Int. J. Engng Ed. Vol. 20, No. 3, pp. 433–439, 2004
Printed in Great Britain.

0949-149X/91 $3.00+0.00
© 2004 TEMPUS Publications.

Decision-Based Design: A Vehicle for Curriculum Integration*

WILLIAM H. WOOD
Department of Mechanical Engineering, University of Maryland, Baltimore, MD 21250, USA.
E-mail: bwood@umbc.edu

Educational activities in engineering are often partitioned into analysis, experimentation, and synthesis. While depth is provided in each of these areas throughout a typical curriculum, little is done to unify them as interdependent components of engineering. This paper proposes an overarching structure that integrates these components under the notion that, at its core, engineering is decision-making. Design research provides the building blocks for decision-based design, providing a unifying framework not only within the design process but also among analysis, experimentation, and synthesis and extending outward to connect engineering to the society of which it is a part. Existing curricula can be integrated with only small changes.

INTRODUCTION

ENGINEERING PRACTICE has evolved tremendously with the development of more and more accurate methods for predicting the behavior of engineered systems. But increased prediction accuracy does not change the role of the engineering designer: to develop design candidates based on the needs of a customer and to project those candidates back into the customer's context to evaluate their efficacy. The uncertainty of this process now lies mainly in the representation of the external world within the modeling environment (e.g. expected loads, customer needs, manufacturing capabilities, etc.) rather than in the models themselves. Decision-based design (DBD) is a response to this shift, focusing attention on the remaining uncertainty toward providing the best possible designs.

The engineering curriculum can be divided roughly into three main components. Synthesis is directed toward generating design options in response to customer needs. Analysis operates over these options, evaluating their performance against well-defined performance measures or failure modes. Experimentation serves dual purposes: to confirm or elaborate existing theoretical models and to parameterize empirical models. The pedagogical question is: how should a curriculum weave these three components together? Historically, analysis leads in the engineering curriculum—students first learn mathematical methods and physical principles. Experimentation convinces students that their analytical results actually obtain in the physical world. Synthesis follows once students have gained a measure of engineering intuition, but it still suffers from an inherent 'fuzziness' that is difficult to reconcile with analysis and experimentation. In the end, students are left with a disintegrated view of engineering.

We propose to integrate the engineering curriculum by casting engineering into a decision framework. If engineers are society's tool-makers, then engineering students should be immersed in a socio-technical process from the very start, where they must decide which projects to consider, how to approach those chosen, how to evaluate their results, and when they have completed their task. Each of these decisions requires the generation of alternatives, their evaluation, the selection of the best alternative, the identification of information that might alter that selection, and consideration of the cost of making a commitment.

DECISION-BASED DESIGN

While often defined as a tool for evaluating a predefined set of design options [1], decision-based design can provide a framework for a rational design process [2, 3]. Under resource constraints, designers must trade deeper exploration of individual design concepts against the development of new concepts; framing this trade-off as a decision leads to a normative design process.

One key to this process is the realization that engineers must make decisions throughout the design process, from the earliest stages, where evaluation models and design concepts are both quite abstract to the final detail stages of design, where alternatives and their effects on design performance are more concrete. Decision-based design thus provides a unifying framework for this process: at any single point in the design process, options exist not only in the form of different concepts under consideration but also in the form of different possible design actions. If all design options are known and their performance

* Accepted 3 November 2003.

predicted to the highest possible precision, then decision-based design is simply a matter of choosing the best option (perhaps under the risk preferences or time-value of money preferences of the designer's enterprise).

Unfortunately, resource bounds prevent either of these conditions [4]. Decision-based design must instead focus on the development and refinement of both the set of designs to be considered and the performance metrics used to predict their value. In the early stages of design, it is useful to regard the performance metrics as a single, uncertain value function. In this case, uncertainty extends not only to the form of function for aggregating these metrics (and weights used within such functions [5]), but also to the set metrics that should be included. Typically, designers focus on one or two main performance measures, tuning the overall performance model as conflicts arise.

Narrowing the design space

Given a set of options and a value function, expected-value decision-making (EVDM) provides a normative method for determining the best option. Each option is evaluated in isolation, using mathematical expectation to resolve uncertainty in option performance as measured by the value function:

$$E(V|dec_i, \mathbf{c}, \mathbf{u}) = \int_{\Omega_\mathbf{u}} obj(dec_i, \mathbf{c}, \mathbf{u}) p(\mathbf{c}, \mathbf{u})\, d\mathbf{u} \tag{1}$$

where V is the value function, dec_i is one of the set of possible decisions, $\mathbf{c}$ are deterministic constraints (e.g. classification variable assignments), $\mathbf{u}$ are uncertain constraints (e.g. performance variable constraints), and $\Omega_\mathbf{u}$ is the state space of the uncertainty.

The designer is then free to choose the option which maximizes expected value. For example, Fig. 1 shows performance along with sets of decision plots for various partitions within a motor selection design space. Here, the overall objective is to minimize motor weight, the expectation of mass given motor length is shown in the upper left panel along with the probability density function for motor length in the design space. Choices for narrowing the permanent magnet, DC motor design space include: rare-earth vs. ferrite magnets, frame vs. frameless motors, and brush vs. brushless commutation. The '*' portions of the plots indicate the decision among all possible partitions that minimizes expected mass.

Refining the evaluation function

As Fig. 1 shows, the decision about which type of motor to select depends on the allowable length required by the design. Within the existing evaluation function (minimize motor weight), motor package length is not considered. For low length values a designer should have a clear preference for framed, brushless, rare-earth motors, perhaps letting motor selection drive significant portions of the design process. If longer lengths are allowed (or required by other constraints), the only clear choice would be for brushless motors; the designer might then look at secondary aspects of the design (e.g. cost) to decide which frame and magnet type to choose. In situations where the optimal choice of a design might change based on the value of an uncertain (or, in this case, ambiguous) parameter, there might be value in resolving the uncertainty:

$$EVPI(u_j) = \int_{\Omega_{u_j}} \{\max_i E(V|dec_i, \mathbf{c}, u_j) - E(V|dec^*, \mathbf{c}, u_j)\} p(u_j|\mathbf{c})\, du_j \tag{2}$$

where u_j is an uncertain design variable, V the value function, dec_i the set of possible decisions, dec^* the current (best) decision, $\mathbf{c}$ the deterministic constraints, and $p(u_j)$ is the probability of the design variable under applied constraints.

This expected value of perfect information (EVPI) is as an upper bound on the value of

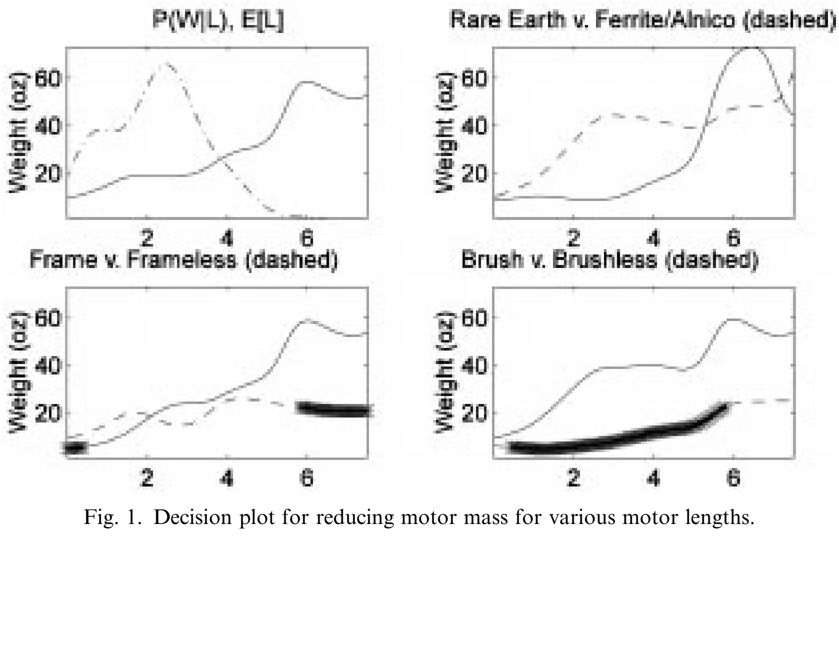

Fig. 1. Decision plot for reducing motor mass for various motor lengths.

determining an uncertain parameter. It can be calculated for all uncertain variables: performance attributes, aggregation weights, even aggregation function parameters. It is expressed in terms of the current value function and so can be compared directly to the expected value of a decision: if a decision is clear-cut (like choosing a brushless motor in Fig. 1), there will be no value in resolving uncertainty. Where resolving uncertainty would change a decision, there is value to be derived. This value can then be traded against the cost of resolving the uncertainty: running tests, doing a more detailed analysis, performing customer surveys, etc.

Design freedom—the cost of making commitments

The above value measures can help identify the current best design choices and their sensitivity to all known sources of uncertainty. The process of design becomes the progressive reduction of uncertainty—making design commitments and refining the evaluation function and performance estimates. While the latter two involve engineering efforts whose cost can easily be estimated, commitment incurs hidden costs. These costs come from the nature of the design process—designers are often forced to make decisions while also fleshing out the design requirements. Attributes that have been explicitly included in the value function are accounted for by EVPI; attributes absent from the value function may prove to be important downstream in the design process. Decisions that reduce the freedom of the designer to select among values of these hidden attributes implied by the current level of design commitment need to be flagged for more careful consideration. This effect is often most pronounced when setting performance targets for a design. Wood [6] proposes a measure of design freedom based on Shannon's information entropy [7]:

$$DF(u) = \frac{\sum_{i=1}^{n} \ln\left(\frac{1}{P(u_i|V)}\right)}{\ln(n)} \tag{3}$$

where V is the value function, u_i is a sample of the uncertainty, and n is the number of samples taken.

Figure 2 shows a plot of design freedom of speed, torque, length, and diameter with respect to motor mass targets. Reducing the allowable mass has a marked effect on motor torque, with respectively less impact on the freedom to choose diameter, length, or speed. A designer wishing to keep options open would not set motor mass targets below about 0.15 (normalized). As this target is reduced (increasing design value), the ability to adapt to other possible constraints by changing any of the four design variables is reduced.

In summation, EVDM provides a means of evaluating possible design decisions. This is most readily applicable to selection among a set of defined, discrete design concepts but can also be used to determine the optimal choice of a continuous valued design parameter as well. EVPI measures the potential impact of reducing uncertainty in the problem, focusing the attention of the designer on actions that could change the current decision context. Finally, design freedom provides a counterpoint to potentially greedy methods that can lead to commitments that reduce a designer's ability to adapt to unforeseen requirements. We now show how these three concepts provide a unifying framework for concepts in modern engineering design. We then extend them into a more general engineering context and explore further

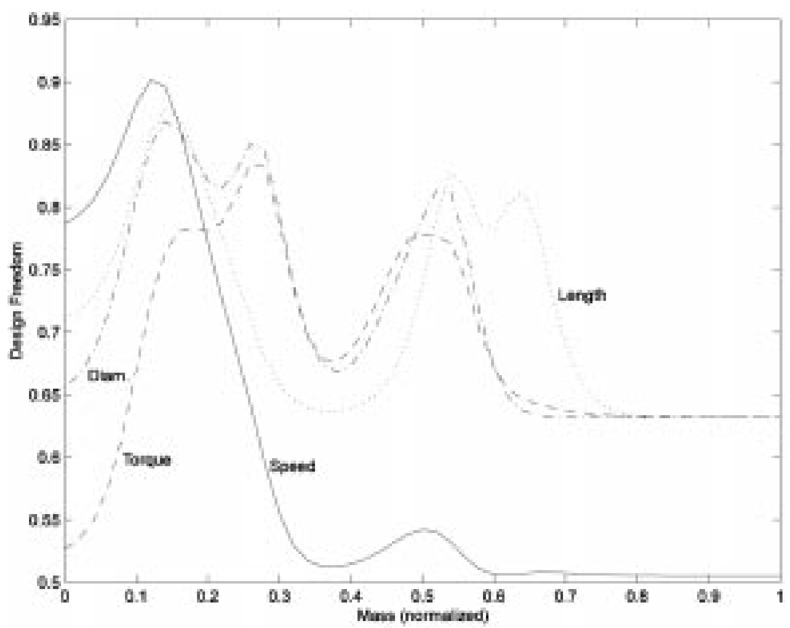

Fig. 2. Design freedom as a function of mass.

extensions beyond traditional engineering decision-making.

DBD: LINKING THE TOOLS OF MODERN ENGINEERING DESIGN

While much has been made of the advances in analysis and modeling methods within the 'hard' center of engineering (primarily through the widespread application of FEA methods), similar advances have been made on the 'softer' boundaries between the customer and engineer. In this case, we take 'customer' to mean any stakeholder in the design life-cycle: marketing, sales, manufacturing, distribution, service, maintenance, retirement, etc. The following is a list of modern design techniques presented in a typical curriculum:

- *Quality functional deployment (QFD)*: These methods focus on the communication process between customer(s) and designers. The aim is to translate the 'voice' of the customer into objective engineering measures of performance. Goals for each of the measures are then set in a semi-rigorous process that trades performance goals against each other, respecting inferred dependencies among them. Difficulties in this process often result in goals set too low, resulting in a set of possible designs (with no clear-cut way of discriminating among them), or set too high, resulting in no acceptable designs [8]. Initial ranges are developed that trade performance goals against design freedom. The designer explores these initial, uncertain trade-offs by reducing goal ranges through the DBD process. Of course, design freedom at the earliest stages is difficult to measure directly; absent initial design concepts, benchmarking methods, can be used to establish the probability densities from which design freedom is calculated. As design concepts evolve and develop through the DBD process, this initial notion of design freedom becomes more accurate and more useful for driving the design process.
- *Robust design*: At the other end of the design spectrum, noise from the realization process must be reflected back into the design process. Rather than just ignoring imperfections in the final product, designers must account for manufacturing variation and its impact on design function and value. Once one takes the DBD stance that it is the designer's job to account for one source of uncertainty, it becomes difficult to ignore other sources of uncertainty. Robust design methods rooted in evaluation under uncertainty fit naturally into a DBD framework oriented toward expected value decision-making. In situations where constraints must be satisfied with robustness, the imposed super-feasible constraints [9] can be reflected in the design freedom calculation [10]. This provides a manufacturing-accurate view of design freedom whose influence can then reach back into the conceptual design stage.
- *Taguchi methods*: From an evaluation standpoint, the overarching quality loss functions characteristic of Taguchi methods also fit in to DBD. The design methods these loss functions underpin can also be cast into the proposed DBD framework: the notion of preserving design space is a key part of solving n-type Taguchi design problems—attributes whose value can reduce performance variation are set to their optimal values, design freedom is reserved for attributes that least affect variation so that they can be used to drive performance back to nominal goals.
- *Six sigma methods*: By including aspects of both QFD and robust resign into a single framework, six sigma attempt to resolve the tension between reducing manufacturing costs (by allowing reduced process capabilities) and satisfying customer needs (as measured by characteristics that are 'critical to quality'). The main tool in six sigma design is the Monte Carlo simulation: designers perform parameter studies of design variables under various process capability and inspection strategies. The goal is to meet an enterprise-wide sigma score (nominally six) with the lowest manufactured cost. Clearly, DBD can help in this process, starting by allowing uncertainty in critical to quality thresholds and manufacturing capability and propagating this uncertainty through the rest of the design process. The substitution of more meaningful value functions for the current sigma score (many companies find six sigma too expensive and so reduce the sigma score goal).
- *Axiomatic design*: Rather than a generic measure of design freedom, axiomatic design [11] values options in which the various aspects of performance can be decoupled. Ideally, adapting to a change in one functional requirement should be done by changing a minimum of design parameters (according to the independence axiom), limiting the propagation of design changes. Among designs that are uncoupled or decoupled, those whose flexibility meets but does not exceed the range of potential requirements are preferred (according to the information axiom). Clearly, axiomatic design places a high value on design freedom; the independence it encourages produces designs that are easy to modify and adapt to changing requirements. However, the penalty for this independence may be high in situations where product specifications are not likely to change after the design process is complete. Preserving post-realization design freedom is different from preserving freedom in conceptual design, where wholesale design changes are still possible.
- *Lean production*: Lean production concepts, like reducing work in process, implementing 'pull'-based, just-in-time manufacturing processes, and reducing line changeover times,

have created a second revolution in mass production. On the surface, none of these methods appears to be closely connected to decision-based design. However, set-based design [12] is at the core of lean production Design commitment is carefully managed here by forcing designers to meet ranges of design goals. Given a range of designs for a range of goals, a lead engineer is free to resolve uncertainty in these goals without incurring additional iteration. The DBD framework provides ways not only of measuring the level of design commitment (through design freedom), but also ways of identifying critical sources of uncertainty.

- *Design for X*: By partitioning the product life cycle into distinct segments (i.e. the Xs), and assigning customer status to each, DfX creates a set of customers, each with specific needs. Successful in isolation, the challenge is aggregating all of the DfX concerns in a single design context. A hallmark of DfX is the establishment of metrics or heuristics that can be used to estimate design performance. In a DBD framework, these metrics can be applied as part of the value function, aggregated with other aspects of performance. Alternatively, the design freedom of these metrics can be monitored and decisions implying poor DfX performance flagged for special consideration.
- *Synthesis*: While not directly addressed by DBD, the introduction of design freedom helps to recognize synthesis as a valid design option. By helping designers to recognize that moving toward high-performing designs can lock in undesired effects, a DBD method that explicitly trades performance for design freedom can help focus attention on needed synthesis. In order to justify greater performance goals, a designer faced with little design freedom might choose to develop more through synthesis. Design freedom is measured relative to the current set of design options; developing new options can restore design freedom that is lost in the quest for greater performance.

Clearly, many of the above modern design methods can be cast within a DBD framework that supports decision-making under uncertainty, can focus designer attention on sources of uncertainty whose resolution will clarify the decision process, and can help designers value the freedom needed to adapt to unforeseen issues. Thus, DBD is an integrating framework for design, but can it integrate the rest of the engineering science curriculum?

INTEGRATING DBD AND ENGINEERING SCIENCE

A typical division in the engineering curriculum is between activities that produce deterministic results (e.g. analysis) or are defined by rigorous processes (e.g. experimentation) and those that are rife with uncertainty and offer only loose process direction (i.e. design/synthesis). Perhaps the first step in integrating the two sides is the realization that the assumptions that lead to deterministic processes (e.g. the load cases for a beam analysis problem) are often uncertain. In the case of a bridge design, one might need to consider possible earthquake or wind loading; to generate economic designs, these are best modeled as probability densities rather than single values. Once the camel's nose of uncertainty is allowed under the tent of engineering science, many more sources of uncertainty can be identified: material properties, geometric tolerances, customer 'requirements', future part costs, model accuracy, etc.

So, the question should not be whether to include uncertainty emanating from these sources, but how to continue to justify the more typical factor of safety approach to engineering analysis. Traditional factors of safety are used to cover all sources of uncertainty. It is difficult to convince students that a design can fail regardless of the factor of safety applied; it is likewise easy to get students to reduce a factor of safety in response to unfavorable design results—as long as the design has a factor of safety of one, there is a belief that it will not fail.

A DBD framework with an emphasis on propagating uncertainty through all design evaluations is more difficult to handle mathematically (at least on a hand calculator). Given access to any computer with a spreadsheet or matrix analysis package, students can easily implement Monte Carlo analyses to propagate input uncertainty through 'design' equations. Conceptually, this shifts the practice of design back to a mode of prediction; algebraic manipulation is replaced with (or augmented by) a search process in which design concept performance is predicted and compared to explicit failure criteria. The factor of safety is replaced by a probability of failure.

Industry has taken the lead here: six sigma design, Taguchi methods, and robust design have replaced more traditional design optimization methods. Easy Monte Carlo simulation provides designers with the ability to more closely predict the real-world response to the nominal design variables they control.

Prerequisites

Perhaps the most pressing roadblock to introducing DBD in the design curriculum is an absence of probability and statistics from the mathematics requirements. When such requirements are in place, the focus is often on the scientific mode of statistical hypothesis testing rather than on the characterization of data and the use of these characteristics for predictive behavior. The concepts of the latter are much simpler than the former: identify distributions and parameters and draw samples from them to model performance of stochastic systems. Because they are orthogonal to the traditional calculus/differential equation curriculum, training in basic statistics and probability can take place very early on.

Engineering analysis

Given a basic probability and statistics foundation, the next step is to integrate uncertainty into the typical engineering analysis courses. Instructors in such classes often bemoan the modern student's use of excessive significant figures, but students are simply transcribing what tools are giving them. Rather than berate the students, we might give them the tools to thoughtfully consider the effects of uncertain inputs to their calculations. They will clearly see the spread in the calculation results. If they still provide too many significant figures, they will more likely be associated with a mean and variance than with a single, nominal value—certainly a step forward. Of course, not all problems need be framed in uncertain terms, but in seeing the effect of common sources of uncertainty students will have a better intuition regarding the accuracy of nominal calculations.

Providing a means for propagating uncertain inputs through models that are approximate opens the door for the characterization of the accuracy of the analysis models themselves. In most cases, modeling error is dominated by uncertainty in the input; this will serve to contextualize the common use of models known to have specific faults. Increasing the physical fidelity of a model often requires the introduction of new, difficult to characterize parameters. The progress of engineering is dominated by the 'pull' of models that improve performance prediction, not by the 'push' of ostensibly more accurate models.

Experimentation

For the most part, experimentation courses require error analysis, prompting students to identify potential sources of uncertainty and deviations from modeling assumptions. So the probabilistic portion of DBD is covered here. What tends to be missing in courses that serve primarily to demonstrate analytical theory is that engineering experimentation plays a role beyond theory verification. Experimentation is not a slave to analysis but a full partner: engineers develop prototypes to test aspects of performance that cannot be adequately modeled. Experimentation is expensive, so it must be undertaken in a rigorous way. How much value is the experiment likely to bring to the design decision process? What will it cost? Are there cheaper experiments (e.g. scale modeling) that could provide reasonable value at much lower cost? This is much closer to the experimentation mode used by engineers in industry; students must be introduced to experimentation in the context of engineering projects.

Electives

Another point of curriculum integration within engineering is manufacturing. Often presented as either a process course or a statistics course, the two are connected in a DBD framework. Process capability is a key part of process selection; manufacturing errors can be propagated through a design and their impact on design value quantified. Investment in more accurate processes or more complete inspection can also be considered. The popularity of six sigma methods as an industrial panacea has, underlying it, a real need to recognize and manage uncertainty throughout the design process, from customer to manufactured product.

INTEGRATING DBD AND THE HUMANITIES

The DBD framework described above lends context and direction to the whole of the engineering curriculum. By formally recognizing the impact of uncertainty on engineering decisions, DBD helps to contextualize more faithfully both analysis and experimentation within the field of engineering. An engineer/decision-maker must forge links outside of the typical engineering curriculum, placing the whole of engineering education into a richer societal context.

Perhaps the most direct link back into society is the realization that designs fail. Once factors of safety have been replaced by probability of failure, an engineer must find ways of determining whether this probability is acceptable. This might involve comparison of the given application to other, similar situations to identify tolerable risk. It might involve experimentation to determine acceptable risk levels based on customer feedback. In cases where loss of human life is a possibility (i.e. most design cases), a student must be connected to society not on just a technical level but also on an ethical level (what is the value that the project will bring to society? what are its potential costs?). A student who has achieved a factor of safety of two need not look further for improving her design. The incommensurability of human life casts DBD as an exercise in ethics—is the risk of loss of life worth the benefit of the design? How can this risk be mitigated? Training engineers in ethics is popular, but this training is only put into practice when design is cast as an active, decision-making process rather than a passive exercise in algebra.

Once engineering is integrated into a societal context, it becomes difficult to divorce it from the more liberal aspects of education. Casting the engineer as a decision-maker does not end with the realization that designs and their failures have societal effects; an engineering decision-maker must also judge when and where a technical solution is the best way to meet society's needs. The knowledge that systems will fail in specific ways gives students the impetus to explore the full context of the application of technical systems. Removing the compartmentalization afforded by a deterministic view of engineering will motivate students to integrate the technical and non-technical sides of their education.

CONCLUSIONS

A decision-based design framework has been presented and related to modern design tools, the engineering curriculum, and the humanities (representing society in general). DBD returns engineering from a technical/scientific exercise to its roots as a means of providing the tools and systems that improve our world. This re-emphasis on the human side of engineering perhaps means shifting some of our teaching emphasis away from better technical analysis and prediction toward methods for managing the uncertainty inherent in the interface between technology and society. In doing this we reconnect engineering with the 'soft' sciences and humanities, integrating engineering with issues from the society it both serves and helps create.

Certainly there are many ways of integrating the engineering curriculum. But no matter how well design is integrated with analysis and experimentation or engineering is integrated into the social sciences and humanities, it is difficult to find a unifying concept within design education. DBD forges a methodology out of the disparate design methods commonly used in industry. A focus on decision-making under uncertainty, on the identification of potentially significant information, and on the careful investment of design freedom characterizes not only the design process but also engineering science. Forcing students to embrace, quantify, and manipulate uncertainty can turn inward-looking technocrats into circumspect world citizens.

REFERENCES

1. G. A. Hazelrigg, On irrationality in engineering design, *Transactions of the ASME*, **119** (1997) pp. 194–196.
2. W. H. Wood and A. Agogino, A prescription for information prospecting, data mining, and design refinement' in *Proceedings of the 10th International Conference on Design Theory & Methodology*, Sacramento, CA, 1997.
3. G. A. Hazelrigg, A framework for decision-based, *Journal of Mechanical Design*, **120**(4) (1998) pp. 653–658.
4. H. A. Simon, *Models of Man*, Wiley, New York (1957).
5. M. J. Scott and E. K. Antonsson, Aggregation functions for engineering design trade-offs, *Fuzzy Sets and Systems*, **99**(3) (1999) pp. 253–264.
6. W. Wood, Quantifying eesign freedom in decision-based conceptual design, in *12th International Conference on Design Theory and Methodology*, ASME, Baltimore, MD (2000).
7. C. E. Shannon, A mathematical theory of communication, *Bell System Technical Journal*, **27** (1948) pp. 379–423, 623–656.
8. R. Ramaswamy and K. Ulrich, A designer's spreadsheet, *Journal of Mechanical Design*, **119** (1997) pp. 48–56.
9. J.-C. Yu and K. Ishii, Design for robustness based on manufacturing variation patterns, *Journal of Mechanical Design*, **120** (1998) pp. 196–202.
10. W. H. Wood, A view of design theory and methodology from the standpoint of design freedom, in *Proceedings of DETC 2001*, Pittsburgh, PA (2001).
11. N. P. Suh, *The Principles of Design*, Oxford University Press, New York (1990).
12. D. K. I. Sobek, Set-based model of design, *Mechanical Engineering*, **118**(7) (1996) pp. 78–81.

William H. Wood is currently on the faculty of the Department of Mechanical Engineering at the University of Maryland, Baltimore County. His research interests include design theory and methodology, the role of information in the design process, design for manufacture, and design education. Dr Wood's teaching interests revolve around delivering an effective capstone design course that balances theory and practice in a consumer- and society-conscious setting. He received his Ph.D. in mechanical engineering from the University of California at Berkeley after receiving B.Sc. and M.Sc. degrees at Duke University.