DATE DOWNLOADED: Mon May 25 18:47:22 2020
SOURCE: Content Downloaded from *HeinOnline*

Citations:

Bluebook 20th ed.
Richard Zeckhauser, Measuring Risks and Benefits of Food Safety Decisions, 38 Vand. L. Rev. 539 (1985).

ALWD 6th ed.
Richard Zeckhauser, Measuring Risks and Benefits of Food Safety Decisions, 38 Vand. L. Rev. 539 (1985).

APA 6th ed.
Zeckhauser, R. (1985). Measuring risks and benefits of food safety decisions. Vanderbilt Law Review, 38(3), 539-582.

Chicago 7th ed.
Richard Zeckhauser, "Measuring Risks and Benefits of Food Safety Decisions," Vanderbilt Law Review 38, no. 3 (April 1985): 539-582

McGill Guide 9th ed.
Richard Zeckhauser, "Measuring Risks and Benefits of Food Safety Decisions" (1985) 38:3 Vand L Rev 539.

MLA 8th ed.
Zeckhauser, Richard. "Measuring Risks and Benefits of Food Safety Decisions." Vanderbilt Law Review, vol. 38, no. 3, April 1985, p. 539-582. HeinOnline.

OSCOLA 4th ed.
Richard Zeckhauser, 'Measuring Risks and Benefits of Food Safety Decisions' (1985) 38 Vand L Rev 539

Provided by:
Harvard Law School Library

-- Your use of this HeinOnline PDF indicates your acceptance of HeinOnline's Terms and Conditions of the license agreement available at *https://heinonline.org/HOL/License*
-- The search text of this PDF is generated from uncorrected OCR text.
-- To obtain permission to use this article beyond the scope of your license, please use: *Copyright Information*

Measuring Risks and Benefits of Food Safety Decisions

*Richard Zeckhauser**

* Professor of Political Economy, John F. Kennedy School of Government, Harvard University.

Initial research was supported by the Food Safety Council. Much of the present document was prepared as a basis for the Social and Economic Committee's discussions, and it appeared as Appendix A of the Committee's report. It is presented here with permission. SOCIAL AND ECONOMIC COMMITTEE, FOOD SAFETY COUNCIL, PRINCIPLES AND PROCESSES FOR MAKING FOOD SAFETY DECISIONS, (1979). Current research is sponsored by the Business and Government Research Center, Harvard University.

Civilized societies cherish the belief that life is priceless, not to be bartered for mere economic returns. If life has infinite value, should we not prohibit even the voluntary acceptance of risks to life?

It now seems clear that this ethic cannot survive in modern society. Risks are inherent in many of the technologies that have so dramatically improved the quality of twentieth-century life and, in the process, extended longevity. To compound the problem, we have developed the ability to detect risks at exceedingly low levels. Many products of modern society once thought safe now would flunk the no-risk test.

With this dissonance revealed, society is increasingly moving to a standard of *acceptable* risk. This test is easily met if the risks created by a product are outweighed by those reduced by it—that is, if the product's net effect is to diminish risks to health. Often, however, it may be difficult to trace such net effects. Moreover, many of our present regulatory standards violate this net risk

rule.[1]

Even if the net risk standard could be promulgated and used effectively, it would not be appropriate. A product or technology may offer other, nonhealth benefits, such as cheaper power or swifter transport, that must be traded off against health risks. Such nonhealth benefits may have positive second-order consequences for health, because the resource savings associated with these benefits can be redeployed to promote health and longevity.[2]

Attempts to define acceptable risk encounter intractable scientific problems. Even if they could be solved—if we could predict the consequences of using any product or technology—we still would have to ask how risk choices should be made. This question is ultimately a philosophical issue, and one to which Samuel Stumpf has contributed not only in his writings, but in testimony, in committee deliberations, and in his work with industry, consumer groups, and the government.

In the late 1970s, I had the privilege to serve with Sam Stumpf on the Social and Economic Committee of the Food Safety Council. The Committee's assignment was to help define the principles that should underlie food safety regulation and to design an appropriate, effective regulatory process.[3]

I. INTRODUCTION

Whereas the risks of nuclear power and lead emissions from gasoline are imposed on the individual by others, food safety risks are associated with voluntary behavior. The principal argument for regulating any voluntarily assumed risk—whether it involves food

1. For example, the Delaney Amendment rules out any food additive shown to be carcinogenic in animals or human beings. This standard would prohibit even a highly beneficial food preservative if it imposed an infinitesimal carcinogenic risk. Food Additives (Delaney) Amendment of 1958, 21 U.S.C. § 348(c)(3)(A) (1982).

2. This view reflects the economic perspective that stresses the transmutability of resources within the economy as a whole. This is not to suggest that a power plant can be turned into a safer automobile, but that dollars measure the value of resources saved, and the saved dollars can be used to buy alternative resources such as highway barriers or car padding, or vaccinations against disease. Assuming that the processes of social choice are roughly rational, if society chooses not to redeploy the saved resources to lifesaving, it is better off than it would be otherwise. For example, assume that policy B could be transformed into policy B′, which offers less risk and more benefits than the status quo policy, A. If society chooses B in preference to B′, then B is preferred to the status quo, A.

3. The Food Safety Council was composed of industry, consumer, public, and academic representatives. For the Food Safety Council's final report, see FOOD SAFETY COUNCIL, A PROPOSED FOOD SAFETY EVALUATION PROCESS: FINAL REPORT OF BOARD OF TRUSTEES (1982) [hereinafter cited as 1982 FOOD SAFETY COUNCIL REPORT].

safety or the use of automobile seat belts—is that it is difficult for individual decision makers to obtain, process, and use information on the consequences of their choices.

The objective of food safety regulation should be to promote the consumption pattern that well-informed consumers would choose for themselves.[4] That premise, widely though not uniformly accepted, underlies this analysis.[5]

The process of making regulatory decisions comprises three critical elements: predicting consequences, organizing information, and choosing among alternative sets of attributes. Prediction concerns such questions as: How much of a product will be consumed; how much a particular use for a product will reduce costs; and what level of risk is entailed by various levels of use. If all predictive information was provided in undigested form, the regulatory process would be overwhelmed. The information must be organized so that rational choices can be made. This Article discusses the assessment of risks and benefits as one approach to organizing information.

The way information is organized should depend on the way it will be valued and used. For example, the decision making authorities within the regulatory process may choose to take different approaches to food substances consumed by young and old, or rich and poor. In that case, information should be organized into those categories. An exquisite breakdown of consumption patterns by counties would do little for an age-regarding regulatory process.

The remainder of this Article is divided into four parts. Part II reviews the general nature of risk and benefit assessment within the context of its somewhat better-understood relatives, cost-benefit analysis and cost-effectiveness analysis. Part III discusses a classification mechanism for the benefits of additives and contaminants in food and outlines the economic approach for converting them to a common measure. Part IV discusses the measurement of risk. Part V outlines a number of common errors that are made when risk and benefit assessments are undertaken.

---

4. The countervailing danger is that regulation may overrule some individuals' preferences, preventing them from consuming products that they would select if fully informed.

5. Though this discussion is directed to food safety, the reader should he able to translate it to such fields as product safety, building codes, air bags, or even areas in which health risk is not concerned, including the entire field of consumer protection.

II. Cost-Benefit, Risk-Benefit, and Cost-Effectiveness

The concepts of cost-benefit, risk-benefit, and cost-effectiveness analysis are valuable decision aids when used appropriately. Because these terms often are used incorrectly, however, it is important to define each of them clearly. These methods all start from one basic premise: When making decisions, particularly complex and difficult decisions, it is desirable to consider all consequences of possible action. Thus, the first step in each type of analysis is to identify the consequences that will ensue. If the area dealt with is one of uncertain outcomes, then an attempt should be made to estimate the probability distribution of these different outcomes.

A. Cost-Benefit

Cost-benefit analysis is the principal analytic framework used in evaluating public expenditure decisions and is also useful in guiding regulatory decisions. The objective of cost-benefit analysis is to assure that resources are put to their most valuable use. Cost-benefit analysis requires careful elaboration of all the benefits and of all the costs that will accrue to all segments of society if a particular project is adopted. The basic principle is to tally the gains from a project and then subtract the losses.[6] If the difference is positive, the project is sound from the standpoint of efficiency, and it is recommended.

In cost-benefit analysis all consequences are converted to a common metric. Most frequently the unit of measurement is dollars, but other measures such as lives saved or units of food quality could be used instead. When applied to the food safety area, cost-benefit analysis runs into the excruciatingly delicate issue of valuing lives. For example, what dollar price tag should be attached to a life at a one percent elevated risk of cancer?[7]

---

6. Losses include the direct costs of undertaking the project.

7. It is important to value the consequences on a probabilistic basis, for that is how they are received by those confronting the risk. One who ingests a low-level carcinogen does not necessarily contract cancer, but incurs a low-probability risk of getting cancer. The valuation depends neither on the consequence alone nor on such frequently proposed aggregate measures of the consequence, such as X number of cases in the nation per year. Rather, one encounters a risk with both a severity and a probability.

This formulation strikes at the traditional distinction between statistical and identified lives. For a variety of predominantly ethical reasons, our society abhors failure to sacrifice huge amounts of resources to save identified lives. However, we are willing to accept substantially larger risks in terms of expected (statistical) lives to save the same level of resources.

Consider the application of cost-benefit analysis to the simplest possible problem in the food safety area. Should we permit additive A on the market or should we ban it? The first step in answering this question is to detail all the costs and benefits that accompany use of the additive. If the costs outweigh the benefits, it should be banned, and vice versa.

A number of difficulties arise when we attempt to employ this approach in food safety. First, rarely will we have the scientific ability to make a full determination of the health costs of permitting the use of an additive or other food substance. Second, even if we could make exact assessments of the health risks, we still would need a way of valuing those risks. With traditional cost-benefit analysis, valuation is less problematic. In deciding whether to construct a dam, for example, the question might be: How much are 100,000 gallons of irrigation water worth in terms of the capital expenditures required to divert them? This question could be answered by looking at market prices, or at economic and engineering analyses to discover the value in use of the water, or by finding out what people would be willing to pay for the water. Risks to life present quite a different situation. They are never traded directly on traditional markets, though they are sometimes one component of other goods or commodities.[8] No engineering or economic study could tell us a life's value in use.[9] Finally, asking individuals about their valuation of life yields little guidance. Few individuals have experienced purchasing the good, "reduced risk to life." An even smaller number can distinguish what the probabilities associated with such risks might mean.

Although the cost-benefit approach is clearly difficult to implement, its spirit must be incorporated in any process of rational

---

Interest in the question of appropriate valuation of lives has increased substantially in recent years. *See, e.g.*, THE VALUE OF LIFE AND SAFETY (M. Jones-Lee ed. 1982); W. Viscusi, *The Valuation of Risks to Life and Health: Guidelines for Policy Analysis*, in EVALUATION OF THE STATE-OF-THE-ART IN BENEFITS ASSESSMENT METHODS FOR PUBLIC POLICY PURPOSES (J. Bentkover, V. Covello, L. Lave, J. Menkes & J. Mumpower eds. 1985); *Special Issue on Risk Analysis*, 30 MGMT. SCI. 395, 395-528 (R. Sarin ed. 1984).

8. In deducing individuals' valuations for lives, a number of researchers have looked to wage premiums paid for more hazardous employment. This is a highly imperfect method, because it is hard to hold constant other aspects of the jobs. Moreover, not everyone is bidding for the risky job. Some would demand a premium 10 times as great as the one generated by the market.

9. Sometimes efforts are made to look at discounted earnings, or earnings less consumption, but these measures hardly indicate the value an individual receives from increasing his probability of survival. *See* Zeckhauser, *Procedures for Valuing Lives*, 23 PUB. POL'Y 419, 419-64 (1975).

decision making. The central tenet of cost-benefit analysis is that the consumer should pay for something no more than what it is worth. For example, if an additional unit of electricity is worth ten dollars, it should be purchased if offered for nine dollars. If the electricity is worth a new shirt and someone offers to trade a shirt for the electricity, that trade should be made.[10]

Consider now the same problem with risks to life.[11] How do individuals decide what a lowered risk to life is worth? Individuals and society continuously accept risks to life to save money or to secure other benefits. As we drive faster on the highway, we save travel time at the expense of greater accident risk. We eat rich desserts though they may increase somewhat our chances of having heart attacks. We build buildings to resist earthquakes that are expected to occur once in 100 years but not those that are expected to occur once in 10,000 years. In each case, we decide that the extra cost of further protection is not worth the sacrifice required. On what basis should these decisions be made? The risk-benefit approach provides guidance for dealing with this class of issues.

B. Risk-Benefit

Risk-benefit analysis starts from the same presumption as cost-benefit analysis: all consequences of an action should be identified and evaluated. However, risk-benefit analysis stops short of fully aggregating the consequences. If the action under consideration makes a resource expenditure to reduce a health risk, the traditional categories for aggregation are health risks and economic benefits. These categories compete against one another in such decisions as those concerning highway safety or occupational safety and health.

Food safety decisions are frequently more complex than many familiar risk-benefit situations. For example, the decision to permit nitrites in bacon may both promote health and entail some risks. Samuel Stumpf contributed to this debate in his article *Social As-*

10. Let us first see the implication of this principle for efficient production of benefits. Suppose there are two different ways to produce a unit of electricity. One costs $12, the other $10. Currently 100 units are being produced with each method. The total cost is $2200. Obviously, money could be saved if more were produced with the second technology and less with the first. Indeed, if we produced 210 units with the second technology, the total cost would be reduced and the total amount of electricity would be increased. This would be an unquestioned gain.

11. We will avoid dealing with lives themselves, because fortunately we rarely have to barter them directly.

pects of Risk/Benefit Analysis of the Food Supply.[12] He concludes by arguing for "the legitimacy and even necessity under certain circumstances, where fundamental human values are concerned, to engage in a reasoned risk/benefit calculation." As examples, Professor Stumpf cites the banning of DDT in Ceylon, which led to deaths due to malaria, and the removal of saccharin from the food supply, which imposes the risks associated with sugar.[13] The present analysis goes further, arguing that nonhealth benefits may trade off against health risks.[14]

Costs may appear on both sides of the equation as well. If removal of an additive eliminates some convenience food, the price may well be lowered, but the time cost for the food may increase. Various possibilities suggest themselves for aggregating. All health benefits and costs[15] could be lumped together in one category, with all other benefits and costs in another. This would provide the traditional risk-benefit formulation. Alternatively, we could have four separate categories: health costs, health benefits, nonhealth costs, and nonhealth benefits. To simplify the following discussion, we shall deal with the two-category formulation in which all health consequences are lumped together as are all nonhealth consequences.

The first objective of risk-benefit analysis is to provide accurate and usable estimates of all the consequences of a regulatory decision. Thus, risk-benefit analysis should help to promote the generation of information. Suppose the decision under consideration is a simple one: should food additive X be allowed on the market or should it be banned from use? Applying the risk-benefit analysis, the first step is to ask what gains consumers and producers would reap if the additive were permitted. The risks they would incur then would be detailed. The result of this calculation would be the computation of a single point in a graph plotting risks against benefits.

If a risk-benefit analysis is pursued, the final decision on how to value the risk or trade it off against benefits is not explicitly resolved. This need not lead to inaction. First, it may be possible to eliminate certain alternatives if the analysis shows that they are inferior to another option in terms of both risk and benefit. Sec-

---

12. 32 FOOD TECH., Aug. 1978, at 68-69.

13. *Id.*

14. Perhaps this difference reflects the perspective of an economist as opposed to a philosopher/law professor.

15. We have employed the term costs here, rather than risks, to achieve parallelism.

ond, by starting with the alternatives that offer the highest net benefit relative to risk incurred and working down among the alternatives, we can be sure that the set of risks accepted offers the highest level of net benefits. This approach also assures that we are not accepting a higher risk in one area in return for a particular level of net benefits than we are in another area. This analysis is made clearer by the use of a diagram.

Figure 1—Geometric Representation of Risks and Benefits

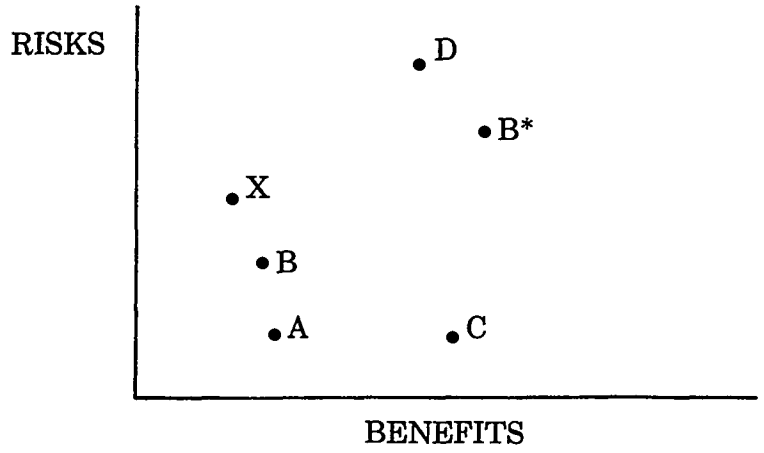

In some decision contexts, matters will be simple. If the choice, for example, were between points B and A in Figure 1, we would have no difficulty selecting A. It offers the same level of benefits at a lower level of risk. Similarly, point C would be preferred over A. We would have difficulty, however, choosing between points D and B. We would have to determine what level of risk to accept to secure the additional benefits that D offers over B.[16] The decision between point B and point D is of the same qualitative nature as the choice between point X and the origin, in effect, the choice between permitting and banning additive X. One alternative offers higher risk (a bad attribute) and higher benefits (a good attribute) than the other. A comparison must entail a weighing of

16. It is important to recognize that point B is not to be preferred simply because it offers a higher ratio of benefits to risks, a commonly heard but fallacious argument. Choices based on ratios are only appropriate if projects can be extended in scale, if, for example, B could be doubled to B*, a point that dominates D.

the two. Consumers typically engage in a weighing of this sort when they have to choose among alternative goods. For example, one automobile may offer greater fuel economy but less comfort than its competitor. The basic decision will depend on whether the sacrifice in comfort is worth the gain in fuel economy.[17]

Microeconomics offers some further principles for rational choice. The central principle relates efficiency to rates of tradeoff. Efficiency requires that the marginal tradeoff between risk and net benefits be the same across all decisions. This is an important finding for food safety. Even if a prescribed use for a food additive is determined to be worth the risk entailed, its use may be modified so that the benefits sacrificed for a reduction in risk (or the benefits gained by accepting a higher risk) more closely accord with the preferences of society or the prescriptions of the regulatory process. Consider a situation with two food additives, A and B. The proposed uses convey the risks and benefits depicted in Figure 2.

Figure 2

	Proposed Use		Alternative Use	
	Risks to Health	Net Benefits Other than Health	Risks to Health	Net Benefits Other than Health
Additive A	.0001	$1,000	.00005	$ 950
Additive B	.0002	$ 800	.00025	$1,000

For the proposed use, Additive A is obviously more attractive than B. It offers higher net benefits and less risk. This might incorrectly suggest that we should increase the use of A and reduce the use of B. This conclusion, however, is inaccurate, as Figure 2 illustrates. The critical question is the ability to trade off at the margin between risk and net benefits. Restricting the use of A entails rela-

17. A detour into symbols may aid exposition. Label the two attributes as x and y and the alternatives as 1 and 2. Alternative 1 possesses the characteristics x_1, y_1. The choice criterion is whether x_1, y_1 is preferred to x_2, y_2. One way of determining which is preferred is to attach a score to each possible combination of x and y that indicates the combination's degree of desirability. If the scoring system accurately reflects consumer preferences, the alternative with the higher score should be chosen. This is the most basic principle of the microeconomics of rational individual choice.

Traditional microeconomic choice problems frequently present the commodities themselves as objects of choice. That may be appropriate when we understand what the items are, such as apples or pears. However, when they are complicated products such as cars or food additives, it helps to break them down into components such as miles per gallon or level of risk.

For a good introduction to microeconomics, see W. NICHOLSON, MICROECONOMIC THEORY (1978).

tively little reduction in benefits ($50) for a .00005 reduction in risk, (.00005/50), while Additive B can offer $200 in additional benefits for a .00005 increase in risk (.00005/200). Because the first quotient is substantially greater than the second, we should expand the use of B and reduce that of A. At an efficient point the two quotients, the rates of tradeoff between risks and benefits, would be the same.[18]

C. *Cost-Effectiveness*

Cost-effectiveness analysis is the general form of which risk-benefit analysis is a special case. Cost-effectiveness analysis aggregates consequences into two categories, but makes no attempt to trade off or compare the two. That process is left to the ultimate decision maker.[19] The underlying principle is that, although it is possible for an analyst to aggregate within categories, it should not be the analyst or scientific expert who feeds in the delicate value judgments that determine just what balance of cost and effectiveness should be achieved.

Applying the cost-effectiveness approach to the food safety area, the two logical categories are the benefits of having the substance in the particular products in the food supply (foregone benefits being cost) and the risk entailed (reduction in risk being effectiveness).

There is no generally accepted procedure for aggregating risks and benefits. Thus, even when it becomes essential to undertake such aggregation, estimates of risk and benefit should be presented separately in the analysis. This enables each analyst, critic, and interested party to make judgments and determinations of his own. The political, regulatory, and judicial institutions can respond to this increase in usable information.

III. Classifying and Evaluating Benefits

A. *Classifying Benefits*

Food substances provide a variety of benefits. Together, they sustain life. Individually, they may enhance health and make life

---

18. It is possible that the optimum would not involve use of one of the additives—what is often referred to as a corner solution. Even if margins are equated, this will only assure efficiency if the ratio for totals is appropriate. There need he no concern over totals if each alternative offers ever increasing ratios of risk/benefits as its use is expanded.

19. For example, military problems frequently are studied with the aid of cost-effectiveness analysis, the two categories being military capability and dollar costs.

more pleasurable in a variety of ways. Most food substances have received general approval for unlimited use in the food supply. Under the present system of food safety regulation, additives and contaminants supposedly are judged for their safety, though the benefits they provide are considered implicitly within the regulatory process.

In 1956 a National Academy of Sciences research council report[20] listed over 550 substances intentionally added to food, approximately 300 of which were flavoring agents. The number now surely is far larger. These additives fall into a number of categories,[21] and they are found throughout the food supply—in cereal products, processed fruits and vegetables, beverages, dairy products, candy, meat products, and fats. In addition to these deliberately added substances, the food supply also contains contaminating substances that occur—either naturally or inadvertently—in many products.[22] These additives and contaminants are placed or allowed to remain in the food supply for a reason. They are designed either to make food products more appealing to consumers, to lower the costs of producing the products, or to provide nutritive or health benefits. We have classified such benefits under four headings: appeal, convenience, resource saving and availability, and health. We shall discuss them in this order, which is the reverse of their likely order of overall importance in influencing food safety decisions. Health is likely to prove most important, convenience and appeal least important.[23]

1. Appeal

Appeal relates to the way the consumer reacts to a food product. It includes flavor, appearance, texture, aroma, and taste. Fla-

---

20. FOOD PROTECTION COMMISSION OF THE FOOD AND NUTRITION BOARD, NATIONAL ACADEMY OF SCIENCES, THE USE OF CHEMICAL ADDITIVES IN FOOD PROCESSING (1956).

21. They include preservatives, antioxidants, sequestrants, surfactants, emulsifiers, stabilizers, thickeners, bleaching and maturing agents, buffers, acids, alkalies, food colors, humectants, anticaking agents, foaming agents, foam inhibitors, solvents, flavoring agents, nonnutritive sweeteners, and nutrient supplements.

22. The presence of aflatoxin in peanuts provides a salient example.

23. Economists are not willing to say one attribute, in general, is more important than another. To make such a statement requires knowledge of how much of the various attributes already is being received. The magnitudes of observed differences are also relevant. For example, someone who says that fuel economy is more important to him than comfort in selecting a car is likely to be proved wrong if offered a choice between two cars that achieve 17.5 and 17.6 miles per gallon, with the 17.5 car offering twice as much leg room. What he probably meant to say is that, with the distribution of cars one generally finds on the market, gas mileage usually turns out to be the determining characteristic.

voring agents are used to preserve or enhance the flavors of processed foods. Flavoring agents may be either naturally occurring spices and oils, such as cloves, ginger, pepper, and citrus oils, or synthetic. Often, naturally occurring flavors do not survive processing procedures as well as their synthetic counterparts. Moreover, natural flavors are in short supply and often are available only during certain times of the year. One can argue that artificial flavors increase food consumption because food lacking good flavor might otherwise remain uneaten.

The extent to which manufactured flavors permeate our food supply is staggering. Aromatic chemicals are used to impart such flavors as pineapple, cherry, walnut, and wintergreen. The consumption of synthesized strawberry flavor, converted to its equivalent in fruit, approximates twice the U.S. production of strawberries for all purposes. The consumption of synthesized Concord grape flavor converted to its equivalent in fruit approximates five times the U.S. production of Concord grapes for all purposes.[24]

Synthetic food colors are used to improve appearance. At least thirty-seven different color additives were in use in 1977.[25] Synthetic agents are used in the coloring of cured meats,[26] soft drinks, candy, puddings, and dairy and bakery products. Antioxidants are used to prevent or delay "enzymatic browning," the darkening of certain fruits and vegetables when exposed to air after being cut, bruised, or allowed to overmature. Consumers now exhibit a strong desire for various colors in food. The extent to which this desire is natural as opposed to artificially created is questionable, however, as is the firmness with which consumers would adhere to these desires in the face of added health risks.

Food additives are also used to create and maintain various textures in foods. For example, emulsifiers are used to smooth the texture of ice cream, while stabilizers are used both to smooth ice cream and to increase its viscosity. Substances used to control pH help create the textures of cheese spreads and processed cheese, and they also control the texture of candy.

Two controversial issues related to the concept of appeal are

---

24. *See* MFG. CHEMISTS' ASS'N, FOOD ADDITIVES: WHAT THEY ARE/HOW THEY ARE USED 21 (1976) [hereinafter cited as MFG. CHEMISTS' ASS'N REPORT].

25. *See* Harkins, *Food Additive Safety Evaluation*, 32 FOOD DRUG COSM. L.J. 182, 187 (1977).

26. For example, nitrites give cured meats their red color. The meats otherwise would turn grayish brown.

created tastes and authenticity. Some critics assert that consumers do not have a basic preference for some of the attributes that they choose to purchase in the food supply. These critics maintain that such preferences are created both through advertising and through continual exposure to products that have the disputed attributes. According to this theory, after a limited period of exposure to colorless gelatin desserts, a wide segment of the population might discover that it no longer had a preference for colored gelatins. To the extent that this theory is true, it would affect the calculation of the benefits deriving from artificial coloring agents. Conceptually, however, the means for computing such benefits would not change.

Some of the additives placed in products to increase their aesthetic appeal are used to replace more expensive or less readily processed natural agents. This substitution raises a critical issue that has been discussed widely in other contexts, principally environmental preservation. To what extent does it matter that a consumer's perception of what he is consuming diverges from what he actually is consuming? If a consumer cannot distinguish between a product made with natural strawberries and one made with synthetic strawberry flavoring and coloring, should we assign any lower level of benefits to the second product? Here again, we would recommend that consumers' preferences serve as the final arbiter. Holding cost, convenience, and health effects fixed, if the consumer does not prefer natural to synthetic agents, we should assign no difference between them on the appeal dimension. On the other hand, if some margin of preference persists, despite consumers' inability to differentiate between the two without outside information, that preference should be respected. If consumers are willing to pay for authenticity, they should be allowed to do so.

2. Convenience

The convenience benefits of food additives come in many forms. Food substances containing additives may be easier to purchase, to store, and to process for final consumption. Antioxidants and other preservative agents allow many foods to be stored easily, resulting in the availability of a wider range of products and also requiring less frequent trips to the store. In recent years, a wide variety of products have become available that allow easier food preparation. These products sometimes make preparation swifter or less subject to human error and sometimes add fewer outside agents. Many of these convenience products require the use of additives.

3. Resource Savings and Availability

The major justification for the presence of some level of contaminants in food is the exceedingly high cost of removing them completely. The cost of removal would raise the cost of food and, in many instances, reduce its abundance. The critical question, then, is what level of contamination should be permissible; the lower the level, the greater the cost. In general, we should expect that further reductions of the same increment will cost more and more to make.[27] Even assuming that the health benefits from reducing contaminants are linear, because the costs of reducing contaminants are nonlinear, the optimal level for contaminants may be positive, though less than the natural level.[28]

In contrast to contaminants, additives are put into, rather than taken out of the food supply. Additives have cost implications in a number of ways. They can assist with the production of the food itself.[29] In addition, an additive may be employed at some stage in the industrial preparation of a food to reduce its cost. Additives as cost-reducing agents receive the most attention in their roles as preservatives. The longer a product keeps in the store or in the home, the less expensive it will be. Finally, additives can reduce costs by replacing a more expensive natural product. Resources will be saved if a cheaper, synthetic flavor is used instead of a natural flavor. Although additive-related products will not be identical to natural products, they will save resources to the extent that they compete with or are substituted for more expensive products.

---

27. For example, it would cost more to reduce the aflatoxin level in peanuts from 15 to 10 parts per million than it would cost to reduce the level from 20 to 15 parts per million. It is not clear that this will always be the case. Are there increasing marginal costs to reduced percentages of fertilizer contamination? The answer would be no if the productivity of fertilizers is fair, but the contamination as a function of level of application increases at a decreasing rate.

In many instances contamination is a probabilistic phenomenon. A standard might set a probability that some fixed level of contamination was exceeded. Here we would expect that the cost of going from a two percent chance of exceeding the level down to a one percent chance would be less than the cost of going from one percent to zero.

28. Most standards impose a level that cannot be exceeded. More sophisticated standards may set a level that can be exceeded only with a certain probability. An ideal standard would define some expected performance measure that would have to be met. Each level would be assigned a score and the probability of exceeding that level would be multiplied times that score to give expected performance. This method is not a recommendation for practice, but rather a theoretical point that should be considered when designing any standards scheme.

29. For example, DES serves as a growth enhancer for cattle.

We have assumed that the perspective for decision should be that of an individual consumer. How can we be assured that reductions in cost achieved through the inclusion of additives and contaminants will be passed along to consumers? The answer depends on the nature of the market in which the product is produced and on the elasticity of demand for the product. If the market is competitive, the full cost savings will be passed along to consumers. At the other extreme, if the market is monopolistic, the extent of price reduction to the consumer will vary directly with the elasticity of demand. Thus, if a monopoly is using a cheaper, additive-related product to increase sales and there is fairly elastic demand, most of the benefit will be passed along to consumers.

4. Health

Additives intended to protect or improve the health of consumers fall into three main categories. One class of additives protects against disease. Nitrites, for example, are used to inhibit the growth of *Clostridic botulinum*, a bacterium capable of producing a potent toxin that, if ingested, can produce a number of neurological manifestations and a fatality rate of twenty to thirty-five percent in this country.[30] Synthetic antioxidants are added to prevent rancidity, to inhibit the carcinogenic action of polycyclic hydrocarbon carcinogens,[31] and to inhibit the formation of free radicals—an action found to prolong life in lab animals.[32]

A second class of additives is used to provide important vitamins and minerals. This practice decreases the incidence of beriberi and pellagra in the United States as well as other diseases related to dietary deficiencies. Vitamin A, vitamin D, thiamine, riboflavin, folic acid, and iron are especially important in this regard. Scientists have estimated that the total cost of providing the Recommended Daily Allowance of ten vitamins plus the necessary

---

30. The use of nitrites is not riskless, however. Nitrites can react with secondary amines to form nitrosamines, which have been demonstrated to be carcinogenic to lab animals at low dosages. Awareness of this risk has led to proposals for the use of substances such as ascorbate and erythorbate, which not only decrease the amounts of nitrites needed to inhibit the growth of *C. botulinum*, but also inhibit the formation of nitrosamines in the digestive tract. *See* B. DAVIS, R. DULBECCO, H. EISEN, H. GINSBERG, W. WOOD & M. MCCARTY, MICROBIOLOGY 834 (2d ed. 1973).

31. "[T]he recent decrease in incidence of stomach cancer in the United States has been attributed, at least in part, to the addition of BHA and BHT (butylated hydroxyanisole and butylated hydroxytoluene) to foods." Jukes, *Current Concepts in Nutrition: Food Additives*, 297 NEW ENG. J. MED. 427, 428 (1977); *see id.* at 427-29.

32. *See id.*

minerals through additives would be less than one dollar per person per year.[33] The third class consists of additives that are substituted for other potentially harmful ingredients. Examples in this category include the use of saccharin instead of sugar and margarine instead of butter.

B. Evaluating Benefits

1. Nonhealth Benefits

Economics suggests the appropriate procedure for evaluating the benefits that food additives yield to consumers: assign the values that well-informed consumers themselves would assign. If the matter is a straightforward reduction in cost, there is no problem. Being able to have a pound of meat available for $1.50 rather than $1.80 is clearly worth $.30. Conceptually, appeal and convenience offer little more complication. If a product is more appealing, easier to prepare, or easier to use, consumers will pay something for that benefit. This is obviously the reason that food companies continually are trying to create more appealing and more convenient products.

Quantifying how much such attributes are worth is more difficult. We might try to value the time the consumer saves when he prepares a product with less effort or energy. The difficulty with this approach is that we cannot find another use of his time that he finds equally valuable. Many individuals, for example, may enjoy a leisurely process of preparing a meal. Therefore, to value the time saving at their wage would be a gross overstatement.

Once again, the answer to the problem of quantifying is the consumer's willingness to pay. How much will consumers pay in the market for more attractive or more convenient products? There are two major difficulties with this approach. First, the market rarely gives unequivocal information. Products that differ in convenience undoubtedly differ in other characteristics as well,[34] one of which may be marketing effort. Merely observing the difference in prices paid in the market, therefore, will not be a good way to estimate the benefits of convenience or improved aesthetics.

---

33. *See* MFG. CHEMISTS' ASS'N REPORT, *supra* note 24.

34. A most salient characteristic on which the products might differ is in their safety. Assume, for example, that the risk level (or perceived risk level) of the more convenient product made a consumer just indifferent between it and its competitor. Therefore, he would pay no more for the convenient product. Although he still might be valuing the convenience quite highly, the risk offsets the convenience.

A second, more technical complication relates to the different valuation of a product by different consumers. Suppose two products differ only in convenience. The more convenient sells for three dollars; the other sells for two dollars.[35] This cost difference tells us only that the marginal consumer, who is otherwise indifferent between the two products values the convenience at one dollar. There may be some people who would pay much more than one dollar extra for the more convenient product, and others, who are not now buying it, who would pay substantially less. The true value of the convenience is the average value to all the individuals who would choose to purchase the product.[36]

2. Health Benefits

Evaluating the health benefits of having a substance in the food supply is often the mirror image of evaluating health risks. We have substantially less experience with the benefit side of the equation, however. We do not know how many cases of botulism the presence of nitrites in meat prevents, nor do we have experience providing answers to such questions. Nevertheless, we believe that some estimates along these lines are essential to make rational food safety decisions. In general, procedures used to evaluate health risks are appropriate for assessing health benefits as well.[37] There is one major difference, however. With risk, it is important to identify and assess every category of consequences. The negative aspects of a particular substance or use could swamp the positive elements. With benefits, however, we need only a level sufficient to outweigh the risks. Therefore, a decision to permit the use of a substance that is based on an incomplete summing-up of benefits is not inherently flawed. The following section develops a methodology for organizing information on health benefits based on health risks. For each class of benefit we look both to its magnitude or severity and to the likelihood that it will be received.

---

35. When there is strong cross-elasticity of demand, consumer surplus cannot be computed merely as areas under demand curves. See W. NICHOLSON, *supra* note 17, for appropriate techniques.

36. Fortunately, we need not be in the uncomfortable situation in which one product must be selected over another. Both can survive on the market simultaneously.

37. See, for example, the proposed evaluation system laid out in the 1982 FOOD SAFETY COUNCIL REPORT, *supra* note 3.

IV. Classifying and Evaluating Risk

A. *Quantifying Risk-Severity and Probability*

The term risk is used imprecisely in most food safety discussions. In part because of the discussion that has grown up around the Delaney Clause,[38] risk is sometimes treated inappropriately as a yes-no variable. The primary question has become: Is there any risk? The honest answer must always be yes, even if all tests up to now have failed to reveal any risk. Tests may be inadequate; errors are possible. We can never prove the absolute absence of risk in the use of a substance or process in the food area or elsewhere. A zero-risk standard often ultimately leads to situations in which we either override our regulatory procedures, as with saccharin, or we gently delude ourselves by treating a risk computed to be below a particular magnitude as if it were zero.

Decisions, however, should be made on the basis of the best possible estimates of risk. Decision processes will inevitably trade off benefits against risks in one way or another, and regulatory processes should be open about the way in which such tradeoffs are conducted. Nevertheless, it is often useful to streamline regulatory procedures by relying on rules of thumb, and it may make sense sometimes to proceed on the basis of a zero level of *perceived* risk. If such an approach is employed, we should recognize that the risk-benefit assessment is being applied not to a particular substance but to alternative regulatory approaches themselves. When evaluating procedures rather than substances, it is important to subject them to reevaluation as circumstances change, for example, as scientific capability to perceive positive risks improves, as the regulatory process itself evolves, as the mix of substances to be regulated changes, or as the values of society transform.[39]

How should risk be discussed and quantified? Assessments of risk must embrace both its probability and its severity. Much scientific attention has been directed to the problem of measuring the probability of an adverse effect. Although no less important, evaluating the severity of an adverse effect has received considerably less attention. Any discussion of a food safety risk should start by

---

38. *See supra* note 1.

39. Concern about reorienting the food safety regulatory process is a response to some such changes. Scientific capabilities bave improved dramatically; low-level risks are much more readily perceived. The overall view of regulation within society is much more skeptical. In particular, it is recognized increasingly that not all risks can be eliminated from society, nor should they be.

identifying the nature of the unfortunate consequence that may result.[40] How serious will be the impairment of function; when will it occur; how many years will it persist? No cardiovascular problem is welcome, but one that will appear at age seventy-five is less dreadful than one likely to become evident at age forty.

Before beginning the actual quantifying process, a significant question must be resolved: What level of resources should be devoted to the process? Finer and finer calibrations of risk levels always will be possible for additional expenditures of resources. However, if a relatively inexpensive, possibly cruder test indicates that a substance is likely to prove unacceptable, it may be worthwhile to truncate the information-gathering process. The regulatory process also may wish to establish standards of minimum benefit and maximum risk that products must meet to win ultimate approval. These benchmarks could provide an early indication that a substance is not likely to prove acceptable.[41] Equally important, a pattern of consistent decisions that follow consistent criteria should enable companies or interest groups to make informed judgments about the acceptability of their proposed products. If a substantial proportion of substances that would not receive ultimate approval are stopped voluntarily before they reach the formal, governmental portion of the regulatory process, the system is working effectively.

In our discussion we will employ concrete terms for measuring the probability of a health risk and the severity of that risk. The probability concept, an RU or risk unit, was created to reflect the probabilistic manner in which risks are incurred. Most health risks that society incurs are low-probability risks spread over a wide population.[42] In order to measure the severity of a risk, we will employ the term QALY, quality-adjusted life year. For example, instantaneous death to a forty-year-old who would otherwise live a normal lifetime of seventy-five years would represent the loss of roughly thirty-five QALYs.

---

40. For example, a food ingredient with a .001 probability of causing a minor rash is surely more acceptable than one yielding equivalent benefits but accompanied by a .00001 chance of liver failure.

41. A "revisit process" might be employed to enable substances that exceed the maximum risk level to seek approval if they also offer a very substantial level of benefits.

42. A situation in which 1,000,000 people face a 1 in 10,000 chance of incurring a disease is far different from one in which 100 identified people out of a population of 1,000,000 will, with certainty, contract the disease.

1. Quality-Adjusted Life Years—A Measure of Severity

Ideally, in order to measure the severity of a risk we would have some readily available measure of health impairment. One possibility, not likely to gain acceptance whatever its theoretical validity, would be to develop a monetary scale for health loss. We suggest an alternative approach that compares health losses directly with each other. One concept that has been employed experimentally in this regard is loss of quality-adjusted life years. Although we do not expect this concept to be employed directly in any regulatory decision, we do feel that attention to the differential consequences of various types of health losses is important. Any regulatory mechanism, therefore, should explicitly identify the nature of any potential health loss.

The QALY concept has been specifically developed to deal with situations of uncertainty. QALYs are calculated on an expected-value basis using a well-accepted technique in decision analysis called von Neumann-Morgenstern utility. Thus, for example, in choosing between one course of action that yields a QALY level of .9 and another that offers a QALY level of 1 with probability .85 and a QALY level of .5 with probability .15, the individual would choose the second course of action. It offers an expected QALY level of (.85 x 1) + (.15 x .5) = .925, and is therefore preferred to the alternative, which offers an expected QALY level of .9.

Risks must be identified in terms of both severity and probability, and within the severity category an effort must be made to distinguish further such aggregate categories as "causes tumor." If regulations can be formulated that disaggregate by these two elements, the evolution of regulatory processes through the actions of companies, individuals, courts, and regulatory agencies will slowly lead to better-informed decisions and, therefore, better decisions. It would be unrealistic to attempt to describe now a methodological approach that could be employed successfully. The outline of the QALY methodology was meant to be merely suggestive.

Some may object that the QALY approach pays excessive attention to quantification. The counterargument is that by not quantifying one may gloss over distinctions that are important in practice, and the difficult process of thinking about numerical values will engage us in an appropriate thought process. Even if the objection to the QALY approach were accepted, it would still be imperative for society to make some attempt to assess the severity

of the identified risks. At the very least, the consequences of the risks should be described in words. The object is not to lose sight of attributes that are not easily quantified.

2. Risk Units—A Measure of Probability

We have stressed throughout this Article the need to focus attention on food risks as they would be incurred and perceived by well-informed consumers. At the outset of this section on risk, we distinguished between severity—the maguitude of the loss should the risky event occur—and probability—the likelihood that the event will occur. We have proposed the use of some concept analogous to QALYs, which measures loss by the preferences of tbe consumer, who naturally would be concerned with the timing and magnitude of his potential loss.

Let us now turn to the probability factor. Fortunately, most high-risk food items have been eliminated from the food supply, and so very few of the remaining risks to health or life associated with food safety are high-probability events. Though commentators sometimes refer to X cases of cancer or Y additional incidents of liver malfunction, there are not X individuals who knowingly will contract cancer because they consume some specific food substance. We believe that the conversion of small risks to a large population into a certainty number is fundamentally misleading. It is quite appropriate for an individual knowingly to take a 1 in 1,000,000 risk of cancer in return for some finite, indeed small, benefit, yet be quite unwilling to accept the certainty of cancer in return for any given amount of money or other benefit.[43] Given that individuals confront risks as low-probability events, and given that there may be nonproportional responses in terms of the amount of money or other goods required to compensate for risks, we think it appropriate to discuss risks according to the probabilities with which they are incurred. Tlus more exacting procedure, unfortunately, tends to bring us into unfamiliar terrain, requiring us to consider such figures as probabilities of 10^{-7} or 10^{-5}.

To measure risk probability we have developed the concept of a risk unit, or RU (pronounced "Roo"). A RU is a probability of loss of 10^{-3}, or 1 in, 1000. The severity umt might be a QALY, a number of QALYs, or a dollar amount. In the discussion that follows, assume that a single QALY is at risk. For example, we miglit

43. *See* Zeckhauser, *Coverage for Catastrophic Illness*, 21 PUB. POL'Y 149, 149-72 (1973).

be thinking of alternative risks to the life of a forty-year-old man. Assume that the decision maker had a variety of ways of avoiding RUs, that is, of reducing his risk of losing a QALY. He could purchase safety devices for his automobile, fireproof his house, or give up rich desserts. Assume that auto safety devices would offer reductions in RUs for ten dollars each, and that fireproofing would provide a reduction of one RU at a cost of twelve dollars. Clearly our protagonist should purchase the auto safety devices before the fireproofing. Similarly, we could calibrate the value of the rich desserts in terms of money—how much he would be willing to pay not to have to give them up—and see whether that was the most economical way to purchase reductions in RUs. Purchasing reductions in RUs in the cheapest available manner provides both lower risk (fewer RUs) and more other benefits.

In practice, matters will not be so simple. Rarely will one technology always provide the cheapest way to purchase reductions in RUs. In general we must expect to experience diminishing returns. For example, a person who equips his car for safety may start by getting new tires to replace his bald ones. Assume that this transaction would reduce his risk by ten RUs for one hundred dollars, or ten dollars per RU. Next he may put in airbags, which reduce his risk by twenty RUs for three hundred dollars or fifteen dollars apiece for a total reduction of twenty RUs. Finally, he may purchase extra heavy bumpers, which reduce his risk by two RUs at a cost of forty dollars each. Obviously before purchasing airbags or bumpers he should purchase some of the reductions in RUs available in home fireproofing for twelve dollars each.[44] The principle of effective purchase ensures that reductions in RUs could not be purchased in some other area at less cost than in the areas where they currently are being purchased. This principle, in essence, is the principle of cost-effectiveness analysis.

Cost-benefit analysis takes matters one step further. Not only should you expend your resources in a manner that secures maximum benefit for the level spent, but you should be guided by your preferences in deciding how much to spend. Consider the purchase of electricity. Assume that as more electricity gets produced it becomes increasingly costly to produce. If the value of a unit of electricity is nine dollars—that is, if you would give up just nine dol-

44. Diminishing returns must be expected here as well. Eliminating clearly faulty wiring may provide RU reductions cheaply. Assuring that the structure will not burn even if the temperature gets to 500°, however, may reduce risk only at exceptionally high cost.

lars' worth of other goods to get the benefits that electricity would produce—then you should continue to purchase electricity until the price rises to this level. Another factor may come into play here. As more electricity is purchased, its value in use may diminish. At first the electricity may provide light to see other individuals or perhaps power to turn on the television set, a significant source of entertainment. However, more marginal uses, such as for a trash masher, may yield smaller benefits. The efficient outcome pays attention both to declining benefits as production goes up and to increasing costs. The efficient equilibrium is reached when the cost of producing an additional unit just equals the value of the benefits it will provide.

Applying the same principles to situations with risks to life reveals some parallels and some differences. The most dramatic difference is that we are not comfortable with talking in terms of valuing concepts such as RUs. In many contexts it will seem downright unpleasant; some would say unethical. We continue to make efforts toward some valuing concept because we know we will sacrifice some welfare if we do not. The second difference actually simplifies matters. Over the relevant range for consideration, the value of an additional reduction of a RU may not change very much. This suggests that the slope of the demand for reductions in RUs is not great. The supply curve, however, may be rather steep. Assume that the present risk level is .01—one chance in 100 of losing a QALY. It might be possible to cut this risk to .009—reduce the RU level from 10 to 9—for an expenditure of $1,000. A further reduction of the same amount, to .008, may cost $5,000.[45] Additional reductions in RUs dramatically increase in price. Such a steep supply curve has two important implications. First, the supply of reductions in RUs will be the principal determinant of the number that is purchased. Second, the valuation placed on a reduction of a RU—the value that normally would be read off the demand curve—will not make much difference for the number of RUs that will be demanded.

Let us see how this discussion would apply to the area of food safety. First, once we had decided how stringently we wished to regulate the food supply, we certainly should employ cost-effectiveness analysis. If reduction in risk is our measure of effectiveness,

45. The demand curve might be downward sloping if an individual bad to pay a substantial portion of his income to purchase reductions in RUs. We do not think that in general such a price will be required.

this analysis is equivalent to risk-benefit analysis. Risk-benefit procedures should be followed so that we do not sacrifice both reduced RU reductions and resources. Second, we should attempt to infer the level of risk that individuals would be willing to incur for a particular level of foregone benefits. Great precision may not be essential in this assessment. Third, if decisions are being made within a political arena, as they certainly will be in the food safety area, then providing the decision makers with relevant information about risks and benefits and letting them make the value judgments is the only intellectually defensible process. We do not require that a full cost-benefit analysis be undertaken, for we know that the types of decisions that will be made ultimately will rely on the valuations of these decision makers, not of analysts. Last, any gross inefficiencies in the process will come not from overvaluing or undervaluing reductions in RUs, but rather from purchasing expensive reductions in RUs before reductions that are less costly.[46] This analysis leads to a recommendation for a strategy of full information, in which decisions are made in light of the best possible estimates of all consequences. The important objective is to avoid the major errors.

Consider the use of an additive or different additives in three different applications. Standardize the quantities for these applications so that each entails a one RU risk of losing a QALY. The nonhealth benefits of application A are ten dollars; for application B one hundred dollars; for application C one dollar per RU of QALY loss. Any rational process of decision would allow application B ahead of application A, which in turn would be ahead of C. We might well "purchase" the reductions in RUs that would be available by keeping A and C off the market while allowing B. In what follows, unless otherwise indicated, the severity unit to which an RU applies is one QALY.

B. A Conceptual Approach to Decision Making

An intellectually defensible regulatory approach may not follow the dictates of cost-benefit or risk-benefit analysis slavishly. Any approach, however, should follow such analysis in spirit, employing ad hoc procedures and rules of thumb when they greatly simplify choice decisions. The committee on which I served with Samuel Stumpf, which was charged with proposing a design for a

46. This could occur, for example, if different criteria are used for old and new substances, for additives and natural contaminants.

food safety regulatory process, advocated principles in that spirit.

1. When regulatory actions can be taken to eliminate a major risk, take that action, unless the benefits lost by taking that action are excessive.
2. When no major risk is entailed, conduct an analysis that identifies both benefits and risks, the latter quantified through the combined RU and QALY approaches. Take appropriate regulatory decisions when the ratio of indexed risk to benefits is clearly high (opt for safety) or clearly low (opt for benefits).
3. For whatever regulatory decision is taken, refine the strategy in any manner that can reduce RUs or enhance benefits cheaply. This refinement might require a selective ban, a warning strategy, a limitation to particular uses, etc.[47]

Figure 3

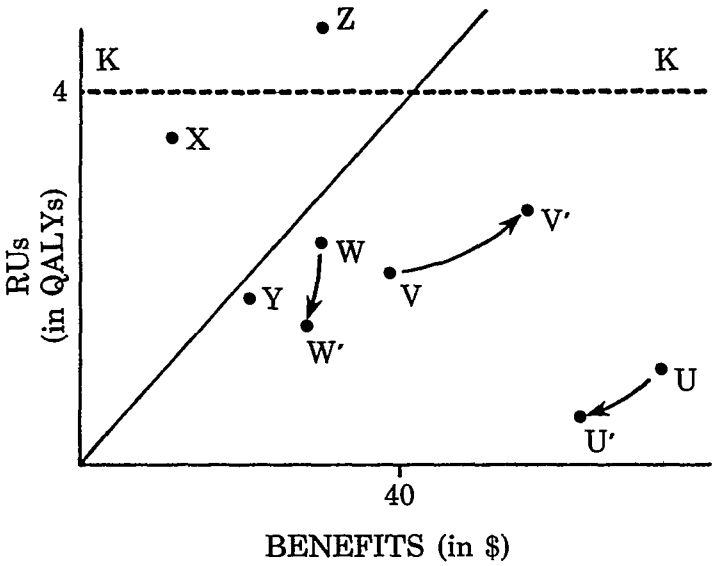

A simple diagram clarifies the recommended procedure. The line KK in Figure 3 depicts the risk level above which a substance automatically is banned.[48] Substances that cluster well above some benchmark line are eliminated. Substances that cluster well below the benchmark line will be approved. A very large percentage of

47. In theory one should undertake marginal adjustments until the cost per reduction in RU purchased just comes into line with the marginal value of a reduced RU, which in turn would be given by the slope of the benchmark line. We argued in the text that the value of a reduced RU probably would not vary significantly with the quantity purchased. To the extent it does vary, it will be from income effects. *See* 1982 FOOD SAFETY COUNCIL REPORT, *supra* note 3.

48. This rule-of-thumb approach can be overruled if dramatic benefits can be demonstrated as well.

substances probably will fall in these rather distinct regions. For each accepted substance and the substances that cluster reasonably close to the benchmark line we examine regulatory refinements that can purchase reductions in RUs at a low price in benefits, or benefits at a low price in RUs.[49]

In the example diagrammed in Figure 3, any substance that entails a risk of more than four RUs for a consumer will be banned automatically, as is point Z. Point X is ruled out by the benchmark line, which suggests that a reduction of one RU should be valued roughly at ten dollars. The substances represented by points U, V, and W are accepted. Note that V would be transformed to V′ in an optimal scheme, and W to W′. V′ accepts additional risk to secure additional benefits; W′ reduces risk at a sacrifice in benefits. Although the movement from U to U′ would improve the risk/benefit ratio, it would be rejected because the reduction in RUs per dollar of benefits sacrificed does not equal or exceed the rate that the consumer would accept. Moreover, there are cheaper reductions in RUs available elsewhere.[50] What about point Y? We have argued that there will be few such points. Assuming that the initial benchmark was drawn with some attention to consumer preferences, it really does not matter very much what is done with Y.

Cost-benefit analysis traditionally is justified as an efficiency-seeking measure. Properly performed, it assures that benefits are maximized for whatever level of expenditure is undertaken. Indeed, this goal can be reached simply through cost-effectiveness analysis, of which risk-benefit analysis is a special case.

The second accomplishment of cost-benefit analysis is that it determines an appropriate level of expenditure for any particular type of benefit-providing program. When lives are at risk, the nature of the tradeoffs will not vary dramatically with the level of reduced RUs purchased. Moreover, the world is unlikely to generate a distribution of RUs that huddles about any benchmark line. Rule-of-thumb techniques are likely to serve well as a mechanism for decision making in such a world.

---

49. Note that we are looking at movements from a particular position. It would not be appropriate to judge a refinement, therefore, by what it does to the RU/benefit ratio.

50. In general, if there are diminishing marginal benefits and costs, the curves will be bowed as shown in Figure 3. At the optimum, the slope of all transformation arrows will be equal.

C. *Risks to Special Groups*

Given the array of potential policy objectives for regulating food safety, a relevant question is whether adopted regulations should be applied uniformly across the population or whether they should be varied according to the market to which a product is aimed.[51] If the objective is either to eliminate or to minimize risk, the proper strategy would seem to be uniform regulation. If the objectives are less extreme, however, this question becomes very relevant.[52]

There are several reasons why one might want different forms of regulation for different segments of the population. First, different societal groups might have different preferences, for example, in the tradeoff between risk and sacrificed resources. Second, different societal groups might have different potential medical reactions to particular substances. Third, different societal groups might have different abilities to make informed choices. Last, different societal groups may experience different levels of externalities.

One approach to this question of the proper scope of regulation is to examine how society deals with similar concerns in other regulatory areas and to look at possible methods that could be used in dealing with food additives. For instance, many products contain toxic chemicals that are intended for external use but are harmful if taken internally. The regulatory strategy here has been to label these products with warnings of risk but not to limit their sale. Thus, regulators have decided not to eliminate risk, but rather to inform consumers of the risk and to tolerate the risk of tragedy associated with the chance that an unattended infant might ingest the product. Similarly, a decision has been made to

---

51. A further relevant question in this context is whether it is reasonable to assume that consumption of a product can be limited to that segment of the market to which it is aimed.

52. This question can be incorporated into a decision tree, as depicted in Figure 4.

Figure 4

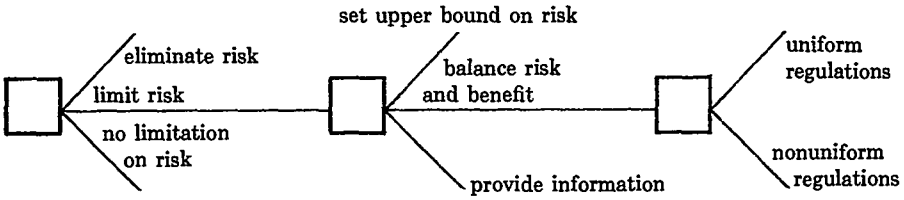

tolerate some risk in over-the-counter drugs. A number of substances, such as aspirin, are made readily accessible to consumers in potentially lethal quantities.

A third, more complicated situation is prescription drugs. Physicians have been given the responsibility of regulating the distribution of these commodities and of knowing their "tolerable" levels of consumption. Doctors must be especially aware of the differential effects of certain substances on various subgroups like fetuses, infants, the elderly, and women of childbearing age. Doctors also are expected to know the interactive effects, both synergistic and antagonistic, associated with giving certain drugs in combination and to advise patients of possible adverse effects of ingesting certain substances while taking particular drugs. The doctor presumably chooses to incorporate his knowledge of differential effects of drugs into his strategy of distribution, rather than choosing to deny drugs inflexibly to all groups because they are potentially toxic to some groups. Examples of this strategy abound in medicine. Eating food containing a high content of tyramine, such as cheese, is dangerous to a patient being treated for depression with a MAO inhibitor but not for a patient being treated with a tricyclic antidepressant drug. A physician should warn his patient accordingly. A second, broader example is the many drugs that have adverse side effects only when administered to a mother in the perinatal period. In this situation, doctors typically refrain from giving these drugs to women in the perinatal period whenever possible, but continue to prescribe the drugs to others.[53]

In many situations, then, society has opted to tolerate some health risks and to apply standards differentially to various population subgroups. What should be done in the case of food additives? A number of factors dealt with in other sections of this Article are relevant to this question. The concerns of this section, as they bear on this question, can be summarized as follows:

(1) Objectives should be outlined with regard to attitudes toward risk.
(2) If appropriate methodology can be devised, the objective of limiting risk

---

53. The complexity of the differential effects may lead doctors to concentrate their attention on a single specialty such as pediatrics or cardiology. Such specialization enables doctors to master the information necessary for dealing effectively with their patients.

Issues of this sort have arisen in the occupational safety and health area. Some risks are incurred primarily in particular groups. An example that has received widespread discussion concerns pregnant women working in plants that make batteries. Critical lawsuits have been fought to determine whether exposures at all johs in such plants must be set so that the work is safe for pregnant women, or whether these women can be rotated off endangering jobs during pregnancy.

is preferable to eliminating or ignoring risk. Eliminating risk seems infeasible, and ignoring risk may lead to significant health loss without commensurate gains in benefits.

(3) If appropriate methodology can be devised at reasonable cost, nonuniform regulation is preferable to uniform regulation.[54]

(4) Whatever strategy is adopted, the methodology should be fiexible enough to adapt to changes in the available knowledge about the risks of consumption of various materials.

(5) Whatever system is adopted and however individualized it might be, we should recognize that when groups differ either on preferences or on risk levels, there will be inevitable tradeoffs between the well-being of different groups.[55] The purpose of a nonuniform strategy is to limit the magnitude of the problem caused by confiicting interests, to allow one group to benefit from a product at the same time that another group can readily limit its use of the product.

D. Use of Information on Risk

Much of the debate within the food safety area starts with the observation that our quantitative assessments of benefits and risks cannot be precise. Rather than throw our hands in the air at this point or stay with the status quo, however, it is worthwhile to inquire what we would have done or should do with complete information. This Article argues that we should use that information to design a regulatory system that will approximate the choice that well-informed individuals would make for themselves. This does not imply, for instance, that we should ban substances that are unacceptable to the median consumer. To do so would deny the right to use the substances to those consumers who would be willing to accept higher risks relative to benefits or to those consumers who are at lower than average risk. Neither does it imply that we should rely solely on an informational strategy. This approach would impose severe decision-making costs on consumers, and they undoubtedly would make some errors as well. The ideal system undoubtedly will employ a mix of strategies—doing nothing, limiting

54. While there are many similarities between the food additive situation and the prescription drug situation, there are many differences as well. The costs of individualizing distribution are much lower in the prescription drug situation, in which tbe doctor must consider each case individually whether drugs are concerned or not.

An alternative and less complicated strategy could be to label products with obvious symbols, such as a red dot, or a syringe, or a baby, as an indication that known or unknown health risks might be associated with consumption of the product by particular groups—each group being clearly specified.

55. To provide the highest level of protection to group A—individuals who are violently allergic to a particular yellow food dye—it may be necessary to limit tbe opportunity of group B to consume the substance. In the example cited, given the potential failures of a labeling system, we may not want to allow the product on the market at all.

use, providing information, and banning.

The design of a regulatory system is essentially an exercise in decision making under uncertainty. No system can expect to reproduce exactly the choices that well-informed consumers themselves would have made. The objective should be to define the system that performs best on the average. In designing such a system, we must also be alert to the divergence between the actual performance of systems and the way they would perform under ideal circumstances. The present regulatory system, responding as we would expect to a host of political pressures, to some extent reflects a risk-benefit orientation. Its operation, however, is extremely haphazard.[56] A somewhat rationalized system could offer both more benefits and less risk to consumers.

E. Uncertainty in Probability Assessments and Updating Estimates

Whatever regulatory system we design must recognize that measures of risk are likely to be imprecise. The theory of decision making under uncertainty has dealt with this problem in a number of arenas in addition to food safety. Two important lessons have emerged.

First, the appropriate measure for an uncertain probability is the mean of the distribution of that probability. Assume that we have a substance that has one chance in ten of conveying a ten-RU risk of losing a QALY and nine chances in ten of conveying a one-RU risk of losing a QALY. The mean risk is then 1.9 RUs. Consider an individual who must choose between consuming this substance or another that incurs a certain two-RU risk. The individual should choose the first substance even though the true value of the probability is an uncertain parameter itself. This principle—the resolution of lotteries on probabilities into their expected values—is the cornerstone of von Neumann-Morgenstern utility the-

56. A cynical view would suggest that the greater the total economic losses to the affected industry, the less likely it is that a product would be severely limited in use or removed from the market. One might argue that this would be rational if consumers' and producers' benefits were roughly proportional, but they are not. The appropriate measure is risk per unit of benefit. The system described would make limitation less likely the greater the total risk. Other factors being equal, the greater the consumption of a product, the less likely it is to be banned.

There are strong forces operating on the risk side of the equation as well. They too do not always push for rationality. In the present state of affairs, it seems much simpler for those who highlight the risk side of the equation to keep products off the market than to reduce the use of those already in widespread use.

ory, which itself serves as the basis for the contemporary theory of rational decision making under uncertainty.[57]

Second, when there is uncertainty on a probability there may be attractive opportunities for sequential strategies for learning and action. Opportunities for learning should push the regulatory system in two directions. First, it should be more flexible. If learning is to be useful, the system must be able to change the way a substance is regulated or the strictness with which it is regulated. For example, an informational approach should be able to evolve into a use-limitation approach and vice versa. Second, the system should be less restrictive in permitting the initial use of substances, on the assumption that experience will generate useful information. When there is uncertainty, there is the possibility that the product will have a better than expected risk-benefit ratio. In this case, its use can be continued or expanded, thereby yielding significant net benefits that offset the expected losses from setting a lower standard at the outset.

If a substance has an acceptable risk-benefit ratio but further testing would produce more information, it may seem appropriate to ban the substance and wait to see if the new information reduces the expected value of the ratio. On an expected-value basis, however, that approach would produce a loss in welfare along the way. Because further study might push the risk-benefit ratio into the unacceptable range, the only analytic argument in favor of banning a substance with an acceptable ratio is that such a ban may be necessary to ensure that the additional information is generated. Ideally, the regulatory system will be sufficiently flexible to allow conditional approval, which will be removed if additional study is not undertaken in some specified time period or if the additional study provides unfavorable information.

In short, the regulatory system should treat substances in a dynamic context. An initial regulatory decision should be recognized as such. Further information should be gathered so that a more refined decision can be made later. Indeed, the design of the regulatory system should encourage or require the collection of such information.

V. Common Errors in Risk-Benefit Assessment

This analysis rests on the assumption that regulatory efforts should provide individuals with the commodities they would select

57. *See* H. Raiffa, Decision Analysis (1968).

for themselves were they well informed and purchasing in well-functioning markets. In addition, proper regulation has an objective function that requires balancing the risks to health with other valued attributes such as nutrition, convenience, economy, and appeal of food products. Fully informed individuals would be expected to select a diet consistent with maximization of their welfare on such an objective function. Many common approaches to risk-benefit analysis address the tradeoff between these sets of attributes, yet fall into logical errors. Here we catalogue a few of those errors that are most commonly observed.

A. Taking Excessive Guidance from Risk Levels in Other Areas

Many analyses of risk suggest that a problem is of secondary importance in some sense because other risks are much greater. Whether we should or should not accept a substance into our food supply is independent of the risks entailed in crossing the street or getting hit by lightning. In the category of food, the argument sometimes is made that, because contamination causes more illness and death than food-induced carcinogenesis, we should be more concerned about eliminating contamination. The inference does not hold. Only if we could show that the same reduction in illness and death could be achieved in the contamination area at a smaller sacrifice in benefits than in the carcinogenesis area would we be correct in pushing for tighter regulation in the first area as opposed to the second. The central principle of effective regulation is that the marginal cost of achieving some given amount of risk reduction should be constant across areas.

B. Failing to Take Appropriate Guidance From Risk Levels in Other Areas

One type of information can be secured from examination of risk levels in other areas. To the extent that those risks are voluntarily accepted and understood, they may indicate the shape of individuals' objective functions. If a worker knowingly will take a one in 10,000 risk of cancer to earn an extra $200 per year, we get a very rough indication of how he feels about the risk of cancer. Suppose such risk premiums were widely deserved in the labor market. Moreover, assume that a food additive was under consideration that would save an individual $100 a year in food costs, but would entail a risk of cancer estimated to be one in 100,000. This occupational risk information would suggest that the food additive risk

level was not excessive. If, however, the food risk turned out to be one in 1,000, the occupational evidence would suggest that the risk should not be taken. Obviously there is a range of in-between values for which the occupational risk comparison would provide no guidance. In those areas, however, in which guidance can be secured—and there are many more than we would expect, given the capricious nature of both the regulatory process and decisions concerning low-probability events—it should be taken.

C. Recognizing When the Absolute Size of a Problem is Relevant

The absolute magnitude of a food safety problem is of consequence when we wish to determine what level of research and regulatory resources should be devoted to its examination and control. Assume that there are two substances on the market, whose benefits and risks per unit under present regulatory strictures are estimated to be the same. Over 100 million units of substance A but only 1 million units of substance B are ingested each year in the United States. This "incidence" information alone would not suggest that strictures on substance A should be any different from those on substance B. It would imply, however, that resources available to find out more about the risks of the two substances would be addressed more profitably to substance A.[58] If the information suggested any change in regulation, a much greater quantity of resources would be at issue with substance A. Therefore, though the expected information gained about the two substances might be equal, the expected value of that information would be greater for substance A.

What would happen if, after the research were undertaken, the estimated risk and benefit levels for substance A were not noticeably changed? This would suggest that present regulations should not be changed either. Regulation should be based on the best available estimates of marginal benefits and risks. Just because the estimates are based on firmer information does not imply that regulation should be tighter or looser.

---

58. This conclusion assumes that the likely scientific returns on research are no greater for one substance than the other. In the language of decision analysis the value of the information gained from a laboratory-month would be equal for the two substances.

D. Failing to Make Benefits Part of the Calculation

The most common error is to ignore benefits when choosing regulated levels of risk. Ignoring benefits would be appropriate only if some rather extreme assumptions were valid: (1) Society legitimately thought that safety and health were its only concerns and, therefore, only a zero level of risk was permissible; and (2) Artificially induced risks associated with food safety affected welfare completely differently from risks inherent in the food products themselves. If the second assumption is not satisfied, a zero risk level for man-made additives would make no sense, because we are already running nonnegligible risks elsewhere. If the first assumption is not satisfied—some nonzero risk level is deemed acceptable—then it would have to be a risk level per some unit. Some of the units that are at times invoked do not make sense unless they are proportional to the benefits at issue. The unit that comes closest to accurately reflecting benefits is probably per adult serving, for if there were equal benefits from a serving of anything, we would have a reasonable surrogate for benefits. However, the benefits one gets from a serving will depend first on how much one values the particular commodity in question, and second on what the alternative is. If the alternative is much less healthful, much less tasty, or much more expensive, the benefits of the particular commodity are high.

The appropriate decision criterion is, as stressed repeatedly, to do as consumers would do for themselves. Consumers naturally look to benefits in deciding whether to accept a particular increase in risks. For example, on an icy morning, many people would drive to get to work, but not just to go downtown to pick up a candy bar. We eat candy bars, which may be bad for us, because we get enjoyment from them. Other equally harmful substances are not ingested because they give less pleasure. This example does not suggest that individuals are finely honed balancing machines that delicately calibrate costs and benefits and make quantitative decisions. Individuals are roughly guided by costs and benefits, however, and they attempt to direct their diets accordingly. When cholesterol became associated in the public consciousness with high risk levels, egg consumption went down. Benefits did not change, but individuals believed that risk levels were higher. Similarly, if some process were discovered to make liver more tasty, we could expect the consumption of this perceived healthful product to increase. From the standpoint of the regulatory process it is not important whether consumers make mistakes or act on the basis of

poor information. The critical objective still should be to help consumers do what they would like to do for themselves.

E. Giving Inadequate Consideration to the Shapes of Benefit and Risk Curves

At various junctures in this Article, we have highlighted the difficulties of making accurate assessments of benefits and the much more extreme problems of attempting to assess risks. At the same time, we have suggested that an orderly inquiry to determine these magnitudes as best as possible would be desirable. Even decisions based on highly imperfect information are superior to those made in a vacuum.

One useful piece of information is the shape of benefit and risk curves. This information would be important when considering whether a ban of a substance is desirable as opposed to some limitation on its use. Though there is widespread disagreement on the form as well as the location of dose-response curves, over the range of likely human consumption, reduction in risk from reduction in dose probably will be fairly constant.[59]

Assuming a roughly linear response, it does not make sense to limit, as opposed to ban, human consumption because of reduction in risk alone. If it is worthwhile to reduce one unit, it would be worthwhile to reduce the next, unless another factor were at play. The other critical factor is benefits. Even if risks are roughly linear over the range of consideration, the level of benefits foregone as intake is reduced may increase significantly. For example, cutting consumption from four units to three units may cut benefits by fifteen, and moving from three to two may reduce benefits by twenty-five. If the benefit curve has steep upward slope, and there is both empirical and theoretical evidence suggesting that it will in some instances, then there is an argument for limiting rather than banning use. We wish to limit use to the more essential areas in which benefits are high—areas in which consumers would be willing to accept the risks because the benefits foregone to reduce them would be even larger.

Consider two possible problem areas: aflatoxins in peanuts and artificial sweeteners. Assume for purposes of argument that the per-unit risks over the relevant range are constant. As we reduce the permissible levels of aflatoxin, thereby increasing the propor-

59. Thus, even if the dose-response curve has an S shape, it is likely that we will be operating primarily on one of the flatter ranges.

tion of the peanut crop that would have to be destroyed, the level of economic dislocation would rise at an increasing rate. Because price would be rising, the value of what was destroyed would rise faster. Therefore, it might be reasonable to have, as we now do, a permissible level for aflatoxins in peanuts, recognizing that we are still running a risk. The marginal benefits of artificial sweeteners are harder to identify. However, a simulated experiment can lend some insight. Assume we were to levy a tax on the sweeteners. As the tax grew larger, the uses of the sweetener would grow more restrictive. At least in broad strokes, this is an indication of the increasing marginal cost, as perceived by consumers, of limiting intake of the sweeteners. Even with a severe tax some uses would persist. This persistence suggests that for those uses, even after taking account of the health risks, the perceived benefits of the sweeteners are exceedingly high.

This argument is as cogent for new additives as for those existing on the market. However, once uses have become established it is much harder politically to prohibit them. This reluctance to ban previously accepted substances undoubtedly explains the violent reaction against a possible saccharin ban. If the FDA had sought merely to reduce the use of saccharin by taxing it, limiting its production, or banning its use in more marginal products, the outcry probably would have been far less severe. With new additives, the individuals who would benefit from the high-valued uses may be unaware of their potential benefits. There is not much of a lobby, for instance, for cheaper meat that would be produced with the assistance of a risky substance. But if a substance now widely employed were to be taken off the market, raising meat prices noticeably, there would be a great deal of opposition.

The point of this section is that regulatory policies should give strong consideration to the shape of benefit curves as well as risk curves. Because benefit curves may be much less linear, it may be reasonable to restrict uses of some additives, but not to ban them.

F. Over Distinguishing Between Natural and Man-Induced Risks

At first glance, it seems reasonable for the regulatory process to distinguish between aflatoxins—a mold that occurs naturally in peanuts—and man-induced additives. Nature has burdened us with the first; we would be burdening ourselves with the latter. To the extent, however, that we can eliminate natural risks from the food supply just as easily as we can avoid taking on new ones, the

distinction diminishes in importance. Obviously, if we rely on estimated benefits and risks, any distinction between natural risks and man-induced risks vanishes unless we feel differently about those two classes of risks. Why might we have such feelings? First, our society feels differently about errors of omission and commission, treating commission much more harshly. Second, any attempt to remove risks currently in the food supply might require substantial dislocation of present resources and alteration of present consumption patterns.[60] Failing to have the opportunity to move toward a "superior" consumption pattern, the sacrifice when a new additive is banned is hardly as severe. Obviously political forces would be strongly responsive to this factor: in-place resources are weighed much more heavily than others.[61]

What should we conclude about the distinction between natural and man-induced risks? First, there will be powerful forces and traditions suggesting that natural risks should be counted less heavily. Second, in terms of benefits foregone and risks incurred, there is no distinction between the two. Third, though somewhat different standards may be warranted in the two areas, gross differences are inappropriate. If we can accept fifteen parts per billion of aflatoxins in corn, we also should be willing to accept some nonzero carcinogenic risk levels in additives, assuming that the benefits are commensurate with those of naturally "contaminated" corn.

G. Drawing the False Implication That High Natural Risk Reduces the Need for Regulatory Stringency

We discussed above the fallacious nature of the argument that because risk A, which entails one chance in 10,000 of death, is present in society, risk B, which entails merely one chance in 30,000 of death, also should be accepted. A closely related error is to say that because some risk due to C is already accepted, indeed perhaps cannot be eliminated, some additional risk due to C, particularly if it is small relative to the present level of risk, is acceptable.[62]

---

60. There are immediate losses in capital value as well as losses to displaced factors of production that must seek their highest return during a transitional period.

61. The disparity in the treatment that present policies give to new versus previously approved substances can be seen as a product of these factors at work.

62. This argument has been made frequently with regard to man-induced radiation. Some contend that the prohlem is less serious because of background radiation and inevitahle radiation from X-rays.

For efficiency, as we have argued above, we wish to look at the benefits foregone for marginal reductions in risk level in each area in which reductions could be achieved. Assume that there is currently no risk due to D, but that naturally caused consequences of risk C are already ten RUs of the loss of one QALY. New foods containing C and D now come up for review. For a given quantity of benefits, perhaps calibrated on a dollar scale to equal $200, the food containing C introduces an additional risk of one RU. For that same quantity of benefits, the food containing D introduces a risk of two RUs. If total risk level and total benefits are the objects of our choice, clearly it would be foolish to approve the food containing D and reject the one containing C. Yet this is frequently proposed. Sometimes the supporting argument is that we should have zero risk level except when nature has made a higher level inevitable. Assuming that this argument is correct, it also would suggest that man-produced additions to the level of risks initially produced by nature also should be zero. If no quantity of benefits is worth any additional risk, it does not matter what the base is, whether zero or larger.

There is one possible argument for accepting different risk levels when a sizable natural risk or previously approved risk is already in place. It is an anxiety argument. Individuals worry about small-probability events almost independently of the size of the probability. Most people probably would worry more if they were told there was one chance in 10,000 of contracting cancer from ice cream consumption and one chance in 10,000 of contracting cancer from bacon consumption than if they were told that bacon was safe, but that ice cream consumption entailed a risk of three in 10,000. If the probabilities are small enough, then the anxiety cost may be large relative to the cost associated with the health loss itself.[63]

---

63. Recognition of the potential importance of anxiety presents a peculiar dilemma for policy. Normally we think that individuals are better off when they are given information about the products they will consume. If, for example, consumers are told that something is less safe than they had previously thought, they will be able to reduce their consumption of that good. This assumes that they can process that information reasonably well at negligible cost. However, if there is evidence that they worry when they confront two separate risks whose total risk probabilities are lower than an alternative single risk, they are not processing in the manner that decision analysis would prescribe. Indeed, it could be argued that the information should not be provided, and the only options would be to ban, limit, or permit the substance.

Consider a slightly different example that makes the same point. Assume that there is a risk that is difficult to avoid but generally unrecognized, such as the risk of being hit by a meteor. It probably would not do the world much good to publicize the existence of that

If the dose-response curve is nonlinear,[64] there would be another reason to distinguish between additions to underlying risk and new risks. Without such attention we might miscompute the increments to risk. Assume, for example, that the dose-response curve has the logistic shape, so that the increment to risk for an additional dose at low levels is substantially smaller than the increment to risk for large doses. Two equally carcinogenic substances are under comparison; one is already prevalent in the environment. If equal additional doses of the two substances are under consideration, someone naively might decide that the increments to risk are the same, but they are not. The substance that is already prevalent would produce a much greater increment to risk because of the nonlinear relationship between dosage and risk. This relationship suggests that, other things equal, ruling out anxiety, and positing a dose-response curve that is bowed toward the horizontal axis, we should be more willing to accept an additional dose of substances that are not currently prevalent.[65]

H. Failing to Recognize That the Mean of a Distribution on a Probability is the Only Appropriate Variable for Basing Decisions

The food safety literature is replete with statements that the estimated risk level is X, but further information might reveal it to be .1X or 10X. Therefore, we should regulate this substance more stringently than another one whose estimated risk level is Y, when Y = X. There are two points to be made about this situation.

First, if X is the best guess, the center of scientific opinion, or the midpoint of the range of subjective probability estimates, X is likely to be an inappropriate variable for decision. The appropriate variable is the mean of this underlying subjective distribution. If

---

risk. Few people optimally would choose to change their lives by staying indoors, and the major consequence might be just that people would worry more.

64. At least two of the leading models on food safety suggest that the dose-response curve is nonlinear.

65. A more sophisticated analysis would look to across-substance complementarities or interactions in producing risks. It may be that substance A becomes a more powerful carcinogen or mutagen in the presence of substance B. If so, even if the dose-response curves for each of the substances, holding the quantity of the other substance fixed, have the required bowed shape, it would be better to concentrate on one substance or the other. Assume, for example, that over the relevant range the true total risk (with K a constant) is

$$R = K(1 + A^2)(1 + B^2).$$

The dose-response curves over a total lifespan for either substance shows increasing bodily damage. The choice is to find a total dosage of A plus B equal to 10. The optimal division would be 9.9 of A and .1 of B or vice versa.

the true value is distributed symmetrically about X, then the median and the mean will coincide. If, however, it turns out that the true value is as likely to be twice as big as X as it is to be half as big as X, then the mean will be substantially greater than the median. Yet we may find it simpler to think about medians of such asymmetric distributions. This is perfectly acceptable; we just have to learn how to translate from medians to means. One commonly employed distribution is the lognormal. It assumes that the log of a variable is distributed normally, which implies that a value twice as high as the best guess is as likely as one half as high as the best guess. This assumption may be appropriate when dealing with small probabilities on a subjective basis. If a variable is distributed lognormally, the relationship between the median or mode and the mean is described in the footnote.[66] This discussion suggests that regulatory procedures should be more conservative when a highly uncertain risk estimate has been provided, because we may have employed an incorrect summary statistic for that estimate, perhaps the median instead of the mean.

If we truly have an estimate of the mean value of a probability, then that is the correct number to employ. This is a central principle of decision analysis: Lotteries on probabilities are resolved on an expected-value basis. Thus, for example, let us assume that a substance has one chance in ten of being at zero risk level, eight chances in ten of having a risk of one in 10,000, and one chance in ten of having a risk of one chance in 1,000. The mean risk would be 1.8 in 1,000. If no further information could be secured, this substance should be treated identically with another substance that was known to bear a risk of 1.8 in 1,000.[67]

The skeptic might argue: This information on variability in risk levels is all very interesting, but it hardly makes sense to attempt to attach any numbers to food safety risks. Our models of extrapolation from high doses to low doses are highly inexact, to say nothing of the problems in extrapolation from consequences on

---

66.

	Lognormal Distribution Standard Deviation as a Factor of Mean				
	1	2	5	10	100
Mean as Multiple of Median and Mode	1	1.27	3.65	14.17	40,287

I would like to thank Donald Shepard for carrying out this calculation. The multiple equals $\exp[\frac{1}{2}\sigma^2]$.

67. *See* H. RAIFFA, *supra* note 57.

animals to consequences for human beings. The answer to this objection is: What other approach is better? It is irresponsible not to take advantage of as much information as is available. Although for some probability estimates the range of error may be significant, a body of thought has built up around decision analysis to deal with hard-to-assess, low-probability events.

Thus, for example, consider two alternative artificial sweeteners that are alike in cost, taste, and so on. Both sweeteners are subjected to animal tests at high dosages. One leads to tumors in X out of 10,000 animals, the other to tumors in 2X out of 10,000. On the basis of this information, it is still possible that the second substance has less tumor-inducing effect in man than in animals. But, if we decide it is worthwhile to take the gamble and permit the marketing of one of the sweeteners, surely it would be preferable to choose the first substance over the second.

This is but one example of a much more general principle. Even if we cannot estimate absolute risk levels accurately, we frequently may have information on comparative *expected* risk levels. We may not know how risky substance A is, but we may know that on average it is safer than B. What do we mean by average if only two substances are under consideration? If 1,000 such pairs of substances were under consideration and if the less risky was selected in each pair to make up a less risky group, then with high probability the total risk associated with the less risky group would be lower than that of the more risky group.

I. Assigning Excessive Credibility to the Choice of Extrapolative Model

There are at least three widely discussed models for extrapolating risk levels back from high to low dosages. Each model has its own theoretical underpinnings. Unfortunately, the three models may produce very different predictions about the risks inherent at low dosage levels.[68] It seems unlikely that substantial evidence will be available soon to enable us to make a solid scientific choice among these models. We must recognize that any regulatory procedure concerning food safety is a decision taken under conditions of significant uncertainty. The lack of sufficient scientific evidence about extrapolative models is one significant factor generating such uncertainty. One should not take false reassurance or ring false alarm through the choice of one model over another.

---

68. *See* 1982 FOOD SAFETY COUNCIL REPORT, *supra* note 3.

Perhaps one of the reasons why the debate over models has gained such prominence is the zero risk level requirement with regard to carcinogenic substances. So long as there is some chance that the one-hit model is correct, any substance that induces any malignancies in any concentrations in animals must be ruled out. There always will be some risk. One part of the problem, of course, is that the zero risk level is probably inappropriate as a criterion. It implies that there is no tradeoff with other valued variables.

VI. Conclusion

This Article presents one approach to the problem of assessing risks and benefits. Much of our discussion focused on health risks and benefits, for we believe that is the less familiar area in which to employ the prediction and bookkeeping procedures we recommend. The RU and QALY concepts outlined here as a means of tabulating the probability and severity associated with health risks represent just one approach. There has been an explosion of proposed procedures for tallying information on health risks. The most controversial aspect of our recommended procedure will be the fact that it aggregates risks received by different individuals. To the extent that it is possible to design the food safety regulatory system so that different resources flow to different individuals, this should create no difficulty. Each category of consumer could be considered separately. Unfortunately, rarely will such fine-tuning be possible. Securing health benefit A for individual 1 will imply that individual 2 suffers health risk B. Such implicit tradeoffs, however uncomfortable, are part of any pattern of government or nongovernment decisions affecting health. Still, they are uncomfortable tradeoffs. To some, the RU and QALY approach will have the virtue of treating all individuals as equals, in effect making health benefits and risks anonymous quantities. Others may feel, however, that this very element renders such methods of aggregation unacceptable. For individuals with this attitude, we recommend that they disaggregate by individuals or groups, but maintain a classification system that looks to both probability and severity.

Risk and benefit considerations already affect regulatory decisions for food, albeit in an unacknowledged and often haphazard fashion. Risk and benefit assessment should be recognized as a vital tool and systematically incorporated into the regulatory process along with considerable safeguards so that its primary purpose—to reproduce the choices the well-informed consumer would make for

himself—is not lost. This Article outlines some methodologies that might be employed and warns of some problems that might be encountered.

Whatever the difficulties inherent in the use of systematic tools for guiding food safety decisions, they are likely to be dwarfed by the irrationalities inherent in a system that pretends not to recognize certain tradeoffs that are a vital part of its working. Though a little knowledge may be dangerous, a pretense to having none is downright reckless.

This Article's analysis has focused on food safety regulation as it is viewed and might be reformed within our own country. The problem of food supply, as the tragic famines of recent years dramatize, is substantially more monumental for less developed nations. Many risk-bearing technologies have the capability to expand world food supply. Incidents such as the Bhopal poison-gas disaster, however, have reinforced the argument that developing nations should employ the same levels of stringency in risk protection as the wealthy, developed nations. The tradeoff is stark. As Samuel Stumpf phrased it, "The use of pesticides, fertilizers, and chemical additives appears to require a cost in the form of probable hazards to life, yet these products provide the most support for survival."[69] In determining how the world is to feed itself, it seems evident that we are confronted ultimately with the question of trading risks for benefits.

69. *See* Stumpf, *The Moral Dimension of the World's Food Supply*, 1 ANN. REV. NUTRITION 1, 25 (1981).