Process Cost Modeling: Strategic Engineering and Economic Evaluation of Materials Technologies

Frank Field, Randolph Kirchain, and Richard Roth

Production cost is a vital performance metric for engineering and management analysis. Despite its obvious relevance throughout the product development cycle, cost analysis has not been a focus of the design engineer. In part, this is because of some key misunderstandings of what cost is—engineers have not been trained in the techniques that tie manufacturing cost to the technical and design parameters with which they are more comfortable and familiar. While there have been many calls for a closer relationship between engineering and economic analysis, these key conceptual obstacles, in conjunction with the limits of the computational tools available, have limited the integration of cost analysis into product and process development. This paper summarizes the conceptual limitations that need to be overcome and presents a basis for revising the notion of process cost analysis. Moreover, it presents a series of cost analysis cases that demonstrate the way in which the notion of "context" lies at the heart of effective use of engineering cost estimates.

INTRODUCTION

One of the attractions of using cost as a basis for decision making is the apparent simplicity of the metric—an economic measure of the resources employed to undertake a set of actions, typically yielding a good or service. The notion of cost is a part of everyone's day-to-day experience, serving as a well-established basis for evaluating alternatives by balancing the costs of acquiring or producing against the value received.

However, engineers are usually far less comfortable with cost when tasked with relating it to a set of specific technical or design changes. When asked to establish the cost consequences of a technical change, engineers typically fall back upon the principles of engineering economics, emphasizing the dynamics of cash flows and the evaluation of their net present value in order to select among alternatives. Unfortunately, these techniques rely upon the availability of cash flow data that, by definition, are not readily available, in that they depend upon an economic assessment of actions that have not yet transpired.

In the face of the need for these data, engineers typically rely upon a handful of simple methods for characterizing the relationship between cost and production, generally driven either by a set of normative methods based upon past experience or by a reliance upon the specific skills of cost estimators who are largely divorced from the day-to-day considerations of engineering practice.[1] In effect, as cost moves toward the center stage of the design decision, it becomes increasingly difficult for an engineer to comfortably work through and rigorously defend a cost analysis.

In large part, this difficulty arises from the fact that cost has traditionally been the focus of the accountant rather than the engineer. The typical instruments for cost assessment are retrospective, directed toward the performance of existing plants and products and tasked with informing the decisions of plant management rather than the decisions of those inhabiting design and R&D offices.[2,3] Because of this focus upon existing facilities, as well as the need to support management and operational, rather than technical, decisions, cost analyses have tended to be narrow in scope and, thus, narrow in utility when considering the kind of technical questions that are of increasing importance to the modern manufacturing firm. Accounting cost systems are largely focused upon how well (or poorly) an operating unit is performing, while the question facing product and project planners is how well a unit might perform with a change in product or process technology. This transition from the assessment of the normative to an assessment of the speculative demands a kind of analysis that engineers are perfectly happy to undertake in the technical domain, but loath to rely heavily upon in the economic one.

As a consequence, engineers have focused on other ways to speak of the benefits of design and process change. Because so many cost instruments are so poorly suited to the questions confronting them, the first response when confronted with cost questions has been to defer those analyses to others, while focusing upon the kind of scientific or engineering metrics that are the lingua franca of the engineers' daily work—and thus, that engineers are more comfortable deriving and defending. Unfortunately, those who adopt this strategy find that their largely unfamiliar technical metrics are set aside by firm management in favor of more poorly supported, but wholly familiar, cost metrics devised by those with only a passing familiarity with the technical issues at the real heart of the issue.

A classic example of this difficulty can be developed from a consideration of the following question: how will cost change if a material substitution is undertaken? While the simplest evaluation would focus upon the price differential between the current and new material, in general there is far more to it. The process engineer knows that a different material will mean different operating conditions for the production line and the possible need for a completely different process. Similarly, the product designer might find that the use of a different material will lead to changes in product performance, changes that may require a redesign in order to ensure that performance targets are met. Comparisons of material prices will miss these important consequences of material change, yet they frequently can dominate the decision-making process, particularly if these other consequences are presented only in technical terms. Making the connection between these technical changes and production cost is key to successful implementation of any process or product change, not to mention identifying the associated opportunities and perils.

A HIERARCHY OF DESIGN PROPERTIES

Designers are routinely confronted with the need to make choices based upon product or process features long before all of the consequences of these choices can be known. Because many of these consequences are important to the success of the product, a considerable amount of effort has gone into devising methods that enable designers to predict these consequences so that designers can understand the implications of the design choices that they are making. Part of the work to date suggests that there is a specific hierarchy that can be used to think about these consequences and their relationship to the designer's choices.[4]

Intrinsic Consequences/Properties

At the simplest level, there are consequences which can be called intrinsic to the designer's choice. An example would be the weight of a specific design. By selecting a material and a geometry, the designer has defined the weight of that design. Moreover, if the designer wishes to change the weight, an appreciation of the trivial relationship between the density of the material selected and the volume of material required to achieve a specified geometry makes it easy for the designer to effect a change, either through a change in geometry or in the specification of a material of a different density.

Resultant Consequences/Properties

While some product characteristics can be predicted with a handful of simple relations based upon basic design parameters, there are other product characteristics whose prediction depends upon much more. For example, while mass depends merely upon a material specification and design geometry, a load limit requires a larger set of parameters, many of which are external to the design itself—the magnitude and direction of desired loads, the number of cycles, and the boundary conditions for the structure.

This class of properties depends not only upon the intrinsic characteristics of the design, but also upon characteristics external to the design and an analytical modeling framework that relates these intrinsic and external characteristics to the property (or properties) of interest. At this level of complexity, these consequences of the designer's choices can be called analytic or resultant,[4] meaning that, while employing intrinsic properties, they are derived from the application of engineering and scientific models to predict the behavior of the design. For example, while geometry and material can be used to generate the intrinsic design property of weight, the application of a set of more elaborate structural analyses would be required to know its failure modes under loading conditions. While designers might have some feel for certain load limits and how changes in the design might affect that performance, they would certainly employ more formal methods of analysis—not only to lend support to their intuition or experience, but also to provide a rigorous basis for any explanation or defense of their design choices to others.

Emergent Consequences/Properties

At the most difficult level of evaluation complexity, there are the consequences of the designer's choices that can be called emergent, meaning that they are only apparent when the product is considered within an even wider realm—the larger context within which it will be employed. Emergent properties are, essentially, properties dependent not only upon characteristics of the product, but also upon the ways in which the product is produced and used, the characteristics of the market, and other, broader domains. Another way of thinking about emergent properties is to think of them as properties that an object clearly has, yet these properties cannot be derived solely from consideration of the object's intrinsic and resultant properties. Such properties could include characteristics like stylishness, reliability, safety—all clearly features of a product, yet ones that can only be estimated by looking beyond the specifics of a product design and its engineering context to consider the interactions between characteristics of the design and those of the other systems—economic, environmental, etc.—that influence and are influenced by the product.

While there are many other product characteristics that are influenced by engineering and design choices, cost is one of the few that leave engineers conflicted about its analysis. On the one hand, cost is viewed as a relatively simple thing, to be resolved in the course of a design or development project, but certainly not one that should be at the center of the engineering effort. On the other hand, cost is also frequently viewed as being far too "messy" a problem to tackle, lying far from the domains and abstractions of engineering science. In either case, the result is that cost is frequently left as a question for other (usually operational or management) teams.

In many ways, the problem arises because a cost analysis is notionally distant from the conventional realms of engineering analysis. While there are hard, technical factors that are fundamental to understanding some elements of cost (e.g., technical relationships between part mass and material prices that can be used to derive the cost of the materials in the part), there are also a host of "softer" features of processing that must be employed (e.g., production volume, operational processing yields and losses, discount rate, etc.) before a meaningful cost analysis can be undertaken. This combination of hard technical analysis with soft, normative, and operational analyses can make the notion of a "cost analysis" seem almost chimerical. To help to resolve this dichotomy, it is useful to reconsider the ways in which the properties of products are developed, and to consider where the property known as "cost" fits into those models of analysis. A review of a basic hierarchy of these properties suggests that cost's position in this hierarchy can help to explain the difficulties associated with this important engineering property. (Read about a hierarchy of design properties in the sidebar.)

While properties from each level of the design hierarchy might be taken into consideration when product or process choices are being made, a designer is confronted with dramatically different sorts of design options, depending upon where the properties of interest might lie in this hierarchy—intrinsic, resultant, or emergent. In part, this is a consequence of the kinds of tools available to assess these properties. Intrinsic properties, for example, are largely developed from the physical properties of the materials and geometries employed in the design, while resultant properties depend not only upon those features, but also upon a set of engineering science models that can be used to frame the prediction of these characteristics. Emergent properties rely not only upon the tools and conceptual bases of intrinsic and resultant properties, but also upon consideration of the ways that other important systems impinge upon and respond to the design being studied.

While the availability of and familiarity with engineering tools to develop each of these classes of properties is vital to the designer, the key to good design lies in the conceptual framework that the designer employs to relate a design's properties to the design goals. A great deal of engineering education is centered upon developing this kind of understanding for intrinsic and resultant characteristics, while effective training in the notions of emergent properties has not been developing at the same rate. As a consequence, there is a temptation to treat cost as if it were one of these more familiar kinds of properties, rather than fully embracing the implications of cost as an emergent property.

COST AS AN EMERGENT PROPERTY

The notion of an emergent property helps to explain the peculiar conceptual difficulties that engineers and designers face when attempting to incorporate economic considerations into their design processes. Traditional economics texts routinely speak of cost in two basic ways. The simplest notion of cost is embedded in the concept of the budget constraint—the sum of the resources required in production, weighted by the price of those resources. In this view, cost is derived by enumerating the resources required, multiplying these resources by their prices, and then summing the results—$\sum_{i=1}^{n} p_i x_i$. The more complex approach is the notion of the supply curve—a function describing the least expensive combination of resources x_i required to produce a desired amount of output Q. This supply curve is developed through optimization, minimizing the budget— min $\sum_{i=1}^{n} p_i x_i$ —subject to a constraint on the minimum required output, Q, defined by a production function—$Q \le f(x_1, x_2, \ldots, x_n)$. The resulting solution to this optimization describes how, at each level of output Q, one can convert resources x_i into Q most efficiently (i.e., the least costly way to produce output Q).

To a certain extent, these two notions, the sum of the cost of factors of production and the minimum cost of producing an amount of output, represent two extremes of the notion of cost. In one case, the notion of cost is a fixed amount, comparable to an intrinsic property. While factor prices might change, the fundamental notion is that factor prices are fixed or at least relatively stable over time. At the other extreme, the supply curve is an incarnation of the notion that cost is the result of a mathematical optimization, establishing that, for a given technology and set of factor prices, there is one "best" way to produce a desired amount of output that results in the lowest possible cost to the producer—a resultant property of the product. In effect, at one extreme, cost is an amount established in the marketplace while, in the other case, cost is an amount established by the technology.

However, both of these notions simplify the realities of cost. Treating cost as a sum of the factors of production, weighted by the price set by the market, renders cost into something that the designer must merely accept. Treating cost as the result of a technological optimization seems to establish that there is something like "a" cost. In both cases, the underlying notions drive one to conclude that cost is a particular property of a good, and that one's effort should be directed toward classifying it. Moreover, there are clear operational difficulties when actually attempting to develop cost under real circumstances. For example, the budget (or the budget constraint) refers to prices p_i. Which prices are to be employed: list price, transactions price, or transfer cost? How does one treat the fact that these factors change with time? And what does that mean about the notion of optimization, whose solutions are demonstrably fragile with changes in such parameters?

In fact, a design has many possible costs. Moreover, the cost of a design cannot be derived solely by considering only the components of a design and their properties. Rather, cost is a consequence of the relationships among these components that are established in the context of the manufacturing process, which itself is influenced by the economic and organizational systems that govern that process. Thus, cost is not a simple intrinsic property (like a factor price) or a resultant property (like an optimized supply curve). Rather, cost is an emergent property, dependent upon the context within which products are designed and manufactured. And being able to capture that nature requires that effective analysis of cost rely upon a consistent and transparent representation of that context.

THE CONTEXT OF COST

If cost is an emergent property, how best to conceptualize the ways in which these properties emerge? While there are many possible ways to frame this discussion, the field of materials science offers a useful framework for consideration that can be readily extended to usefully treat this question.

In materials science, the working paradigm for evaluating and predicting a material's performance is based upon the realization that the physical behavior of a manufactured good is dependent upon the interrelationships between the properties of the material, physical (micro)structure of that material, and the processes employed to generate that structure. (See, for example, the Carnegie Mellon University Department of Material Science and Engineering website *http://neon.mems.cmu.edu/ugrad/index.html*, which states "The overarching paradigm of materials science and engineering is thus to exploit the connection between processing, microstructure and the properties of a material in order to choose a material that will fit the performance criteria for a given application." Also, see the National Academy of Sciences[5] for the now-classical presentation of this four-part interrelationship in the form of a tetrahedron.) Only by taking all of these elements into consideration can a materials scientist be able to predict the performance of a manufactured good, and only by understanding these interrelationships can a materials engineer act to develop a new manufactured good to meet a desired set of performance targets.

This set of notions is directly applicable to the problem of developing a manufacturing cost. Like the question of materials performance, a product cost is dependent upon the architecture and composition of the product, the properties of the elements employed in that composition, and the processes whereby those elements are shaped to yield that desired architecture.

The challenge, then, is how to represent those interrelationships in a way that can serve as a tool for the designer. (See the sidebar for a description of the basics of cost modeling.)

BENEFITS OF MODELING

As suggested at the outset of this article, the primary rationale for developing a model of production cost is to be able to translate the complex and interrelated consequences of design and/or process technology changes into a cost metric. Moreover, a cost model undertakes that translation in a fashion that is technologically defensible while retaining the transparency necessary to support that defense.

Once such models are in hand, it rapidly becomes clear that there are many more uses for them than merely estimating costs. Two broad applications of these tools have proven to be particularly valuable: as tools for improving communication and discussion of cost with diverse groups, and as tools for strategic analysis of design, material, and process choice. Although the second application is the most common justification for developing a cost model, substantial benefits frequently derive from the use of a cost model as a platform for creating a working dialog about the underlying factors that drive production cost.

Cost Models as Instruments for Dialog

As cited earlier in this paper, one of the most difficult things about cost is that it is such a mutable notion within the engineering and business community. The propagation of terminology for different ways of accounting cost gives evidence to this confusion: “life cycle cost,” “fully accounted cost,” “operating cost,” and “sunk costs,” just to name a few. As a consequence, meaningful use of cost information for decision making is drastically impeded, largely because many different definitions of cost are being used by each participant in the discussion. Even quite simple questions of cost can become almost impossible to resolve, particularly when these discussions cross organizational boundaries.

For example, developing the cost of a series of automobile bumpers was a part of the dissertation research of one of the authors.[14] The cost of the shock absorbers that were used to mount these bumpers to the automobile frame was surprisingly difficult to establish. Despite the fact that the part was listed in the supplier’s manual, conversations with design and manufacturing engineers in

BASICS OF COST MODELING

While there are many possible ways to develop a model of production cost,[6–8] the Massachusetts Institute of Technology (MIT) Materials Systems Laboratory’s research over the past 20 years has demonstrated that an effective way to structure a cost model is to think of it as composed of three interrelated and interdependent models: a technical process model, a production operations model, and a financial accounting model. The decomposition of the problem of cost estimation into these three categories helps not only to establish the scope and boundaries of the task, but also to structure the task of model development so that analyses of otherwise bewildering complexity can be handled. More importantly, this structure explicitly offers the modeler the opportunity to incorporate engineering and operational functionality into estimates of resource requirements, whose technical interactions are largely ignored or oversimplified in other modeling approaches. The notional interrelationship of these three components is illustrated in Figure A.

Process Model

At the heart of any manufacturing process is a set of technologies that are employed to accomplish production. These technologies can range from something as basic as the application of forces to reshape a material to something as complex as a combination of thermal and chemical processes to create fine structures of specific materials on the surface of a crystal. For any of these processes, there is a set of fundamental operations taking place that can be characterized according to scientific and engineering principles.

The field of chemical engineering has probably done the most to characterize the relationship between process definition and production costs. However, many analysts have recognized that there are key insights into cost that can be gained through a careful assessment of the ways in which engineering principles can be applied to the consideration of the economics of production processes. Fundamental engineering principles that are routinely employed in the process model include materials, energy and mass balances, the implications of process characteristics upon the rate at which processes can take place (e.g., cycle times), and the ways in which these principles of engineering science constrain and structure the ways in which the necessary resources for production can be employed. Classical examples of these kinds of treatments can be found in References 6, 9, and 10. The literature of chemical engineering (for example, References 2, 11, and 12) also has employed elements of process modeling to establish equipment and energy requirements as well as typical labor and materials requirements.

Operations Model

While the process model helps to structure the problem of cost estimation based upon the technical and scientific principles that underlie the manufacturing process, the process alone is insufficient to completely specify the costs of production. In large part, the cost of production depends upon how the technical process is physically implemented and how the actual operation of the physical plant is organized. The desired scale of operation is also a fundamental parameter in any cost model, because so many technical and operational decisions are predicated upon satisfying production targets.

In many respects, the modeling of operations helps to define how the manufacturer acts to optimize its use of the one resource that is most difficult to obtain—time. As a consequence, much of the modeling of operation is centered upon the ways in which rates of machine operation (typically determined by features of the process model) are managed by the plant operator to achieve the most cost-effective way to allocate the capital resources of the plant.[13]

This interplay among the technical constraints underlying the manufacturing process and the capital costs of the physical artifacts required to implement that process are the bumper development and production teams yielded cost estimates ranging from roughly $3 to $60—on a part with a listed price of about $15 in the supplier's catalog.

Creating a basis for meaningful discussion of the cost implications of product and process choices can be a powerful application of a cost model. The success of such an application depends heavily upon the ways in which the cost modeler decides to implement those elements described in the preceding section. In practice, cost modeling is more than merely internalizing these notions. Like any modeling methodology, it is the way in which the problem is framed that determines the applicability and utility of the method. This is particularly true in the case of cost modeling, because the nature of the resulting cost analysis tool is intimately tied up in the set of questions the modeler has in mind when applying the method.

primary examples of the way in which manufacturing cost is an emergent property. In order to make the most effective use of that capital, the plant operator must make key choices about how to balance operating time against maintenance time, how to allocate capital resources among the various operations of the plant, how to source the factors of production (labor, materials, energy, etc.), and other equally mundane, yet vitally important decisions. Where there are many possible outcomes of those decisions that will result in successful production, only a handful (or less) will result in an efficient use of all the available resources and, thus, a cost-effective product.

Capturing that decision process, structured and informed by the process model, yields a set of operational parameters that can be used to characterize the full set of resources required to achieve a desired production output (see, for example, the detailed methods listed in Reference 1). Once that set of resources is fully determined, the modeler can turn to the final stage of the analysis.

Financial Model

Once one has successfully developed a complete enumeration of the resources required to produce a product, and established costs of those resources, there remains the task of converting resource requirements into their economic costs. In the case of factors of production directly employed in production (such as energy, materials, and labor), the task is fairly straightforward. A simple accounting of the factors required, weighted by their purchase price, is sufficient.

In the case of less direct factors, more indirect allocation strategies must be employed. The fundamental questions tend to center upon the ways in which capital costs should be allocated to units of production. Classical finance models provide standard methods for performing this distribution, centered upon the notion of the opportunity cost of capital—essentially requiring that the use of capital goods in production must yield a financial return that must equal the return that would otherwise be required by the firm. (Many cost analysts erroneously conflate the notion of opportunity cost of capital with the notion of depreciation. Although depreciation is typically carried on a company's books as a cost, it is no such thing. Rather, depreciation is an artifact of government tax policies, giving profitable firms a way to shelter income from taxation by creating a new kind of cost item that is based upon the cost and age of a firm's capital stock. While policymakers tend to justify the parameters of the equations used to calculate this income shelter using language suggesting that this capital stock "wears out," these calculations only rarely are reflective of the actual longevity of most capital goods, and it is important to avoid the confusion that can arise from this approach. The simplest way to avoid it is to remember that depreciation is not a cost, but rather a tax credit and should be treated accordingly.) This required return can be viewed as another operations parameter, although its value is reflective of a broader strategy than resource efficiency. The theory of finance suggests that there are many possible ways to estimate and employ this opportunity cost. Thus, a good cost model should generally offer full transparency in this analysis so that alternative strategies can be readily incorporated.

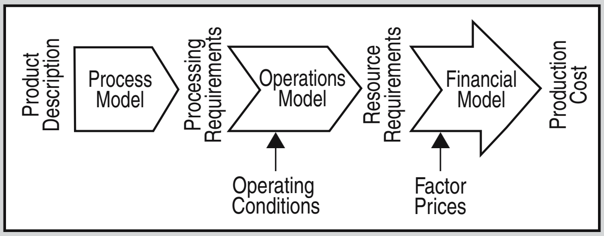

Figure A. Conceptual decomposition of a technical cost model.

Most cost models ultimately become either ones that focus upon process or ones that focus upon product. This evolution is a consequence of the way in which the modeler strives to sharpen the cost tool to treat the family of problems of interest. The model designer is always torn between the desire for maximum model precision and flexibility and the set of parameters required of the model user; an increase in precision and flexibility almost always requires a concomitant growth in the parameters required. The decision to concentrate on one class or the other is one approach to managing this problem. However, it is important to remember that any effort to reduce or control the number and complexity of model parameters should be tempered by a consideration of the ways in which a cost model is employed to explain (and defend) a cost estimate.

In short, the need to establish the credibility of a cost estimate challenges the cost modeler to make transparency a primary design goal. And, in turn, this transparency can then make the cost model an instrument for dialog among designers, developers, and business people.

One case, or rather a family of models, demonstrates this ability of a good technical cost model to engender a complex dialog about the economics consequences of technological and design innovations. In order to promote the use of advanced steels and steel-forming technologies in the automobile, an international consortium of steel companies set up a program known as the Ultra Light Steel Auto Body (ULSAB) program and its complementary projects (see the project information at *www.worldautosteel.org*). A fundamental element of this project was not merely to conduct a "paper study," but to conduct a full demonstration of design and processing technology innovations by actually constructing several automobile bodies. While this effort would serve to prove the technical feasibility of their efforts, a demonstration of the economic feasibility of the ULSAB design could not be accomplished so directly (doing so would have required setting up and executing full-scale production of the design—an unjustifiable expense). However, enough technical information about both the design and process were on hand from the prototype production that a technical cost model of the production of the ULSAB could be constructed.

The model was developed for the ULSAB consortium with a specific goal, beyond the development of a production cost estimate for the ULSAB: to expose all of the processing and operating parameters that drove the economics of the ULSAB production so that anyone who wished to challenge the results of the cost model would have to frame that challenge in terms of the operational or processing parameters of the model, rather than merely declaring that the resulting cost was "wrong."[15] Several studies have continued to employ this approach, using a cost model not merely as a tool for cost estimation, but also as a basis for discussion of the basic notion of manufacturing costs and their drivers.

Rather than present the model (some of whose results accompany the following sections), it is more instructive to examine the overall structure of the model design and how that structure helps support the effective discussion of cost. The four panels in Figure 1 summarize key elements of this model's structure. The two panels on the left describe the kinds of inputs that a typical cost model requires, mapping the three key structural elements of any cost model: technical modeling inputs like fabrication and other process and equipment data, operating data like production volumes, and financial data like factor prices. The "Body Inputs" panel lists the part-specific inputs that are required, such as the description of the part, the kind of process steps that will (or will not) be undertaken, and the ways in which subassemblies will be built up from the individual parts.

The heart of the model is summarized in the panel labeled "Fab & Assy Calculations." This portion of the model yields each of the individual cost elements that are eventually brought together to produce the total part cost, but the key element of this part of the model is the fact that the model also presents processing and operational results that can be reviewed by the model user. While the cost breakdowns are important, it is these processing and operational results, derived from the process engineering and the resource constraints, that are most useful in cost and technical discussions. Processing rates, equipment requirements, material yields—these results and others like them are what determine the final production cost and, when cost results are challenged, these are the results that are actually being questioned. A basic cost estimate tends to bury these important intermediate results in the calculations, and the delay required to regenerate cost estimates based on revised assumptions can frustrate the discussants.

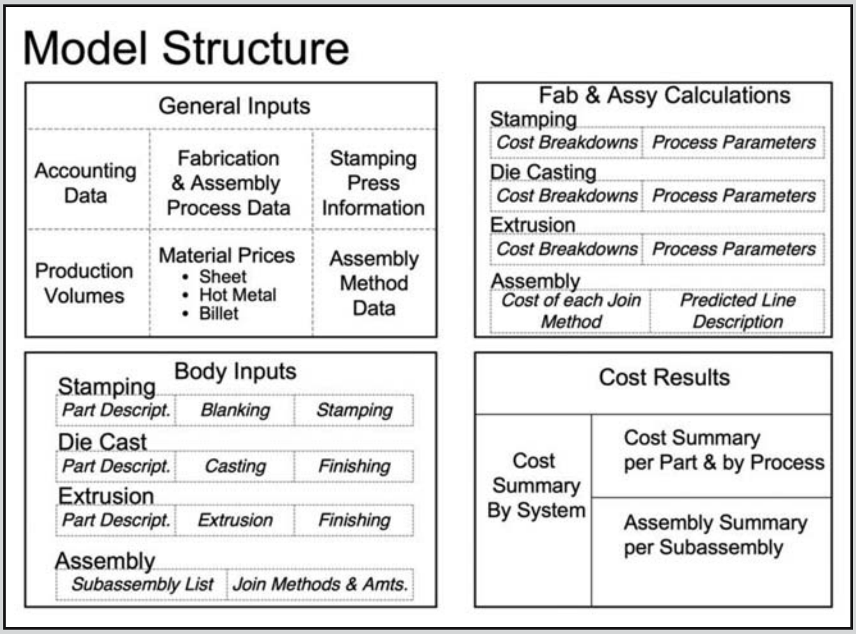

Figure 1. Ultra Light Steel Auto Body model family input and output structure.

With these important intermediate results immediately available, and directly tied to the generated cost estimates, it becomes possible not only to defend the basis for cost results in technical terms (e.g., "the low yield means we're throwing away thousands of dollars"), but also to question the basis for unsupportable claims (e.g., "the cycle time would have to fall to two seconds in order to achieve that cost"). This ability to directly translate processing and operating parameters into costs, as well as to explore the ways in which one changes or demands changes of the other, can lead to a kind of discussion of cost credibility that is otherwise almost impossible to achieve. The credibility of the cost analysis presented in the ULSAB and related programs is a direct consequence of being able to undertake exactly these sorts of discussions.

Cost Models as Instruments of Analysis

Of course, the fundamental justification for developing a cost model is the ability to estimate production costs. However, with a well-constructed cost model on hand, a host of related analyses are also possible. By building up the cost estimate through the systematic consideration of the processes that underlie it, the analyst is then able to explore which elements contribute most (and least) to the total cost. In many respects, a firm can most effectively drive down production cost by concentrating its efforts first on those technology choices that consume the greatest number of resources and therefore represent the greatest contribution to production cost. Armed with knowledge of these key cost drivers, appropriate stakeholders can focus development effort on their improvement or elimination. Moreover, because these models tie the technological, operational, and financial with the economic, analyses of the economic consequences of changes in operating parameters deriving from technological or operational improvements can be readily undertaken.

Another opportunity that cost modeling affords is the ability to develop a coherent basis for making comparisons across multiple design and processing alternatives. A process-based cost model provides a common platform that development and design teams can use to discuss cost issues. Because of the different perspectives of the design, engineering, and operational units within a firm (not to mention differing perspectives across firms), participants in multidivision or multifirm development programs may find it difficult to discuss alternatives because of different (and, typically unexamined) accounting perspectives. Because cost models can be constructed in a transparent manner, it is possible for all involved to focus design discussions on the underlying drivers and implications of cost rather than on the method of computing those values. By ensuring that a common set of assumptions underlies each set of estimates (as well as achieving the transparency necessary to explore those assumptions), comparisons among multiple designs can be developed and subjected to the broad range of analyses listed here.

In fact, with the ability to develop this level of insight into the underpinnings of cost, it becomes possible to undertake a wide range of strategic analyses of technological opportunities. Because process-based models build cost up from technological bases, it is possible to use these models to explore the impact of changing technological conditions. For example, a systematic exploration of the sensitivity of cost estimates to technical and operational parameters can reveal the key drivers of cost, particularly when combined with a realistic assessment of the range of reasonable variation in those parameters. However, it is also possible to use this information to establish the levels of technical or operational performance that would be required to achieve targeted cost levels. While a researcher might be exploring ways in which a process might be improved, a cost model can quickly establish which technical changes will lead to significant economic benefits. Alternatively, with a sufficiently robust model, the technological hurdles to achieving a particular production cost target can be established, not only in general but in terms of specific technological, operational, or financial factors. Moreover, such assessments can be tied to forecasting and planning efforts to consider the possible future evolution of the technology and its implications for the firm.

Fundamental Cost Tradeoffs In Production

Development	Material	Fabrication	Assembly
• Time & complexity • Sharing among related systems • Novelty vs. incremental refinement • …	• Raw material cost • Ability to reuse scrap and process materials • Properties & performance • …	• Feasible geometries & complexity • Final product properties • Rates of production, loss, defects, rework • Tooling & equipment requirements • …	• Number of parts • Diversity of processes, materials • Rates of production, loss, defects, rework • Tooling & equipment requirements • …

Figure 2. Key tradeoffs across product development and production that lead to product cost.

Central to all such analyses is the availability of a tool that directly couples design and process specifics with an estimate of production costs. With such a tool, an analyst is able to examine the connections between the technical and the economic domains that will govern the competitiveness and the feasibility of product decisions.

Rather than focusing upon the abstract, the power and utility of technical cost models can best be illustrated through brief reviews of some applications of these tools to problems in materials, design, and process choice. (Because of space constraints, only certain highlights will be presented here. Interested readers are urged to review the reports cited in the references.)

SAMPLE COST ANALYSES FROM THE AUTOMOBILE INDUSTRY

The many challenges facing the automobile industry can be used to develop case studies that illustrate the scope and depth of insight that can be gained from the application of technical cost modeling. While material strategy is not generally perceived as an obvious dimension of product strategy, it is frequently one of the fundamental features of success in this (and many other) industries.

It is instructive to step back and to explore briefly why this might be the case. Automobile product designers, like those in many others industries, are routinely challenged to develop new products that satisfy an ever-more-demanding set of customer needs and wants at a price that ensures economic viability. Moreover, the industry players are seeking to satisfy a mass, rather than a niche, market, implying that these products must be manufactured in volume. While “mass production” is a familiar notion, the implications of being able to produce 100,000 units (or more) per year are not as widely appreciated. Successful production at such rates depends upon a host of engineering tradeoffs, largely dominated by the way in which materials choice and the consequential manufacturing process options are balanced to yield an acceptable (and economical) product.

The automobile is a particularly dramatic platform upon which to exercise these engineering decisions. Unlike many other mass-produced goods, the automobile is composed of a wildly diverse set of components, serving a broad range of consumer needs. This diversity means that there frequently is more than one way to achieve desired product performance, but rarely more than a handful of ways to do so economically, particularly when facing such a competitive marketplace. And the consequences of making the wrong choice are not generally apparent until long after the point of no return has

Table I. IP Beam Part Data[19]

Name	Process	Mass (kg)	Reject Rate (%)	Trim Loss (%)	Melt Loss (%)	Cycle Time (s)	Total Investment (rel.)
Magnesium IP Beam Parts							
Main IP Structure	Die casting	8.1	1.0	2	3	142	1.00
Average Bracket (4 total)	Stamping	0.2	1.0	20	0	2	0.13
Steel IP Beam Parts							
IP Reinforcement, Upper	Tube bending	2.0	0.2	5	0	70	0.11
IP Reinforcement, Lower 1	Tube bending	0.4	0.2	5	0	49	0.04
IP Reinforcement, Lower 2	Purchased	0.3	N/A	N/A	N/A	N/A	N/A
Average Bracket (27 total)	Stamping	0.3	1.0	20	0	2	3.83

been passed. Thus, there are substantial needs to be able to better anticipate the consequences of inappropriate choices in a timely fashion.

Presently, the automobile industry is striving to develop cost-effective ways to reduce the mass of primary automobile structures. This is a continuation of the long-standing effort by the automobile industry to improve vehicle performance, and weight reduction is one approach to improving vehicle efficiency. Even without the impetus of rising fuel costs, there are important reasons to find ways to reduce vehicle mass, including improvements in vehicle performance and handling as well as the opportunity to add vehicle features that are heavy, most notably safety systems.

Material and design alternatives are regularly suggested as ways to achieve desired weight reduction targets. In the case of structural automobile applications, these alternatives routinely target component systems that are largely composed of stamped steel components. Stamped steel has been the dominant structural material/process combination in the automobile for essentially as long as there have been mass-produced automobiles.[16] (See References 16 and 17 for somewhat more nuanced takes on steel's centrality to this manufacturing transition.) Steel has occupied this position because of its relatively low cost and high performance, as well as its compatibility with high-speed forming and processing technologies. However, in an era of weight reduction, steel has been targeted because of its relatively high density. As compared to light metals and reinforced plastics, steel's density has become a perceived weakness of the material. However, these lighter materials also have weaknesses, not only in terms of their engineering performance (strength, stiffness, etc.), but also in terms of their cost. In general, on a per-pound basis, these materials are substantially more expensive.

However, material cost is only one of many factors that contribute to the cost that really matters to an automobile maker—the total cost of the final product. In that regard, there are other equally important contributors to final cost that must be taken into account when evaluating material and process alternatives. These include the costs of developing a design that can make the most effective use of the alternative material, the costs of forming the components that make up the component system and the costs of assembling those components. Figure 2 presents some of the key tradeoffs that have to be taken into account when evaluating production costs. A new production process might lead to parts consolidation and reduced assembly costs, but at the expense of a more complex fabrication process, a more demanding design effort, or a more expensive material. Only by considering these related and interconnected costs together can one make a reliable comparison among alternatives—an analysis that depends heavily upon the development of a complete picture of the interrelationships among all of these stages leading to the final cost of the component system.

Four cases of design and process selection in the automobile are presented in the following sections to illustrate the utility of technical cost modeling as a methodology for strategic evaluation of alternatives.

Instrument Panel Beams in Magnesium and Steel

A material option receiving consideration for light weighting is magnesium, for use in vehicle structures. While magnesium is significantly more expensive on a per-kilogram basis than the steel that it is typically substituting for, its advantages lie in the fact that substantial parts consolidation (and thus reduced tooling and assembly costs) is possible through its use in die castings. One specific application that has received consideration is in instrument panel (IP) structures, specifically the IP beam. In steel designs, the IP beam is typically a beam structure that forms a structural base to which a large number

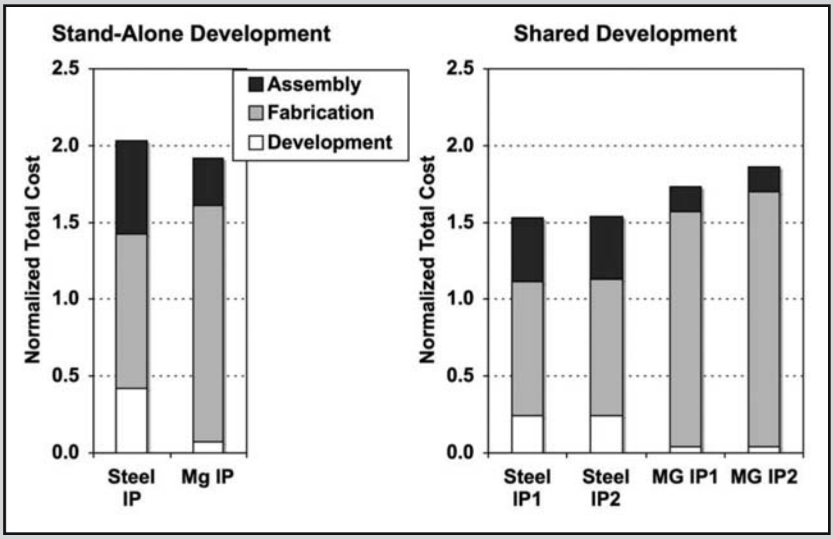

Figure 3. Comparison of total costs of development and production between steel and magnesium instrument panel (IP) beams at 75,000 units/year production over 5 years.[19]

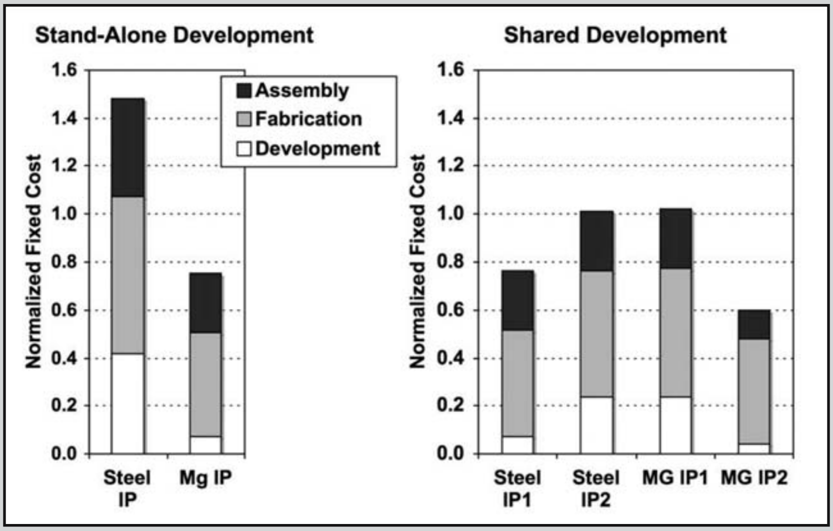

Figure 4. Comparison of fixed costs of development and production between steel and magnesium instrument panel (IP) beams at 75,000 units/year production over five years.[19]

of attachment brackets are assembled so that the rest of the instrument panel can be attached. Designers seeking to use magnesium in the IP beam have tried to minimize the number of attachment brackets required by devising designs where these mounting positions can be molded into the die cast component. Because of its casting properties, the degree of complexity that can be feasibly molded can enable a dramatic reduction in the amount of IP bracketry.

This forming complexity comes at a cost, however. As already indicated, magnesium is substantially more expensive than steel on a per-kilogram basis. However, that cost is offset in part because magnesium has a lower density, leading to a lower mass component. Another key processing difficulty associated with magnesium die casting is that its processing rate is substantially slower than that of steel stamping; in fact, the production rate of a die casting with the kind of complex geometry required for this application can mean a production rate ranging between one-half to one-quarter that of a conventional steel IP beam.

A study of the relative costs of a specific steel IP beam and its magnesium competitor reveals a third consideration after higher material cost and slower processing rates: the cost of developing the product and its associated manufacturing process.[18] Table I shows the comparative level of complexity of the two alternative designs and draws a clear picture of the relative tooling cost advantage that the magnesium alternative supplies. Moreover, this tooling cost advantage derives from a reduction in total parts, which also indicates that assembly costs for the magnesium IP beam should be reduced, helping to offset the disadvantageous cycle time.

Offsetting these advantages in tooling and assembly costs is the cost of raw materials. Here, the magnesium IP beam is at a significant disadvantage. In fact, the material cost for the magnesium IP beam is roughly three times the material cost of the steel IP beam, despite the fact that the magnesium system offers a roughly 18% reduction in weight. ($3.10/kg magnesium, $0.79/kg steel).

However, the cost analysis revealed that there was a subtler aspect to the cost comparison that had to be taken into account—the role of development cost in total production cost. One of the key advantages that an established manufacturing process affords a firm is the fact that familiarity and facility with the process and materials enables the firm to find economies across products that might otherwise be unavailable. In the case of steel stamping, where complex geometries are produced by assembling relatively simple stamped shapes, economies can be achieved by exploiting opportunities for component or design sharing. In this way, it may be possible to offset the cost penalties of multi-part assembly (which part consolidation is generally pursued to reduce) by savings in sharing designs and fabrication across components.

In the case of the IP beam, exactly this kind of economy can be achieved. Figure 3 shows that, while the total cost of the magnesium IP beam is less that that of the steel beam when the designs are developed wholly independently of all other IP beams, the cost advantage shifts to the steel IP beam when developers can exploit sharing of elements of design and production across two IP beam systems. Figure 4 shows that this cost advantage is largely derived from the significant fixed cost reductions that can be achieved through this sharing. This result illustrates the fact that established materials can have real economic advantages, and that the observed difficulty of introducing new materials into automotive systems arises from more than mere inertia on the part of the automaker.

ULSAB (AVC) -Innovative Processing and Design

This notion of material inertia suggests that finding approaches to weight reduction that rely upon incremental innovations within the broader context of an established development and production systems can lead to improvements in both vehicle mass and manufacturing cost. The case of the Ultralight Steel Auto

Table II. Design and Cost Summary, ULSAB Project[20]

	Reference Vehicle	ULSAB	ULSAB-AVC*
Steel Cost	$369	$416	$468
Forming Cost	$282	$250	$213
Assembly Cost	$328	$281	$291
Total Body Cost	**$979**	**$947**	**$972**
Tooling Investment	$68.0 million	$51.2 million	$40.3 million
Part Count	135	96	81
Body Mass	270 kg	203 kg	218 kg

Note: ULSAB-AVC was designed to meet more stringent safety (e.g., crash) criteria than the listed alternatives.

Table III. Aluminum vs. Steel Front End Analytical Statistics

	StFE	AlFE
Stamped Parts		
Steel	93	43
Aluminum	—	52
Stamping Dies		
Steel	72	25
Aluminum	—	44
Extrusions	—	7
Casting	—	3
Total Parts	**93**	**105**
Total Dies	**72**	**79**

Body (ULSAB) project shows that it is possible to introduce a new, more expensive material solution, provided there are sufficient processing and assembly cost benefits to offset the cost of using a more advanced material.[19,20] The ULSAB project was largely focused upon exploring whether innovative designs in steel could achieve significant vehicle weight reductions without significant cost increases.

The fundamental question at the heart of this analysis was whether high-performance steels could be employed in automobile manufacturing in such as way as to reduce the mass of the vehicle without compromising the performance of the body and the economics of its construction. A deeper question underlying this project was related to the question of the relative merits of incremental versus radical changes in design, materials and processes. Because the automobile is a largely ferrous structure today, this approach to vehicle redesign for lightweighting would build upon an existing industry expertise in steel sheet metal by demonstrating that, with careful process and design choice, more aggressive material strategies might not necessarily be needed to achieve significant vehicle mass reduction.

Because of the increased unit cost of the high-performance steel, it would be important to exploit possible processing and design opportunities that were available, but not conventional, in current automobile production and design. The program objective led to the use of advanced high-strength steels, complemented by the use of advanced processing technologies including tailor-welded blanks as well as sheet and tubular hydroforming. These manufacturing techniques were largely introduced to find ways to reduce part count, so that savings in assembly costs and reductions in the number of forming tools would offset the increased cost of the material employed as well as the increased complexity of some of the forming processes that were required.

The ULSAB study was conducted in multiple stages, with the ULSAB phase II effort completing in 1998 (ULSAB2) and the ULSAB Advanced Vehicle Concepts study finishing in 2002 (ULSAB-AVC). Additional studies focused upon specific subassemblies of the vehicle, including suspensions and closures. The ULSAB body design was developed in ULSAB2 and modified in ULSAB-AVC to treat two "sizes"—a C-class and a Partnership for a New Generation of Vehicles-class body. A notable modification of the ULSAB-AVC was its redesign to meet more stringent safety criteria than those set for the ULSAB and the reference vehicle.

While the ULSAB-AVC study developed cost estimates for a complete automobile, this discussion will only treat the body studies, which focused upon the performance and costs of using of advanced high strength steels in automobile body structures, as well as advanced forming technologies, particularly tailor-welded blanks.

Table II summarizes the results of the design and cost studies undertaken in the project. The immediate conclusion that can be drawn is that there are opportunities for weight savings in the steel car, provided the developer is willing to undertake substantial design effort. Moreover, that design effort must be tied to a careful consideration of the technological opportunities afforded by advanced steel forming technologies.

The ULSAB vehicles employ more expensive steel grades in more complex processes, but the resulting reduction in part count and material used means that savings in assembly can offset these other costs. While these results do not conclusively resolve the deeper question of relative merits of design radicalism versus incrementalism, they do tend to defend the merits of the incremental approach of limiting design and process changes in the absence of a quantum leap in performance. As the following cases suggest, while such leaps may be achieved with more radical design and process choices, balancing the value of these changes against their cost may be more problematic.

Table IV. Cost Breakdown (in dollars) for StFEand AlFE at Annual Production Volume of 225,000 Units

	StFE	AlFE
Material Costs	74	139
Forming Costs	120	240
Assembly Costs	161	200
Total Cost	**355**	**579**

Table V. Assembly Content and Materials Costs for StFE and AlFE Assembly at Annual Production Volume of 225,000 Units

	Assembly Content		Materials Costs	
	StFE	AlFE	StFE	AlFE
Spot Welds	1,274	292	$0.07	$0.02
Rivets	—	638	—	$25.52
MIG Weld (m)	1.40	1.72	$0.00	$0.00
Laser Weld (m)	3.15	0.98	$0.02	$0.01
Adhesive Bond (m)	2.45	14.21	$1.13	$6.54
Total Cost	**$161**	**$200**	**$1.22**	**$32.09**

Front End Studies—Innovative Processes and Design

Vehicle front end systems are a key subsystem for lightweighting innovation in the automobile. Aside from the general desire to reduce overall vehicle weight, there is also the fact that a more even weight distribution between the front and rear of the vehicle helps to improve vehicle performance and handling. Since the engine and major drive train components are toward the front of the vehicle, the design problem is about finding ways to reduce the structural components required to support these systems without compromising structural performance or manufacturing economics.

A Steel and Aluminum Front End Comparison

In one study, the economics and performance of two production vehicle front end systems, one all-steel and one with a significant amount of aluminum structures, were assessed. Tear-down studies were employed to determine the part and assembly characteristics of each system. The steel system (StFE) was composed entirely of steel stampings. The aluminum system (AlFE) employed

Table VII. Modeled Cost Breakdown at 225,000 Units per Annum Production Volume[22]

	Base ($)	Stamped ($)	Hydroform ($)
Material Cost	49	56	43
Ordinary Stamping Cost	58	28	18
Tailored Blank Stamping Cost	13	46	—
Tubular Hydroforming Cost	—	—	47
Rollforming Cost	1	—	—
Assembly Costs	94	74	97
Total Cost	**214**	**203**	**205**

Note: Columns may not sum to totals due to rounding.

both aluminum and steel stampings, as well as aluminum extrusions and die castings. The overall objective in the AlFE design was to achieve a front end weight reduction through the use of aluminum structures, containing costs through the use of manufacturing processes that would lead to parts consolidation. In particular, extrusions and castings were employed to reduce the stamped part (and tool) count as well as assembly steps. The part, process and die counts for the two systems are presented in Table III.

Despite the design goals, the fact that the AlFE aluminum structure was to be made a part of an overall steel vehicle caused substantial assembly complications. While the die-cast and extruded parts did mean a reduction in part count for those component subassemblies, additional steel stampings were required to accomplish workable mating of the aluminum front end system to the steel vehicle body. This requirement led to an overall increase in the total number of tools for the AlFE, as well as increased assembly complexity. In effect, the desire to offset the increased cost of the aluminum material through parts consolidation was defeated by the need to attach the aluminum system to a steel vehicle. Tables IV and V present the overall and assembly-specific cost breakdowns for the two systems.

These results suggest that, while aluminum can yield lightweight structures, putting it onto a steel body introduces complications in assembly that can challenge the economic feasibility of such a design choice. In particular, the high cost of the rivets and the large number of assembly steps required to manage the part count increase in the mixed material design (note that one-third of the parts in the AlFE design were steel) indicate that, while the AlFE design might offer significant performance advantages over the StFE design, the automakers clearly decided that they were prepared to accept a cost penalty in order to get that performance gain. It is reasonable to assume that, while the aluminum components helped to reduce the front end mass, its cost penalty has been a substantial incentive to further design refinements.

Table VI. Process and Assembly Characteristics of Front-End Structure Design Alternatives[22]

	Base		Stamped		Hydroform	
	No. of Parts	Mass (kg)	No. of Parts	Mass (kg)	No. of Parts	Mass (kg)
Stampings	32	30.77	12	12.75	10	8.25
Tailored Blanks	2	5.36	6	18.90	—	—
Roll Formings	1	6.39	—	—	—	—
Hydroformings	—	—	—	—	2	20.60
Totals	**35**	**42.52**	**20**	**31.65**	**12**	**28.85**
Spot Welds	796		675		143	
MIG Welds (m)	0.40		0.40		9.13	
Bolts	8		8		8	
Rivets	—		—		10	

Front Rail and Bumper—A-SP Study

Another front end study examined the front rail and bumper assemblies from a mid-size vehicle.[21] In comparison with a baseline design composed of steel stampings and roll formings, two alternatives were developed and evaluated for technical and economic feasibility. The first of these alternatives employed laser-welded blanks as a way to achieve parts consolidation and thereby reduce both forming and assembly steps, at the expense of introducing a more complex blank preparation process. The second alternative strove for further consolidation through the use of a tailor-welded tubular hydroformed component. In the case of both design alternatives, the goal was to achieve weight reduction and economic feasibility through the use of more specialized materials and processes. The fundamental economic question, once technical feasibility had been demonstrated, was whether the additional costs of process complexity and specialized materials would be offset by simplified assembly operations and related parts production activities.

Table VI demonstrates the overall trends of the design and complexity tradeoffs implicit in these alternatives. Note that, for purposes of achieving a one-to-one comparison among the alternatives, all systems required a certain number of steel stamping for the purposes of bumper and engine attachment. As the forming complexity increases in each case (moving left to right) the part count and mass declines, although the type of assembly technology becomes more complex (for example, the increase in metal inert gas welding and riveting required for the tubular hydroforming design over simpler spot-welding technologies). The cost analysis question thus becomes whether the cost advantages of part reduction will outweigh the cost penalties of specialized part forming and assembly.

The production and assembly costs of these three systems were modeled. Two key set of results are presented in Table VII and Figure 5. The table shows that all three systems have comparable costs, although the distribution among material, forming, and assembly costs varies significantly. Moreover, as Figure 5 shows, the relative sum of these costs does not vary widely over the range of reasonable production volumes.

These results suggest that there can be circumstances under which parts consolidation can pay for itself. More

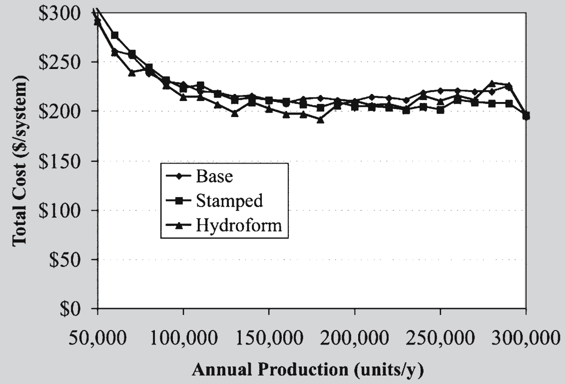

Figure 5. Total cost as a function of production volume.[22]

interestingly, these results also show that assembly savings through parts consolidation are not the only economic defense for the use of complex forming technologies and the associated expensive materials. In this case, the economic advantage of the hydroformed design derives from materials cost savings achieved through improved process efficiencies and processing cost savings resulting from a reduction in the total number of parts to be formed.

CONCLUSIONS

The modeling of production cost is a vital instrument of engineering and management analysis. Despite its obvious relevance throughout the product development cycle, cost analysis has not been a focus on the technological side of the house. In part, this is because of some key misunderstandings of what cost is; in part, this is because engineers have not been trained in the techniques that tie manufacturing cost to the technical and design parameters with which they are more comfortable and familiar.

However, these techniques do exist, and have been successfully employed to examine a wide range of technical options and design alternatives. The difficulty of accomplishing meaningful cost estimates derives from the fact that cost is a context-dependent metric of performance of a manufactured system. More precisely, cost is an emergent property, suggesting that it cannot be successfully evaluated without consideration not only of the specifics of the design or process under consideration, but also the ways in which these design objects are influenced by the systems that impinge upon them (markets, etc.).

While there have been many calls for a more effective partnership between engineering and economics, the development of techniques bridging these two domains has depended in part upon advances in simulation methods and computational availability, and in part upon a wider appreciation of the capabilities and promise of such a partnership. While practitioners continue to develop and refine these methods, the audience for these tools has been limited by the wider perception of cost as an inevitable constant, rather than a metric of performance that designers can strive to engineer. The examples presented here demonstrate that there are real opportunities for such an approach to product and process development, which can lead not only to more cost-effective designs, but also to a more nuanced approach to material and process innovations. Such an approach can move designers away from a "one size fits all" approach to technology choice and toward a set of engineering and design practices that explicitly recognizes that the economic competitiveness of every technology option is a function of the context of its application.

References

1. John Victor Busch and Frank Remsen Field III, "Technical Cost Modeling," *Blow Molding Handbook,* ed. Donald V. Rosato and Dominick Rosato (New York: Hanser Publishers,1989), pp. 839–871.
2. Donald E. Garrett, *Chemical Engineering Economics* (New York: Van Nostrand Reinhold, 1989).
3. Randolph Kirchain and Frank Field III, "Process Based Cost Modeling: Understanding the Economics of Technical Decisions," *Encyclopedia of Materials: Science and Technology,* ed. K.H. Jürgen Buschow et al. (Dordrecht: Elsevier, 2001), pp. 1718–1727.
4. Jaegwon Kim, "Making Sense of Emergence," *Philosophical Studies,* 95 (1-2) (August 1999), pp. 3–36.
5. *Materials Science and Engineering for the 1990s* (Washington, DC: National Academy Press, 1989); *www.nap.edu/catalog/758.html.*
6. Hollis B. Chenery, "Engineering Production Functions," *The Quarterly Journal of Economics,* 63 (4) (1949), pp. 507–531.
7. Soren Wibe, "Engineering Production Functions: A Survey," *Economica,* 51 (204) (1984), pp. 401–411.
8. V. Kerry Smith, "Another View of the State of Engineering Production Functions," *Economica,* 53 (212) (1986), pp. 529–532.
9. Anne P. Grosse, "The Technological Structure of the Cotton Textile Industry," *Studies in the Structure of the American Economy: Theoretical and Empirical Exploration in Input-Output Analysis,* ed. Wassily Leontief (New York: Oxford University Press, 1953), pp. 360–420.
10. Allen R. Ferguson, "Commercial Air Transportation in the United States," *Studies in the Structure of the American Economy: Theoretical and Empirical Exploration in Input-Output Analysis,* ed. Wassily Leontief (New York: Oxford University Press, 1953), pp. 421–447.
11. Frank C. Vilbrandt, *Chemical Engineering Plant Design* (New York: McGraw-Hill Book Company, 1942).
12. Cecil H. Chilton, editor, *Cost Engineering in the Process Industries* (New York: McGraw-Hill Book Company, 1960).
13. John Victor Busch, "Primary Fabrication Methods and Costs in Polymer Processing for Automotive Applications" (Master's thesis, Massachusetts Institute of Technology, Cambridge, MA, 1983).
14. Frank Remsen Field III, "Application of Multi-Attribute Utility Analysis to Problems in Materials Selection" (Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, MA, June 1985).
15. Richard Roth, Jeff Dieffenbach, and Jacqueline A. Isaacs, "Economic Analysis of the Ultra Light Steel Auto Body," SAE International Technical Paper No. 982399 (Warrendale, PA: SAE International, 1998).
16. David A. Hounshell, *From the American System to Mass Production, 1800-1932: The Development of Manufacturing Technology in the United States* (Baltimore, MD: The Johns Hopkins University Press, 1984).
17. Ramon Casadesus-Masanell and Daniel F. Spulber, "The Fable of Fisher Body," *Journal of Law and Economics,* 43 (1) (April 2000), pp. 67–104.
18. Richard N. Langlois and Paul L. Robertson, "Explaining Vertical Integration: Lessons from the American Automobile Industry," *The Journal of Economic History,* 49 (2) (1989), pp. 361–375.
19. Michael DeShawn Johnson, "A Methodology for Determining Engineering Costs and Their Effects on the Development of Product Families" (Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, MA, 7 May 2004).
20. Porsche Engineering Services, "Ultra Light Steel Auto Body: Phase 2 Findings," Tech. Rep. 1.0.2 (Ultra Light Steel Auto Body Consortium, 1998).
21. Automotive Applications Committee. "ULSAB-AVC (Advanced Vehicle Concepts)," Tech. rep. (Southfield, MI: American Iron and Steel Institute, January 2002), *www.steel.org/AM/Template.cfm?Section=ULSAB_AVC1.* [Accessed August 15, 2006.]
22. Auto/Steel Partnership, "Lightweight Front End Structure: Phase I & II—Final Report," Tech. Rep. (Auto/Steel Partnership, October 2006); *http://www.a-sp.org/database/ custom/ASP_Lightweight_Frontend_%20Phase1_and2_Final_Report.pdf.* [Accessed August 10, 2006.]

Frank Field, Randolph Kirchain, and Richard Roth are with the Materials Systems Laboratory, Massachusetts Institute of Technology, Cambridge, MA, USA. Dr. Field can be reached at (617) 253-2146; e-mail furd@mit.edu.