constraints among the secondary modules, and only one constraint between a pair of the primary modules.

To use this partition in a market simulation, or in an algorithm for optimal product design, we begin by constructing a composite suspension/motor compartment module; its designs are the feasible combinations of the front suspension (module 3) and motor compartment (module 5). We select a design for this composite suspension/motor compartment module and module 1 (floorpan); propagates the constraints to each of the secondary modules (2, 4, 6, 7 and 8), restricting choices to only their feasible designs; and then map a feasible engineering profile (specified by one design for each module) onto a conjoint product profile.

The constraints thus specified encode, in a sense, the current technological frontier. Further technological innovation may lead to certain new module designs becoming feasible (i.e., new levels of attributes associated with a module, or new combination of attribute levels associated with a module). If product designers wanted to study the potential market impact of such innovations, it can be implemented within our methodology in a straightforward manner, simply by adjusting the constraints. For instance, in the MCD study, product designers wanted to study the impact of a new type of engine technology. This was implemented by specifying an additional level of the technology attribute associated with the engine module.

Costing

Maximizing a profit function is often difficult in conjoint models of optimal product design because there are problems in estimating cost as a function of the conjoint attributes. For example, it is not easy to specify the marginal cost of an extra mile per gallon of fuel efficiency. The reason is that there is no direct relationship between such conjoint attributes and manufacturing costs. Instead, costs are more naturally associated with the engineering modules.

We estimate the total cost of a product as the sum of the costs of its constituent modules. The module costs are determined using a method described below. Once these cost estimates are obtained, we use the mapping procedure described above to identify a conjoint profile with the product. Then we simulate consumer preferences for the product at a given price, and estimate the demand, revenue and profit for the product. In this way, we obtain a model in which costs, estimated in the engineering domain, are combined with the revenue predicted by consumer choice simulations that are conducted in the conjoint domain. The simulations do not have to be done explicitly, one product at a time. Instead, we can use algorithms for selecting optimal products and product lines, similar to those in the literature. The only difference in the use of these algorithms for the present problem is that products and costs are defined in terms of engineering modules, while the market simulations are performed over conjoint attributes and profiles. The link between the two domains is achieved by the procedure described in section 3.2.

Product designers at GM often use the following method for estimating the variable costs for various module designs. They start with an existing car in a product class, for which the variable costs are known for the entire car and for the various module designs. We will call this a base car, and will refer to the associated costs as base costs in the following discussion. The costs of other module designs are specified as increments to the base cost. In a problem with m modules, where module j is associated with n_j levels, we need to estimate $(n_1 - 1) + \ldots + (n_m - 1) = -m + \sum_j n_j$ incremental costs. For example, a product defined using $m = 5$ modules, each with $n_j = 5$ designs per modules, requires the specification of $-5 + 25 = 20$ incremental variable costs, relative to the baseline cost. Some of these costs are easily estimated, for example when there are existing module designs that are produced for other products, or that can be purchased from outside vendors.

In some cases, as for example with hybrid engines, there is substantial uncertainty about costs because of lack of prior experience within the firm. GM uses two methods, often in combination, in these cases. First, when possible, a vehicle teardown center reverse engineers competing products in which these technologies are used, and then builds up to a cost estimate. Second, individuals from several different groups, such as engineering, finance and manufacturing, collaborate to identify the cost drivers and estimate variable costs. In certain cases, fixed costs also need to be estimated, if for example a new engine is to be built for a line of cars. The same inter-functional group then performs a detailed feasibility and cost analysis before projecting a fixed cost for plant and equipment. Estimates of this kind are the least accurate, and there is at present no quick substitute for the detailed analysis required to obtain cost projections.

Product designers at GM typically use the following iterative procedure when using the proposed method. First, they assume the lower end of the range of fixed and variable costs for a new module. If it is picked by the optimizer, they then re-estimate costs using the projected sales volume for a car, or for a line of cars. The optimizer is rerun, and the costs updated if the module is again selected by the optimizer. A few iterations are normally sufficient to get a reasonably accurate estimate of costs.

4 Using the methodology

The integrated engineering and marketing model provides product planners with a learning laboratory where they can conduct hundreds of experiments on the impact of specific design choices on market acceptance and financial returns. We describe below a number of different ways in which this machinery can be put to use, and illustrate these applications with learnings from the MCD portfolio design project at GM.

Engineering choices and customer decisions

The linking of engineering attributes, module designs, and conjoint attributes provides a method for assessing the sales and profit impact of different choices of engineering designs. Product designers and marketers can collaborate to explore this integrated landscape and develop greater intuition on how certain engineering decisions increase or decrease the product's appeal to customers. While some of these effects are directly apparent — such as the increased customer appeal of a superior engine technology — others are based on complex interactions in the model and can yield counter-intuitive results.

For instance, in the MCD project, the design team was surprised to find that a car with a larger cargo capacity did not perform as well in market simulations as another design involving a smaller cargo space. Sensitivity and root-cause analysis revealed the reason to be a tradeoff between cargo capacity and turn circle that was not directly evident to the designers. The reason for this is that cargo capacity imposes constraints on the dimensions of a car's platform. These dimensions, in turn, affect the turn circle. From the consumer standpoint, a smaller turn circle is more important than a larger cargo capacity. Similar analysis involving other attributes and linkages are now done commonly by engineers and marketers, allowing them to better communicate how and why the ideas and objectives proposed by one group affect the feasibility and impose constraints on the other group. For example, the idea of designing more environmentally friendly cars has been proposed several times, by different groups over the years. These cars tend not to be selected by the conjoint optimizers because the simulations suggest that consumers are not willing to pay the incremental price that would be necessary to make the design economically viable. As one of the designers puts it, people are "green" only as long as they don't have to pay. Ethanol engines are another example. These engines are attractive because Ethanol is made from biomaterial (like corn or grain products), helping reduce dependence on petroleum and greenhouse gas emissions. However, evaluations of these engines using the proposed methodology suggest that the cost of Ethanol is a key factor in consumer acceptance of such cars, not the cost of the car itself. GM thus began working to make Ethanol competitive with gasoline. Today, about 2 million GM cars on the road have the capability of running on gasoline or E85, a blend of 85% ethanol and 15% gasoline.[5]

Sometimes, products have been developed despite the proposed method predicting poor market performance. An example is a two-seater electric car. It was developed, launched, and then withdrawn after it failed in the market. There are also cases where the proposed methodology has identified important market opportunities. In these cases, the key decision is not market demand but the feasibility of developing the required technology at a cost that can be justified by management. For hybrid cars, R&D and production costs, in addition to demand projections, have been major factors for management to consider.

[5] For details, see http://www.gm.com/company/onlygm/energy_flexfuel.html.

Product-line differentiation

Product planners constantly wrestle with the following dilemma. On the one hand, they want to standardize their product line along a common platform, architecture, and set of modules, since this reduces manufacturing and logistical complexity and often yields substantial cost savings. On the other hand, they want to offer a differentiated product line in order to better meet the needs of different segments of customers. To assess the tradeoff between costs and customization in the MCD study, GM designers established three test wells. A test well was a proposed design for each of the six GM mid-size cars. The three test wells were meant to represent the designers judgments on what may be optimal portfolio designs. Each test well was then compared to the designs identified by the proposed integrated model. The extent of differentiation was assessed by comparing the attribute bandwidth, which is defined as the range over which attribute levels vary across a product portfolio.[6] As shown in Table 4, the proposed method recommended a product line with a wider bandwidth — greater differentiation — than the test wells. It also recommended a product line that was predicted by the conjoint simulator to yield significantly higher revenues and profits. The results suggested that the designers had overemphasized the cost of product variety over its benefits in coming up with the test wells.

Product designers can also utilize the proposed methodology to identify those attributes for which an increased bandwidth yields the greatest gain in market share or profits. For instance, the following analyses were conducted in the MCD study. First, a variety of optimal mid-size car portfolio recommendations, corresponding to different starting conditions or competitive assumptions, were investigated. The objective of this investigation was to highlight those attributes that consis-

[6] For a discrete attribute, this range is simply the number of levels that appear across products in the portfolio.

tently appear with a larger bandwidth. Next, simplified versions of these recommended product portfolios were evaluated. These simplified portfolios involved adjustments to the recommended portfolios where all cars were constrained to be similar along one or more attributes. This analysis provided additional richness to the team's understanding on the impact of differentiation along each attribute. Among the attributes in the MCD study where differentiation was identified to be important are cargo capacity, turn circle, passing acceleration, fuel economy and towing capability.

Over time, GM has combined many platforms. It now uses a common platform architecture across a large number of cars. Platforms have also migrated across product lines, for example from mid-size sedans to small SUVs like the Saturn Vue. The most recent thrust in this regard is the consolidation to a single architecture of related product lines across many countries across the world. Similar efforts are known to be taking place across the automobile industry; Toyota is considered to be leading the way in this effort.

Like all conjoint-based models for product-line design, the present approach accounts for cannibalization in sales, profit and market share when selecting product lines. The main difference is that the proposed model allows this assessment simultaneously in the engineering space and in the market space. In practice, it seems to be able to differentiate the low margin products from high margin products, while at the same time seeking to establish common technology platforms that reduce investment requirements and R&D complexity.

Impact of technological innovations

New technologies are often associated with substantial levels of uncertainty in terms of cost and performance. These technological and financial uncertainties need to be factored into the product design process. The proposed methodology can be used to evaluate technology development options and to establish target performance and cost ranges within which a technology would create value for the business. The essence of this approach to evaluating technological choices involves two steps. First, designers need to establish linkages between the technology and the engineering modules and attributes/levels. The technology may, for instance, trigger changes in the way engineering attributes and levels translate into conjoint attributes and levels — for instance, by increasing the reliability of a car. Second, if there are uncertainties in performance or cost, designers can create multiple scenarios with alternative performance levels and costs for the technology. By running the market simulation under these alternative scenarios, designers can map out the range of possible financial outcomes from deploying the technology, and then determine its suitability and target performance/cost levels.

In the MCD study, designers faced a decision on whether or not GM should invest in a new engine technology. To evaluate this decision, a number of scenarios were constructed involving a range of performance levels — in terms of fuel economy and passing acceleration — and variable costs for the technology. Running the market simulation model under these alternative scenarios provided rich learnings on the conditions under which it was attractive to introduce this engine technology into GM's mid-size car portfolio. Such analyses have led to the investments in such technologies as Ethanol engines.

Highlighting Key Engineering and Customer Attributes

Given the combinatorial complexity of evaluating new, multi-attribute products, and particularly portfolios of products, a mechanism that efficiently highlights the likely subspace of the overall feasible region where optimal solutions may lie is of substantial value to designers. One such mechanism that is facilitated by the integrated methodology is the identification of specific engineering attributes that are particularly important in driving business performance.

In the MCD study, the methodology provided the project team with specific insights into the selection of several key engineering attributes such as motor compartment and engine, and key customer attributes such as fuel economy and turning circle. In particular, the methodology consistently identified vehicles with higher fuel economy, slower acceleration and smaller turning circles relative to vehicles in the three test wells. By further analyzing the interactions across engineering attributes, the team was able to link the recommended smaller turning circles with specific motor compartments that facilitated a number of benefits to the customer (of which turning circle was one), discovering through the process the importance of certain motor compartments.

Product designers also look for designs that are difficult for competitors to copy. For example, new seat configurations or new floorpans are less likely to be copied by competitors because these can have high design costs. On the other hand, smaller engines are easy to copy and build and are therefore considered less attractive by designers. There are also cases where the conjoint simulations might make recommendations that are not accepted by the designers because they are untried and require large investments. An example is the design of folding seats, which was considered but rejected for a line of SUVs because it required a large investment. Another example was the decision to not introduce new colors in a product line because the cost of building a new paint shop can be very high.

5 Conclusion

The methodology we have described has gained wide adoption at GM. After its initial application to designing mid-size cars, the methodology has been extended to design small SUVs like the Chevy Equinox, and subsequently, full-sized SUVs like the Chevy Tahoe and Escalade. Design teams for several additional lines of

GM cars have also used these methods. Product engineers and marketers have benefited by developing a common understanding of the tradeoffs involved between engineering and marketing decisions. GM focuses on share-maximization and profit maximization and actively explore the trade-offs between the two.

GM designers have learned several important lessons using the proposed methodology over the last ten years. One is that conjoint simulations may tend to understate the price sensitivity of consumers. Designers now make subjective adjustments for this in their models, and are currently working on ways to make these adjustments more formal in the product optimizer. Prices in most simulations are typically restricted to realistic ranges that depend on models and brands. Closer to market launch, the actual costs of modules and designs tend to be higher than the earlier cost estimates. Designers at GM now try to build in a scaling factor in the conjoint simulations to represent this likelihood.

Experience suggests that it is better to not rely on one optimization run with all the constraints specified ahead of time. It is easier to start with unconstrained optimization, and then add constraints one at a time to produce a series of solutions. The advantage is that there is a lot of learning in this process. It is in effect a structured approach to exploring the design space. GM designers know when to call in experts, and the entire development team learns about why certain technological or cost constraints are necessary.

The focus on common platforms tends to make unique architectures less viable from an economic standpoint. And there are instances where seemingly simple proposals for changes in a product lead to a cascading sequence of inter-related changes in design that are otherwise not evident to designers. Certain types of designs are favored by GM designers if these appear in the simulations. These are designs that are difficult for competitors to copy; and designs that can be used across product lines and countries. GM designers try to keep away from easily copied changes, especially if these require large investments.

The largest impact of the method is that GM has been able to better differentiate models like Oldsmobile and Chevrolet while using a common platform architecture. The methodology is used to evaluate new investment decision, to target R&D efforts and to identify entirely new ways of introducing competing products. Navigation systems are one example. Nonreflective flat screens have been used by many of GM's competitors in car navigation systems. These screens are very expensive. They can also be difficult for consumers to use. The methods described in this paper led to the development of Onstar. GM was looking to find an alternative that would compete well against flat screens and would also be much less expensive to produce and maintain. The model described in this paper recommended a radio (text) screen, accompanied by verbal commands for directions and map downloads. Design engineers developed this option in the Onstar system. It was first offered on select models of Buick, Pontiac and Cadillac cars in the late 1990's. Today, General Motors offers OnStar on more than 50 models.

References

Balakrishnan, P.V. and V.S. Jacob (1996), "Genetic algorithms for product design," *Management Science,* 42 (1), 1105–1117.

Boothroyd, G. and P. Dewhurst (1983), *Design for assembly: a designer's handbook,* Amherst, MA.: Boothroyd and Dewhurst.

Chen, K.D. and W.H. Hausman (2000), "Mathematical properties of the optimal product line selection problem using choice-based conjoint analysis," *Management Science,* 46 (2), 327–332.

Dobson, G. and S. Kalish (1988), "Positioning and pricing a product line," *Marketing Science,* 7 (2), 107–125.

Dobson, G. and S. Kalish (1993), "Heuristics for pricing and positioning a product-line using conjoint and cost data," *Management Science,* 39 (2), 160–175.

Fisher, M., K. Ramdas and K. Ulrich (1999), "Component sharing in the management of product variety: a study of automotive braking systems," *Management Science,* 45 (3), 297–315.

Frieze, A. and M. Jerrum (1997), "Improved approximation algorithms for MAX k-CUT and MAX BISECTION," *Algorithmica,* 18, 61-77.

Garey, M.R. and D.S. Johnson (1979), *Computers and intractability: a guide to the theory of NP-completeness,* San Francisco: Freeman.

Goemans, M.X. and D.P. Williamson (1995), "Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming," *Journal of the Association for Computing Machinery,* 42, 1115–1145.

Green, P.E., J.D. Carroll and S.M. Goldberg (1981), "A general approach to product design optimization via conjoint analysis," *Journal of Marketing,* 45, (Summer), 17–37.

Green, P.E. and A.M. Krieger (1985), "Models and heuristics for product line selection," *Marketing Science,* 4, 1–19.

Green, P.E. and A.M. Krieger (1989), "Recent contributions to optimal product positioning and buyer segmentation," *European Journal of Operational Research,* 41, 127–141 .

Green, P.E. and A.M. Krieger (1991), "Product design strategies for target-market positioning," *Journal of Product innovation Management,* 189 (8), 189–202.

Griffin, A. (1992), "Evaluating QFD's use in U.S. firms as a process for developing products," *Journal of Product Innovation Management,* 9 (3), 171-187.

Griffin, A. and J.R. Hauser (1993), "The Voice of the Customer," *Marketing Science,* 12 (1), 1-27.

Hauser, J.R. and D. Clausing (1988), "The House of Quality," *Harvard Business Review,* 3 (May-June), 63–73.

Hauser, J.R. (1993), “How Puritan-Bennett used the House of Quality,” *Sloan Management Review,* 34 (3), 61–70.

Kohli, R. and R. Krishnamurti (1987), “A heuristic approach to product design,” *Management Science,* 33,1523–1533.

Kohli, R. and V. Mahajan (1991), “A reservation price model for optimal pricing of multiattribute products in conjoint analysis,” *Journal of Marketing Research,* 28 (August), 347-354.

Kohli, R. and R. Sukumar (1990), “Heuristics for product-line design using conjoint analysis,” *Management Science,* 36, 1464–1478.

Krishnan, V. and K. T. Ulrich (2001), “Product development decisions: a review of the literature,” *Management Science,* 47 (1), 1–21.

McBride, R.D. and F.S. Zufryden (1988), “An integer programming approach to the optimal product line selection problem,” *Marketing Science,* 7,126–140.

Nair, S.K., Thakur, L.S. and Wen, K. (1995), “Near optimal solutions for product line design and selection: Beam search heuristics,” *Management Science,* 41 (5), 767–785.

Ramdas, K., M. Fisher and K. Ulrich (2003), “Managing variety for assembled products: modeling component systems sharing,” *Manufacturing and Service Operations Management,* 5 (2), 142–156.

Srinivasan, V., W. S. Lovejoy and D. Beach (1997), “Integrating product design for marketability and manufacturing,” *Journal of Marketing Research,* 34 (1), 154–163.

Taguchi, G., E.A. Elsayed and T. Hsiang (1989), *Quality engineering in production systems,* New York: McGraw Hill.

Ulrich, K.T. (1995), “The role of product architecture in the manufacturing firm,” *Research Policy,* 24, 419–440.

Ulrich, K.T. and W.P. Seering (1990), “Function sharing in mechanical design,” *Design Studies,* 11 (4), 223–234.

Ulrich, K.T. and S.D. Eppinger (2004), *Product design and development,* 3rd edition, New York: Irwin/McGraw-Hill.

Appendix

Maxcut representation of the module partitioning problem

The separation of modules into master and secondary categories can be achieved by solving the following maxcut problem. Consider a dependency graph $G(V, E)$ in which the vertices $i \in V$ correspond to modules. An edge $e \in E$ connects a pair of vertices $i, j \in V$ if the associated modules are not independent; i.e., if the selection of one model constrains the choices of designs for the other module. It is sometimes advantageous to associate weights w_e for all edges $e \in E$ (notationally, we will interchangeably write w_e and w_{ij} with the understanding that edge e connects the pair of vertices i and j). The weights reflect the relative degree of dependence between pairs of modules. For example, one can use as weights the number of pairs of module designs that are mutually infeasible for a pair of modules; these values then reflect the extent to which a pair of modules is mutually constraining. The special case $w_e = 1$ for all edges $e \in E$ implies no difference in the importance associated with one or another pair of attribute dependencies.

A cut comprises the edges between a partition of the vertices into two subsets $V_1, V_2 = V \setminus V_1$. A cut in which there are no edges within either subset (i.e., all edges are between subsets) has the property that all modules within a partition are independent: choosing one module design places no feasibility constraints on the choice of other modules in the subset. Such a solution may not exist; one then wishes to find a partition of vertices that maximizes the weighted sum of the edge weights across a cut. This is the maxcut problem. It can be written as the following integer quadratic program (see, e.g., Goemans and Williamson 1995):

$$\begin{array}{rl} \text{Maximize} & \frac{1}{2}\sum_{i<j} w_{ij}(1 - y_i y_j) \\ \text{subject to} & y_i \in \{-1, 1\}, \text{ for all } i \in V. \end{array}$$

In the above formulation, $y_i = -1$ if vertex i belongs to (say) V_1; and $y_i = 1$ if it belongs to the other subset of vertices. Note that $y_i = y_j$ if both modules i and j are in the same subset and then the associated term in the objective function has a value of zero; but if i and j are modules in different subsets, then $y_i y_j = -1$ and the associated term in the objective function is

$$\frac{1}{2} w_{ij}(1 - y_i y_j) = w_{ij}.$$

The optimal solution to the maxcut problem is a partition of the modules into two sets so that the weighted number of dependencies is maximized across subsets. As the total number of edges in the graph is fixed, this is equivalent to minimizing the weighted number of dependencies within each subset of modules. The maxcut problem is NP-Hard (Garey and Johnson 1979) and so it has no exact polynomial time solution procedures unless P=NP; but there are well known approximation algorithms (e.g., Goemann and Williamson's (1995) randomized algorithm that uses a semidefinite programming approach).

Given a maxcut solution for a specific set of modules, we define one subset as the collection of primary modules; and the other set as a collection of secondary modules. We then propagate constraints by first selecting designs for primary modules and then imposing the implied constraints on the designs for the secondary modules. The labeling of maxcut partitions as primary and secondary is based on the judgments of engineers and designers.

Table 1: Conjoint attributes used in MCD study

Attribute	Number of levels
1. Fuel economy	4
2. Make	2
3. Cargo capacity	4
4. Turning circle	4
5. Drive type	4
6. Passing acceleration	4
7. Towing capability	4
8. Transmission type	4
9. Front interior package	4
10. Rear interior package	4
11. Exteriors	4
12. Engine technology	4
13. Price	4

Table 2: Engineering modules and module designs in MCD study

Module	Number of designs
1. Floorpan	12
2. Cowl	4
3. Front suspension	9
4. Rear compartment plan	18
5. Motor compartment	16
6. Rear suspension	9
7. Engine	35
8. Transmission	11
9. Sideframe	18

Table 3: Linking fuel efficiency and module designs in MCD study

Gallons per 100 miles = $a + b$ Curb weight $+ \sum_i c_i x_i$
Curb weight = $k \cdot$ Floor area
Floor area = Overall width $\times$ Overall length
Overall width = max$\{(d_1$+Front shoulder room), (d_2+Track$\})$
Overall length = (Front overhang) + (Wheelbase) + (Rear overhang)
Wheelbase = (Axle to Ball-of-Foot) + (Ball-of-Foot to Rear H-Point)
+ (Rear H-Point to Axle)
Axle to BOF = (Dash to Bumper) $-$ (Front overhang) + (Ball-of-foot to Dash)

— — — — — — — — — — — — — — — — — — — -

Note: a, b and c_i are model parameters; k, d_1, and d_2 are positive constants; x_i denote dummy variables representing the alternative engine, transmission and the final drive modules.

Table 4: Attribute bandwidths for the test wells and the proposed model's recommendations

Attribute	Test well bandwidth(s)	Proposed model bandwidth
1. Transmission	1 level	$2-3$ levels
2. Drive	1 level	2 levels
3. Passing acceleration	0.5 sec. and 1 sec.	1.5 sec.
4. Fuel economy	1 mpg and 3 mpg	$2-3$ mpg

Figure 1: Overview Of Product Portfolio Optimization System (PROPOSE)

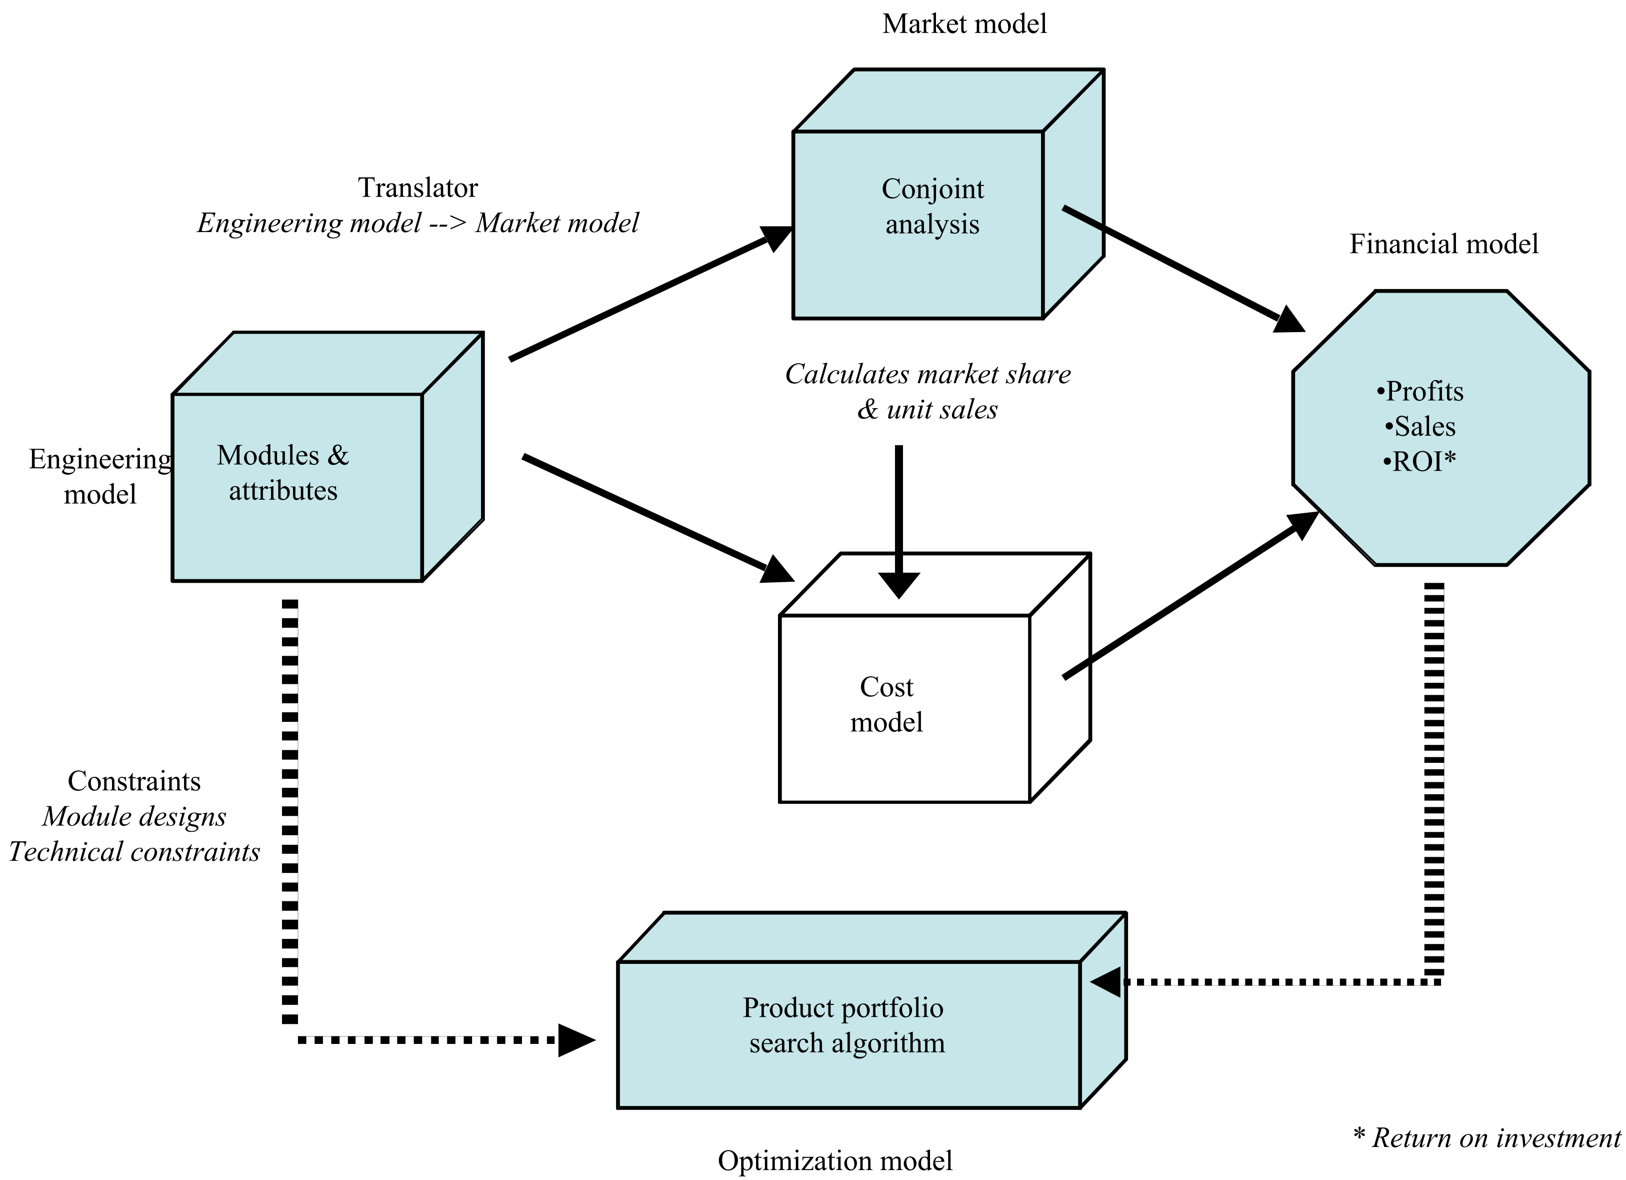

Figure 2: Hierarchical Representation of a Product as a Collection Of Module Designs, Engineering Attributes, And Levels of Engineering Attributes

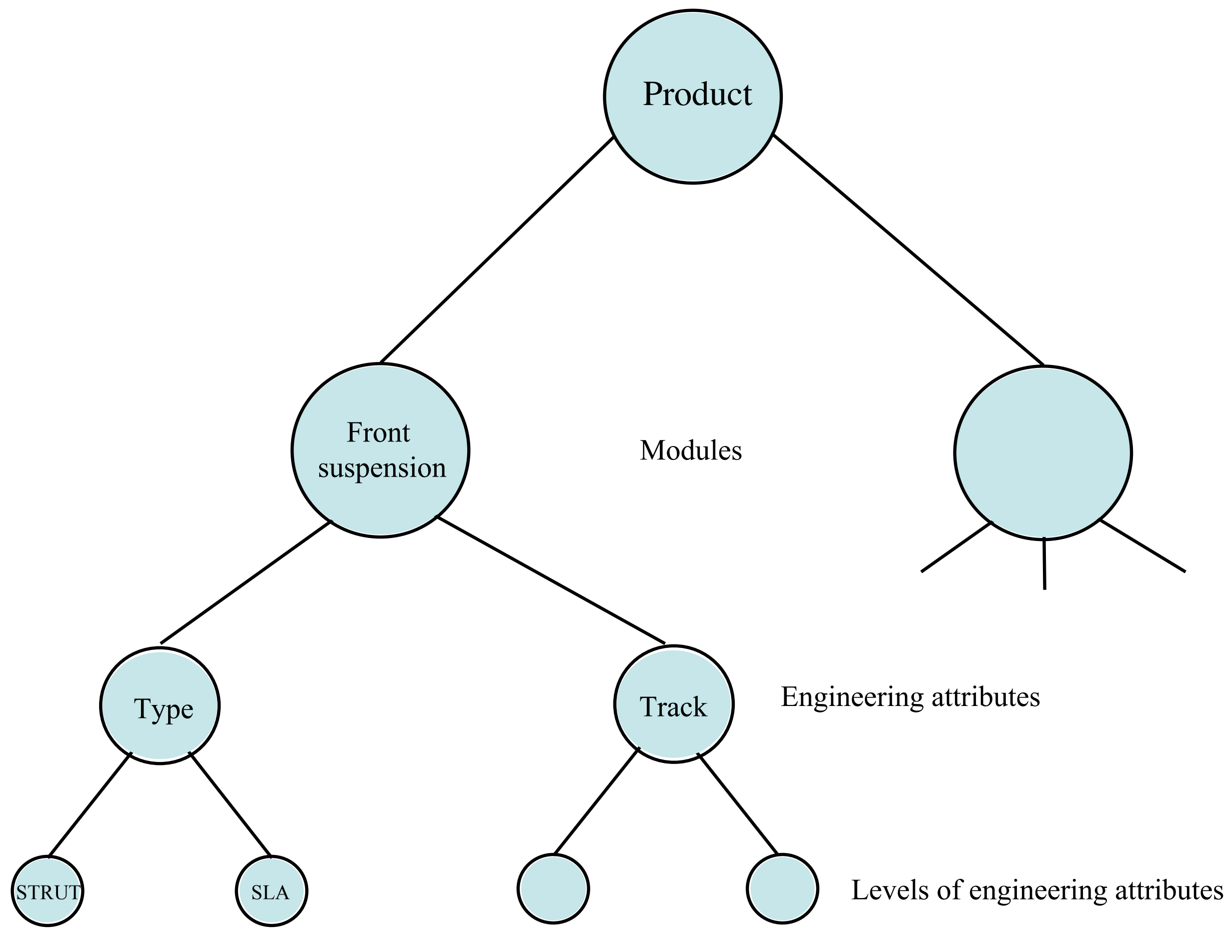

Each combination of engineering attributes levels describes a possible module design; not all such combinations may be feasible. A product is described by a collection of module designs, one design per module.

Figure 3: An Illustration of Decomposition for Fuel Economy

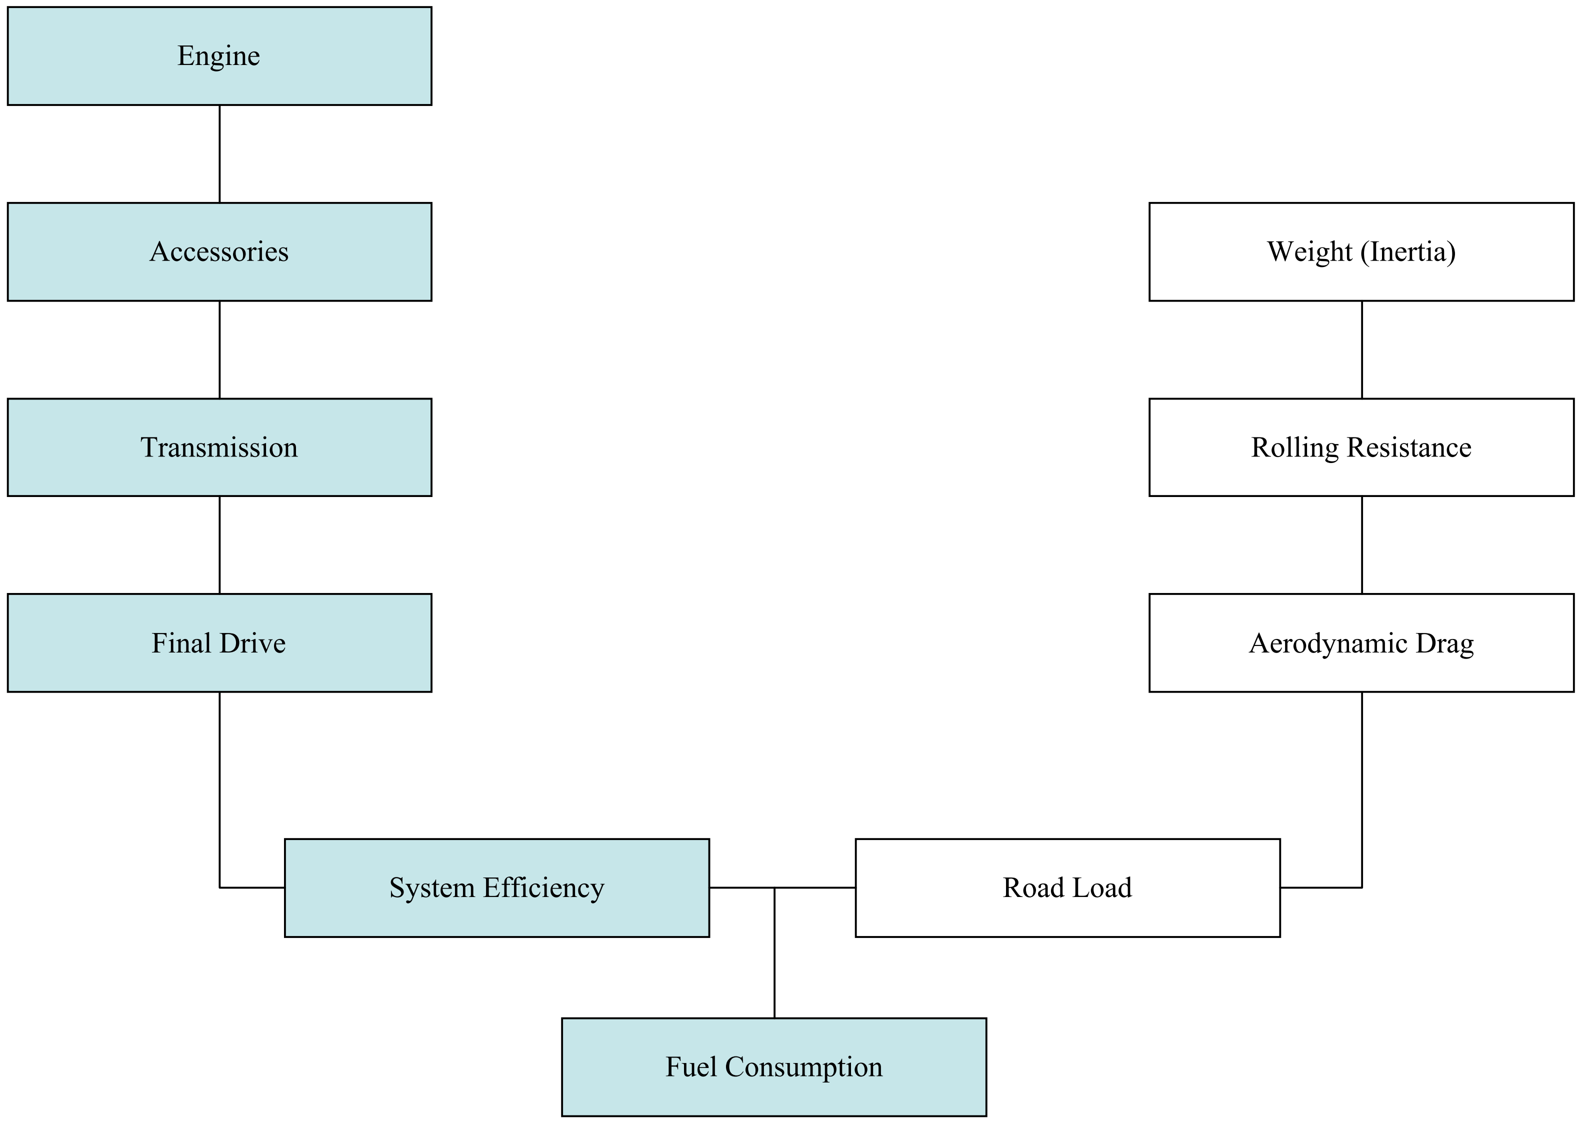

Source: Effectiveness and Impact of Corporate Average Fuel Economy (CAFÉ) Standards (2002)

Figure 4: Empirical Relationship Between Automobile Fuel Efficiency And Weight

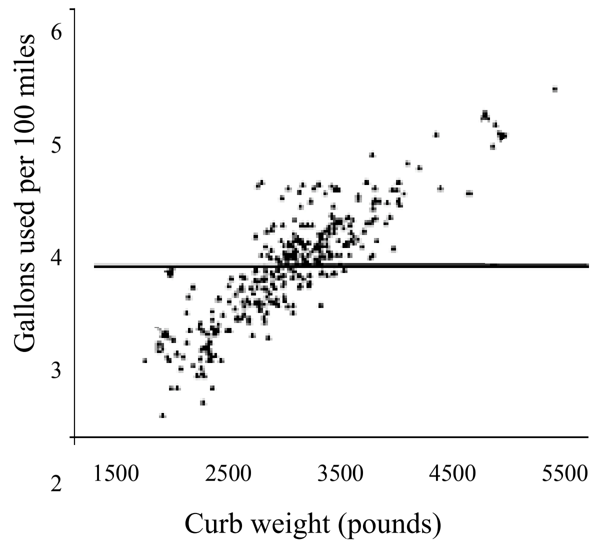

Source: Effectiveness and Impact of Corporate Average Fuel Economy (CAFÉ) Standards (2002)

Figure 5: Graph Representing Modules (Vertices) And Their Dependencies (Edges) For The GM Mid-Car Design Study

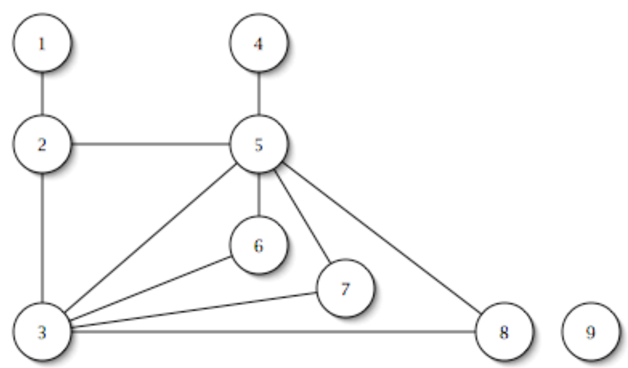

The labels for the vertices reflect the numbering of engineering attributes in Table 2

Figure 6: Partition Of Modules (Vertices) Into Independent, Primary And Secondary Modules
For The GM Mid-Car Design Study