

Arrow's Theorem and Engineering Design Decision Making^{*†}

Michael J. Scott, Ph.D.[‡]
Department of Mechanical Engineering
University of Illinois at Chicago
Chicago, Illinois, USA

Erik K. Antonsson, Ph.D., P.E.[§]
Engineering Design Research Laboratory
Division of Engineering and Applied Science
California Institute of Technology
Pasadena, California, USA

January 25, 2001

Abstract

This article establishes that Arrow's General Possibility Theorem has only indirect application to engineering design. Arrow's Theorem states that there can be no consistent, equitable method for social choice. Many engineering design decisions are based on the aggregation of preferences. The foundation of many engineering decision methods is the explicit comparison of degrees of preference, a comparison that is not available in the social choice problem. This explicit comparison of preference levels is coupled with the choice of an aggregation method, and some forms of aggregation may be inadequate or inappropriate in engineering design.

Keywords: Design decision-making; Engineering design; Multi-criteria analysis

**Research in Engineering Design*, Volume 11, Number 4 (2000), pages 218-228.

[†]Word Count: 7,564 words.

[‡]Assistant Professor of Mechanical Engineering

[§]Professor of Mechanical Engineering; corresponding author.

1 Introduction

Important tasks in engineering design are to generate and refine design alternatives, and then to select a single design, or a set of designs, to fulfill a particular need. In informal terms, the latter problem may be stated, "Find the best alternative(s)." Sometimes the determination of "best" is relatively straightforward and unambiguous, as in: "Find the lightest alternative," or, "Find the stiffest alternative." Sometimes it is less so, as in: "Find the lightest, stiffest alternative." This simple directive is inadequate to choose between one alternative that is stiff but heavy, and another alternative that is light but compliant.

This engineering design decision problem is a problem of decision with multiple criteria, and can be stated as follows:

Given several performance criteria which are to be simultaneously optimized, determine a method for comparing any two design alternatives that depends only the values of the individual criteria for each alternative.

The multiple criteria engineering design decision problem has been addressed by various decision-making systems, such as Quality Function Deployment (QFD) [Hauser & Clausing 1988], the Analytic Hierarchy Process (AHP) [Saaty 1980], Pugh charts [Pugh 1990], and multi-criteria optimization [Papalambros & Wilde 1988], all of which help guide designers in choosing designs to meet a global performance criterion based on the aggregation of performances on distinct criteria. It is generally assumed that evaluation on the basis of any of the individual criteria is straightforward. In practice, the calculation of the individual criteria may require considerable engineering judgement or resources. It is also assumed that any two alternatives could be compared directly, without resorting to individual criteria. In practice, engineers working on the designs do not make these direct comparisons, but instead work to meet performance specifications. The expertise to assess individual performance criteria is often distributed throughout an engineering enterprise. A method to compare alternatives based on separate criteria is an efficient and feasible alternative to direct comparison.

In addition to the engineering decision methods listed above, there is an extensive literature in decision theory on the topic of multiple criteria decision-making by an individual decision maker (see, for instance, Luce & Raiffa [1957], French [1988], Keeney & Raiffa [1993]). The idealized decision maker of decision theory need not be an individual person. Rather, it is assumed that there is a single entity that renders decisions; the inner workings of that entity may safely be ignored.

The problem of decision with multiple criteria is similar to the problem of social choice, or group decision-making, in which the rankings of several alternatives

by individuals are to be combined into a single, "social," ranking (*e.g.*, selecting candidates in a multi-party election). In social choice, of course, there is no longer a single decision maker, and the goal is to arrive at rational decisions that respect the sovereignty of the individual citizens involved in the decision. In the theory of group decision-making, a well-known objection to the validity of combining separate weak orders into a single social order is Kenneth J. Arrow's General Possibility Theorem [Arrow 1951, Arrow & Raynaud 1986].

Various engineering design decision methods, such as those mentioned above, depend on the aggregation of several weak orders into a single order. Furthermore, it is common for many individuals to participate in an engineering design decision. Recent discussions in the design research community have raised the question of the applicability of Arrow's Theorem to decision-making in engineering design [Hazelrigg 1996]. The relevance of Arrow's Theorem to the engineering design decision problem depends on whether making a decision in engineering design is sufficiently similar to a social choice problem, or to a decision with a single decision maker but with multiple objectives.

This paper shows how Arrow's Theorem does not apply to the multi-criteria engineering decision problem. It further shows that, despite the common participation of many individuals in engineering design decisions, engineering design is closer to multiple criteria decision-making than it is to social choice. Thus, engineering design decision-making occupies a middle ground between decision with an idealized decision maker and decision by groups of fully autonomous citizens, and on this middle ground Arrow's Theorem has no detrimental consequences.

Also related to multiple criteria analysis is decision under uncertainty, in which the problem is to determine the best alternative when the consequences of each alternative are probability distributions over possible outcomes (*e.g.*, investment decisions). Both multiple criteria analysis and decision under uncertainty usually overlay the weak ordering of alternatives with a numerical scale.

1.1 Preliminaries: Notation and Definitions

Design alternatives are distinguished by uppercase Roman characters:

A, B, C, ...

Design criteria are denoted by lowercase italics: *x*, *y*, *z*, ...

Individuals or voters are denoted by Arabic numerals: 1, 2, 3, ...

The most basic concept in the ranking of alternatives is simple comparison, in which there is no association of numbers with alternatives, but only the notion that

one alternative A is preferred to another alternative B. A ranking that depends only upon simple comparison is called a *weak order*:

Definition 1 A weak order on a set of alternatives $\mathbf{X} = \{A, B, C, \dots\}$ is a transitive binary relation \succeq such that for any two elements A and B, either $A \succeq B$ (A is at least as preferable as B), or $B \succeq A$ (B is at least as preferable as A). Indifference is possible: if $A \succeq B$ and $B \succeq A$, then one writes $A \sim B$ (A is indifferent to B). If $A \succeq B$ but $B \not\succeq A$, then A is (strictly) preferred to B, written $A \succ B$.

A weak order is an ordinal ranking: it orders the alternatives without assigning numerical values. Any computational method for decision-making requires the further structure of a numerical scale that ranks alternatives. Such a numerical scale is called a *value function*. The familiar $>$ and \geq on the real numbers of the value function correspond to the preference relations \succ and \succeq among alternatives:

Definition 2 A value function is an assignment of real numbers to alternatives that preserves a weak order of acceptability of those alternatives. A value function maps a set together with a weak order $\{\mathbf{X}, \succeq\}$ to the real numbers with its usual ordering $\{\mathbb{R}, \geq\}$. For a value function v , $v(A) \geq v(B)$ iff. $A \succeq B$, with equality for indifference.

While it is always possible to construct a value function from a weak order (see Krantz, Luce, Suppes & Tversky [1971]), the mere assignment of a value function does not imply a measure of *degree* of acceptability. A value function is of greater use in a computational decision method if the numerical scale can be used to compare levels of acceptability. The weak order $A \succeq B \succeq C$ is captured both by v_1 , where $v_1(A) = 10$, $v_1(B) = 9$, and $v_1(C) = 1$, and by v_2 , where $v_2(A) = 10$, $v_2(B) = 2$, and $v_2(C) = 1$, but the relative preference for A, B, and C implied by the two value functions is quite different. It shall be shown below that the correct specification of the numerical scale is crucial to the satisfactory resolution of both the multi-criteria decision problem and the problem of decision under uncertainty.

2 Arrow's Impossibility Theorem and its implications for the aggregation of preference

Kenneth J. Arrow's General Possibility Theorem, now commonly known as Arrow's Impossibility Theorem or simply Arrow's Theorem, is an important and powerful result in the theory of social choice. For that reason, and because a thorough understanding of Arrow's Theorem will facilitate a comparison between the social

Table 1: Weak orders of three voters

	1st	2nd	3rd
Voter 1	A	B	C
Voter 2	B	C	A
Voter 3	C	A	B

choice and multi-criteria decision problems, the Impossibility Theorem will be presented here. The treatment refers mainly to Arrow [1951] and Arrow & Raynaud [1986].

2.1 The motivating paradox

The context of Arrow’s work is politico-economic. Political scientists are interested in determining a “fair” method of reconciling the potentially conflicting interests of individuals in a society. Economists seek the most “satisfactory” distribution of a set of commodities throughout a society. The similarities between the two problems are evident, and indeed both can be formalized in the same way; the notions of “fair” and “satisfactory” are explored through this formalization.

The *majority method of decision-making* is one possible answer to the loosely formulated question of fair social choice, and one that is sufficiently obvious that a contradiction that arises from its employment motivates the Impossibility Theorem. For an odd number of people and two options to choose among, a simple vote is guaranteed to satisfy the most people. But when there are three alternatives, a paradox arises: say that there are three voters, one who orders the options $A \succeq B \succeq C$, another who holds $B \succeq C \succeq A$, and a third who holds $C \succeq A \succeq B$. These preferences are shown in Table 1. All three voters have rationally ordered preferences, but a pairwise vote shows that as a group, these three prefer A to B, and prefer B to C, yet also prefer C to A. The resulting social order is not rational, and provides no basis on which to make a decision. This paradox is called the *failure to ensure the transitivity of the majority method*, or the *paradox of voting* [Arrow 1951, p. 2]. It is in the context of this paradox that economists and political scientists explore the limits of “fair” and “satisfactory” social choice: is there any procedure for aggregating social preferences that can avoid this paradox? Arrow’s Impossibility Theorem shows that, given a particular set of axioms that define fair and satisfactory, there is no procedure that can (always) fulfill them all. The formal proof proceeds from the description of the problem with a set of axioms.

2.2 Axioms for the social choice problem

By introducing axioms to define any decision problem, two ends are accomplished. Primarily, the problem is modelled so that conclusions about the problem can be derived mathematically. Results are certain with respect to the axiomatic model; their certainty with respect to real problems depends on the validity of the axioms. Additionally, axioms cast such vague descriptions as "fair" and "satisfactory" in precise terms.

The axioms which model the social choice problem make a formal statement of requirements for any procedure that purports to solve the social choice problem, in particular two high-level requirements: consistency of the result, or rationality, and autonomy, or sovereignty, of the individual voters whose preferences are to be combined.

The social choice problem considers decision cases where all options are known, mutually exclusive, and ordered by individuals, and where the task is to produce a single social order yielding the greatest overall benefit while respecting the equal worth of each individual (as with the idealized decision maker, an "individual" is a single decision-making entity, not necessarily an individual person). Note that in the social choice context, each weak order corresponds to the wishes of an autonomous individual; in multiple criteria decision, each order corresponds to a single criterion. In engineering design, there may be many people involved, but decisions still depend upon the aggregation of engineering *criteria*. To formalize a desirable decision situation for social choice, Arrow introduces five axioms [Arrow 1951, pp. 24-30]:

Axiom 1 (unrestricted domain) *Each individual is free to order the alternatives in any way.*

Restricting Axiom 1 is one way to address the paradox, and methods that guarantee the transitivity of the majority method can be ranked by how severely they restrict this freedom. It is not at all obvious that Axiom 1 is reasonable for design decisions, as will be discussed below.

Axiom 2 (positive response) *If a set of orders ranks A before B and a second set of orders is identical except that individuals who ranked B before A are entitled to switch, then A is before B in the second set of orders.*

Axiom 2 is an ordinal version of monotonicity, and is a consistency or rationality requirement.

Axiom 3 (independence of irrelevant alternatives) *If A is before B in a social order, then A is still before B if a third alternative C is ignored or added.*

Note that Axiom 3 is violated in the motivating paradox, where the relative rankings of A and B are influenced by the addition of the alternative C. Thus, this axiom implicitly enforces the transitivity of the social order.

Axiom 4 (not imposed) *An order is called imposed if some A is before some B in all possible social orders. The social choice problem must not be imposed.*

As Axiom 1 states that individuals can hold any preferences they like, Axiom 4 says that they have a reasonable expectation that their preferences are not excluded from being chosen as the social order. This is a fairness requirement.

Axiom 5 (not dictatorial) *An order is called dictatorial if there is one individual whose decisions dictate the social order. This is likewise not allowed.*

Perhaps the simplest way to guarantee consistency of results is to violate Axiom 5: by choosing an individual whose preferences exactly determine the social order, the other four axioms are trivially satisfied. Such a solution is intuitively unfair in a social choice context, and is not allowed. While there is no need in engineering to respect all attributes equally, it is wasteful at best to evaluate attributes that have no impact on the decision. Dictatorship by one evaluation *criterion* is not a rational solution for engineering design. Some engineering cultures may appear to have a dictator in the form of a single decision maker, perhaps a manager with ultimate responsibility for all decisions; however, decisions will still be made by considering several criteria. Rather than violating Axiom 5, this changes the problem to decision with multiple criteria and a single decision maker.

In addition to these five axioms, there are implicit assumptions. Some are merely technical: there are at least three alternatives, and there is an odd and finite number of alternatives. One assumption, however, is substantive and crucial. The social order must be chosen based only on the weak orders supplied by the individuals; it is not permitted to ask for any further information to determine strength of preference, as any comparison between individuals is held to be meaningless.

2.3 The resulting contradiction

Arrow's Theorem shows that a social choice function satisfying all five conditions is an impossibility:

Theorem 1 *Any social choice function satisfying Axioms 1–3 must be either imposed or dictatorial.*

The reader is referred to Arrow [1951, pp. 51-59] or Arrow & Raynaud [1986, pp. 20-21] for details of the proof, but the basic line of reasoning is as follows.

A *decisive* set of individuals for A over B is a set who guarantee that A will be preferred to B whenever they unanimously agree so; any decisive set must contain a smaller decisive set; there is always a decisive set; any set that is decisive for A over B is decisive for A over anything else and for anything else over B, and thus for all A over all B; thus there must be a dictator. The only way to avoid this dictatorship is to impose some preferences, violating Axiom 4.

Thus, the paradox of the intransitivity of the majority method is a manifestation of a difficulty so deeply embedded in the social choice problem that it cannot be resolved without compromising the defining axioms. It is an intuition-building exercise to take the three orders shown in Table 1, and attempt to combine them using a general procedure. It quickly becomes clear that transitivity can be insured by dictatorship. Arrow's Theorem proves essentially that transitivity can only be insured by dictatorship.

2.4 Ways around the contradiction

Arrow's Theorem shows that there is no method of aggregating social choice that is guaranteed to satisfy all five axioms. Are all socio-political systems then fundamentally irrational? Or are there systems that, implicitly or explicitly, operate with restricted versions of one or more of the axioms, and thus avoid gross inconsistency?

Arrow and others have attempted to resolve the paradox by weakening the first condition, arguing that in real political, economic, and even moral¹ systems participants tacitly agree to structure their choices in a "logical" way, *i.e.*, in a way that keeps contradictions from arising. Thus Arrow introduces the notion of *single-peakedness* as a way around the dilemma. A single-peaked set of alternatives is ordered on some (one-dimensional) external scale, so that each individual is free to choose a favorite alternative, but then must hold descending regard for the other alternatives to the two sides of his first choice. The example of a political spectrum is given: each voter has a preferred, or ideal party, and each step away from the ideal party, whether to the left or to the right, is an ever less desirable alternative. This condition says nothing about comparison between parties to the left and parties to the right of ideal. If a condition of single-peakedness is substituted for the axiom of unrestricted domain, then the Impossibility Theorem no longer holds. An abstracted parliamentary system thus avoids the difficulties of Arrow's Theorem. Of course, in a two-party system there is no contradiction, as the two-alternative situation is not paradoxical.

In general, the difficulty of the Impossibility Theorem can be overcome by

¹Arrow goes so far as to quote Kant [Arrow 1951, pp. 81–82]

restricting the freedom of individuals participating in the process by structuring their preferences in some way. Ranking all alternatives on an external scale as discussed above is one form of structure; allowing limited veto power is another.

2.5 Numerical scales, or comparing strength of feeling

Numerical scales for comparison of different attributes are key to the resolution of the multi-criteria decision problem. Their introduction into the social choice problem does not obviate the problems raised by Arrow's Theorem, but a discussion of why they do not solve the social choice problem will facilitate a later comparison with the multi-criteria decision problem. Consider the example of the motivating paradox discussed above, with the individual voters' weak orders from Table 1. A majority pairwise vote to combine these three weak orders leads to an intransitive, and thus untenable answer. Consider adding a numerical measure of strength of feeling to the problem, by supposing that each voter is given 10 points (or votes) to distribute among the three alternatives. A representative combination with a weighted sum is shown here:

Table 2: Three voter's preferences

	A	B	C
Voter 1	6	3	1
Voter 2	1	7	2
Voter 3	2	0	8
Total	9	10	11

In this case, there is no ambiguity: B is clearly preferred to A, and C is clearly preferred to both A and B.

This solution, however, is accidental. If Voter 3, for instance, holds the slightly different preferences shown here, preferences which are still consistent with the weak order in Table 1, there is no longer a clear choice between the three alternatives:

Table 3: Three voter's alternate preferences

	A	B	C
Voter 1	6	3	1
Voter 2	1	7	2
Voter 3	3	0	7
Total	10	10	10

Indeed, any possible ordering of A, B, and C, including indifferences, can be achieved with numerical preferences consistent with the weak orders in Table 1.

It is unclear whether the procedure of allotting ten votes to each individual violates Axiom 3, the axiom of independence of irrelevant alternatives. Say one of the three options is ignored. Then the comparison between the remaining two is the same if the vote distributions remain identical. If, on the other hand, the full ten votes must now be distributed over only two alternatives, then the final order can change. This procedure does not address the difficulty raised by Arrow's Theorem, where the *only* information allowed in the formation of the social order is the set of individual orders. Also, since any possible social order may result from this method of combination, every individual has a strong incentive to assign all ten points to their most preferred option, thus returning the problem to the precise statement of Arrow's Theorem.

Nor is the difficulty overcome by using different arithmetic (by letting each individual rank each alternative on a scale of 1 to 10, say, or by aggregating with something other than a weighted sum). A numerical scale cannot work in the social choice problem. Because interpersonal comparison is not allowed, the scale can only be assigned arbitrarily. While a particular scale may appear to address an inconsistency for a particular problem, it is merely a matter of arithmetic to recast the problem so that it is directly subject to Arrow's Theorem. The inability of numerical scales to address the social choice problem is discussed further in the Appendix.

3 Decision with multiple criteria

The engineering design decision problem is not a social choice problem, but instead is a decision with multiple criteria, that is: rank a number of alternatives, each of which is ranked separately by several ranking criteria. Although this problem appears superficially similar to the social choice problem, since it seeks to combine several individual rankings into one, it is a distinct problem. Two differences are:

- In the social choice problem, all orderings are accorded equal worth. In the multiple criteria problem, it is desirable to be able to assign importance weightings to criteria. While it is natural to accord all human voters equal worth, there is no obvious reason to require equal weighting of the different engineering criteria that describe a device or system.

This difference may disappear if the weighted problem can be recast as an unweighted problem with more individuals.

- The social choice problem permits no interpersonal comparison of preferences, and is thus limited to the discussion of weak orders. The heart of the

multi-criteria engineering problem is the inter-attribute comparison of preferences. When considering many design goals, it is crucial to understand their relative importance and the way in which they interact. Again, what is natural to require when modelling the sovereignty of individual citizens is not necessarily applicable to separate engineering design criteria.

This difference makes decision with multiple criteria structurally different from social choice, and has deep implications for the applicability of the Impossibility Theorem to engineering design decision-making.

Even the informal motivating paradox for the Impossibility Theorem (where a majority vote ranked A before B before C before A) loses much of its power if cast in the framework of multi-criteria decision-making. Consider the analogous example of a design or a product that is to be judged on the basis of three criteria: x , y , and z . It is certainly plausible to assume that the designer may be faced with a choice of three candidate designs, A, B, and C, such that A is better than B is better than C with respect to criterion x , B is better than C is better than A with respect to y , and C is better than A is better than B with respect to z . The analogous "paradox" here is that giving x , y , and z one vote each as a method to determine the best design yields no obvious answer. In other words, if all that is known about a design is a weak order among alternatives for each of the three criteria x , y , and z , then there is not enough information to decide upon an overall best design. This "paradox" is resolved in the multi-criteria problem by more careful consideration of preferences for x , y , and z , and consideration of how those preferences interact. Note that even if this decision is made by a group of three individuals, each responsible for one criterion, there is still no paradox; the involvement of more than one person does not by itself make group decision-making.

A crucial assumption was made at the outset about the engineering design decision problem. In general, in a real design situation, there *is* a rational weak order among A, B, and C.² Furthermore, that order could in principle be found directly. The "paradox" is merely that additional information beyond the weak orders on x , y , and z is required to recognize the overall order. The question asked here is whether and when it is possible to find consistent, rational techniques to discover the ranking among A, B, and C. If so, it will necessarily be with a slightly different set of assumptions from the ones of the Impossibility Theorem, assumptions more appropriate for engineering decision-making than for social choice.

²This does not mean that there is one optimal solution that fits all situations; a design situation includes the attitudes and preferences of the people and corporate entities involved. Given the need for personal transportation, some auto manufacturers choose to make luxury cars, while others choose to make economy cars.

3.1 Comparison of axioms for the two decision problems

A careful examination of the axioms is necessary before considering the Impossibility Theorem in the context of the design decision problem. When combining engineering criteria (the multi-criteria problem), rather than individual orderings (the social choice problem), the axioms of positive response, independence of irrelevant alternatives, and inadmissibility of dictatorial solutions still appear to hold. However, it is not obvious that domains must be unrestricted or that orders must not be imposed.

The social choice problem must respect individuals by affording them the freedom to order alternatives as they choose; in a design situation, cultural, customer, or managerial structure is almost always imposed on the orders. For instance, if the three candidate vehicle structure designs A, B, and C have bending stiffnesses of 3000, 3200, and 3400 N/mm respectively, Axiom 1 states that any individual is free to express the preference $C \succ A \succ B$. A vehicle structures group, however, which proposed this order to management, would be criticized for "irrational" preferences *over bending stiffness*. The order that ranks $C \succ A \succ B$ is transitive, and any transitive order must be considered rational in the social choice problem; it would be an acceptable final order of candidate designs. With respect to the particular evaluation criterion of bending stiffness, however, it is not rational; many such transitive orders would be considered irrational in an engineering context. No individual is given veto power in the social choice context; almost any attribute of an engineered design has a level so unacceptable as to veto the entire design.

Recall that Arrow proposed *single-peaked* preferences as a way to resolve the contradictions of the Impossibility Theorem: preferences are single-peaked when all options are ordered on an external scale, and each individual has one preferred option and holds descending regard for alternatives to the two sides of that preferred option. Engineering variables are almost always ordered on an external scale, and preferences for engineering requirements are commonly single-peaked around an ideal target. Indeed, nearly all engineering requirements are of one of three forms: less is better, more is better, or closer to a particular target is better [Byrne & Taguchi 1986]. All three of these forms are single-peaked.

There are (rare) evaluation criteria that do not behave in a single-peaked manner. A design preference for availability of a particular material stock may be one criterion for a design, and it may change over time and take on any order. The preference for the frequency of the first acoustic mode is often to *avoid* a particular unpleasant range. However, designers can and do restrict criteria that are not globally single-peaked to regions of local single-peakedness. The vehicle structure designer seeking to avoid a particular range of frequencies of the first acoustic mode, for example, chooses to target either higher or lower frequencies, thus con-

sidering only a range over which the criterion is single-peaked. Thus, while the completely generic design decision problem should obey Arrow's axiom of unrestricted domain, designers strive to avoid the generic problem, and rather to cast each problem so that domains, rather than being unrestricted, are single-peaked along the obvious external scales provided by the design parameterization. Indeed, in terms of the decision problem, the parameterization of a design serves to restrict domains. For the multi-criteria decision problem, the axiom of unrestricted domain is replaced by an exhortation to the designer to verify that criteria are single-peaked, or restrict the problem until they are. It is understood that this may not always be possible, but when it is not, the designer realizes that the design problem is not completely well-conditioned. Such problems are difficult for formal methods and informal methods alike.

The axioms of the multi-criteria decision problem are thus not identical to those of the social choice problem. Axiom 1 is crucial to the social choice problem, while in the engineering design problem, it appears in a modified form that is known to overcome Arrow's Theorem. This difference in Axiom 1 allows multi-criteria methods to operate on some large classes of problems without violating Arrow's Theorem, even without the use of a numerical scale to compare preferences across attributes. However, the differences in axiomatic structure are minor compared to the difference in one fundamental assumption: the social choice problem does not admit interpersonal comparison, while the multi-criteria decision problem would be meaningless without inter-attribute comparisons. The next section will discuss how a numerical scale for the inter-attribute comparison of preference is used in two contexts: the aforementioned multi-criteria decision problem, and the related problem of decision-making under uncertainty.

There is a further, practical, difference between social choice and engineering design: designs have constraints. A maximum stress indicates the point at which a design breaks and fails; government regulations must be fulfilled or a design is not allowed on the market. However, it is usually far from obvious, *a priori*, which candidate designs violate constraints, and the engineer must be free to consider all potentially viable candidates. In decision theory, the decision maker chooses among a neatly defined set of viable alternatives; in engineering design, deciding which are the viable alternatives is a major task, often involving considerably more resources than the final decision. Furthermore, the early consideration of infeasible designs is often a crucial part of the refinement process.

A decision method that captures this special feature of engineering design would, in principle, violate Axiom 4. The positions of alternatives that fail on the basis of a single criterion could be seen as imposed (to be last) in any aggregated order. Constraints in engineering design, if translated into social choice terms, are a sort of veto that individual criteria may exercise over the entire decision, and thus

violate the axiom of no imposed orders. However, when the design is acceptable on the basis of each individual criterion, there is no reason to abandon the axiom of no imposed orders. For the multi-criteria engineering decision problem, the axiom of no imposed orders is weakened:

Axiom 4.4a (limited imposed orders) *Axiom 4 holds, with the exception that some alternatives may be declared unacceptable, and thus last in any combined order, on the basis of an unacceptable ranking on a single criterion. All unacceptable alternatives are equally unacceptable.*

This modification of Axiom 4 is not a significant theoretical objection to the application of Arrow's Theorem in engineering design, as the same end can be achieved by simply removing alternatives, after analysis, from consideration. It does, however, provide a framework for decision methods for engineering design that are capable of considering alternatives that may turn out later to be infeasible.

4 Inter-attribute comparison of preference

Decision with multiple criteria differs empirically from social choice in an important way. In the former, there is always a well-defined aggregated order among alternatives, which is available to anyone with the time and resources to query a decision maker directly about all possible combinations; in the latter, Arrow's Theorem calls into question the very existence of a well-defined aggregated order. A direct specification of preference in many dimensions in the multi-criteria problem presents no more theoretical difficulties than a direct specification of preferences in one dimension; the practical implementation, however, can present great difficulty.

There is more than one way to assign a value function when a weak order among alternatives is given. In this section, two distinct approaches to the assignment of a value function are discussed: the *utilities* (or *von Neumann–Morgenstern utilities*) [von Neumann & Morgenstern 1953, Keeney & Raiffa 1993] that are used for decision-making under uncertainty, and directly specified *preferences* such as those used in many multi-criteria decision systems. Both techniques use inter-attribute comparison to guarantee consistency (and avoid the pitfalls of the Impossibility Theorem), but the emphasis in each is different.

4.1 Utility

The specification of utility depends upon a weak order among alternatives, and on the mathematics of expectation. To determine a utility function, the so-called

lottery question must have an answer: "Given that A is preferred to B, and B is preferred to C, at which probability p is there indifference between the two choices 'B with probability 1' and 'a lottery that yields A with probability p and C with probability $(1 - p)$ '?" (Note that the question need not have a *direct* answer; see Wakker & Deneffe [1996], for instance, for a discussion of the elicitation of von Neumann–Morgenstern utilities with little or no probability information.) Von Neumann and Morgenstern [1953] show that, given the assumption that utilities combine with the mathematics of expectation, the numerical utility scale is determined up to an affine transformation.

The assumption of the use of mathematical expectation arises because utility theory is intended to treat questions of decision-making under probabilistic uncertainty, such as those that are germane to gambling. This makes the specification of *relative* utilities with probabilities natural. The lottery method provides an elegant method to determine not only that A is preferred to B and B is preferred to C, but also how great the preference is for A over B relative to the preference for B over C.

However, the development of utility specifically excluded the notion of interpersonal comparison of utility as too difficult to address:

We do not undertake to fix an absolute zero and an absolute unit of utility. [von Neumann & Morgenstern 1953, footnote, p. 25]

Utility theory is intended for use in decision-making under uncertainty or risk, rather than as a solution of the multi-criteria decision-making problem. Decision under uncertainty is often required in multi-criteria settings. Keeney & Raiffa [1993] present conditions under which a decision maker's overall utility is a simple function of individual, independent utility functions. The overall utility function is determined by direct comparison of several alternatives in order to determine appropriate weighting factors for the individual utility scales determined with regard to the separate attributes. Arrow's Theorem presents no difficulty whatsoever to this sort of calculation; even if several people are involved in the decision, there is an idealized single decision maker.

The success of the von Neumann–Morgenstern utility paradigm, and the ease of its application in terms of quantified risk, have led to a situation in which many decision problems are treated as problems of (economic) decision-making under uncertainty. The lottery question seems natural, and so it is assumed that the lottery question is the right way to impose a numerical scale on preferences. Nevertheless, engineering design may *not* be best classified as decision-making under risk and uncertainty. Utility theory is one paradigm for decision-making, appropriate for a particular set of problems, those where the "estimation of expectations for each option" is the most pertinent information. When design reaches the manufacturing

stage, and probability distributions over manufacturing variations are usually the most relevant uncertainties, the design decision problem is indeed similar to the problem addressed by utility theory. Earlier in the design process, where uncertainty will be resolved by refinement of a design alternative, rather than by random selection from a perceived distribution among alternatives, a utility model is less appropriate. Direct preference specification, for uncertainty other than quantified risk, will be discussed in the next section.

From the point of view of classical utility theory, the design decision problem examined here is a case of decision-making under certainty. In the classical utility theory view, the construction of the utility function and the choice of a "best" solution with limited computation are uninteresting problems. The construction of the utility function is only support for the subsequent question of which course of action to select in the face of uncertain consequences.

4.2 Direct specification of preference

The engineering design decision problem that is the subject of this paper is not primarily a problem of decision-making under uncertainty. While methods such as those proposed by Keeney & Raiffa [1993] could be applied to the problem, and the resulting utility function would indeed provide a basis for comparison of any two alternatives, the relative complexity of these methods is not justified in cases when no probabilistic uncertainty is to be incorporated in the problem. This is particularly true as the number of individual performance criteria increases (the number of coefficients to be determined is $2^n - 1$, where n is the number of criteria). Thus methods intended to solve the multi-criteria decision problem as stated at the beginning of this paper, without the incorporation of uncertainty, commonly use direct specification of preference (*e.g.*, QFD [Hauser & Clausing 1988], AHP [Saaty 1980]). Direct preferences, unlike utilities, are expressed on an absolute scale. For example, a preference of $\mu = 1$ could indicate a completely acceptable value, and $\mu = 0$ a completely unacceptable value. Alternatively, each criterion may be judged on a discrete scale (such as 1–3–9). The individual numerical orderings associated with each criterion are then aggregated into an overall numerical ordering of all alternatives based on all criteria simultaneously. The functions used for aggregation of the orderings vary; a weighted arithmetic mean is a common choice.

It was made clear earlier that Arrow's Theorem does not apply to the engineering design decision problem, whose chief concern is to compare preferences based on multiple attributes. Nevertheless, Arrow's Theorem does bear indirectly on the legitimacy of these methods: if the comparison of preferences is effected by the arbitrary assignment of numbers to alternatives, then those preferences contain no

more information than weak orders, as discussed in Section 2.5 and the Appendix. In that case no aggregation method is adequate. A method that delivers an overall preference as a function of individual preferences must explicitly, and justifiably, compare individual preference values.

The individual preference orders may be generated by, or associated with, different people or groups involved in a design. One of the motivations for the engineering design decision problem as described in this paper is the need to assess design alternatives for overall worthiness when the most readily accessible information comes from many independent performance measures, each of which is the responsibility of a different portion of a design team. If these groups or individuals must be treated as autonomous citizens, then there is no fair and rational way to combine their preferences. However, team design is not driven by the need to respect the sovereignty of the individuals involved, but rather by a desire to create superior products.

The team engineering decision problem thus has two aspects: the multiple criteria problem, to which Arrow's Theorem does not apply, and the problem of collaboration, to which Arrow's Theorem would apply if it were necessary to guarantee sovereignty of team members. The mere involvement of groups does not make it group decision making; in engineering, rather than social systems, attributes, not people, must be reconciled. There are certainly difficult and interesting issues in the management of groups in an engineering context, but Arrow's Theorem does not make such management impossible.

5 Conclusions

This paper establishes the legitimacy of constructing aggregated preferences in engineering design, even though such aggregation is not permissible in the social choice problem. It does so by a thorough examination of the context and results of Kenneth J. Arrow's General Possibility Theorem. The design decision problem has structural (axiomatic) peculiarities which make comparisons in the manner of social choice acceptable in engineering design; more fundamentally, however, the aggregation relies on a procedure for comparing the relative importance and interaction of individual preferences that violates an assumption of citizens' sovereignty in the social choice milieu.

Decision with multiple criteria and a single, idealized decision maker is one problem in decision theory; social choice, in which each idealized individual's sovereignty must be guaranteed, is another. Engineering design often falls somewhere in the middle, but even when many people are involved, the group nature of the decision is subordinate to the multi-criteria aspect, as the sovereignty of individ-

uals is not necessary. Interesting parallels are evident between the ideas expressed in this paper and different management styles [Goold & Campbell 1987]. A top-down management style would correspond to forcing a single decision maker on the problem, whereas distributed control would correspond to the group decision environment described here; in either case, the decision problem is framed in terms of engineering requirements rather than individuals' desires.

The aggregation of preference in engineering design is often discussed as the assessment of a utility function, though utility is a particular form of preference assessment that is useful in decision under uncertainty and risk. The multi-criteria decision problem discussed in this paper does not include the incorporation of probabilistic uncertainty. The specification of a multi-attribute utility function for decision under uncertainty, and the use of direct specification of preferences for the resolution of the multi-criteria decision problem, both avoid the problems raised by Arrow's Theorem, but their emphases are different.

While Arrow's Theorem does not prove that the engineering design decision problem treated here cannot have a rational answer, it does offer a caution to methods that attempt to solve the multi-criteria decision problem: comparison of preferences must be made explicit. The arbitrary assignment of numbers to alternatives can lead to undesired conclusions. Similarly, the arbitrary assignment of an aggregation method can lead to undesired conclusions. A decision method must have an explicit procedure for assigning values to alternatives, and for combining those values into a single performance function, and the two must agree. If a decision method gives inadequate answers, it is not because Arrow's Theorem declares that comparison is impossible; it is because the particular implementation is flawed — the representation and comparison of preferences in the method are insufficiently rich.

Finally, the need in engineering design to consider and evaluate alternatives that may turn out to be infeasible is not well captured by the classic formulation of the decision problem, where only feasible alternatives are considered. This difficulty can be overcome by reformulating the decision problem whenever an infeasible candidate arises. Alternatively, a decision methodology can handle such constraints by a modification of the underlying decision axioms.

Acknowledgements

This material is based upon work supported, in part, by the National Science Foundation under NSF Grant Numbers DMI-9523232 and DMI-9813121. Any opinions, findings, conclusions, or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the sponsors.

References

- Arrow, K. J. [1951], *Social Choice and Individual Values*, 1st edn, J. Wiley, New York, NY.
- Arrow, K. J. & Raynaud, H. [1986], *Social Choice and Multicriterion Decision-Making*, The MIT Press, Cambridge, MA.
- Byrne, D. M. & Taguchi, S. [1986], The Taguchi approach to parameter design, in 'Quality Congress Transaction – Anaheim', ASQC, pp. 168–177.
- French, S. [1988], *Decision Theory: An Introduction to the Mathematics of Rationality*, Ellis Horwood, Chichester.
- Goold, M. & Campbell, A. [1987], 'Many best ways to make strategy', *Harvard Business Review* **65**(6), 70–76.
- Hauser, J. & Clausing, D. [1988], 'The House of Quality', *Harvard Business Review* **66**(3), 63–73.
- Hazelrigg, G. A. [1996], 'The implication of arrow's impossibility theorem on approaches to optimal engineering design', *ASME Journal of Mechanical Design* **118**(2), 161–164.
- Keeney, R. & Raiffa, H. [1993], *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*, Cambridge University Press, Cambridge, U.K.
- Krantz, D., Luce, R., Suppes, P. & Tversky, A. [1971], *Foundations of Measurement*, Vol. I, Academic Press, New York, NY.
- Luce, R. & Raiffa, H. [1957], *Games and Decisions*, J. Wiley and Sons, New York, NY. Dover Edition, 1985.
- Papalambros, P. & Wilde, D. [1988], *Principles of Optimal Design*, Cambridge University Press, New York, NY.
- Pugh, S. [1990], *Total Design*, Addison-Wesley, New York, NY.
- Saaty, T. L. [1980], *The Analytic Hierarchy Process: Planning, Priority Setting, Resource Allocation*, McGraw Hill, New York, NY.
- von Neumann, J. & Morgenstern, O. [1953], *Theory of Games and Economic Behavior*, 2nd edn, Princeton University Press.
- Wakker, P. & Deneffe, D. [1996], 'Eliciting von Neumann-Morgenstern utilities when probabilities are distorted or unknown', *Management Science* **42**(8), 1131–1150.

Appendix

In Section 2.5 the allotment of ten votes to each individual was proposed as a potential way around Arrow's Theorem, and was shown to be inadequate. That point is discussed in more detail here.

Any of the sets of numerical preferences attributed to individual voters in Section 2.5 can be recast as weak orders held by more voters. Combining all the weak orders of the fictitious voters using the majority method of decision-making will then yield the same social order as a weighted sum aggregation of the numerical rankings, with ties in the majority method corresponding to equal overall numerical rankings. To express 10 preference points requires 20 individual weak orders: for example, if alternative A received 1 of the 10 possible points, that is expressed by one weak order $A \succ B \succ C$, and one $A \succ C \succ B$. (Two orders are required so that B and C are indifferent.) Voter 2's numerical preferences, for instance, are equivalent to the following 20 weak orders:

1st	2nd	3rd	instances of this order
A	B	C	1
A	C	B	1
B	C	A	7
B	A	C	7
C	A	B	2
C	B	A	2

The entire weighted sum aggregation is equivalent to the majority method applied to 60 individual voters. However, these 60 individual orders do not obey Axiom 1, as they are actually 30 dependent pairs. This clearly demonstrates that the assignment of ten votes to each individual imposes restrictions that are not in the original problem, which by definition must obey all the axioms of Arrow's Theorem. Indeed, this numerical method can be seen as an attempted end run around Arrow's Theorem, by the means of eliminating Axiom 1. Thus, Arrow's Theorem still applies: the use of a numerical method does not by itself guarantee both fairness and consistency.