Default and Inference Options: Use in Recurrent and Ordinary Risk Decisions

James D. Wilson

Discussion Paper 98-17

February 1998

1616 P Street, NW
Washington, DC 20036
Telephone 202-328-5000
Fax 202-939-3460

Default and Inference Options:
Use in Recurrent and Ordinary Risk Decisions

James D. Wilson

Abstract

How "default options" should be used in health risk assessment divides the risk analysis profession. Some argue that these should be "hard": set by policy, generally biased to be "health protective" and requiring a substantial body of evidence to replace by decision-specific alternatives. Others argue that they should be science-based, identified by consensus of the professional community, replaced by whatever decision-specific information may be available to the analyst. This paper shows that both positions have validity, and that both are incomplete. Each kind of construct has a useful role to play, but in different kinds of decisions. Because the two are different, we suggest giving them different names, "default option" being assigned to the policy-based construct, "inference option" (NRC, 1983) assigned to science-based assumptions, etc. We develop a theory that explains why these two different kinds of construct exist, and comment on some of the implications.

"Inference options" constitute an integral part of human health risk assessment, providing practitioners with consensus theories, models, or parameters that can be used to bridge knowledge gaps in specific analyses. Because human health risk assessment is both considered "scientific" and employs scientific reasoning, inference options must be treated as priors in an empirical-Bayesian inference process. Decision- or case-specific information modifies each prior according to the reliability of this information, with conflicts resolved by a scientific, weight-of-evidence process. Inference options are science-based "best estimates" and evolve through consensus within the professional community.

"Default options" constitute policy-derived components of particular kinds of decisions, serving as instructions to analysts. Such use is appropriate when many very similar, nontrivial decisions are to be made by a particular agency. In these decisions, which we suggest calling "recurrent," value-judgments are prescribed in advance, usually by legislation; generally only two decision options exist and the decisions usually turn on judgments made by experts. Authority to make these decisions is often delegated (sometimes tacitly) to permanent staff members who have the requisite expertise.

Policy-based default options exist in part because delegation of decision authority carries risks for organizations; those to whom it is delegated may unwittingly make decisions differently from senior officials, and may thus in some way harm the organization. Thus, delegation of authority is always conditioned by various forms of controls, including limits on the authority. We postulate that defaults serve as one means to control delegation risk. (They also simplify decision-making and make it more consistent.) These default options need to be tailored to the policy ends served, which generally means that some or all will be biased. They must be "hard," with "departure from" them requiring a high standard of evidence and also assurance that choosing an alternative will still satisfy the policy ends of the decision process in which these are a part. Default options need to be developed in the same way as any other policy-implementation practices, including deliberations that engage those who will be affected by the decisions. Such deliberations do not always take place.

Key Words: risk, regulation, decision making, default options, EPA, FDA, science, policy, public participation, risk analysis

Acknowledgments

Special thanks are due to Ms. Ann Holbrow for deciphering for me the Superfund decision process, the understanding of which played an important role in the results reported here, and for editorial assistance. Thanks also to Robert Scheuplein, Richard Merrill, Antonia Mattia, Paul Kuznesof, David Hattan, and Mike Bolger for tutoring me on FDA decision making, Alwynelle Ahl for explaining exotic food-plant import decisions, Dick Hill for keeping me accurate about EPA's practices, and Bob Barnard, Dan Byrd, Terry Davies, Adam Finkel, Dick Hill, Fred Hoerger, Sheila Jasanoff, Lester Lave, Warner North, Kate Probst, Resha Putzrath, Joe Rodricks and Paul Stern for enlarging my horizons, challenging my partly-baked ideas, sharing their wisdom, and encouragement.

Table of Contents

Introduction 1
Theory: Nature of Scientific Inference; Recurrent Decisions and Their Analytical Support Needs 3
Results 12
Discussion 16
Conclusion 17
Appendix: Characteristics of Some Decisions 19
References 23

List of Tables

Table 1. Some "Components" and Corresponding Default Options 2
Table 2. Comparison of Food Additive and Food Contaminant Decisions 8
Table 3. Recurrent Decisions 8
Table 4. Default Options in FDA's Food Additive Approvals Decision Process 11
Table 5. Some Inference Options and Related Default Options 15

DEFAULT AND INFERENCE OPTIONS: USE IN RECURRENT AND ORDINARY RISK DECISIONS

James D. Wilson[1]

INTRODUCTION

One of the more contentious and difficult issues in human health risk assessment has concerned the nature and employment of "default options" (Barnard, 1995; NRC, 1994). This term denotes a kind of construct commonly used by risk assessors to speed and simplify analyses. While such constructs are generally agreed to be useful, even essential to the practice of risk assessment, certain uses are controversial.

Barnard observed that defaults actually come in two kinds:

> "– generic default assumptions chosen in advance as a policy matter on the basis of conservatism;
> – assumptions selected for consistency in a particular case to bridge a data inadequacy." (Barnard, 1995, p. 436)

We suggest here that reality is a little more complex than his formulation suggests, and agree that defaults are of two different kinds. We present below theories accounting for the origin and utilities of both kinds of defaults, together with some implications for policy making that flow from those theories.

Attention was first called to the existence of policy-based constraints by the 1983 NAS "red book" (NRC, 1983). This report called these constructs "inference options" and described them as choices made to bridge chasms of scientific ignorance:

> "The uncertainties inherent in risk assessment can be grouped in two general categories: missing or ambiguous information on a particular substance and gaps in current scientific theory. When scientific uncertainty is encountered in the risk assessment process, inferential bridges are needed to allow the process to continue. The Committee has defined the points in the risk assessment process where such inferences to be made as components. The judgments made by the scientist/risk assessor for each component of risk assessment often entail a choice among several scientifically plausible options; the Committee has designated these inference options." (NRC, 1983, p. 28; emphasis in the original)

---

[1] Senior Fellow, Center for Risk Management, Resources for the Future.

This NRC report goes on to describe how the options to be chosen may -- but need not -- be decided as a matter of science policy. It did not otherwise distinguish between consensus assumptions and policy-based inference options. Later the term "default option" came to be used instead of "inference option." Table 1 lists several "components" and, where they exist, corresponding default options. When there is no default option, it is left to the judgment of the expert analyst to arrive at an appropriate resolution of the questions raised in each component. Policy directives constrain only a fraction of these choices.

TABLE 1. Some "Components" and Corresponding Default Options

"Component" [a]	***Default***
• How should evidence of different metabolic pathways . . . be factored into a risk assessment?	• Responses observed in animals predict how humans will react.
• What mathematical models should be used with carcinogens to extrapolate from experimental doses to human exposures?	• The Linearized Multistage Model is to be used.
• How should exposure to more than one toxicant be taken into consideration in doing a epidemiology study?	• (no default)
• The responses from which dose group should be used as the basis for setting a Reference Dose (or other exposure standard)?	• If the valid study exhibiting the lowest LOAEL has a NOAEL, that NOAEL should be chosen.
• Should a particular study be considered valid?	• (no default)
• What value should be taken as the body mass of people in a particular exposure assessment?	• The mean body mass is 70 kg.
• What should be taken as the daily intake of water?	• Average water intake is 2 liters per day
• How long, on average, do people reside in a rental apartment?	• Mean residence time in rental apartments is 3 years.

[a] National Research Council. 1983. *Risk Assessment in the Federal Government: Managing the Process* (Washington, D.C., NAS Press).

Within the risk analysis profession, there is disagreement on two fundamental aspects of these default options:

- The input data and theories may be intentionally biased or they may be "best estimates."
- The options may serve as point of departure for further analysis, in which case they are expected to be modified by whatever case-specific information is available to the analyst, or they may be "hard" instructions -- meaning that the value, procedure, method of calculation, or whatever is to be used unless there is strong reason to do otherwise.

A recent National Academy of Sciences (NAS) Committee could not come to agreement on these two issues. This Committee's report (NRC, 1994) included as appendices two strongly-worded essays that described the poles of this disagreement. We conclude that neither position is correct, and that both positions have validity under different circumstances.

THEORY: NATURE OF SCIENTIFIC INFERENCE; RECURRENT DECISIONS AND THEIR ANALYTICAL SUPPORT NEEDS

The conclusions presented in this Discussion Paper derive from observations on the natures of the risk analysis profession and scientific reasoning, and on results of recent research into the characteristics of regulatory decisions in which defaults are used. We demonstrate that, as Barnard suggests, policy-based default options must receive different treatment from the consensus assumptions commonly used by risk assessors. These two kinds of constructs are different, and are to be used differently. In general, the latter (which we suggest be called "inference options"[2]) must be based on the best estimate of a theory, model, or parameter currently available; inference options function as points of departure for specific assessments, to be modified in accord with whatever situation-specific information is available to the analyst. (It is not appropriate to think in terms of "departing from" inference options.)

Further, we show that in certain cases, specifically in a kind of decision described below and termed "recurrent," it is appropriate to use *policy*-based "default options." Default options may be biased (depending on the needs of the policy being implemented) and should be "hard:" For analysts to "depart from" default options requires a high degree of certainty on the analyst's part that an alternative theory, model, or parameter would satisfy the policy needs as well as or better than the default option.

These conclusions follow from six major premises, the first three of which are well enough established to be granted axiomatic status. The latter three statements are presented here as postulates; some observations support them but they have not been extensively tested and their generality has not yet been demonstrated.

1) Risk analyses are done in order to organize information for policy decisions. In general, for each of these policy decisions many possible decision options will exist.
2) Complex risk analyses are often carried out by professional risk analysts; the members of this profession are usually trained as scientists, and the situations they analyze are commonly characterized by high degrees of uncertainty concerning the future consequences of decisions being contemplated.
3) Scientific reasoning is accurately described as empirical Bayesian inference.
4) Risk analysis is based on scientific reasoning.

---

[2] In accord with standard taxonomic practice we retain the first published name for this species of intellectual construct. Thus "inference option" is used here to describe the general phenomenon, and "default option" the policy-specific uses that will be illuminated below.

5) Some policy decisions are made through simplified processes, with decision authority delegated to staff who possess critical knowledge and skills. We call these decisions "recurrent." Truncated risk analyses are used in (some) such simplified processes.
6) Delegation of authority poses political risks to the delegator (e.g., Congress, senior agency managers) which are managed by constraining the options of those to whom authority is delegated.

For discussion of the first premise, see NRC (1996) and Morgan and Henrion (1989). The second comes from observation of risk analysts and comparison of their behavior with that described by Schön (1983) as characteristic of professionals. The remaining premises are discussed immediately below.

The Methods of Health Risk Assessment Derive from Scientific Reasoning; Scientific Reasoning Employs Empirical Bayesian Inference

Modern health risk assessment is almost always described as "scientific." Whether or not that encomium is entirely appropriate, much of the credibility of risk assessments derives from their being considered "scientific." Further, it is demonstrable that the founders and most current practitioners of human health risk assessment received their formal training in one of the sciences (usually biochemistry, pharmacology, or pathology). Many of the folkways of the profession, especially their accepted ways for drawing inferences and making arguments, come from the customs of chemistry and biology (rather than, for instance, law or the social sciences). This mode of thought can be seen in the guidance documents prepared by several regulatory agencies, including the U.S. Food and Drug Administration (FDA, 1982) and Environmental Protection Agency (EPA, 1986, 1996), the State of California (California, 1985), and others.

Scientific reasoning is characterized by building specific conclusions from general theories (Ziman, 1978). If an existing theory fails to rationalize certain observations, a scientist will develop an hypothesis that modifies the existing theory in order to accommodate to these data, and test the newly modified theory against the relevant observations (Howson and Urbach, 1989). When observations conflict, as is often the case, scientists weight the evidence from each according to their evaluation of the likelihood that each observation is correct. The result is a weighted sum of the observations.

This mode of reasoning is described as inference, specifically empirical Bayesian inference. In Bayesian inference, one starts with a theory or other mental construct or model of some phenomenon of interest (a "prior" to Bayesian logicians), and uses observations to modify that theory to arrive at a better description of the phenomenon. Howson and Urbach have described (1989) the identity of scientific reasoning and empirical Bayesian inference.

This process is flexible; it is subjective, and it proceeds, ideally in public, according to behavioral norms imbued during scientists' training[3] (Ziman, 1978).

When applied to specific instances, a scientist will use generic information -- broad theories, generally applicable data, etc. -- as a basis for inferring characteristics of the particular case. The particular information is used to modify the general.

For instance, asked to predict the response of humans to a particular chemical for which there is no direct human experience, scientists will first consider how humans respond to closely-similar chemicals, then modify that preliminary conclusion by considering results from tests of how animals respond when treated. The generic information in this case consists of the theory that relates chemical structure to effects in humans and the theory that describes the similarities and differences between humans and test animal species. The specific information used to modify the general consists of the chemical structure of the chemical at issue and the results of tests in animals. Depending on the degree of similarity between the structures of the instant chemical and others for which responses are known, and the nature of the animal responses, the analysts may be more or less confident in the predictions made. This confidence can be assigned a numerical value, although in practice it is seldom carried beyond subjective ordinal ranking (*cf.* NRC, 1996). Thus, not being able to dose people with poisons, we dose animals and use the results to predict how humans may respond, using a formula based on other general knowledge.

Origin of Inference Options as Consensus Scientific Assumptions

Risk assessors, as most professionals do, find themselves often doing analyses that closely resemble previous jobs. The profession, as do most, has developed practices that streamline these repetitive tasks. One of these concerns assumptions made during the course of an analysis. Information is always missing; commonly, the same or very similar information will be lacking in all similar analyses. To cope with this situation, the profession developed the science-based consensus-accepted assumptions that we call "inference options." In general, analysts using inference options need not justify their use, whereas use of other assumptions must be justified.

It is useful to think of inference options as a folkway characteristic of this profession. Their usage is very much a human construct, developed for a purpose. It evolved, and continues to evolve, from within the profession. No one sat down and invented the particular use of inference options in risk assessment, and no one sat down and decided what assumptions are to be used. These custom-sanctioned choices appear to have been arrived at by a consensus of those working actively in the field. Some of them date back at least fifty years.

Because the customs of science dictate that assumptions be "best estimates," inference options generally are and have been just that -- representations of the description of the world in which the greatest degree of confidence can be placed. (Thus, if scientists believe that a

---

[3] Competition to receive credit for being both first and correct, and the norm that all results be made public, keep this process reasonably open and honest -- but not perfectly so!

"probit" mathematical model best represents all exposure-response curves, that would generally be the consensus assumption; similarly, if the mean body weight for all Americans were 74.3 kilograms, that would be the assumption.) However, there is evidence that from the beginning of the profession, in the early 1950s, some of these assumptions customarily were not "best estimates" but biased estimators (Doull, 1995).

I believe that biased estimators were used because analysts at that time recognized they were in fact making policy decisions -- protecting public health -- and not just analyzing the probable outcomes of alternative decision options. It would be very difficult, now, to document the motives of professionals acting in a milieu in which the separation of science from policy was not an issue; for one thing, almost no one active at that time remains alive. However, the legal and statutory context in which they operated, based on the 1938 amendments to the Food, Drug and Cosmetic Act (Merrill, 1988) and a close working relationship between senior technical people in FDA and relevant Congressmen (Doull, 1995; Scheuplein,1996), would have encouraged these professionals to take policy initiatives. It appears that the 1958 food additives amendments, which really gave FDA the authority to approve new food additives (other than food colors) before they were marketed (as opposed to removing unsafe ones from the market), wrote into law the practices that had evolved informally during the previous decade.

Yet at this time the use of biased assumptions was customary and not required. As noted, science creates new knowledge by using new information to modify existing knowledge. The early professionals felt quite comfortable abandoning customary usage as they perceived the data required (Doull, 1995).

We have shown that the development of customary inference options is a natural consequence of a science-based profession making repetitive analyses in which decision-specific information is often consistently missing. We know that sometime between the appearance of the first inference options, in the mid-1950s, and about 1980, some of these inference options changed from being customary guides to policy-based, required procedures -- "default options" (FDA, 1982, NRC, 1983). In the following section we present a theory that accounts for this "hardening" of inference options.

Recurrent Decisions

The recent National Research Council report, *Understanding Risk: Informing Decisions in a Democracy* (NRC, 1996**)** included a brief discussion of different kinds of risk analyses, divided by the nature of the work product and scope of the anticipated decision, called a "risk landscape." Two of the kinds identified there have relevance for understanding inference and default options: those called "unique, wide-impact" and "routine, narrow-impact." To simplify terms a bit, this paper uses the phrase "ordinary risk assessment" to describe the kind of analysis appropriate for unique, wide-impact decisions.[4]

[4] In the U.S. Federal government, this kind of decision often results in a "major rule" -- a regulation imposing costs of more than $100M on the U.S. economy. This kind of decision has received extensive study. Executive Order 12886 requires an analyses of costs and benefits anticipated from any regulations consequent to such decisions; often the projected benefits are based on analyses of anticipated health impacts that are estimated using health risk assessment.

Decisions identified by the NRC report (1996) as "routine, narrow-impact" have received little attention from students of decision processes. The name chosen is not apt: the decisions themselves are frequent but not necessarily routine, and often have substantial impact. One example: approval of a food additive such as a broadly-effective artificial sweetener may affect hundreds of millions of people and can lead to very large transfers of income from one industry to another (in this case, from sugar growers to the additive's manufacturer). Yet the decision process and the supporting analyses are strongly truncated, compared to other major regulatory decisions (see below). We suggest "recurrent" as a more apt term for this kind of decision.

Recurrent decisions share a critical characteristic which enables their supporting analyses to be truncated. Those making these decisions do not try to resolve conflicting societal values -- and in fact are not permitted the option of considering such values. Any conflicts over values were resolved in establishing the legal and regulatory context in which the recurrent decisions are made. In addition, in all such decisions only pre-determined options are available; usually there are only two such options.

For example, as part of Superfund the U.S. Congress instructed the Environmental Protection Agency that hazardous waste sites which, on examination, could be judged safe were to receive no "remedial action." Thus the decision process for selecting remedies under Superfund includes an early "no further action" decision node. (Few decisions have gone down this path.) This decision is based on a "Baseline Risk Assessment" that shares many characteristics, including use of default options, with food additives decisions. Similarly, Congress instructed EPA, in the Toxic Substances Control Act, not to act on "Pre-manufacturing Notices" sent to inform the Agency of an intent to introduce a new chemical into commerce, unless the intended use of that new substance was not judged safe. Congress -- at least implicitly -- has judged food additives to have value, and has left to the marketplace the task of calibrating that value; most food additives have the effect of reducing the cost of food, through reducing spoilage or preserving palatability. (Those that affect only esthetics, notably food colors, are treated differently.)

To illustrate further the concept of a recurrent decision process, Table 2 summarizes some characteristics of two closely-related decisions carried out within one organization, the FDA's Center for Food Safety and Applied Nutrition. These differ in a number of important regards, including the kinds of risk and decision analyses done to support each decision, and the explicit weighing of values (cost and abundance of the food supply) against uncertainty of safety in the case of contaminants.[5]

[5] Acceptable levels of contaminants are largely set by balancing the increasing cost of smaller levels against the uncertainty that such levels are safe. If it were possible to ascertain an absolutely certain level of safety, all contaminant action levels would be set there; in practice, such certainty does not exist, and the uncertainty in the judgment of what constitutes safety is weighed against cost of further reduction.

TABLE 2. Comparison of Food Additive and Food Contaminant Decisions

Additives	*Contaminants*
Numerous (>100/y)	Few (< 5/y)
Decision authority effectively delegated	Decision authority not delegated
Simple decision logic	Complex decision
Information limited	Information not limited
Values not considered	Values considered

In all these decisions, the choice between approval (or "no action") and rejection (or "do further analysis") turns on an analyst's considered judgment about whether or not a projected exposure can be considered safe. For each of them, the agency responsible for the decision employs decision and analysis procedures that simplify the process, reducing both the alternatives available to decision makers and the information they may use in making decisions. We have identified seven decision processes that fit the definition implicit in the preceding discussion; they are listed in Table 3.[6]

TABLE 3. Recurrent Decisions

Decision	*Agency*
• Application to register new food additives	• U. S. Food and Drug Administration
• Review of Pre-Manufacturing Notices under TSCA	• U.S. Environmental Protection Agency
• "No action needed" under CERCLA (Superfund)	• EPA
• Setting pesticide residue tolerances	• EPA
• Establishing "safe harbor" values under "Proposition 65"	• California EPA
• Identifying "priority substances" under the Canada Environmental Protection Act	• Health Canada
• Submission of toxicity information under TSCA §8(e) (practiced 1992-1996).	• Monsanto Company
Some possible or potential recurrent decisions	
• Applications to approve import of new kinds of produce	• Animal and Plant Health Inspection Service
• Establishing Permissible Exposure Levels (PELs) in workplaces (potentially a recurrent process)	• Occupational Safety and Health Administration
• Applications to register new securities	• Securities and Exchange Commission

[6] Note that in the Monsanto Company process identified in Table 3, the decision did not turn on a judgment of safety but on a judgment that a particular test result was valid and met submission criteria.

Recurrent Decision Processes Require Constraints on Decision Makers and Supporting Analysts

In essentially all organizations those who legally hold decision authority delegate some of it to others within the organization. We observe that this authority is almost never delegated without conditions, and suggest that default assumptions form part of the structure used to condition delegation of authority in recurrent decision situations.

Senior executive time and attention constitutes one of the most scarce of resources in almost all organizations. Successful senior executives almost always delegate to others authority to make those decisions that do not require his or her special knowledge and perspective.[7] However, delegation of authority carries risks: others with less knowledge and a different perspective may make decisions that harm the organizations. Senior executives must balance this possibility against the desire to focus their attention on issues where their knowledge and perspective are essential. Commonly, the means used to manage the risk and obtain the desired balance is application of controls on the decisions. For instance, low-level managers may purchase supplies, but only of less than some set value, and specified records of such purchases must be kept.

In for-profit organizations, financially-based controls are almost universally used to manage delegation risks. The practices are highly evolved, forming part of the practices used to track the organization's overall performance. In general, delegation is hierarchical, with authority posing the greater risks being delegated to fewer people. Traditional management control practices focus almost exclusively on operations for which output (performance) can be quantified.

Of interest here is governmental organizations. We observe that in such organizations authority to make policy decisions is commonly delegated. Almost the entire edifice of the modern executive branch represents a delegation of Constitutional authority from the Congress to the President, or directly to some executive-branch agency. These various agencies are charged with correcting various "market failures" such as pollution from manufacturing operations. Congress has consistently limited the discretion of those to whom it delegated this authority, both in setting conditions on the nature of decisions that can be made, and in establishing general decision-making procedures under the Administrative Procedures Act (US Code, 1966).[8]

In the Executive Branch, the risks to senior executives ("political appointees") and to the agencies they manage posed by delegation are largely political rather than financial. Inappropriate decisions by delegatees may harm the organization's public approval, the political standing of the executive, etc. The public bears financial risks that also accompany this delegation, and these risks are managed with conventional management controls. But

---

[7] Inability to delegate often leads to failure of senior executives in both public and private organizations.

[8] In recent years, Congress has shown a tendency increasingly to limit authority and constrain discretion of the regulatory agencies. *Cf.* the Food Quality Protection Act of 1996 for an example relevant to this discussion.

conventional controls have little ability to manage political risks. Other delegation-risk management practices are required.

Often much of the delegation risk is managed by restricting delegated authority to officials who serve at the pleasure of the chief executive and can thus be readily dismissed. The threat of removal provides some control, even if it is only *ex post facto*. However, an *ex post facto* strategy does not provide useful management of delegation risks when the delegatees are career staff. These people can be removed from their positions only by high-cost means.

Because political positions are few in modern bureaucracies, it is logical to postulate that organizations which face a large number of similar decisions will delegate authority to make those decisions to career staff. Further, these organizations will also seek and find means to condition delegation of this authority, imposing controls and constraints to minimize risks associated with this delegation. Among these means are instructions to buck sensitive decisions back upwards or some decision review process; another is formal decision criteria (some of which are imposed by Congress).[9]

Default Options Form a Means of Constraining Decision Authority

We suggest here that another way to condition decision making takes the form of default options -- instructions as to kinds of information, assumptions, and evaluative models that may be used. Obviously these constructs exist (NRC, 1983). As an example, consider the constraints built into the food additives approval process. Petitions for use of food additives are to be approved if FDA finds that such use meets the test, "reasonable certainty of no harm when used as intended" (Merrill, 1988). This decision criterion clearly anticipates exercise of expert judgment.[10] However, the framework for exercising this judgment is quite rigid. An "acceptable daily intake" (ADI) value is compared with an "estimated daily intake" (EDI); EDI less than ADI implies that the prospective use is safe.

Development of the ADI begins with a standard set of animal test data; from this set one particular test datum is to be chosen as the basis for calculating the ADI. That datum consists of the largest "no observed adverse effect level" (NOAEL) from a valid and reliable test, generally the one which exhibited the smallest "lowest observed adverse effect level" (LOAEL). The judgment on reliability is made based on the suite of all results. (In the rare cases in which no test datum meets these conditions, another procedure is invoked.) This particular test NOAEL is divided by 100 to yield the ADI (Hattan, 1996). The estimated daily intake (EDI) is derived by taking the 90th percentile of intake estimated from a dietary profile

---

9 Note also that implicit delegation occurs very often, creating ambiguity that can be both frustrating and useful to the delegatees.

10 Note that FDA commonly acts as though "no harm" means "no discernible harm." For instance, action is taken on botulism whenever there appears a cluster of cases large enough to be reliably detected (> 3); in this case the significant risk level leading to action is near 10^{-8} per year. It is a standard tenet of management that action should be taken only when the result of that action can be discerned.

analysis built on an extensive data base that describes distributions of intake in the entire U.S. population for a wide variety of foods[11] (Kuznesof, 1996).

The default options in this process -- quite clearly policy-based instructions -- are listed in Table 4 together with some of the most important of the many science-based assumptions which undergird them.[12]

TABLE 4. Default Options in FDA's Food Additive Approvals Decision Process

Default Option	*Scientific Bases*
1. Use the largest "No Observed Adverse Effect Level" (NOAEL) from the reliable test with the smallest "Lowest Observed Adverse Effect Level" as the basis for judging safety.[a,b]	♦ Responses in animals predict how humans will respond, in general and with reasonable accuracy. ♦ The doses at which different kinds of response occur are highly correlated (i.e., it is unlikely that any unobserved response would occur at exposures smaller than the lowest observed adverse response.)
2. Obtain the Acceptable Daily Intake (ADI) by dividing the indicator NOAEL by 100.	♦ The spread of inter-individual variability in humans lies between 1000- and 10,000-fold. ♦ Choosing the "most sensitive test" provides a sufficient margin of error to account for uncertainties in interspecies extrapolation.
3. Use the standard database of food consumption[c] to predict intake of the proposed additive.	♦ The standard database reliably represents actual consumption by the U. S. population, including significant subpopulations. ♦ The standard tools for drawing inferences from frequency distributions are valid.
4. Take the 90th percentile of the predicted consumption as the "reasonable worst case" intake to use in judging safety (the "Estimated Daily Intake").	♦ The maximal consumption rate occurs over a relatively short period (i.e., in young males between ages 16 and 22), diminishing thereafter, so that any individual's actual lifetime consumption rate will be less than is predicted by the 90th percentile of the intake distribution.

FOOTNOTES TO TABLE 4

a. Note that by statute no substance found "to induce cancer when ingested" can be approved as a food additive.
b. If no such NOAEL is found in an otherwise satisfactory series of animal tests, an ADI may be derived from the smallest LOAEL from an appropriate study, by dividing that value by 1000. The experts who evaluate the animal test data enjoy considerable latitude in judging what constitutes an appropriate study.
c. This database is "TASDIET" (maintained by Technology Assessment Systems, Washington, D.C.).

[11] The system, TAS-DIET, was developed and is maintained by Technical Assessment Systems, Inc., Washington, D.C., and includes results of surveys of food consumption conducted by USDA and a proprietary data base of food contents. FDA (and others) license its use from Technical Assessment Systems.

[12] Not all of these are widely recognized or understood. In particular, the tenfold "uncertainty factor" employed to "account for variability within the human population" is often interpreted as implying that this variance spans only an order of magnitude. In fact, since this tenfold divisor is applied to an exposure level corresponding to <4% of the population,(Gaylor, 1989) and typically NOAELs are observed to be less than one-third the median (50%) response, if the distribution of susceptibilities is symmetrical then this implies the width of this distribution to be >1000-fold. Also not widely understood is the assumption that choosing the most sensitive response from a spectrum of observed responses is very likely to protect against all responses.

Note that these default options are conditional, in the sense that analysts responsible for evaluating data packages may choose alternate instructions in certain cases. For instance, the ADI may be defined as LOAEL / 1000, if in the analyst's judgment this provides adequate protection.

RESULTS

Four results follow from the theory and observations described above:

1. In general, inference options are to be scientific "best estimates".
2. In general, inference options are to be modified by whatever case- or decision-specific information is available.
3. There exist situations in which an agency may find it useful formally to establish truncated decision-making processes for repetitive decisions ("recurrent decision procedures"), as part of which it may be appropriate to employ policy-based default options in place of some inference options.
4. Default options may be biased, if doing so is appropriate to policy needs, and "departing from" these options must be constrained so that policy objectives will be met.

We examine each of these in turn.

1. In general, inference options must be scientific best estimates

We use the term "inference option" to denote the consensus assumptions customarily used within the risk analysis profession, as bridges over common information gaps. Inference options evolve from professionals' practice, and, in that sense, belong to the profession. That these should be "best estimates" follows from the premises that the purpose of the analysis is to inform decision making and that risk assessment is based on science, and the observation that, in general, policy decisions always have many possible decision options. Health risk assessment relies on scientific modes of thought for the reliability of results and credibility that obtain when these modes are used.

This paper is concerned with policy decisions that have been delegated by a legislature (i.e., Congress or an equivalent) to an executive branch. While this delegation is observed to be conditional, in the sense that the statutes under which this delegation occurs place some limits on the behavior of the delegatees, we observe that in general some range of decision options is available to the executive-branch people who make decisions. The "no action" option is always available, and the range of other possible options will be limited primarily by the ingenuity of the decision maker and those employed to assist her or him. Thus in the general case, only rarely will the conditions required for recurrent decision processes be met. That is, very few decisions will have only two options, one of which is a default "safe" condition. Further, almost all such decision situations include value conflicts that must be considered (if not balanced) (NRC, 1996).

In this general case, a risk analysis must describe for the decision maker the dangers to health and the environment that either may be diminished by action taken or arise as a consequence of these actions. The possible effects of the subject hazards can seldom be the sole considerations affecting the choice of decision options; decision makers must weigh the gains and losses expected to follow from the different choices. Thus, obviously, the decision maker needs to know as accurately as possible the differences resulting from different options. Only by accident can analyses based on and deriving from policy-biased assumptions provide the accuracy necessary for well-informed decision making.

Whenever scientific information is available, it provides the most reliable description of the tangible world useful to a decision maker (Ziman, 1978). Thus to the extent possible risk analyses are based on science. It follows directly that for risk analyses generally, initial assumptions and other inference options must represent the best scientific estimate of the theory, model, or parameter to be used to bridge a chasm of ignorance. No other option will provide as accurate a result for decision makers to use. (By "best" we mean that parameter, formulation of a theory, or other proposition in which the relevant expert scientific community has the greatest confidence.)

Note that sometimes analyses of different decision alternatives will be similar enough that much of the input data and many of the starting assumptions will be common to all. In such cases, the accuracy of the common inputs, inference options, and so on will not matter to the decision maker. In such cases, the change *consequent to the decision* will be what is important. In these cases, biased or otherwise inaccurate inference options can often be employed without affecting the utility of the results. However, analysts can not assume that biased inference options are always appropriate. Each case, each situation must be examined and appropriate choices made.

2. In general, inference options are to be modified by whatever case- or decision-specific information is available

This result follows from the previous result and the premise that scientific reasoning is appropriately considered to be empirical Bayesian inference. The choice of input data, model or theory represents only part of the "defaults" problem: the issue of "departing from" defaults -- use of decision-specific information in place of the normal assumption -- often looms larger. As with the choice of options, the proper course in using decision-specific information depends on the type of decision to be made. Here we continue to describe the situation regarding ordinary risk analyses.

Scientific thought is Bayesian, i.e., new information modifies existing knowledge, new theories incorporate much of the content of theories they replace[13] (Howson and Urbach, 1989). It follows that any analysis done for policy purposes that purports to be "scientific"

[13] Thomas Kuhn wrote about scientific "revolutions," (1949) but his paradigmatic revolution, Einstein's Relativity Theory supplanting Newton's theory of planetary motion, actually illustrates Howson and Urbach's point: Newtonian dynamics are included in relativity as a special case.

must employ similar reasoning. In fact, the prototype risk assessment processes, including FDA's evaluations of food additive safety, employ exactly this kind of reasoning. Decision-specific or situation-specific information is used to modify both scientific and policy defaults if experts' judgment suggests that doing so is appropriate, given all the information in hand. True, this application of expert judgment represents qualitative Bayesian reasoning, not its quantitative application. However, there exists no reason in principle why much of it could not be quantified, were the investment in necessary tools and training to be made.

Thus in the general case, inference options should be treated as Bayesian priors -- general, science-based descriptions of phenomena that are to modified in specific cases by whatever case-specific information is available. If apparently conflicting decision-specific data are at hand, then the analyst must weight them according to his judgment about their relative reliability. The result of the analysis is then the prior altered by the weighted-average of all decision-specific data.

Plainly, in this conclusion lies the resolution to the conundrum posed by "departing from defaults," raised and left unresolved by *Science and Judgment in Risk Assessment* (NRC, 1994). At issue here: if defaults are set by policy, choosing to use an alternative is a policy decision. The NRC Committee responsible for *Science and Judgment* got itself all tangled up in trying to assure separation of "science" from "policy." It ended up recommending that EPA adopt formal policy criteria for deciding how much evidence is enough to "depart from a default," but was unable to suggest what criteria might be appropriate! In fact, this suggestion violates both sense and reason. It ignores why experts are hired to evaluate scientific information in the first place: Their expertise allows them to apply reasonable weights to different pieces of data use to update inference options. In hiring such people, policy makers choose to delegate some of their authority in order to obtain the benefit of specialized experience and knowledge. They take some risk in doing so: In some cases, the personal values that inescapably color experts' evaluations may lead these experts to conclusions implying policy choices different from those believed appropriate by the policy makers. The way to manage this risk is to require full disclosure from analysts; risk characterizations need to include discussion of the impact from alternate choices (NRC, 1996).

3. There exist situations in which an agency may find it useful formally to establish truncated decision-making processes for repetitive decisions ("recurrent decision processes"), as part of which it may be appropriate to employ policy-based default options in place of some inference options

This result follows from the first one and the identification of recurrent decision processes as including policy-based limits on inputs, etc. We suggest that this circumstance emphasizes the utility of defining "default options" as a special case of inference options (NRC, 1983). Since default options are inference options or other similar constructs that have been fixed by a policy decision, their use both simplifies and constrains the analysis, permitting delegation of decision authority to people not susceptible to common executive sanctions while minimizing the risk that their decisions will deviate from the policy goals of the executive.

4. Default options may be biased, if doing so is appropriate to policy needs, and "departing from" hard defaults must be constrained so that policy objectives will be met

This result follows from the previous and the observation that successful application of truncated analyses to recurrent decision processes requires that the process consistently yield decisions congruent with policy aims of the legislature and executive.

Of interest here are decisions that are recurrent: frequent, usually limited by policy to only two options, one of which is a default "safe" condition, and in which the decision maker need not resolve conflicted values. We noted that government organizations facing recurrent decisions may conditionally delegate authority to career employees, establishing limits on the decision options that may be considered, and the models and information that may be used in the course of a risk analysis.

Under these circumstances, it becomes entirely appropriate for inference options to become policy-based, "hard," and often biased. It is these options that deserve to be called "defaults" as we now use this term. Table 5 lists a number of inference options and corresponding defaults. Two characteristics distinguish default options from ordinary inference options. Defaults are often biased, whereas inference options need to be "best estimates" of whatever theory, model, or parameter is of moment. Defaults must be "hard": other values, theories, etc., used only if a strong body of evidence supports an alternative, *and the policy objectives of the process will still be met*. Inference options must be modified by whatever decision-specific information is available in carrying out the analysis.

TABLE 5. Some Inference Options and Related Default Options

Inference Option	***Default Option***
• Observing an adverse response at some exposure in an animal experiment implies that it is more likely than not that humans will respond similarly at a corresponding exposure.	• Observing an adverse response at some exposure in an animal experiment implies that humans will always respond at a corresponding exposure.
• Compared to the exposure eliciting a particular response in some animal species, humans will respond at an exposure equal to the 3/4 power of the relative body masses of the two species.	• Humans will respond at 1/10th the exposure observed to cause a response in rodents. (For mice, this implies no bias; for rats, it overstates toxicity by about 3-fold.)
• With data available from tests of a substance in several animal species, the data most likely to be predictive of human response will be that from the species whose metabolism of the substance most closely resembles that of humans.	• Choose the response observed to occur at lowest exposure as the basis for setting exposure standards and advisories.
• If data are available on a particular response rate for a series of tests on one substance in one species of animal, the group median exposure is the best predictor of that at which humans will respond.	• From a series of tests, choose as the basis for further analysis the response observed to occur at the smallest exposure.
• Effects of concurrent exposure to two or more substances can be predicted only if detailed knowledge is available; at very small exposures, adding individual responses provides an upper bound on the actual response.	• Responses of concurrent exposures are taken to be additive unless there exists strong information to the contrary.

DISCUSSION

Biased (default) options should not be employed when the results will be inputs for cost-benefit calculations

We concluded that biased default options may be appropriate as parts of recurrent decision processes, but should be used as part of a general analysis only if it is very clear that this use will not affect the result. On occasion, results of risk assessments are intended as input to economic analyses. Here, defaults or biased inference options must be resolutely eschewed, lest the economic analysis be distorted. In addition, it is essential in such cases that a thorough understanding of the uncertainties and effects of alternative assumptions be part of the analysis.

Default options will not be the same for all recurrent decisions

It should be obvious that assessments tailored to different recurrent decision processes may well, indeed perhaps should, differ one from another. Ideally, the suite of all defaults in a process should be tuned so that the aims of the decision are met, within the limits of uncertainty set by the statutory and other policy objectives of the regulatory organization. There exists one bit of evidence suggesting that this tuning occurs. Compare the processes used by FDA to approve food additives with those used by EPA Pesticides Office to set pesticide-residue tolerances. In both instances, all the underlying assumptions are identical; both FDA and EPA rely on a standard method for estimating exposure, which is expressed as a distribution of mean daily intakes over the population. However, FDA chooses the 90th percentile of this distribution as its estimate of the reasonably-worst-case exposure, while EPA chooses the 99th percentile. Thus, it is more difficult to prove safety for a pesticide residue than for a food additive. Also differing, between these two, are the values placed by the legislature and executive and thus, implicitly, by the public, on the two different kinds of "contaminants" in food. Additives have direct value, improving keeping, processability, or otherwise reducing the cost of food. Pesticide residues, however, have no direct benefit *in* food; the pesticides themselves also lower the cost of food and are considered valuable, but society clearly would prefer to gain the benefits of smaller food production costs without paying the price of residues. One can plausibly and defensibly interpret this difference as reflecting the different values imputed to additives and pesticide residues. (However, there is no indication at all that that this difference represents a conscious, well-analyzed decision on the part of the two agencies!)

Other decisions and the defaults appropriate to them will also differ, because decisions differ in statutory context, legal decision criteria, information available, hazard of concern, dominating uncertainties, agency custom and culture, and so on. For another example, the defaults appropriate for decisions on Pre-manufacturing Notice (PMN) review carried out by the EPA Office of Toxic Substances cannot be the same as those which would be appropriate for establishing Permissible Exposure Limit (PEL) values by OSHA, even though protection of workers from damaging chemical exposure is a primary concern in both cases. The PMN decision is "approve / require further review" and typically very little information is available.

The PEL decision establishes a legal criterion for evaluation compliance with health and safety standards, and typically considerable substance-specific information is available. Defaults appropriate for the PMN process should be more "conservative" than those which would be appropriate for a PEL decision because of the greater uncertainties.

Finally, note that in the PMN review one default option used is a particular structure-activity computer program to predict the toxicological properties of substances. Such a default would be anathema to the food additives or pesticide residue tolerance decision process.

Despite this clear evidence that different decision processes demand different defaults, it is commonly believed that one size of defaults will fit all. For instance, the sections discussing inference / default options in both the National Research Council's "red book" (NRC, 1983) and *Science and Judgment in Risk Assessment* (NRC, 1994) very clearly assume that risk assessment is essentially the same for all kinds of decisions. Both recommend government-wide adoption of a uniform set of defaults. Yet neither NRC Committee examined this assumption. This same assumption pervades EPA's writings (*cf*. EPA, 1986).

This belief is wrong. Rigid consistency across different decision types is not only not necessary, it may compromise achievement of policy goals. Some who would reform regulatory processes want to drive government agency practices toward identity and consistency. In the context being considered here, this drive seems misplaced.

Developing default options for use in recurrent processes

This observation raises an important issue regarding the origin and derivation of default options. Theory (NRC, 1996) quite clearly states that policy decisions such as these represent should be taken only following appropriate deliberation including those who will be affected by the outcome. Both FDA Center for Food Safety and Applied Nutrition and the EPA Office of Pesticides Programs at least arguably engage those affected as policies potentially impacting their decision processes are developed (although both have much yet to learn about doing so effectively and efficiently).

CONCLUSION

Within the risk analysis profession disagreements about default options have been framed in terms of a single proper or appropriate means of their development and use. The results presented here demonstrate that this framing distorts and retards understanding of the nature and utility of default and inference options. Clearly these two sorts of construct are different, and should be employed differently. Equally clearly, default options must be chosen so as to shape analyses so that policy goals are achieved. Instead of there being just a single appropriate set of default assumptions and options, there are many -- as many as the decisions in whose analyses they are employed.

The analysis presented here shows that when defaults comprise part of the constraints employed to manage risks in delegated policy implementation decisions, it may well be appropriate for some or all to be biased, with substantial information hurdles required for

"departure from" them. (Scientifically, this might be better considered as these policy defaults being intentionally-biased Bayesian priors, with the hurdles identified with the degree of dispersion in each prior.) Results from the National Research Council's *Understanding Risk* (NRC, 1996) study suggest quite cogently that agencies which set up decision processes employing policy defaults should assure that parties affected by the decisions flowing from such processes be involved in adequate deliberations during their design, and that the degree of bias is appropriate to the decisions' needs.

Conversely, the inference options commonly used to bridge information gaps in general risk analytical practice should not be biased. The analyses which use general practice guides commonly inform singular decisions. Although authority to make these decisions may be delegated, the delegation is not to career technical staff but to dismissable appointees. The risks posed by this delegation can be managed by more conventional political means. These people's political judgment and acumen constitutes important input to these decisions, and biasing the technical information available to them is likely to lead to less-satisfactory decisions. More practically, it is impossible to know in advance how much bias will always be appropriate for all the many scientific default assumptions which exist; thus arriving at a single set useful for all general analyses represents a classic exercise in futility.

To sum: policy-based, biased, "hard" inference options, called "default options," may form useful components of truncated decision structures for recurrent decisions. They both simplify the analyses needed for these decisions and provide some measure of control over the process, thus helping to manage the significant delegation risk that goes along with these decisions. In developing or modifying recurrent decision structures, agencies should take pains to involve affected people in deliberations. Experience suggests that doing so leads to better decision structures, more readily accepted by those affected.

APPENDIX: CHARACTERISTICS OF SOME DECISIONS

Registering new food additives (U. S. FDA)[14]

Congress has granted the Food and Drug Administration authority to sanction use of chemical substances called "food additives" if their use is judged safe. FDA has adopted a set of standardized procedures and practices to judge safety. These procedures rely on standard methods of analytical chemistry to demonstrate that a prospective producer of an additive knows what is to be produced and how to determine its presence and estimate its concentration in food. From this information FDA can estimate how much of an additive will be ingested by consumers. FDA relies on data taken from tests in experimental animals (mainly rats and mice) to identify a level of intake of the prospective additive that can be considered safe. Although decision authority formally resides in the office of the FDA Commissioner, in practice effective authority has been largely delegated to a scientist heading the office responsible for safety evaluation. These decisions are frequent, numbering a few dozen to a few hundred per year over the past forty years. Only two decision options are available: if use of the additive can be found safe, FDA is to approve the petition for approval; if not, it must be rejected.[15] The decisions made by FDA are reasonably effective (concerns are raised about their efficiency); in general, those interested in and potentially affected by the decisions seem to be satisfied with the process.

Review of Pre-manufacturing Notices (PMNs) under TSCA §6 (U.S. EPA)

Section 6 of the Toxic Substances Control Act (TSCA) gives EPA authority to prohibit manufacture (or import) of "new chemical substances" if it finds that such manufacture "poses an unreasonable risk to human health or the environment." The statute also gives EPA the authority to maintain an Inventory of Chemical Substances in Production; substances not listed therein are considered "new" for purposes of the Act. Manufacturers who anticipate beginning manufacture of a new substance are required to notify EPA 90 days before manufacture may commence. They must also submit, as part of the notification, information identifying the substance, describing its manufacturing process, and any information they may possess that bears on its propensity to threaten human health or the environment. EPA may decide, within the 90-day period that the substance may pose an unreasonable risk, based either on information available or on the lack of such information, in which case it may prohibit manufacture. Unless the Agency makes such a decision, manufacture may commence. The nature of the health and safety information submitted varies widely; in general, little is submitted for substances that are high polymers, more for

---

[14] Closely similar decisions are made in many OECD countries, and by the United Nations FAO/WHO Joint Expert Committee on Food Additives.

[15] FDA has no authority to judge or otherwise consider the benefits offered by a prospective additive. The decision criterion on safety is not impossibly rigid: "reasonable certainty of no harm when used as intended," and any proposed additive that meets this hurdle is presumed to offer sufficient benefits to justify approval. In essence, Congress relies on the market to judge the benefits of each safe additive.

substances expected to enter the environment. EPA relies heavily on sophisticated correlative methods to predict both possible consequences of manufacture (adverse effects from exposure to the substance and the likelihood that exposures leading to such effects will occur). The authority to make the decisions on PMNs is formally delegated to career staff; in practice, decisions are made by individuals well down in the formal hierarchy. The regulated community generally evinces satisfaction with the overall process (although individual decisions are often argued); outside of the Agency and that community there seems to be very little awareness of how these decisions are made, or their overall effectiveness.

Superfund "no action required"

The Comprehensive Environmental Release and Liability Act of 1980, "Superfund," gives EPA the authority to require or take actions necessary to prevent releases from industrial activities, past or present, that may harm human health or the environment. Included are provisions for "remedies" for dangerous contamination. At an early stage in the "remedy selection" process, a decision is made whether or not further cleanup is needed. (Frequently, "removal actions" are conducted before the remedy selection process begins; these deal with what are considered imminent hazards.) Authority to make this decision is formally delegated to Regional Administrators; informally, the decisions are effectively delegated to Remedial Project Managers. In total, EPA will make more than 1000 remedy-selection decisions. At the "no further action" stage, only two decision options exist: to declare cleanup unnecessary or to proceed. The primary technical input to this decision comes from a "baseline risk assessment." The Project Managers generally do not conduct these assessments, and few appear to possess the technical qualifications to understand them in any detail. Surprisingly few conditions are imposed on the Manager, considering that expenditures at sites commonly total several tens of millions of dollars. Almost no one evinces satisfaction with the process (although since this decision is embedded in a very controversial larger decision process, it is not clear how much dissatisfaction attaches to this step).

Establishing pesticide residue tolerances (EPA)

Amendments to the Food, Drug, and Cosmetic Act give EPA the authority to establish tolerances for residues of agricultural pesticides in foods. (That authority was clarified by the 1996 statute called "Food Quality Protection Act".) The process used by EPA's Office of Pesticide Programs, assigned authority for developing tolerances, resembles very closely that used by FDA for food additives. However, because of authority contained in the Act establishing regulation of pesticides, EPA can and does require submission of more animal test data and information on predicted exposures than is sometimes available to FDA. The 1996 FQP Act changed the requirements in ways whose implications are, as yet, little understood. The Pesticides Office is undertaking a lengthy process of deliberations with interested and affected parties as it develops practices and procedures to implement those new provisions.

Establishing "safe harbor" values under California "Proposition 65"

This initiative statute, "The Safe Drinking Water and Toxics Enforcement Act of 1986," commonly called "Proposition 65", requires private firms to warn people if their operation cause them to be exposed to specifically identified substances at levels large enough to constitute a "significant risk." As originally intended, the statute relied for enforcement on individuals bringing suit for damages against the offending firms; the penalties to be assessed are substantial, and one-quarter of the damages are paid to those who sue; the state Attorney General can also bring such suits. Early in its implementation, the California Department of Health Services to begin issuing "safe harbor" exposure values for substances on the list. These numbers identify exposure levels such that exposures smaller than them the State Attorney General will not support any suits against putative offenders. Having such exposures does not prevent private actions against firms, but, in practice, as was intended, private actions don't occur. It is presumed that such exposures are lawful. Since a government reorganization the Office of Environmental Health Hazard Assessment within the California Department of the Environment develops and issues these values. This Office use methods to develop "safe harbor" numbers that are generally very similar to those employed for establishing ADIs for food additives, although in cases where requisite data are available, the Office has pioneered use of scientifically advanced assessments. The regulated community evinces satisfaction with the process in general, although, not unexpectedly, some individual cases have raised controversy.

Monsanto Company 's reporting information to EPA under §8(e) of TSCA

Manufacturers of chemicals have a duty, under provisions of §8(e) of the 1976 Toxic Substances Control Act, to report to EPA any information that comes into their possession which bears on potential hazards posed by exposure to the chemicals they produce. Penalties for failing to make such reports are substantial, and fall onto the shoulders of individuals. EPA has provided a mechanism by which corporations can assume the burden from its employees, by establishing and following procedures which assure that top management becomes aware of whatever its employees learn. The Monsanto Company adopted a procedure which includes formal delegation to its chief medical officer authority to act on behalf of the Company, and provisions for reporting of new data and information, and review of decisions. With experience, decision rules have been developed for most of the kinds of information that arise (particularly results of tests of chemicals in animals); with these rules, decision authority was effectively delegated to the principal product toxicologist assigned responsibility to review this information. Also included in the conditions of this effective delegation were instructions on how to proceed whenever the information arising was not of the kind for which rules had been derived, and procedures to document how decisions had been reached. These decisions are numerous, ranging from a few dozen to three hundred per year. Only two decision options are available: to report the new information or not to do so. Some of the criteria for the choice are clear and rigid: the information need be reported if it

concerns a substance manufactured or imported by the firm, and if it indicates the possible existence of "an unreasonable risk" unless it is already known to EPA (e.g., published in the scientific literature or reported by someone else). However, both of the latter two criteria involve some expert judgment: issues of reliability and interpretation of test results, for instance, arise in drawing conclusions about hazard posed by exposure, and a particular test may only confirm results already publicly available and not constitute new information. The procedures and practices that evolved over fifteen years' experience resulted in an efficient and effective decision process that produced predictable results, and generally satisfied the various interested and affected parties (within Monsanto and the EPA).

Clean Water Act "Best practicable technology" Regulation[16]

The Water Pollution Control Act Amendments of 1972 required EPA to establish "best practicable technology" (BPT) standards for 30 industries by 1977. These standards were supposed to reduce release of contaminants into surface waters from manufacturing and processing firms, and to incorporate considerations of cost and technological feasibility in determining what was "practicable." EPA contracted with consultants to develop and provide much of the technical information; it specified what information was to be developed and how this information was to be presented to the Agency but did not constrain the inputs. Authority to manage this process and recommend standards was delegated, but approval authority was retained by senior management. (A lawsuit had forced a very rapid decision process.) Economic information on the industries and the impact of the BPT regulations was developed and analyzed within the Agency. The decision process was basically that of "informal rule making" specified by the Administrative Procedures Act. At least some of the stakeholders were satisfied with the process, but not with most outcomes: some two-thirds of the BPT standards were challenged in court, and many changed as a result.

16 W. A. Magat, A. J. Krupnick, and W. Harrington, *Rules in the Making: A Statistical Analysis of Regulatory Agency Behavior*. Washington, Resources for the Future, 1986.

REFERENCES

Barnard, R. C. 1995. "Risk assessment: the default conservatism controversy," *Reg. Toxicol Pharmacol,* 21, pp. 431-438.

Byrd, D., and L. B. Lave. 1988. " 'Significant risk' is not the antonym of '*de minimus* risk'," in C. Whipple, ed., *De Minimus Risk* (Plenum Press).

California, State of. 1985. "Guidelines for Chemical Carcinogen Risk Assessments and Their Scientific Rationale," Department of Health Services, Health and Welfare Agency, Berkeley, Calif.

Doull, J. 1995. Private communication.

EPA (U. S. Environmental Protection Agency). 1986. "Guidelines for Carcinogen Risk Assessment," *Fed. Register*, 51, pp. 33992-34003.

EPA (U. S. Environmental Protection Agency). 1996. "Proposed Guidelines for Ecological Risk Assessment," EPA/630/R-95/002B, August.

Food and Drug Administration. 1982. *Toxicological Principles for the Safety Assessment of Direct Food Additives and Color Additives Used in Food*, Washington, D.C. (Available from the National Technical Information Service, order code PB83170696.)

Gaylor, D. W. 1989. "Quantitative risk analysis for quantal reproductive and developmental effects," *Environmental Health Perspectives*, 79, pp. 229-241. (Provides evidence that for a series of teratogen studies, the upper bound on the "no observed adverse effect level" lies between 2 percent and 6 percent.)

Hattan, D. G. 1996. Interview with Dr. D. G. Hattan, Director, Division of Health Effects Evaluation, Center for Food Safety and Applied Nutrition, Food and Drug Administration.

Howson, C., and P. Urbach. 1989. *Scientific Reasoning: The Bayesian Approach* (La Salle, Ill.: Open Court Publishing Co.).

Kuhn, T. L. 1949. *The Structure of Scientific Revolutions* (Chicago, Ill.: University of Chicago Press).

Kuznesof, P. 1996. Interview with Dr. P. Kuznesof, Chief, Chemistry Review Branch, Office of Pre-Market Approval, Center for Food Safety and Applied Nutrition, Food and Drug Administration.

Lowrance, W. W. 1976. *Of Acceptable Risk: Science and the Determination of Safety* (Los Altos, Calif.: William Kaufmann Co.).

Merrill, R. A. 1988. "FDA's implementation of the Delaney Clause: repudiation of Congressional choice or reasoned adaptation to scientific progress?" *Yale Journal of Regulation*, 5, pp. 1-n.

Morgan, M. G., and M. Henrion. 1989. *Uncertainty: A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis* (Cambridge University Press).

National Research Council. 1983. *Risk Assessment in the Federal Government: Managing the Process* (Washington, D.C.: NAS Press).

National Research Council. 1994. *Science and Judgment in Risk Assessment* (Washington, D.C.: NAS Press).

National Research Council. 1996. *Understanding Risk: Informing Decisions in a Democracy* (Washington, D.C.: NAS Press).

Office of Science and Technology Policy. 1985. "Chemical carcinogens: a review of the science and its associated principles," *Federal Register*, 50, pp. 10371-10442.

Scheuplein, R. A. 1996. Private communication.

Schön, D. A. 1982. *The Reflective Practitioner: How Professionals Think in Action* (New York: Basic Books).

United States. 1966. 5 U. S. Code §§551-557.

Ziman, J. 1978. *Reliable Knowledge: An Exploration of the Grounds for Belief in Science* (Cambridge University Press).