APPLIED SYSTEMS ANALYSIS

Engineering Planning and Technology Management

Richard de Neufville
Massachusetts Institute of Technology

McGraw-Hill Publishing Company
New York St. Louis San Francisco Auckland Bogotá Caracas Hamburg
Lisbon London Madrid Mexico Milan Montreal New Delhi Oklahoma City
Paris San Juan São Paulo Singapore Sydney Tokyo Toronto

CHAPTER 21

COLLECTIVE DECISIONMAKING

21.1 OBJECTIVE

Many different groups of people must often jointly agree on the design of a system. The towns in a metropolitan area, for example, may have to select a subway design that best meets their various needs for service; or a consortium of companies may have to select an industry standard for their products. This kind of collective decisionmaking represents the most difficult form of systems analysis and design.

This final chapter addresses this ultimate question. It indicates how planners can most effectively organize the analysis to assist their clients in selecting the system that will be best for them. This procedure, given in Section 21.6, pulls together the important elements of both system optimization and system evaluation, Parts I and II of this text. An application illustrates the recommended procedure.

The design that is best overall for groups with different interests cannot be found analytically, unfortunately. As a matter of principle, no such procedure will ever be universally valid or acceptable. Section 21.2 covers this theory. This fact does not mean, of course, that analysts do not propose many analytic procedures that purport to define the optimum collective choice for groups. These procedures are, at best, only acceptable under quite limited conditions. Whatever their appeal, they are ultimately flawed by resting on unacceptable assumptions. Section 21.3 makes the point with reference to the most common approach proposed as a solution for collective decisionmaking in design: the naïve idea of product maximization.

The reason there cannot be any universally acceptable analytic solution is simply that the way different people or groups value products differs enormously. As indicated in Chapters 1 and 2, the fact that people have values means that purely technical considerations cannot define the optimum. Sections 21.2 and 21.4 discuss these issues in the context of collective decisionmaking.

The approach recommended for group design of systems is collaborative negotiation (Section 21.5). Negotiation permits the different groups to progress toward mutually beneficial improvements, as must almost inevitably exist. Collaboration further facilitates the exploration of economies of scale (see Chapter 4) for mutual advantage.

21.2 THE ISSUE

The basic difficulty in collective decisionmaking is that there is no way a universally acceptable utility function can be defined for all groups in a decision. If such a function were available then the selection of the optimal design would be obvious: we would simply choose the system that maximizes the expected utility for the group. This section discusses the concept of a group utility function, examines some of its difficulties, and then proceeds to the demonstration why such a function cannot be defined in any uniquely acceptable way.

Concept. The utility function for a group is known formally as either a group utility function (in operations research) or a *social welfare function* (in economics). We use the second term because it ties in most directly with the standard theoretical literature on the nature of choice for group.

> **Semantic caution:** "Social Welfare" does not refer to charity in the sense used currently. It simply means the common good of society.

The social welfare function, SWF, is usefully thought of as some kind of expression of the utilities of different individuals or members of the groups, U_i:

$$SWF = f(U_1, \ldots, U_n)$$

The immediate question is what this aggregation should look like.

A little thought indicates that the nature of the social welfare function should depend on the ethical concepts of the group. If we were to suppose that society valued every individual equally, one might imagine that the social welfare function were some kind of summation. Is it then

$$SWF = \sum U_i \ ?$$

But a society might, as often happens, value some people more than others. The ancient Roman republic, for example, explicitly gave its aristocrats more votes than the plebians in elections. Closer to home, we might discount the desires of some individuals; children for instance typically get no votes and their childish desires are substantially ignored. These ideas would imply an alternative social welfare function, in which individual utilities were somehow weighted. Could it be

$$SWF = \sum w_i U_i \ ?$$

We need not suppose, however, that the social welfare function is necessarily a summation of individual utility, in which the great joy of some counterbalances the misery of others. A sense of social justice may lead us to conclude that the social welfare function should at least reflect some concept of equality; it may also naturally give some weight to ensuring that everyone achieves some minimal level of satisfaction, and thus discount greater individual levels of utility. Would a more realistic version then be

$$SWF = \sum w_i e^{-(U_i - U_{\text{average}})} \ ?$$

This discussion has two implications. First, no simple addition of individual utilities is likely to define a plausible group utility function. Secondly, it is very difficult to imagine what a satisfactory functional form might be.

For discussion purposes, the shape of the social welfare function for most groups can be imagined as in Figure 21.1. Were we able to define it, it would be something rather like that of an individual valuing several benefits. It would favor some kind of balance, reflect the group's diminishing marginal utility for excessive benefits to individuals, and thus curve inward toward the origin.

Interpersonal comparisons. In thinking about the social welfare function, we need also to ask how we can combine the utilities of different people. We can certainly add them up; but to what extent does any such addition make sense?

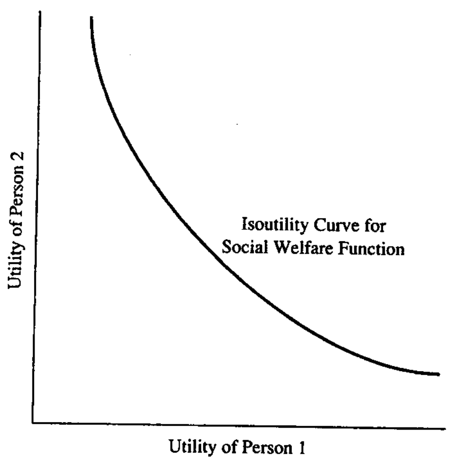

FIGURE 21.1
Plausible nature of a social welfare function.

Are individual utilities measured in comparable units that can legitimately be summed?

Economists and philosophers have thought hard about this question of interpersonal comparison of utility. Their conclusion is that such comparisons simply do not make sense. The problem is that there is no way to transform everyone's utility into a common denominator. Even for an individual, utility is a relative measure; we have no way of measuring any person's absolute degree of happiness. By extension there is no way we can establish comparable degrees of happiness among individuals.

The short of it is that even if we could persuade ourselves that we have devised the right mathematical form for the social welfare function, we could not calculate it in any meaningful way from individual utilities. We would not know how to define the personal utilities on a common scale, and thus could not use the function in practice.

Arrow paradox. Logically, there is an alternative to constructing a social welfare function from the utilities of individuals. We could think of identifying a group's utility from the choices it makes, just as we construct an individual's utility from the choices a person makes. This aggregate approach has the appeal of being direct. Unfortunately, this approach is not satisfactory either. Kenneth Arrow demonstrated this fact and received the Nobel Prize for his investigations of this issue. His essential results are that

1. The choices a group makes depend on its internal rules of decisionmaking; for example, its voting rules.
2. No one voting rule or decisionmaking process is intrinsically best.
3. The choices made by a group are therefore necessarily an ambiguous reflection of its preferences, so that we cannot rely on a group's choices to construct its social welfare function.

The "Arrow Paradox" illustrates the situation. It demonstrates by simple example the kind of ambiguity of choice that may result from voting procedures. Consider a group of three persons evaluating three different design options. Their individual preferences are quite different, as Table 21.1 shows.

Suppose that the group agrees to select its final design by successively comparing pairs of options until it has ranked them all. This is a natural way to search for dominant choices; it is also rather like selecting the president of the United States after a set of primaries.

Let us proceed to apply the rule to the case at hand, starting by comparing A and B alone. We find

$$A > B \quad \text{(C excluded)}$$

by a 2:1 majority. If A is thus retained as a preferred option and compared to C, we find

$$C > A \quad \text{(B excluded)}$$

TABLE 21.1
Individual Rankings of Options A, B, and C for Arrow Paradox

Ranking of option	Individual: Tom	Dick	Harry
First	A	C	B
Second	B	A	C
Third	C	B	A

also by a 2:1 majority. Having compared all three options, can we then conclude

$$C > A > B \quad ?$$

The Arrow Paradox is precisely that this ranking is not valid. We can check the result by comparing B and C alone, and find

$$B > C \quad \text{(A excluded)}$$

also by a 2:1 majority. The net result is intransitive. If we chain the results we obtain

$$C > A > B > C \quad !$$

The interpretation is that we cannot rely on the choices expressed by a group to reflect its social welfare function. The actual choice may depend critically on the precise way a voting procedure is applied. It may depend even more on the voting rules themselves. We thus cannot presume that we can interpret the choices expressed by a group as their clear preferences.

Impossibility of the social welfare function. The point of this discussion has been that it is simply not possible to construct a valid group utility function that we can apply to collective decisionmaking. Recognizing this impossibility provides the right starting point for dealing with collective decisionmaking.

21.3 DISTRIBUTION PROBLEM

What makes system design for a group so difficult is that it involves a complex problem entirely different from the kind we have been discussing throughout the text. In collective decisionmaking we must consider the distribution of the products of a system to each member of the group. This distribution defines the utility of each member and, thus, the social welfare of the group.

The text has so far treated system design as a *production problem*. The issue has been to specify how the system will be configured, and thus what it produces. The entire analysis is based on a combination of good technical design (the technically efficient production functions in Chapter 2), and the prevailing values that may either be interpreted as costs (as in Chapters 4 to 14) or as utilities.

System design for a group could also be viewed simply as a production problem if we could define the utility of the group, the social welfare function. That is impossible, however, as the previous section indicates. The only way to define value for the group is to consider the values of its members. That implies that we have to focus on the distribution of products of the system. This is the *distribution problem*, the determination of what each member of the group will get.

Naïve approach. A popular procedure for dealing with system design for a group is simply to deal with the production and distribution problems separately. The idea is "first, make the pie as big as possible, and then worry about how it's divided." Specifically, this is interpreted in practice as maximizing economic value and letting politicians worry about the rest. This is the underlying concept when development economists focus on maximizing the gross national product of a country, or engineers justify a new highway on the basis of its economic stimulus. This approach has considerable superficial appeal, but unfortunately is naïve and deficient.

At first, it does seem logical to concentrate on "maximizing the pie." It would seem obvious that the more pie we had, the easier it would be to make everyone happy. The more we had to distribute, the easier it would be to give everyone enough. The separation of the problems also is professionally appealing: the production problem seems like a nice technical issue suitable for systems analysts, whereas distribution involves social and political issues one might prefer to leave to others.

The approach is naïve because no system produces only one product that we can unambiguously maximize. Life is not just a bowl of cherries we can only make into a pie. Every system involves multiple products and impacts. Without a unique way to compare them, without a social welfare function, we cannot define what it means to maximize product. In short we have apples and oranges in addition to cherries, and do not know how to compare them. Economic development of a country, for example, typically implies urbanization, disruption of the traditional social structure, and other effects not reflected in the statistics on gross national product. A new highway likewise has pollution and other impacts in addition to its economic benefits.

A simple example illustrates the issue. Consider a country that produces two types of products: exports and fresh food. We may suppose its best combinations of design at any time can be represented by a curve. These are the noninferior solutions described in Chapter 8, also known as the production possibility frontier in economics. If there are diminishing marginal returns in each product, the curve will bow away from the origin as in Figure 21.2. The issue is, what is the best design of the system; specifically, what is the best combination of exports and fresh food to produce?

Maximizing the "pie" consists in practice of maximizing economic value. This depends on the relative price of the products. Graphically, the maximum economic value, Z^*, is defined by the point on the production possibility frontier

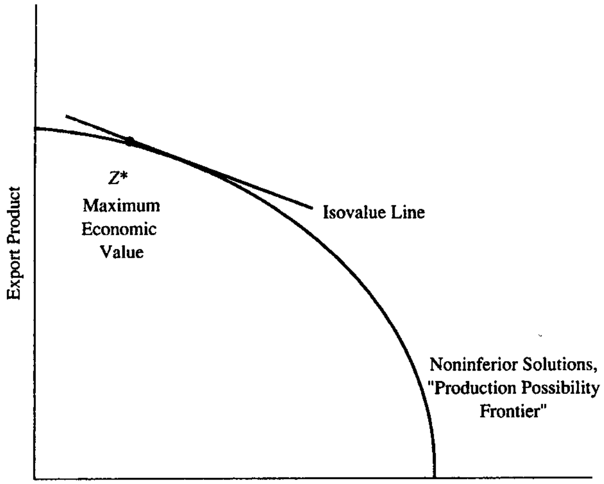

FIGURE 21.2
Maximizing the economic "pie" for example country.

that is tangent to the highest isovalue curve. If we assume, as is most often the case in developing countries, that the relative price of exports is high, maximum economic production will favor exports against local agriculture. This has been a typical result in Nigeria (petroleum, palm oil, and peanuts), Central America (bananas), Malaysia (tin and rubber), and so on. In each of these cases, focus on maximization of economic value has meant that some products, particularly local agriculture, have been neglected.

Maximization of economic value often does not and cannot lead to the best solution for the group. This is most obviously the case if the group's welfare requires minimum levels of products which have not been produced. Suppose for example that our country requires a minimum level of fresh food for its health, as shown in Figure 21.3. We may then imagine that its social welfare function, could it be drawn, would be asymptotic to this minimum as shown. The design that maximizes the social welfare, SWF^*, is quite different from the design that maximizes economic value, Z^*. Moreover, getting the economic maximum may make it impossible to attain the social welfare maximum insofar as money cannot buy some products—such as fresh food—if they have not been produced. In fact, malnutrition is often a common phenomenon of Central American and other countries that export cash crops. Maximizing the economic pie may thus not only be quite different from the social welfare optimum, but in fact preclude it.

The difference between the economic and social optimum is not just a phenomenon of simple examples or foreign countries. The same kind of results

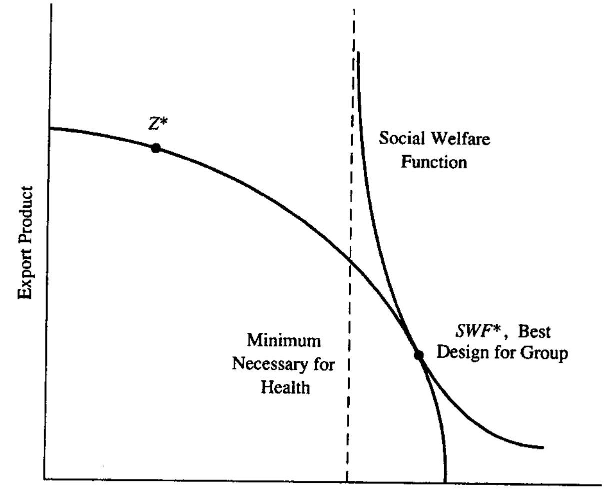

FIGURE 21.3
The economic maximum may be different from and preclude the best design for the group.

can obtain in industrialized nations. We may admit that every country requires an educated, healthy population to ensure its future. But if it focuses narrowly on immediate financial benefits, and neglects the training of teachers, doctors, and researchers, it may not be able to buy them when they are needed. Some attribute the decline in America's international competitiveness precisely to this focus on immediate rewards that neglected investments in the future.

Effect of prior distribution. The price of goods reflects the existing distribution of wealth in society, in the world. The price of 20 carat diamond rings reflects the number of multimillionaires, for example. If there are many, as in the capitalist world, the prices may be high because they can afford to pay a lot for these amusements. If there are no millionaires, for instance in an egalitarian socialist society, then the price of jewels cannot be high—no one can pay for them.

This is an important observation. It means that what we determine as the economic maximum must also depend on the existing distribution of wealth, and must be good or bad depending on how we view this distribution. If you believe that the current situation is unjust, then you will reasonably conclude that the economic maximum does not define an acceptable social optimum.

Put another way, the rich have more economic votes. Having more money, they can call for products to satisfy their desires more effectively than the poor. To the extent that we allow prices to determine the design of a system for a group, we imply that the social welfare function weights the utilities of the rich highly. Many may find this unacceptable.

Consider Figure 21.3 again. The high price of the export product, such as bananas from Central America or coffee from Brazil, depends on the existing wealth of the industrialized countries. Their good fortune gives them the economic power to define what is produced in the exporting country. Does this arrangement necessarily lead to the optimum for all? Only if you believe that the rich somehow have a right to their wealth, that they have earned it in some way.

The possible inequity of the prior distribution of wealth is thus another reason why the economic maximum may not be the social maximum. This reinforces the conclusion that we cannot solve the problem of systems design for a group by the naïve approach that separates the production and distribution problems. These issues must be addressed together.

21.4 ANALYSIS OF DISTRIBUTION

We must first understand how to analyze the distribution problem before we integrate it with production to arrive at the best design for a group. We do this in two stages. First we consider the simpler problem of distributing a fixed set of products to members of a group. Then we examine the larger question of deciding both what to produce and how to share it.

Fixed product. The analysis of the distribution of a fixed set of products to members of a group is best explained graphically. The device for doing this is the *Edgeworth Box*. This is a diagram that indicates the effect of any allocation of two products to two individuals. The lessons drawn from this case apply directly to the more complex situation of multiple products and many members of the group.

The Edgeworth Box displays the two products to be distributed on a rectangle in the X-Y plane. To make it easy to remember the arrangement, think of these products as two chemicals, Xenon and Yttrium. Starting from the origin, any point in the rectangle shows the amount of X and Y distributed to one individual. This varies from none of each at the origin, to all the available products, (X^*, Y^*), at the opposite corner of the rectangle.

The Edgeworth Box is special because it has two origins, one at each opposing corner. As Figure 21.4 shows, any point A in the space simultaneously shows the allocation to individual 1, (X_1, Y_1) and, viewed from the other origin, the allocation to individual 2, (X_2, Y_2). Naturally, these allocations must be complementary, and sum to the total amounts available:

$$X_1 + X_2 = X^* \qquad Y_1 + Y_2 = Y^*$$

Consider how each individual views the allocation of products. Assuming that these have been represented as benefits with diminishing marginal utility (see Section 18.3), we may presume that each individuals's utility function can be represented by isovalue lines bowed away from the origin, rather as shown

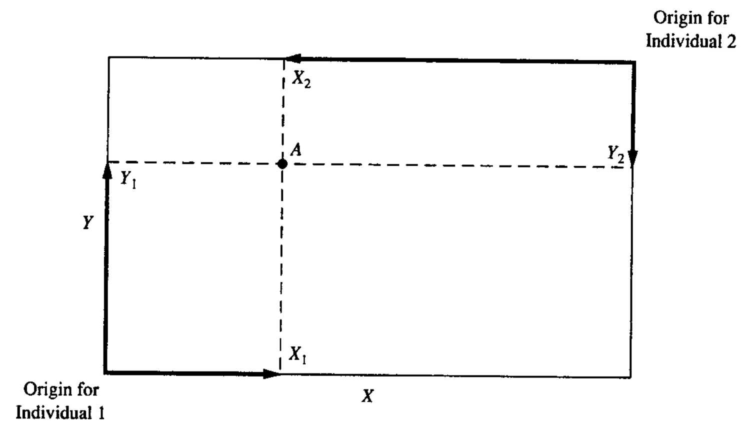

FIGURE 21.4
The Edgeworth Box: any point in the rectangle indicates the allocation of X and Y to two individuals.

in Figure 21.1 and Figure 18.7. Any allocation represented by point A will simultaneously be on two intersecting isovalue lines, one each for individuals 1 and 2, as Figure 21.5 shows. The Edgeworth Box thus also indicates the utility any individual gets from the distribution.

Now the important result: whatever the initial allocation of products to individuals, we should almost always be able to improve the allocation to *everyone's* benefit! To see this, consider the effect of changing the allocation from A to B, as shown in Figure 21.5. Given that each individual's isovalue curves bow away from their origin, then the move from A to B provides both individuals with higher utilities.

What is remarkable about this result is that we achieve *mutually* beneficial improvements in the distribution although the product, the "pie," is fixed. This deserves special attention because it is counterintuitive to most people. The general—but wrong—assumption is that since the "pie" is fixed then any advantage one person gets must be to the disadvantage of the other. (Such a situation is called a *zero-sum game* in operations research.) The simple analysis would be correct if we were only dealing with a single product. But we are generally not dealing with a zero-sum game when we have multiple products because the difference between these products provides mutually beneficial opportunities.

The reason we can usually achieve mutually beneficial improvements is that, at any point, each individual values each product differently. What one person values little, another may really want. There should thus be opportunities for trades that benefit both.

Figure 21.5 illustrates the situation. At allocation A, individual 1 gets little value from more Y, as shown by the steep isovalue curve. Individual 2, on the

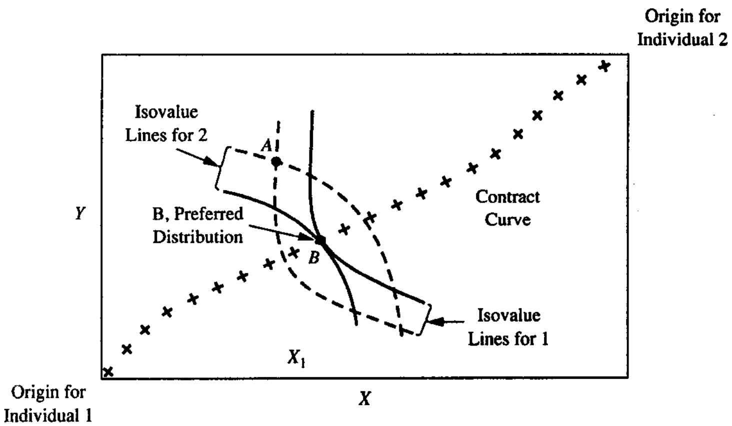

FIGURE 21.5
Any allocation *A* can, in general, be improved to the mutual benefit of all concerned.

other hand, values *Y* relatively highly, and is quite willing to trade some *X* to obtain it. This is what happens as they trade and change the allocation from *A* to *B*. Note that in this analysis we assume, as is generally the case, that any person's utility for a product is nonlinear, and depends on what is in short supply or surplus. Trading between individuals allows them to achieve balanced distributions that maximize their mutual utility. (See box for an example.)

There is a limit to the mutual improvements that can be achieved. In Figure 21.5, *B* is such a point. Any change in distribution from *B* necessarily decreases the utility of one or both individuals. Graphically, this is indicated by the tangency of the isovalue curves. Any movement away from this point must be to a curve of lower value.

The Edgeworth Box

For this example, the products are M and N, with the total amount to be shared, $(M^*, N^*) = (5, 5)$. Each individual values M and N differently. Assuming linear values for simplicity

$$U_1 = 4M + N$$

$$U_2 = M + 2N$$

Suppose the M's and N's are shared so that individual 1 gets (3, 2). Individual 2 would thus get the complement: [(5 − 3), (5 − 2)] or (2, 3). This is point *A* in Figure 21.6.

For the allocation represented by *A*, the value to each individual is:

$$U_1 = 4(3) + 2 = 14$$

$$U_2 = 2 + 2(3) = 8$$

The isovalue line for individual 1 can be expressed as

$$N = 14 - 4M$$

and for individual 2

$$N = (8 - M) / 2$$

These two lines pass through *A* as in Figure 21.6.

The fact that the isovalue lines intersect indicates that it is possible to reallocate the products to mutual advantage. Any allocation represented by the points in the shaded area would lead to increases in utility for both individuals, being further away from the origin than the isovalue line for each individual. For example, consider the allocation represented by *B*: this gives (4, 1) to individual 1 and (1, 4) to individual 2. The resultant benefits to each are greater than those of the initial allocation *A*:

$$U_1 = 4(4) + 1 = 17 > 14$$

$$U_2 = 1 + 2(4) = 9 > 8$$

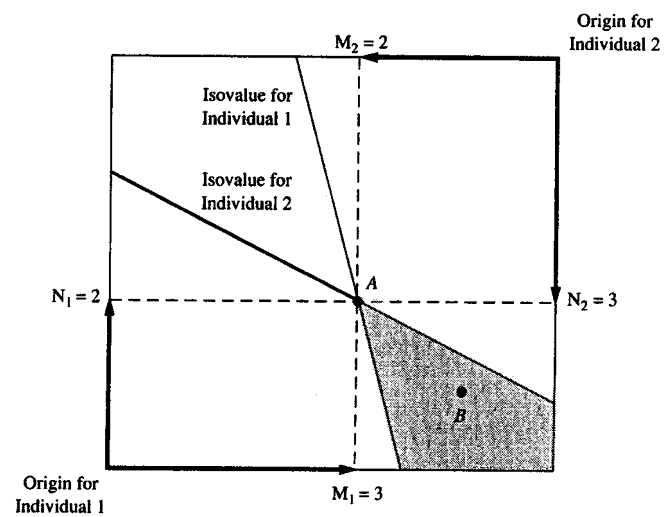

FIGURE 21.6
Illustration for Edgeworth Box analysis.

The locus of all allocations that represent the limit of mutually advantageous changes in distribution is the *contract curve*. Graphically, this is the locus of all points of tangency of the isovalue curves. Because we cannot effect the interpersonal comparison of utilities, we cannot define this curve numerically.

The important point about the contract curve is that it defines only a minuscule fraction of all the possible allocations of a fixed set of products. The probability that any initial distribution is on the contract curve is thus essentially zero. Consequently, we may assume as a practical matter that members of a group can always change any initial allocation of products to their mutual advantage.

Choice of product. The analysis of distribution when the amounts to be produced are not fixed is also best done graphically. This is because, given the impossibility of acceptable interpersonal comparisons of utility and social welfare functions, no analytic solution is possible.

We focus on noninferior designs, the ones on the production possibility frontier that define the maximum combinations of products we can achieve. Each particular noninferior design, Z_i, provides benefits that can be shared by members of the group. The maximum utility to these individuals is defined by the contract curve, so that each design implies a maximum level of utilities to the members. The case for two individuals can be plotted as in Figure 21.7.

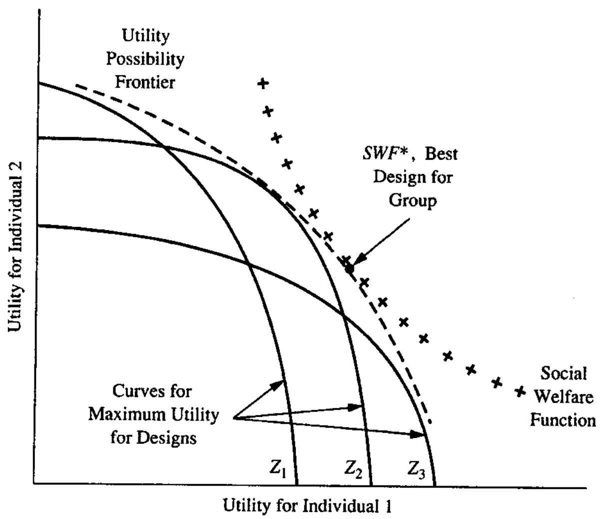

FIGURE 21.7
The best design for a group requires finding both the best system and the best way to allocate its benefits.

The set of all noninferior designs defines a *utility possibility frontier*. This is the maximum utility that individuals can receive from all possible productions and allocations, and is the outer bound of all the maximum utility curves for a fixed production, as Figure 21.7 shows.

The best design for the group, SWF^*, is the one that achieves the highest social welfare. Graphically it would be defined (if we could actually draw acceptable curves) as the point of tangency of the utility possibility frontier and the highest isoutility curve for the social welfare function.

This discussion indicates that achieving the best design for the group of members with different values requires us to do two things:

1. Move along the production possibility frontier of noninferior alternatives to find the design that can give all members the highest utility.
2. Find the allocation of products of this system which exploits all opportunities for increasing mutual advantage.

21.5 COLLABORATIVE NEGOTIATION

The best procedure for collective decisionmaking is collaborative negotiation. Negotiation is essential because there is no analytical way to proceed, given the impossibility of making interpersonal comparisons of utility. Some kind of bargaining, or exploration of each other's preferences and trade-offs has to be worked out. Collaboration is equally important because of the numerous possibilities for mutual gain.

This section discusses some of the characteristics of the negotiation process that system analysts ought to consider. Naturally, this is only an introduction. Negotiation is a subtle subject at the limit of where systems analysis leaves off and psychology, not to say diplomacy, begins.

Trade-offs. Changes in allocation are central to mutually beneficial improvements, as the previous section shows. These are achieved by exchanges between individuals. A key aspect of negotiation thus consists of finding opportunities for these trade-offs.

Trade-offs can most obviously be between products of the system. They can also be between their characteristics, such as cost, reliability, or quality. They can even be between characteristics of the system itself: its site, its configuration, its accessibility.

More subtly, trade-offs can be made between different levels of risk. Some members of a group may be quite reluctant to accept risks when others are not; this provides the opportunity for exchanges between the two. One group can provide insurance to the others. For example, if a community is extremely averse to the possibility of dioxin emissions from a proposed municipal incinerator that an entrepreneur proposes to build, they can come to an arrangement whereby the developer guarantees to close the plant if it does not perform as expected. The developer may then be able to build, at the cost of more risk; the community bears less risk but permits the facility.

Economies of scale. Economies of scale in a system provide an enormous additional motivation for collaborative negotiation. As indicated in Section 4.4, the cost function for a system is often

$$\text{Cost} \sim \text{Constant (Capacity)}^{0.5}$$

This means that when many entities get together to establish a system, their total cost can be very much cheaper than if they each attempt to provide the services for themselves. If nine towns of equal size got together to build a metropolitan incinerator, for example, their total cost would be one-third (!) of what it would be if they each built an individual facility.

The tremendous savings possible by exploiting economies of scale can be used effectively to secure agreement for a large system, such as a metropolitan incinerator, an airport, or a power plant, that might otherwise be rejected by the community. Some of the financial benefits of the larger system can be traded off with other members of the group who may care more about other factors, such as safety, noise, and risk.

Exploitation of great by small. An almost inevitable characteristic of collective decisionmaking is that the smaller members of the group get a larger distribution of benefits than their proportional share according to size. This fact is counterintuitive; one might expect that the powerful members get their own way and exploit the weak.

The strong do dominate the weak when results are determined merely by brute force. But the reverse is true when we consider collective decisionmaking in which each member of a group has some kind of vote on the final design.

The reason the bigger members get exploited is that they have a great deal to gain from collective action, precisely because they are big. They therefore are often willing to trade these savings relatively generously to the smaller members, who may not have as much to gain, to ensure that a collective decision is made. A salient example of this phenomenon is the OPEC oil cartel. Saudi Arabia is by far the greatest producer and yet has consistently got far smaller production quotas relative to its potential than the other nations. Saudi Arabia has, in effect, paid other nations to maintain the cartel.

A numerical example of the exploitation of the great by the small is tabulated in Table 21.2. Suppose a potential group consists of four members, A, B, C, and D. A is the largest, and requires six units of product. The others need only one unit apiece. Assume that cost is proportional to the square root of capacity, the approximate situation with economies of scale. Their total costs are 5.45 if the members act independently, but only 3 if they collaborate, producing a savings of 2.45 or 45%. There is no way to predict how the members of the group would in fact distribute these savings, but the allocation shown in Table 21.2 would be a representative pattern. In this case the larger member, A, gives up some of its proportional share of the savings to ensure that they exist. The smaller members, B, C, and D, do get less than A, but more than their proportional share.

TABLE 21.2
Example of Exploitation of Great by Small

		Cost if acting		Group savings	
Member	Capacity required	Alone	As group	Total	%
A	6	2.45	1.8	0.65	25
B	1	1	0.4	0.6	60
C	1	1	0.4	0.6	60
D	1	1	0.4	0.6	60
All	9	5.45	3	2.45	45

This result may or may not appear equitable to different people. The important fact is that it is frequently what happens normally. It should thus be anticipated.

Process. The process of negotiation is most important to its outcome, to the result of collective decisionmaking. Process is important at two levels. First, the result depends on the voting rules for the group; different arrangements translate the desires of individual members differently. Simple majority rule is different from proportional representation. The United States Senate is surely different, with its rule of two senators per state regardless of size, than it would be if senators represented equal districts.

Secondly, small procedural features can make a substantial difference to the outcome. The discussion of the Arrow Paradox in Section 21.3 illustrates this: the option that emerges as preferred from the paired comparisons depends absolutely on which pairs you consider first. This kind of phenomenon explains why experienced diplomats and negotiators pay great attention to the details of procedure.

Finally, personal styles influence the result of bargaining. Abrasive, aggressive approaches may seem to be effective but often fail utterly in achieving the significant gains that can result from collaborative negotiation. A full treatment of the subject being beyond this text, the interested reader is referred to an excellent introductory manual to effective bargaining: *Getting to Yes*, by Fisher and Ury, (see References).

21.6 RECOMMENDED PROCEDURE

As this chapter stresses, there is no analytic way to define the design, the system that best meets the desires of a group of different members. We can only recommend a process for dealing with the problem.

The recommended procedure consists of five steps:

1. Model the physical system
2. De[illegible] noninferior options, the production possibility frontier
3. Determine individual preferences
4. Explore the possible trade-offs
5. Negotiate toward a collectively satisfactory solution

A model of the system is essential to understand the technical possibilities. This consists of developing either an analytic or a simulation model of the production function, as described in Chapter 2.

The production possibility frontier focuses our attention of the designs that clearly produce the most benefits. These are only a very small fraction of the total number of options. Considering only these keeps the analysis to a minimum. These noninferior solutions can be determined as described in Chapter 8.

Individual utility functions indicate what products or characteristics of the system are most and least important to each member of the group. This information provides the basis for exploring trade-offs. The appropriate multiattribute utility functions can be determined as described in Chapters 19 and 20.

The exploration of possible trade-offs is an information gathering activity. In order to find opportunities for trades we need to understand which products or characteristics are valued differently by different members, and which thus offer the possibility of mutually beneficial exchanges.

Finally, the individual members must negotiate between themselves. In doing this they proceed along two dimensions: the product and its distribution. In terms of the product, they move along the production possibility frontier, from the designs they would prefer themselves, to compromise designs that are collectively acceptable. Simultaneously, they should explore how the production itself should be shared. This process is not necessarily one in which the members lose; as shown in Section 21.4, this process may be mutually advantageous as they change the allocation of benefits. The final result should be an effective collective decision on the design of the system.

21.7 APPLICATION

At MIT we have successfully used the SUMIT case study of a hypothetical electric power system to illustrate the issues of collective decisionmaking. It has been applied in a wide range of situations, in graduate subjects, in summer institutes, and short seminars for senior professionals. Its details are available on request. SUMIT concerns a hypothetical electric company that must meet the demands of its clients while satisfying political and environmental constraints. The company must devise a plan to provide sufficient additional power for its service area, keeping costs within tolerable bounds, yet ensuring environmental quality.

The case is organized as a participatory exercise among the members of the class or seminar. They are divided into small subgroups representing: the company itself; its industrial customers; environmentally conscious consumer advocates; and the state agency that regulates the price of electricity. Each subgroup is given materials enabling it to play its role, and is provided with a diskette that enables it to calculate the consequences of any design. The subgroups are told that collective decisions can only be made when three out of the four subgroups agree. The entire group is then charged with developing a plan for the future.

The SUMIT case provides the participants with the model of the computer physical system, and so automatically fulfills step 1 of the recommended procedure. This is realistic as such models are, in fact, quite common.

The exercise of finding the noninferior solutions is often most revealing, notably to experienced professionals who participate. This is because the usual way to approach the design is to look for the optimum solution, not to define the frontier that includes all kinds of designs that are not optimal to them. The revelation is associated with the expanded set of solutions. It leads all participants to recognize that their preferred design is contingent on their own values, which are probably not representative of the group.

When the participants determine their preferences, they are usually surprised to find how complex they are. The nonlinearity of their utility functions forces them to recognize that, at some points, they are in fact prepared to sacrifice some of their preferred objectives in order to obtain a reasonable overall balance. This step both prepares them to consider trade-offs and helps them understand what kind of trade-offs might be acceptable.

At this point, in running the SUMIT case, we encourage the participants to engage in a lot of informal discussions between themselves. It is necessary for them, as it is in real life, to explore possible trade-offs in preparation for more formal negotiations.

The specific outcome of the final negotiation for any group is unpredictable. It depends on many factors, not the least of which are the skill and personalities of the participants.

Some general results of the SUMIT case stand out consistently. These are:

1. A collectively satisfactory design is possible. This is a big lesson for most participants who, when presented with the problem, assume that resolution can only come through power and at the expense of others.
2. The emergence of the satisfactory collective decision is not a matter of chance, it is systematically facilitated by the recommended procedure.
3. Collaborative efforts are necessary. Participants who adopt combative, unyielding attitudes miss the opportunities for exploiting mutual benefits and systematically end up with inferior solutions as compared to other groups.
4. Working on the distribution problem is essential. It is virtually impossible to determine the acceptable solution without simultaneously defining how the benefits will be shared.

Overall, we have found that the SUMIT case study provides an excellent integration of the range of issues of systems analysis.

REFERENCES

Arrow, K., (1963). *Social Choice and Individual Values*, Yale University Press, New Haven, CT.
Fisher, R., and W. Ury, (1981). *Getting to Yes*, Penguin, New York.
Olson, M., (1971). *The Logic of Collective Action: Public Goods and the Theory of Groups*, Harvard University Press, Cambridge, MA.

PROBLEMS

21.1. *Sharing the Crop*
Peter and Paul jointly grow rye and wheat (R, W) for their own use. They each have different tastes for these products:

$$U(\text{Peter}) = 4R^{0.5} + W^{0.5} \qquad U(\text{Paul}) = 2R^{0.5} + 3W^{0.5}$$

(*a*) Calculate Peter's utility if his share is (16, 16). Plot the isovalue line through this point. You may find it helpful to express W = *f*(R).
(*b*) Repeat (*a*) for Paul when he also gets (16, 16).
(*c*) Suppose the crop yields (32, 32). Compare their utility if they share the crop evenly with that if Paul gets (9, 25) and Peter (23, 7). Are we robbing Peter to pay Paul? That is, does the advantage to one come at the disadvantage of the other?
(*d*) Draw the Edgeworth Box that illustrates the unequal distribution in part (*c*). Define the region in which Peter and Paul can mutually improve their positions from that defined by the "fair" distribution, in which each gets half of each product.

21.2. *Dividing the Pie*
Mother Evening has to divide her four apple and four blueberry pies between her children, Sam and Janet. She knows they have different tastes:

$$U_S = 10A^{0.5} + 2B^{0.5} \qquad U_J = 4A^{0.5} + 8B^{0.5}$$

(*a*) Calculate Sam's utility if he gets half of each. Plot his isovalue curve through this point.
(*b*) Repeat (*a*) for Janet.
(*c*) Draw the Edgeworth Box for Sam and Janet, showing their isovalue lines through an equal division of the pies. Identify the possibilities for a mutually more advantageous way for them to share the pies.
(*d*) Propose a better allocation and indicate how it increases the pleasure of Sam and Janet Evening.

21.3. *Share Cropping*
On a plantation that can grow either cotton or food, the owners have decided—in view of prevailing export prices—to concentrate on growing cotton for export, and only let the workers cultivate small plots for food. The resultant production is (C, F) = (40, 1). The utilities of the owners and workers for the products are:

$$U_o = C + F^{0.5} \qquad U_w = C + 10(F^{0.5})$$

(*a*) Calculate the utility of the owners and workers if the owners get all the cotton and the workers get all the food.
(*b*) Draw the Edgeworth Box for the allocation of (40, 1) and determine its optimal distribution between owners and workers.
(*c*) Suppose that it is possible to produce (30, 9) on the land: discuss whether you think this would be better? What do you mean by better? How do you weigh the importance of the preferences of the two groups?

APPENDIX

NOTATION

Symbol	Meaning	Chapter
$\mathbf{A}$	Matrix of Constraint Parameters	5
A_0	Parameter of Cost Function	4
AC_Y	Average Cost of Product	4
a	Constant in Utility Function	18
	Returns to Scale Parameter	4
a_i, a_{ij}	Parameters of Constraint Equations	5
	Parameters of Production Function	2
$\mathbf{B}$	Vector of Constraints	5
B	Discounted Benefits	13
b	Constant in Utility Function	18
b_j	Constraint	5
$\mathbf{C}$	Vector of c_i	5
C	Discounted costs	13
°C	Degree Centigrade	18
$C(Y)$	Cost Function for Product	4
CE	Certainty Equivalent	19

C_c	Closure Costs	13
C_f	Fixed Costs	14
C_k	Initial Capital Costs	6; 12
C_r	Annual Recurring Costs	12
C_v	Variable Costs	6; 14
CLR_j	Conditional Likelihood Ratio for O_j	15
c	Constant exponent	19
c_i, c_{ij}	Parameters of Objective Function (often costs)	5
c_o	Fixed Charge	5
$c(\mathbf{X})$	Input Cost Function	4
crf	Capital Recovery Factor	15
D_i	A Possible Decision	16
E	Event	15
$EV(\bullet)$	Expected Value of $(\bullet)$	15; 16
e	Naperian number, 2.7183	11
F	Future Amount	11
°F	Degree Fahrenheit	18
$f(\bullet)$	Function of $(\bullet)$	5
$f_s(K)$	Cumulative Return Function for State S	7
$G(\mathbf{X})$	Objective Function for Dynamic Programming	7
$g(\bullet)$	Production Function	2
$g_i(X_i)$	Return Function	7
$h(\bullet)$	Constraint Equation	3
I_i	Inventory in Period or Stage i	7
IRR	Internal Rate of Return	13
i	Interest Rate	13
K	State, in Dynamic Programming Normalizing Parameter for Multiattribute Utility Function	7 20
k	Number of Observations in a Test	17
k_i	Scaling Factors for Individual Dimensions of Multiattribute Utility	20

L	Lagrangean Equation Number of Locations Lottery	3 16 17
LEP	Lottery Equivalent/Probability	19
LI	Lottery Revised by Information	17
LR_N	Likelihood Ratio after N Obserbations	15
M	Number of Points in Utility Assessment	20
MC_i	Marginal Cost	4
$MF(\bullet)$	Monotonic Function of $(\bullet)$	18
MP_i	Marginal Product	2
MRS_i	Marginal Ratio of Substitution of Product X_i for X_j	2
N	Number of Dimensions; Periods; Observations	4; 11; 15; 17
NPV	Net Present Value	11
O	Observation	15
O_{ij}	Outcome Conditional on D_i and E_j	16
OC_K	Opportunity Cost	6
OF	Objective Function	6
P	Present Amount	11
$P(\bullet)$	Probability of $(\bullet)$	15–20
P_e	Response in Lottery Equivalent Method	19
$PF(\bullet)$	Preference Function	18
$PLT(\bullet)$	Positive Linear Transformation of $(\bullet)$	18
PM	Preference Measure	18
P_j	Probability of E_j	16
P_k	Probability of TR_k	17
P(E/O)	Posterior Probability of Event given Observation	15
P(O/E)	Conditional Probability of Observation given that event has occurred	15
p_i	Price of Input X_i	4
R	Equal Annual Payment	11

R_i	Requirement for Period or Stage i	7
RTS	Returns to Scale	2
r	Returns to Scale parameter Discount Rate	4 11
r_{irr}	Internal Rate of Return	13
r_p	Return on Investment	12
S	Future Amount Number of Sizes	11 16
S_j	Slack Variable	3
SP	Vector of Shadow Prices	6
SV	Vector of Slack Variables	6
SWF	Social Welfare Function	21
s	Scaling factor	2
T	Number of Periods	16
TR_k	Test Result	17
$U(\cdot)$	Utility of $(\cdot)$	18
U_i	Utility to Individual i	21
u	Learning Curve Exponent	4
$V(\cdot)$	Value of $(\cdot)$	18
W	Vector of Dual Variables	5
w_i	Relative weight	2; 21
X	Vector of Resources Used in Design	2
X_b	Buying Price of a Lottery	19
X_s	Selling Price of a Lottery	19
X_i	Inputs to a Design	2
Y	Product of a System	2
Y_i	Product in Period i	7
Z	Dual Objective Function	5

Greek Symbol

λ_j	Lagrangean Multiplier	3

Special Mathematical Notation

$(\cdot)^*$	Best Value of $(\cdot)$
$(\cdot)_*$	Worst Value of $(\cdot)$
$\sim$	Is Indifferent to
$>$	Is Greater Than
$<$	Is Less Than
$\succ$	Is Preferred To
$\prec$	Is Preferred By