PROLOGUE
DESIGN IN GENERAL

The design process is widely used in solving problems and achieving desired ends, even by people with little or no training or interest in the general process. So it makes good sense to begin our exploration of the idea of design with a commonplace example familiar to most people: designing a garden. The garden example enables us to then consider attributes of design in general and the more-or-less sequential set of stages that commonly occur in the process. In Chapter 1, the ideas presented here are applied to the particular case of curriculum design. Chapter 2 considers the basic dimensions of curricula that lend themselves to design.

Design

• To create, plan, or calculate for serving a predetermined end.
• To draw, lay out, or otherwise prepare a design or designs.
• The result of a process of designing.
• The process of selecting the means and contriving the elements, steps, and procedures for what will adequately satisfy some need.
—*Webster's Third International Unabridged Dictionary*

An Introductory Example

Suppose our desired end is to have a backyard garden. One does not have to be an expert to know how to design a home garden, though it is necessary to know something about plants—or about how to find out about them. Our approach need not be orderly, one careful step at a time, but among the things we would surely do are these:

- We would gradually become clearer on what we want from a garden. Will we grow vegetables or flowers or both? Or will we use the garden to hold parties, keep bees, or just enjoy working the earth? And how will we judge the success of our design—will it result in more nutritious meals, honey, reduced florist bills, social prominence, peace of mind, or some combination of those goals? At the same time, we would begin to identify any physical, financial, and legal constraints on what we plan to do.

- As our goals and constraints become clear, we would identify some alternative design concepts to help us focus our thinking about design possibilities. A design concept for a garden may be to provide a seasonal succession of vegetables or flowers, imitate an English country garden, attract (or repel) certain wildlife, or simply have a backyard that requires weekends-only maintenance. We would study model gardens in books and magazines, search the Internet, talk to professional gardeners, and look at the gardens of our friends and neighbors—on the chance of finding possibilities that might not have occurred to us.
- We would narrow the possibilities down to a few appealing design ideas that would work within the constraints we face, think over our desired end, and choose an approach that would seem to be the best bet.
- Then we would develop the idea in enough detail to get started actually planning the garden. During this stage, trade-offs would have to be considered—a choice, for instance, between the desire for large shade trees and for sun-loving plants. Our plan would very likely be in the form of a sketch showing how the plants and other features of the garden would be placed. In developing the final design, we could call on experts for advice, or use a commercially available computer-assisted garden design program.
- As the actual work in the backyard progressed, we would come up against unexpected difficulties, forcing us to modify the original design—or even to choose an alternative design concept altogether.
- Even with the garden in place, the design challenge would not be over, as any home gardener knows. Maybe the actual garden would not look exactly like the design, or it would be just like it but would not please us. Maybe mistakes made in implementing the design would now show up. In other words, we would discover or decide that modifications were needed. We would have to make allowances for the fact that even if the garden were entirely satisfactory at first, it could turn out that as the plants matured, the relations among them would change enough to require still other modifications.

As our discussion of the garden example shows, there is nothing mysterious about design—architectural, engineering, horticultural, or any other kind. But as straightforward as it is, design does have certain features that people sometimes overlook in solving problems and achieving desired ends. It is worth, therefore, exploring the design process more generally.

Attributes of Design

The story above would have been much the same had the desired end been the Brooklyn Bridge rather than a home garden. Each particular design undertaking has its own special features depending on traditions and circumstances, but in general the process applies equally to the design of any object, process, or system. As our garden example illustrates, design has these attributes:

Design is purposeful. The purpose (desired end) may be to improve a curriculum's effectiveness, solve a problem of traffic congestion, replace pesticides with crop diversification or rote learning with understanding, exploit some existing technology in new ways, or create a new product or service. In practice, there may be many purposes in a design undertaking. They may be in harmony or in conflict, explicit or hidden, immediate or long-range, political as well as technical. Whatever the mix, the designers of a project are better off if they know all of the purposes at the outset, so that they can respond accordingly.

Design is deliberate. The Brooklyn Bridge did not just appear one day in all its glory, nor did it evolve over the decades from a pontoon bridge, nor did it result from a lot of workers showing up each day and deciding what to do next. That may seem obvious when thinking of a bridge, but not always when thinking of a school curriculum.

Leonardo da Vinci's diagram of ribbed wing for a flying machine.

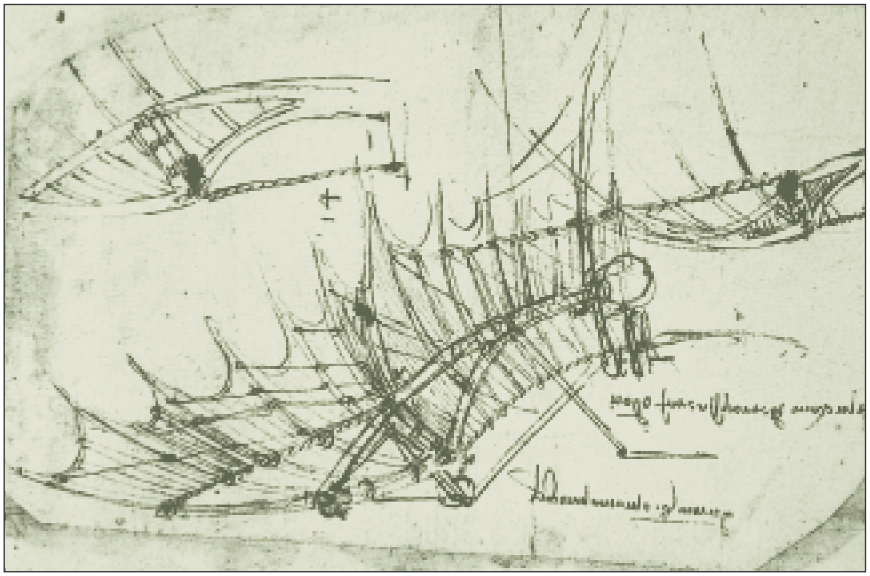

Design is a conscious and deliberate effort to plan something—an object, event, process, or system. A curriculum is one such thing and just as susceptible to deliberate planning as a garden, ship, weapon, banquet, traffic pattern, ballet, assembly line, or scientific experiment.

Design is creative. For all its practical focus, design is not some mechanical process that invariably leads to success. Like science itself, which is often misrepresented as a fixed sequence of steps ("the scientific method"), the design process is a highly variable and creative process. Nevertheless, like scientific inquiry, it has certain features that show up again and again. At every stage of a design undertaking, whether it happens to be the design of a new hospital or a telephone routing system, there are opportunities for innovative thinking, novel concepts, and invention to be introduced. Perhaps it is fair to say that design is powerful precisely because it is at once systematic and creative, feet-on-the-ground and head-in-the-clouds.

For an engrossing account of the importance of creativity in design, read the story of the invention of the chronometer in *Longitude: The True Story of a Lone Genius Who Solved the Greatest Scientific Problem of His Time*, by Dava Sobel (1995)

Design operates on many levels. There was a design for the Brooklyn Bridge as a whole. But there also had to be a design for each of its parts and construction operations. Those designs covered such activities as cutting and transporting the stone for the towers, fabricating the cables, laying the roadway, and even creating special tools. An essential design principle is that the design decisions at one level must be compatible with those at the higher levels. (In the case of the Brooklyn Bridge, a design for a rivet that did not match the expected tension on a main girder could eventually have led to the collapse of the whole bridge.) Many common engineering tasks involve selecting among parts that already exist, so what parts are available may affect the larger-scale design. In other situations, the larger-scale design may also require the design of special new parts and processes.

Design requires compromise. Design is not the pursuit of truth or perfection, but it has to get the job done. Architects and other designers are expected to come up with practical solutions that work well enough in the circumstances. In reaching decisions leading to such practical solutions, trade-offs are made among benefits, costs, constraints, and risks. Therefore, no matter how careful the planning or how inventive the thinking, designs always end up having shortcomings when viewed from one perspective or another, but they are knowingly accepted as reasonable compromises. Some shortcomings, however, may be unanticipated and may not show up in the design process until the designed object or process is put into use, so the design should include provisions for keeping an eye on its success.

Design can fail. There are many reasons why designs may fail. Stereoscopic movies were a marketing failure, even though technically they worked as planned. Leonardo da Vinci's flapping-wing aircraft didn't work for mechanical and conceptual reasons (because no adequate source of power was invented until 300 years later, and because in any case he based his design on the wrong model—birds). Hydrogen-filled dirigibles failed because they were unsafe, as dramatically demonstrated by the *Hindenburg* disaster. A designed system can fail because one or more components fail or because the components do not work well together, even though each works well enough by itself. In most cases, however, designs are neither wholly satisfactory nor abject failures, and so a key element in design is the provision for continuing correction, assessment, and improvement, both in the initial design process and after.

A 1950's design for 3-D movies did not survive.

"There is no perfect design. Accommodating one constraint well can often lead to conflict with others. For example, the lightest material may not be the strongest, or the most efficient shape may not be the safest or the most aesthetically pleasing. Therefore every design problem lends itself to many alternative solutions, depending on what values people place on the various constraints. For example, is strength more desirable than lightness, and is appearance more important than safety? The task is to arrive at a design that reasonably balances the many trade-offs, with the understanding that no single design is ever simultaneously the safest, the most reliable, the most efficient, the most inexpensive, and so on."

—*Science for All Americans, p. 28*

Another version of design stages is outlined in *Design and Problem Solving in Technology* by John Hutchinson and John Karsnitz (1994): (1) Identifying problems and opportunities; (2) Framing a design brief; (3) Investigating and researching; (4) Generating alternative solutions; (5) Choosing a solution; (6) Developmental work; (7) Modeling and prototyping; (8) Testing and evaluating; and (9) Redesigning and improving.

Design has stages. Design is a systematic way of going about planning. While it does not consist of some inflexible set of steps to be followed in strict order, design does involve certain stages that take place at one time or another in the process, or at several times in the process. Practical design is likely to include looping back between stages. Decisions made at one stage may require reconsideration of how they interact with decisions made at other stages. As a result, a design almost inevitably evolves somewhat in the process. The original desired end itself is likely to be clarified as the design progresses, and it may even have to be modified as constraints and costs are discovered. In brief, design involves the following four stages:

- Getting as clear as possible in the design specifications: precisely what is to be achieved and what constraints must be accepted.
- Conceptualizing several alternative design possibilities and thinking about each enough to be able to choose one as a best bet to develop further.
- Developing a complete design, testing aspects of the emerging design along the way, and making adjustments as needed. This stage almost always requires trade-offs involving goals, constraints, benefits, costs, risks, and desirable design features.
- Refining the complete designed product on the basis of experience and feedback from users.

The remaining sections of this chapter elaborate on these four stages.

Establishing Design Specifications

Being purposeful, design is an effort to achieve something specific—devise a means to reach a desired end, satisfy a need, or take advantage of an opportunity. Thus, one of the first steps in design is becoming clear on just what is to be achieved. At the same time, consideration must be given to the essential fact that design is always confronted with constraints on what it can do—time and money limitations, legal restrictions, political considerations, cultural traditions, the laws of nature, and more. Together, goals and constraints determine the specifications a design is expected to meet.

Moreover, design specifications can be technical, aesthetic, financial, political, or moral. The design requirement for the Boeing 777 that it be able to reach an airport safely if one of its two engines fails while over the ocean was both technical and moral. Design specifications can be attributes desired by the client or designer, or they can be demands imposed on them from the outside. In fact, much of the work in developing a design is figuring out how to respond to—or get around—constraints while still reaching the desired goals and serving the purposes that initiated the design undertaking in the first place.

Goals

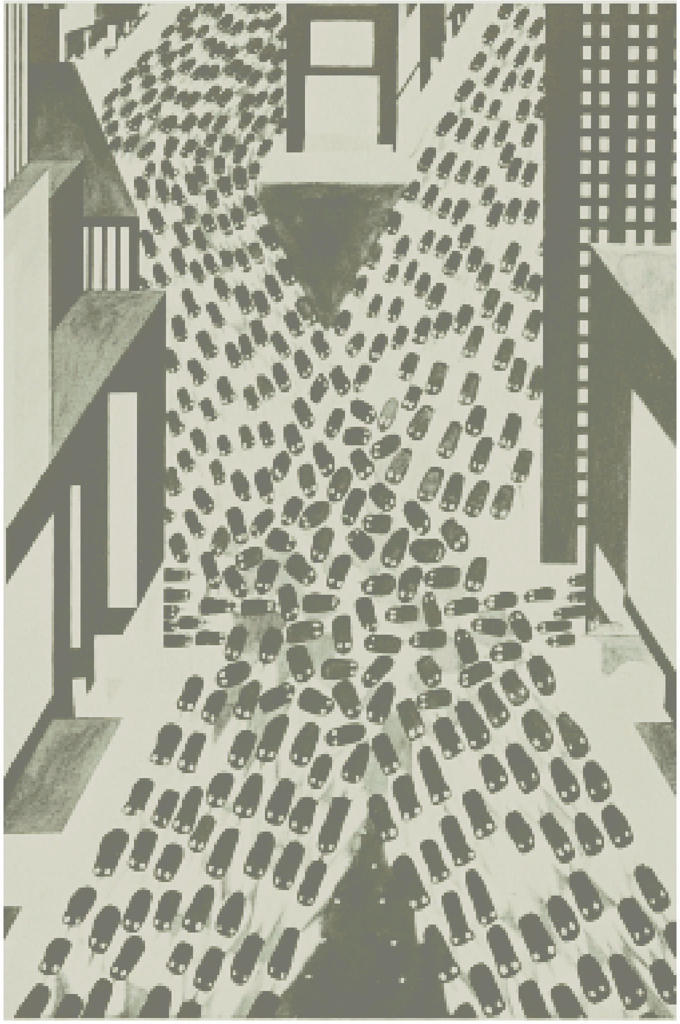

Purposes are usually couched in sweeping language, such as to create a health-care delivery system that is more cost-effective than current ones, an electric vehicle that can travel long distances without recharging, a beautiful backyard garden, or a K-12 curriculum that enables all students to become science literate. Such general purposes then have to be transformed into more concrete goals. Design can get under way before all of the goals have been clearly defined, but there must be more specificity than is provided by the usual statement of purpose. Once the design process is started, the clarification of the goals continues, and they become more and more specific. Occasionally new goals may be added, but for the most part goals are progressively derived as expressions of higher-order goals.

It is important to establish just how specific the goals have to be. Suppose, for example, the mayor of New York City orders city officials to come up with a way to speed up crosstown (east-west) automobile traffic. In response, the officials could try one thing or another and see if crosstown traffic speeds up. But that could easily make matters worse, and so, before tackling the problem, traffic-pattern designers would want to know more specifically what they are expected to achieve in "speeding up crosstown traffic." Are all crosstown streets to be included, or only the main ones? Are all the boroughs to be speeded up, or only Manhattan? Does it mean crosstown traffic all day long or only at certain peak periods? How much faster will do—10 percent? Are we talking about an equal increase for each street or an average increase for all streets? And so on.

Setting goal specifications is rarely as simple as it may seem, however. Often, as goals become clearer, they evoke tensions among groups that do not share the same interests or beliefs. Scientists may be in accord, for instance, with the general proposition that the

United States should create and sustain a vigorous program of space exploration. Yet if Congress translates that idea into increased funding for a space station, scientists who believe that unmanned space exploration is more productive may rise up in opposition, as may scientists in other fields who see their own funding put in jeopardy. Or speeding up crosstown traffic may reasonably be expected to slow down traffic going in other directions. To get one thing, we decide to—or have to—sacrifice another, or to get more of something, we agree to settle for less of something else.

And goals run into constraints. In fact, some goals can become constraints on others. The mayor's dictum could be read as "Speed up traffic in one direction in a way that doesn't slow down traffic in the other direction." The mayor may also stipulate other requirements, such as that the new traffic-flow design must not put citizens and institutions at risk, increase city expenses or decrease city revenues, make it difficult for deliveries to be made to stores, or otherwise impede businesses. Somewhere in the design process, something will have to give, for it simply may not be possible to create a design that fulfills the mayor's purposes and meets all of the conditions he has imposed. However, before informed trade-offs can be made, the constraints and goals have to be specified more precisely.

One can expect goals to be modified in the design process as knowledge of the situation grows and as constraints appear. If, for instance, traffic-flow studies revealed that crosstown traffic is unacceptably slow on only a third of the downtown streets, it may be possible to redefine the goal to bring those streets up to speed without improving the others. Whether this redefinition of the goal would be acceptable to the mayor may have more to do with political considerations than with strictly technical ones.

There is often more to goals than what appears on the surface. In the traffic case, the fictional mayor's real purpose may be to increase his chances of reelection by showing a readiness to deal with a long-standing city problem and to deflect attention from his lackluster performance on other city problems, such as housing and crime. Our real goal in creating a backyard garden may be to keep up with the Joneses or to provide a worthwhile activity for a retired spouse, goals that may not be met if they are made too evident. To take an example from the schools, a goal such as reducing truancy, usually expressed in educational terms, also has as a silent partner—a community goal of keeping unsupervised young people off the street during the day. Designers need to be aware of the possibility that success requires taking into account some goals that have not or cannot be made public.

Constraints

Clarifying constraints is as much a part of design as delineating goals—the goals saying what is to be accomplished, the constraints specifying the limits. In a sense, much of what is said above about goals applies to constraints. For instance, at the beginning of the design process, constraints to be dealt with are often expressed in general terms, such as “must not exclude handicapped customers,” “be able to use existing runways in international airports,” “be environmentally benign,” or “do not increase instructional costs.” None of these particular limiting conditions are stated well enough for design purposes, and so they must be transformed into specifics. For instance, the runway constraint will need early translation into specifications limiting the total permissible weight, brake performance, and maximum take-off distance of the aircraft based on knowledge of the physical properties of runways at international airports.

Some of the most stringent limits on a design are those imposed from the outside. For example, a new museum must meet all of the building codes in the community where it will be located; a new jet aircraft must comply with the safety requirements of the Federal Aviation Administration; a new medicine must meet the effectiveness and safety demands of the Food and Drug Administration. Interestingly, not all such external requirements emanate from government agencies or have the authority of law. Professional and trade associations set expectations and sometimes explicit standards that influence or limit design possibilities.

Just as goals may be modified during the design process because they conflict with one another, so too may constraints. As constraints become more and more precisely defined, it sometimes becomes clear that

“Every engineering design operates within constraints that must be identified and taken into account. One type of constraint is absolute—for example, physical laws such as the conservation of energy or physical properties such as limits of flexibility, electrical conductivity, and friction. Other types have some flexibility: economic (only so much money is available for this purpose), political (local, state, and national regulations), social (public opposition), ecological (likely disruption of the natural environment), and ethical (disadvantages to some people, risk to subsequent generations). An optimum design takes into account all the constraints and strikes some reasonable compromise among them.”
—*Science for All Americans, p. 28*

they block all possibilities for achieving the stated goals. This situation can lead to a renegotiation to eliminate or modify some of the constraints. That may sometimes be difficult or even impossible. With regard to the runway requirement, for example, the manufacturer may try to persuade the authorities that technological innovations like a new kind of landing gear justify raising the maximum permissible landing weight. On the other hand, constraints derived from the laws of physics such as the force of gravity at the earth's surface, are simply not negotiable.

In short, just as goals may turn out to be so unrealistic that they have to be abridged in the design process, constraints also may be so paralyzing that they are challenged along the way: new materials, technologies, and processes can be invented, legal regulations can be overturned, public opinion can be molded, expectations can be transformed, funding priorities can be changed, and so forth.

Conceptualizing a Design

Behind every interesting design—the Panama Canal, the *Whole Earth Catalog*, or the Cannes Film Festival—there is an interesting idea, or several interesting ideas, often likely to be sketchy, rarely precise. Such ideas—"design concepts"—come before designs, and although they rarely survive the design process intact, they are essential for getting started.

The mythical griffin blends aspects of different design concepts.

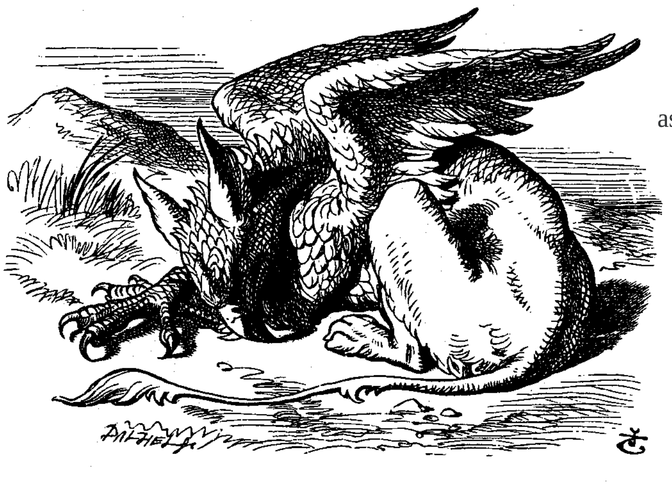

A design concept is any overarching idea, or set of ideas, that suggests the character of the thing to be designed. The designers of the United Nations Headquarters complex in New York City, for instance, set out to create a facility that would proclaim the dignity and significance of the infant organization, yet serve as a practical "workshop for peace." This metaphorical workshop for peace would be international in spirit but still live in harmony with its surroundings, and would point to the future rather than recall the past. Observers differ about how well the final design embodies those concepts, and surely those same guiding concepts might well have led to other designs. But for the designers themselves, the concepts provided a powerful unifying theme. (Design concepts can, incidentally, also play an important role in generating enthusiasm and funding for the project.)

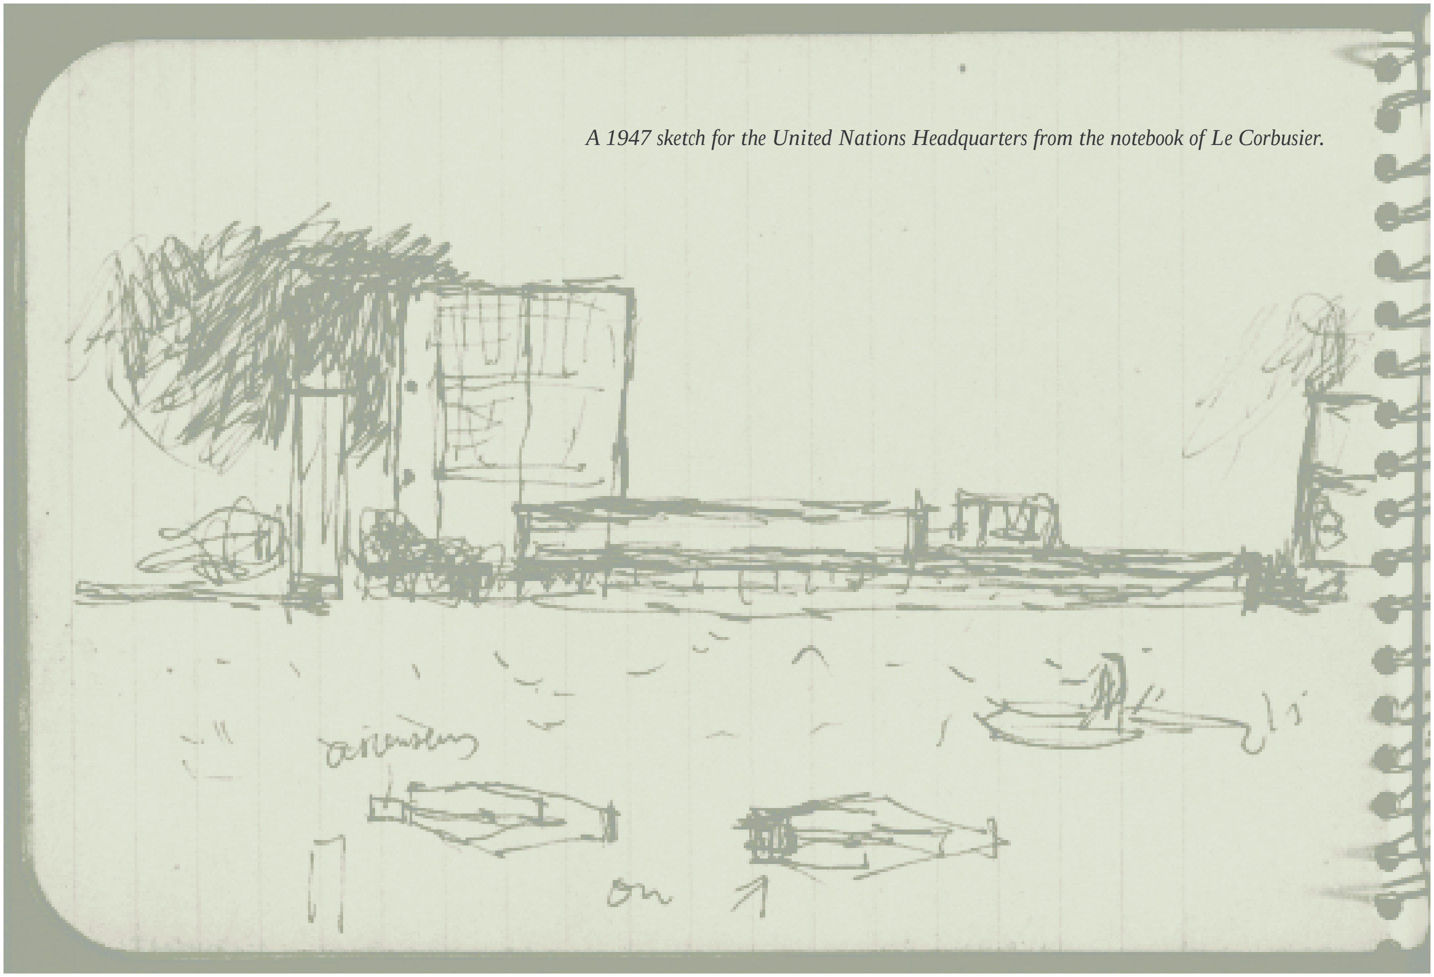

A 1947 sketch for the United Nations Headquarters from the notebook of Le Corbusier.

A "Program of Requirements" for the building specified the staffing and space requirements for the meeting halls, the Secretariat, and other service areas.

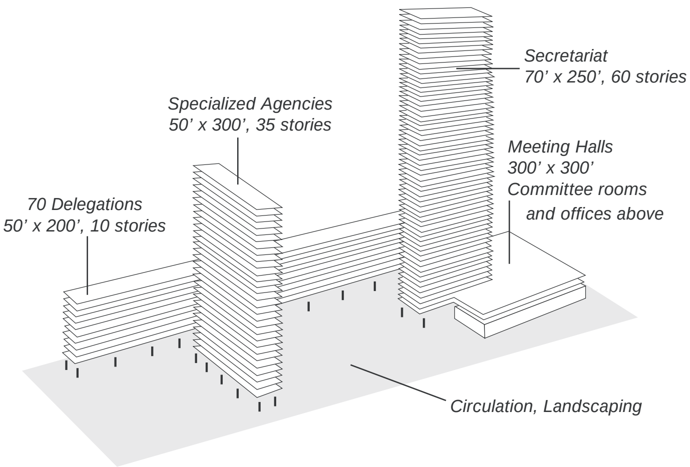

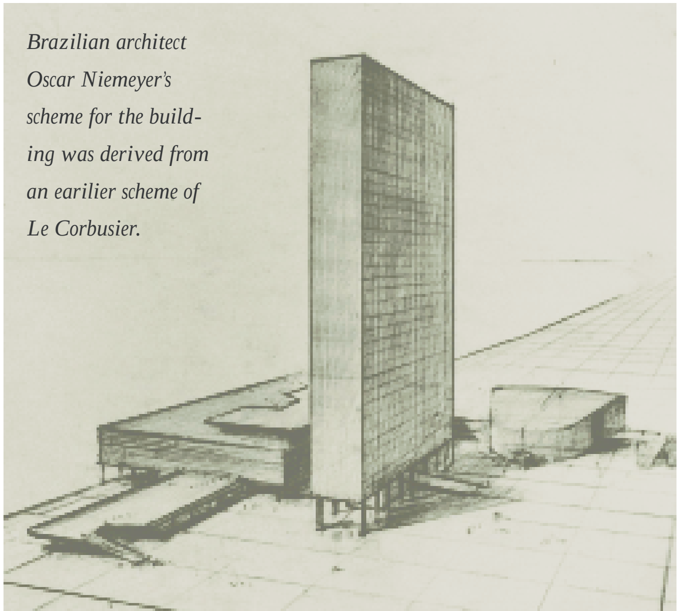

Brazilian architect Oscar Niemeyer's scheme for the building was derived from an earilier scheme of Le Corbusier.

For an interesting account of different design concepts applied to a major undertaking, see *The Path Between the Seas: the Creation of the Panama Canal, 1870-1914* by David McCullough (1978).

Often, the design concept can be captured in a visual sketch, but sometimes it is expressed in prose or a combination of words and images. In more mundane situations, however, there are existing designs so well worked out and widely known that we can just adopt one (a station wagon, a Cape Cod house, a college-preparatory curriculum) or adapt one (a station wagon with bucket seats, a Cape Cod house with a solarium, a college-preparatory curriculum with community projects). Sometimes there is no attractive precedent that will serve as a design concept, and instead the designers emphasize some aspect of an already developing design—a purpose, component, feature, or effect—to provide a character for the design and guide its further development.

In design, there are always alternatives. There is rarely just one right way to do things. Generating and considering competing possibilities is a fundamental step in successful design. Choosing among alternative design concepts may not be easy, but the need to consider and make conscious choices forces designers to think on a grand scale before turning to the details.

How to decide among possible design concepts? Formal analysis of the relative benefits and limitations, costs, dependence on other systems, and possible side effects may help, but intuition based on experience and knowledge of the territory also come into play. Naval architects know a lot about ships, movie directors about cinema, and so forth. But just as the consideration of alternative design concepts is important, it is equally important that design not stall out at the stage of considering alternatives. It is simply too costly to develop many competing designs simultaneously.

Developing a Design

Once progress has been made toward setting goals, identifying constraints, and selecting a design concept, the main task of developing a full-fledged design can proceed. For a monumental example, consider the Brooklyn Bridge. The goal was to find a way of moving large numbers of people back and forth across the East River between Brooklyn and Manhattan. That led to a design concept of a two-tower suspension bridge (rather than any other kind of bridge, or a tunnel or more ferry boats) high enough for oceangoing ships to pass underneath. Then the specifications for every feature of the bridge were formulated—where the towers would be located, what they would be made of, how high they would be, what the wire cable would be made of and how it would be installed and anchored, what the dimensions and slope of the spans would have to be, and so on.

In developing a design for a product or system, new opportunities and ideas may emerge, along with the inevitable unanticipated impediments. Not all of the ideas can be exploited and not all of the impediments can be overcome, so choices have to be made *as the design effort proceeds*. Estimating relative benefits, costs, and risks provides a basis for making trade-offs among the possibilities. For example, construction of the towers of the Brooklyn Bridge involved so many worker injuries and deaths that the chief engineer halted construction at a depth considerably less than that called for in the design; he felt that saving the workers' lives outweighed the perpetual risk of the bridge one day collapsing. So far, the trade-off has been successful; the Brooklyn Bridge is now over a century old. However, there are extreme cases of designs in which none of the trade-offs is desirable, acceptable, or even tolerable, and it may be necessary to reject the selected design concept and take up another one for development.

Some Helpful Strategies

Actually developing a design can be rather easy—or it can be daunting; it depends on the complexity of the challenge. To design our garden may take only a few days or weeks and require little help, whereas designing a space station takes years and involves a cast of thousands. Most design challenges—including curriculum—fall between those two extremes of scale. In even moderately complicated design undertakings, there are some strategies that can help to deal with the complexity. One of them is to copy or modify an existing design; a second is to divide the design task into component parts that are individually more manageable than the whole thing; and a third is to plan on testing the maturing design repeatedly during design development. Following is a brief look at these three strategies.

Copying or modifying an existing design innovation. In turning to an existing design for guidance, the presumption is that it represents a successful design. Because the actual Brooklyn Bridge did work (in the sense of doing what was expected of it and not falling down), many other bridges have used very similar designs. However, design failures can occur when modifications of basically successful designs are carried too far. In his 1994 book *Design Paradigms: Case Histories of Error and Judgment in Engineering*, Henry Petroski claims that the history of bridge building is littered with the stories of designs that have been extrapolated too far from known successes.

Since it is rarely possible to copy an existing design down to the last detail in new circumstances, modifications are usually necessary, although they may entail risk of not working as well. But then there is usually an even greater risk in not making modifica-

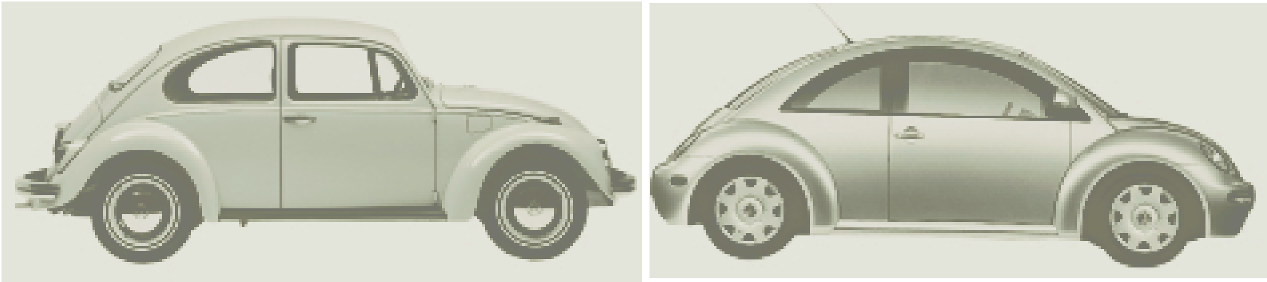

The size and shape of the early Volkswagen is recaptured in a fresh 1998 design.

tions when fitting a given design, no matter how successful, to a different situation. In adapting an existing design, designers must call upon accurate information, relevant experience, and known principles—which may not always be available. And they must be alert to the fact that parts from different designs may not work well together.

When they are known, established principles about what works can be extremely important. To create a design for spanning a 25 percent wider river, it is not simply a matter of planning to make the bridge 25 percent longer. Rather, the new design must be guided by physical principles that relate the strength of beams to their length and cross section, or that relate the wind drag on a cable to how high it is off the ground. Consider taking a curriculum design that has been found to be successful in middle-class suburbs and modifying it for inner-city schools: How much more confidently that could be done if there were known principles for how students in those two types of schools differ in how they can best learn. Curriculum design is often limited because the underlying principles about teaching and learning are not known well.

Copying an existing design could involve identifying a completed, operating object or system (the Paris Opera House, Yale University, the United States Constitution, the Pony Express) and studying it to derive its design—that is, describe it in a way that makes possible fashioning another like it. If, for example, a school district's operating K-12 instructional program were to be copied, it would have to be studied and described to be able to provide a curriculum design for another school district.

Compartmentalizing the design components. Development of a complex design can be greatly simplified if it can be divided into parts, each of which then becomes a separate design challenge. Suppose the concept we adopted for our backyard garden called for an area dedicated to easy-care perennials and an area dedicated to vegetables for the family, and that the borders of each had been set. It would then be easy to concentrate on the design of each component separately, perhaps with each being handled by different family members. When Boeing designs a new aircraft, it typically contracts with other companies to design and manufacture subsystems (for power, navigation, communications, etc.) after writing the specifications (goals and constraints) for each. And it may very well plan to use some "off the shelf" parts in its design, rather than designing new ones. The specs for each subsystem, of course, must include relationships to

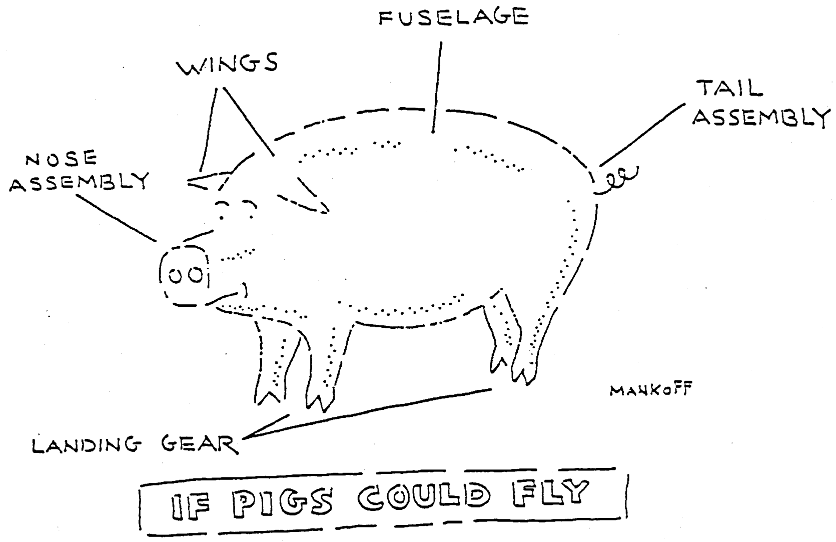

other parts. For example, all electrical components will likely be expected to run on the same voltage. If the specifications for one part have neglected the requirements of other parts it must connect to, it may work poorly or not at all. In the end, however, Boeing is responsible for making sure the whole design works—which is not a foregone conclusion just because all the subsystems individually pass muster.

Although curriculum design is addressed at length later, it provides an excellent example right here for mismatched parts. To the degree that curriculum design for the whole K-12 range is done at all, the task is usually divided into nearly independent parts. Sometimes the division is by subject-matter domain (the reading, mathematics, science, history curricula, etc.), sometimes by grade level (elementary, middle, and high school, or even grade by grade), sometimes by track (vocational, general, college preparatory, advanced placement), and sometimes by combinations of these (the college-prep foreign-language curriculum). The trouble seems to be that, whatever the quality of the design for each component, the parts usually do not get put back together to form a coherent whole that optimizes students' learning over their whole K-12 range of instruction. Good curriculum design should attempt to optimize learning across the entire curriculum, not just unit by unit, subject by subject, or grade by grade.

"Large changes in scale typically are accompanied by changes in the kind of phenomena that occur.... Buildings, animals, and social organizations cannot be made significantly larger or smaller without experiencing fundamental changes in their structure or behavior."
—*Science for All Americans, pp. 179-80*

The Tacoma Narrows Bridge provides a famous example of the terrible consequences that can result from ignoring a known principle. Designers of the unprecedentedly long span correctly extrapolated the increased stiffness required to control up-and-down vibration, but neglected to consider the twisting oscillations that were previously unimportant; when built in 1940, the bridge disintegrated in the first high wind.

"Designs almost always require testing, especially when the design is unusual or complicated, when the final product or process is likely to be expensive or dangerous, or when failure has a very high cost. Performance tests of a design may be conducted by using complete products, but doing so may be prohibitively difficult or expensive. So testing is often done by using small-scale...simulations...or testing of separate components only."
—*Science for All Americans, p. 29*

Testing in the development stage. No matter how good the existing design model is, how thoughtfully the design challenge is subdivided, and what sound principles and storehouses of information are drawn upon, the process of designing a complex system sooner or later (usually sooner) is beset with uncertainties. Designers want to know if they are on the right track before they get irrevocably committed to a design concept, and then they need to find out whether the various components really will perform as required by the design specifications. They can accomplish this in the development stage of design by testing components of the maturing design frequently. Before building a new middle-school curriculum around parent volunteers, for example, educators should see how many volunteers could be turned up in the community. Or before basing a self-paced mathematics curriculum on computer tutorials, educators should test how local students learn from tutorial software.

The best way to test the component of a design is to make a prototype of it and see how well it does what it is supposed to. In some cases, this can be done by using small-scale trials. In designing new aircraft, for instance, scaled-down wing shapes can be

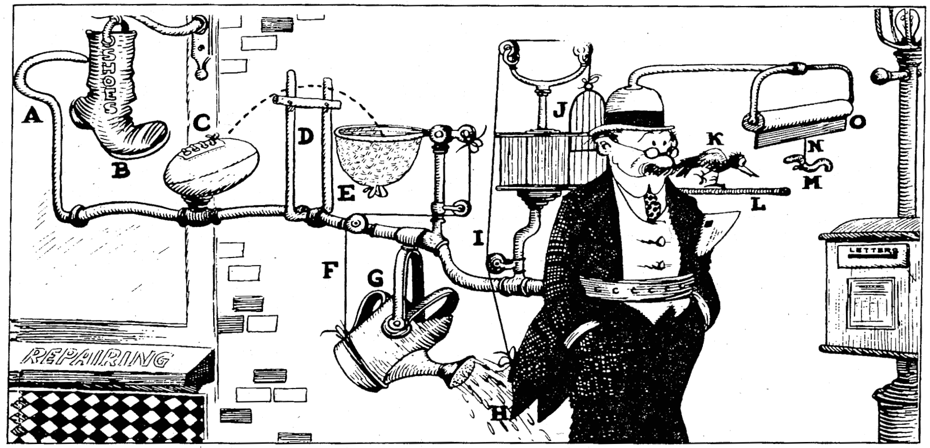

"Keep you from forgetting to mail your wife's letter"—Rube Goldberg™

tested in a wind tunnel; in designing new skyscrapers, computer simulations can be used to answer "what if" questions about the effects of wind shear on a building of various dimensions and orientations. In other cases, it is possible to test a process on a small sample. For instance, aspects of the design for the next round of the national ten-year census of the entire population are tested on a representative sample of some hundreds of households to see how people will respond to the questions. Testing all of the components of a design may not be necessary or even feasible, but the practice is to test those components for which the greatest uncertainty exists or that are most crucial to the success of the design. The question remains open, of course, as to whether the components will all work together when the time comes, and thus the design as a whole will eventually have to be tested. In a home-entertainment system, some top-rated amplifiers may not work well with some of the top-rated loudspeakers. In the school example, the volunteer parents, however numerous, may or may not be effective coaches for the computer tutorials that worked well under experienced teachers.

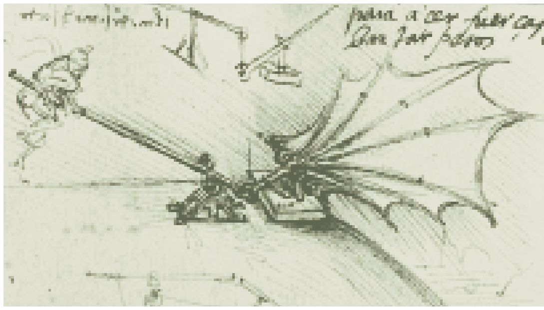

Testing the lifting power of a wing.

Design Decisions

Ideally, one would like to be able to create objects, processes, or systems that would perfectly serve all the identified goals and do so at low cost and without any risk. Alas, needed resources are not always available; insurmountable constraints get in the way, things cost too much, and, like it or not, there is always risk. Because the design process is the pursuit of acceptable solutions, compromises are expected between what is desired and what is feasible. Making design decisions often comes down to agreeing on trade-offs among desired features on the basis of estimates of their relative benefits, costs, risks, and the associated trade-offs.

Benefits and costs. It is natural for the proponents of a new design for a product or process to emphasize its possible benefits: the disease it will cure, the faster it will get people from one place to another, the greater grain yields that will result, the more that students will learn in a year. It may not be known until sometime after the product exists or the process is in effect whether those benefits actually will accrue—although for some few things, such as medical drugs and procedures, stringent testing is required before they can be put on the market.

But potential users are likely to respond to projected costs as well as to claimed benefits. Hence, cost-effectiveness is not far from the minds of designers in any field. The following questions suggest that there are often social costs, as well as immediate and long-term financial costs, to consider:

These questions are based on similar ones posed in the section on technological decision making in Chapter 3 of *Science for All Americans.*

- Who are the main beneficiaries of the proposed design? Who will receive few or no benefits? Will people other than the beneficiaries have to bear the costs? Who will suffer if the design is implemented? Who will suffer if it is not? How long will the benefits last? Will the design have other applications?
- What will the proposed design cost to build and operate? How does that compare to the cost of alternatives?

"The best laid schemes o' mice an' men
Gang aft agley,
An' lea'e us nought but grief an' pain
For promised joy."
—Robert Burns, *To a Mouse*

• What people, materials, tools, and know-how will be needed to build, install, and operate the proposed new system? Are they available? If not, how will they be obtained, from where, and at what cost? What energy sources will be needed for construction or manufacture, and also for operation, and at what cost? What will it cost to maintain, update, and repair the design if it is implemented?

Risk. Seldom are all the effects of a design reliably predictable. This means, of course, that we can never count on getting all the benefits that we hoped for from a design, although it happens sometimes that totally unexpected benefits arise to surpass those intended by the design. But also weighing heavily on the mind of the designer is the specter of unwanted outcomes. There is always risk—the issue is never risk versus no risk—and so statisticians and engineers have worked hard to develop reliable ways of estimating risk. A formal analysis of risk involves estimating a probability of occurrence for every undesirable outcome that can be foreseen and also estimating a measure of the harm that would be done if that outcome did occur. The risk of each undesirable outcome is the product of its probability and its measure of harm. The sum of these risk estimates (perhaps with adjustment for correlation among them) then constitutes the total risk of the design. But such elaborate theoretical risk estimates are always difficult and sometimes impossible to make. Instead, it is usually necessary to settle for rough estimates of how the risks associated with a proposed design compare to those of other possible designs, including the design it is intended to replace (which, however well established, is likely to carry risks of its own). Whatever the outcome, it is still the case that the risk associated with a particular design can never be reduced to zero, and so designers take steps to minimize risk at reasonable cost.

One hedge against failure is what is called overdesign—for example, making something stronger or bigger than is likely to be necessary. Another hedge is redundancy—building in one or more backup systems to take over in case the primary one fails. On NASA missions, the onboard backup computer also has a backup. In an education context, some provision for remediation should always be available if the first-line instruction does not succeed.

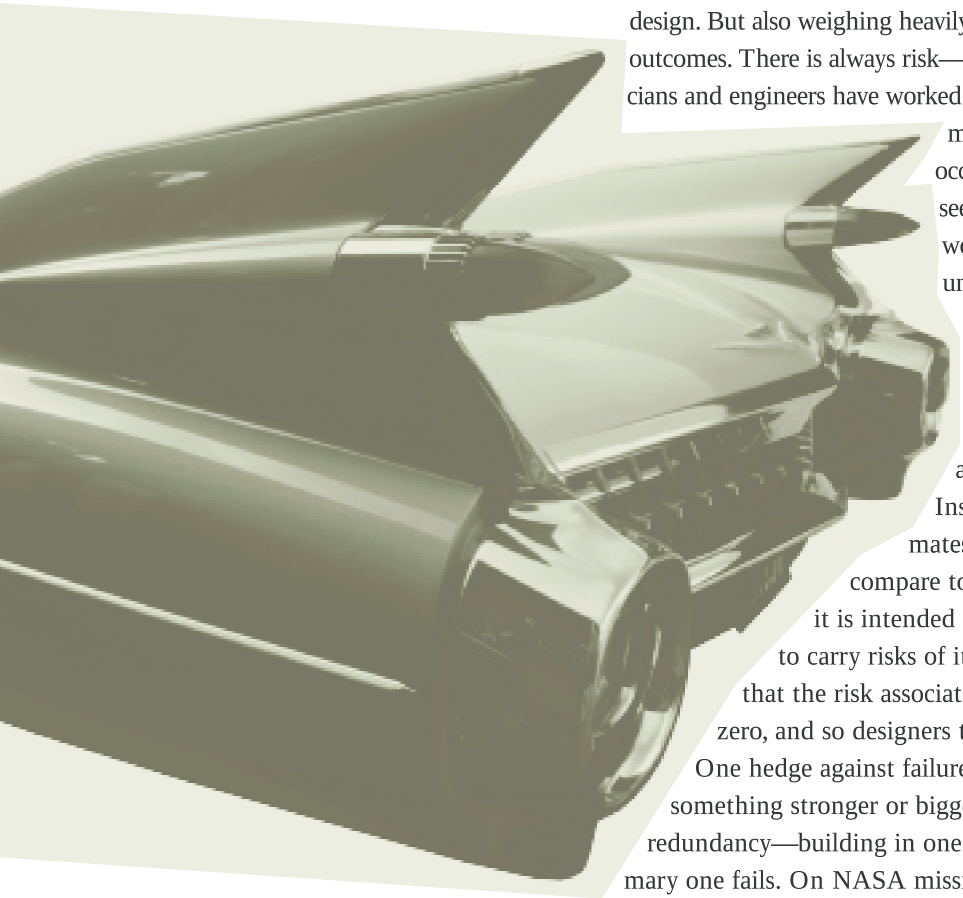

This vintage, finned automobile is an example of carrying one design characteristic too far.

If failure of a system would have very costly consequences, the system may be designed so that its most likely way of failing would do the least harm. This requires a value judgment of the kind, “If there is a possibility of some failures, let’s try to err on the side of bad outcome X rather than the even worse outcome Y.” One example of “fail-safe” design is the on/off switch of an electric lawn mower—if the switch breaks, better it should get stuck in the off position than in the on position. Another example is the U.S. legal policy under which uncertainty about guilt in criminal cases leads to acquittal rather than to conviction (on the value judgment that it is better to free the guilty than to punish the innocent). Of course, not everyone may agree just what the risks and the costlier consequences are. For example, debate persists in education policy about whether the risks of promoting students with dubious achievement should be preferred to the risks of holding them back. Currently it is most common to see greater risk in holding marginal students back than in promoting them, and so there is a preference to err on the side of being too optimistic rather than too pessimistic (a literally “fail-safe” policy). Also, as suggested earlier, the likelihood of failure in either direction is reduced by doing more testing to develop a more robust design.

Trade-offs. The term “trade-off” has come into vogue only in recent years, but the idea behind it is venerable. It points to the age-old practice in human affairs of making compromises—of sacrificing in one way to gain in another. A trade-off is based on the common-sense notion that only very rarely can one have everything one wants—espe-

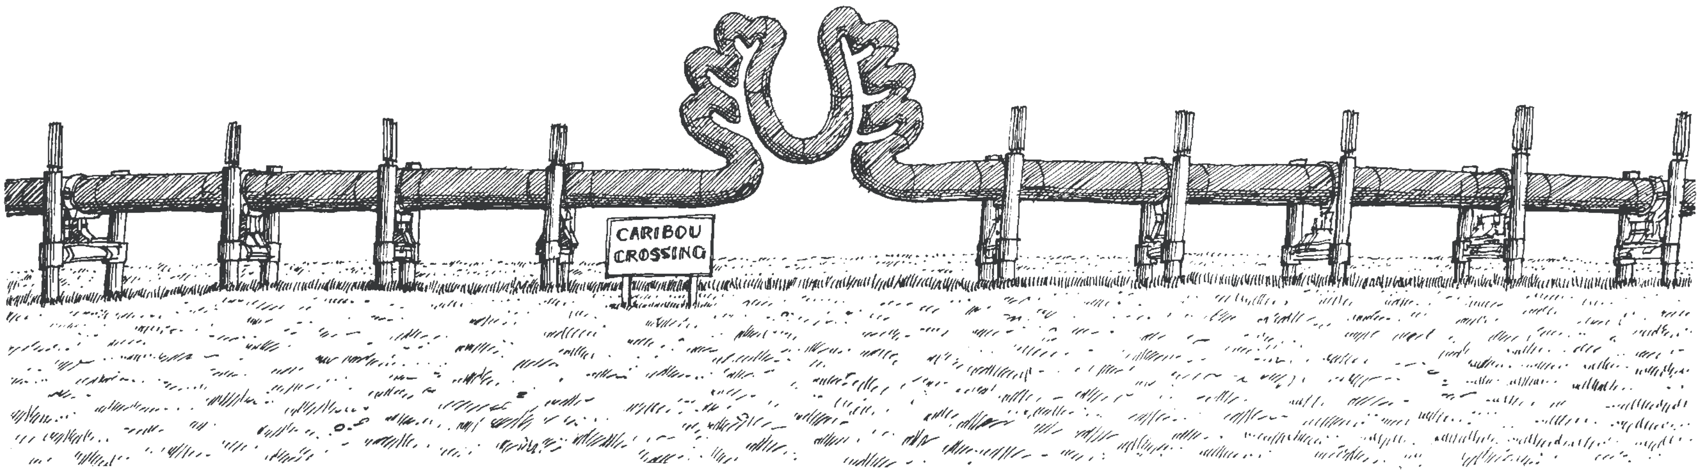

“An Environmental Consideration”

cially if others are also to get what they want. Trade-offs are solutions, and they result from such things as people needing to share a fixed total amount of some available resource, such as money or time or raw materials, or from their being in conflict over incompatible values, as in exercising both generosity and frugality, or in their wanting to treat all people equally and yet respond to individual needs. Trade-offs also occur because individuals and groups do not all see things the same way. When they are at odds with regard to something in particular, each party may have to give ground to come up with something that is agreeable (or at least tolerable) to all parties.

The practice of making trade-offs is central in design. Trade-offs can be used to settle differences in goals (for example, give up the shade trees, or some of the shade trees, in the backyard garden design, to ensure having more sunlight for the flower beds and vegetable plot); negotiate constraints (a downtown office building design may be permitted to be higher than the code stipulates in return for using less of the ground area than is the rule); or to balance goals against constraints (the mayor will settle for less improvement in crosstown traffic than he would like so that north-south traffic will not be slowed down).

The time given to the design process itself is also subject to trade-offs. At some point, it may seem better to go ahead with an imperfect product rather than to wait until the design is improved further.

Making trade-offs in these ways can be simple or complicated depending on what the design undertaking is. In many situations, it may not be possible to base design trade-offs on rigorous benefit-cost-risk analyses because sufficient information is not available, applicable principles are not known, or the complexity is simply too great. In other situations, it may not make sense to back every trade-off with such analysis, if for no other reason than that the cost in time and money would become prohibitive. Good judgment needs to prevail in this as in other aspects of design. However, there is little doubt that good design is fostered by (1) making trade-offs deliberately and (2) doing so in the light of what is known about benefits and costs and risks. Benefit-cost-risk analysis requires some way of assigning relative magnitudes to benefits, of quantifying costs, and of estimating risks. CHAPTER 1: CURRICULUM DESIGN includes a discussion of some of the difficulties of doing such analysis in an education context.

REFINING THE DESIGNED PRODUCT

A design is not the product itself. Indeed, the thing itself—the actual garden, canal, jet aircraft, chair, dress, building, curriculum—often does not perfectly match a design. The difference may result from some flaw in the logic of the design, poor information on which the design is based, or the unexpected influence of factors not covered by the design. If discrepancies become too great, the production may be aborted and, in a now-famous phrase, sent "back to the drawing board." On top of that, designs can usually be interpreted in different ways, and hence what the client believes the design intends may not perfectly match what the designer intends.

When components of a product are tested separately before they are implemented together, feedback and refining have already begun. But no matter how closely the product matches the design, and no matter how well the designed components performed when tested during development, no designed product turns out to work exactly as intended when it is all put together and used in the real world. Even products that are outstanding at first often become less so as new demands are placed on them—the air traffic control system that was designed to track a certain number of planes but that now must track many times that number is a current example. Other products continue to perform well enough, but new technology and policies may make them obsolescent. Anyone who has bought a computer recently can attest to that phenomenon. For all these reasons and others, design is never complete.

"Well, back to the old drawing board."

Another reason that designs need to be refined even after they become actual things is that they are used or affected by human beings. Assumptions are made, often implicitly, about the interaction of the product and people—that all pilots, not just test pilots, will be able to use the controls properly; that drivers will obey the signal lights; that office workers will receive training in the use of the new computer system; that we will tend our new garden faithfully; and so on. In many cases, what later gets labeled as "human error" may be the result of making naive or unwarranted assumptions about how people will actually use the designed product.

Given all these ways in which a finished product can fall short of expectations, good design practice makes provisions for systematic retesting. Performance feedback on a design may come from instruments (say, for monitoring stresses on a bridge), from direct observation of operations, or from testimony from users. In the case of air traffic control, inspectors can measure traffic delays, look over the shoulders of operators and pilots, or interview them. In the case of curriculum, supervisors can observe classrooms or interview teachers or students.

Feedback can also come from assessment of whether the initial goals for the design are being met and the constraints adhered to. In air traffic control, do annual records show that traffic flow is sustained with a tolerable number of accidents? In curriculum, do student achievement records indicate at least short-term learning of concepts and skills—and better, do subsequent studies of graduates show long-term retention? Periodic assessments can identify unanticipated shortcomings and sometimes lead to suggestions for modifications. Because of the incremental changes made in its design based on feedback from users, the Boeing 777 that rolls off the assembly line today is not identical in design to the first one that emerged. The design, under which additional products will be produced, is progressively modified on the basis of real-world experience with the early products. On the other hand, when there is a single, unique product—say, the United Nations Headquarters—fix-ups have to be made in the actual product. The initial *design* followed in producing it is not likely to be modified—on paper—unless similar products are planned. One could see the changes in a product as implying a modified "design" in a descriptive sense. Sometimes, but not often, feedback from actual use is so negative that it becomes necessary to rethink the entire design. But in most cases, designs just replace a piece at a time with new and/or better ones.

Since our garden will not turn out to look exactly like our design for it, why bother with design at all? We could, instead, just start planting things at random in the backyard and eventually a garden would exist—perhaps even a beautiful one. But the odds are very much against it, especially if we expect the different parts of the garden to relate harmoniously to one another. Even though there is some difference between our actual product and the design on which it is based, we are more likely to be satisfied with the product than if we simply forged ahead without careful planning. Obviously, that is even more true for more complicated systems having criteria for success that go beyond pleasing us subjectively.

Looking Ahead

Design does not automatically and inevitably result in successful products. Nor has everything admirable that exists necessarily been designed. But among those things that can be and ought to be designed as whole systems are curricula. In this Prologue, we have given little attention to how design relates to curriculum—in order to emphasize what is common in almost all design in the world outside education. The first chapter in what follows considers how the principles in this Prologue can be applied to the particular case of curriculum, and the chapter after that elaborates on what properties of curriculum are appropriate to design.

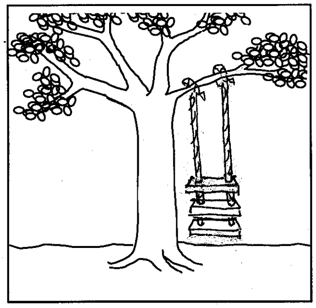
What the state framework specified

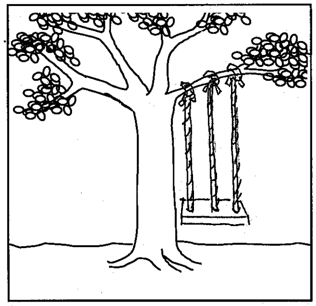
What the curriculum committee designed

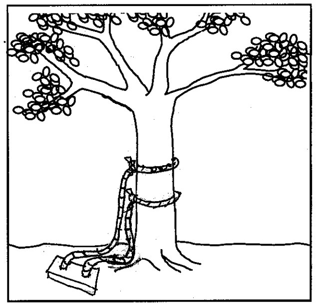
What the superintendent requested

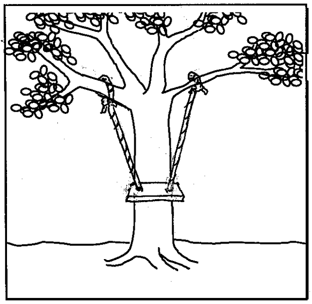
What the board of education approved

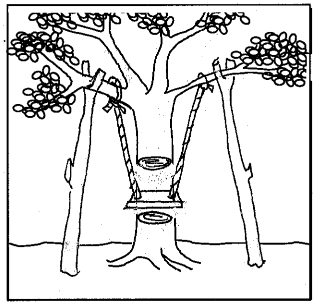
What the publisher actually produced

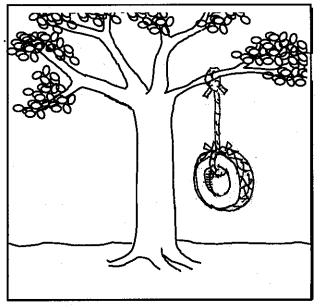
What the literacy goals had intended