STRATEGIES FOR INFORMATION DEVELOPMENT AND UTILIZATION FOR TOXIC CHEMICALS[1]

Mark Sharefkin
"Resources for the Future
1755 Massachusetts Avenue, N.W.
Washington, D.C. 20036

April 1982
CP-82-12

INTERNATIONAL INSTITUTE FOR APPLIED SYSTEMS ANALYSIS
2361 Laxenburg, Austria

CONTENTS

STRATEGIES FOR INFORMATION DEVELOPMENT AND UTILIZATION FOR TOXIC CHEMICALS[1]

Mark Sharefkin

1. INTRODUCTION AND BACKGROUND

1.1. Characteristics of the Problem

In the last five years, several countries have expanded and strengthened their arrangements for regulating existing chemicals, and for controlling the introduction into commerce of new chemicals. One simple, uncontroversial fact has led many to believe that those controls are necessary: the natural environment is now contaminated with many synthetic organic chemicals, some of which are believed to be carcinogens, mutagens, teratogens, or several of the above.[2] Technological

[1]Preparation of this document was supported in part by the Andrew W. Mellon Foundation and in part by the International Institute for Applied Systems Analysis largely through funding by the Bundesministerium für Forschung und Technologie, FRG.

[2]The trihalomethanes, produced as byproducts in the chlorination step of drinking purification, are ubiquitous in American surface water; so are many common organic solvents, like trichlorethelyne. Unfortunately, the same is true of much American groundwater, as we

progress in organic analysis over the last three decades has been so rapid that relatively quick and inexpensive identification of such contaminants at parts per billion concentrations is now possible. A one cubic centimeter sample of drinking water can be processed, essentially instantaneously, into a computer listing of synthetic organic contaminants.

So much is now uncontroversial, but controversy begins with interpretation, and with the subsequent policy proposals. The broad range of positions already occupied can be demarcated by noting the extremes. At one such extreme are advocates of benign, if cautious, neglect. This argument runs: there is no clear evidence that age-adjusted rates of cancer incidence or mortality have increased with the expansion of, production of, and commerce in synthetic organic chemicals.[3] The implication generally drawn, that no radical changes in current toxic chemicals management policy are warranted, is immediate.

At the other extreme are warnings of potential disaster.[4] The latency period for cancer induction, this argument runs, is typically in the order of decades. Thus, the consequences of the enormous post-World War II increase in synthetic organic chemicals production and commerce, and in the implied human exposures, have not yet appeared in cancer incidence and mortality rates. Continuing, if all dose-response relationships are linear, existing measurements of synthetic organic chemical

---

have learned from the survey work of the United States Geological Survey. A much more systematic understanding of the dimensions of the problem will soon be possible, thanks in part to the data gathering requirements of the Interim Primary Drinking Water Regulations of the Safe Drinking Water Act.

[3]See, for example, Sir Richard Doll (1977).

[4]See, for example, R.H. Harris, et al. (1977)

pollutant concentrations in drinking water and food can, in combination with potency estimates, be used to forecast future cancer incidence and mortality rates. Some forecasts constructed in this way imply significantly higher future cancer incidence rates. The policy implication generally drawn: stronger control efforts should put in place "with deliberate speed" are warranted.

Given what we know now, neither of two extreme positions can be confidently rejected. For that reason, toxic chemicals management poses a dilemma for public policy. That dilemma remains even if we ask what would be good policy in a world in which we knew with certainty that the second (or "pessimistic") extreme view was correct--for special characteristics of the toxic chemicals problems make design of a good management strategy particularly difficult.

Three of those special characteristics should be kept in mind in everything that follows. The first is the "large numbers problem". There are a great many chemicals, and thus many potentially hazardous ones. The relevant information about those chemicals is widely dispersed among final users, distributors, and manufacturers and their employees. Second, there are several very different ways of buying additional information about particular chemicals; those alternatives differ in cost, and in the character and quality of the information they yield. This might be called "many alternative tests." Third is the "testing budget constraint" characteristic. The number of potentially hazardous chemicals is so large, and some test alternatives so expensive, that exhaustive testing--subjecting chemicals to all conceivably warranted tests--is economically impossible. Let us take up each of these characteristics briefly and in

turn; later we will have to look more carefully and analytically at each.

1.1.1. **The large numbers problem**

That there are a great many potentially troublesome chemicals is by now widely appreciated. Though the use of such numbers is inevitably open to misinterpretation, the listing of existing chemicals in American commerce, the "inventory" prepared by the United States Environmental Protection Agency (EPA), is about 55,000 chemicals long.[5] The number of new chemicals entering into commerce in the United States in any given year may be as high as several hundred to one thousand.

Compilation of the "inventory" of existing chemicals was a monumental task, in part for reasons already noted. Information on chemicals is widely dispersed among final users, distributors, and manufacturers and their employees. The seriousness of this problem can be illustrated by the difficulty of assembling information on current levels of occupational and general population exposures to a particular chemical.[6] Exposure estimates are critical to health effect estimates. The crudest measure of health impact is simply the product of exposure and potency reasonably accurate estimates of potency can be obtained from laboratory tests.

Though extrapolation from laboratory test to humans remains contentious, those tests do provide some quantitative measures of potency. But where is exposure data to be gotten? Suppose, for example, that the

[5]Office of Pesticides and Toxic Substance, United States Environmental Protection Agency, 1979 Toxic Substances Control Act Chemical Substance Inventory, Initial Inventory.

[6]See, for example, the selection of information required of reporting firms by United States Environmental Protection Agency proposed form, Premanufacture Notice, Part II, Human Exposure and Environmental Release, Federal Register, Vol. 44, No. 201, October 16, 1979. The difficulties of accurately estimating human exposures from the requested information are apparent.

chemical in question first sees the light of day in the production area of some chemical manufacturing plant. Initial human exposures to that chemical are then exposures of plant personnel. But occupational exposures, which are typically easier and less costly to measure than general population exposures, are difficult to measure and have been estimated for very few chemicals. Any such measurement requires excellent knowledge of the particular process, and perhaps even of the particular plant configuration; such knowledge will typically be held only by plant employees and management. The same point can be made for many, if not most, of the other kinds of information that are germane to toxic chemicals management. Information on health and environmental effects, no less than information on production processes and occupational exposures is dispersed among many institutions and individuals.[7] At least some of those institutions and individuals have "more and better" information than can be "centralized"--delivered to some government agency charged with toxic chemicals management policy.

1.1.2. The many alternative tests "problem"

Suppose, nevertheless, that some such agency does decide to gather information on some particular chemical. There will be no lack of possible ways in which to expend that agency's scarce resources. First, the scientific literature can be searched for reported work on the chemical in question, or on chemicals thought to be closely related in chemical structure, and thus perhaps in biological activity.[8] Calculating "simple

[7]For the diversity of sources of information on the environmental effects of, and the persistence of, candidate toxic chemicals, see the first report of the American Interagency Testing Committee (1977). It is almost certain that the available information varies widely in quality.

correlations" of the structure of the chemical under investigation (with structural indicators of other chemicals) is inexpensive, but interpreting the results is difficult. Second, a growing arsenal of so-called short-term, or bacterial tests, is available.[9] Those tests can typically be run in a day, and at a cost of about $500 to $1,000. Third, there is the traditional last resort: long-term, or animal bioassay, testing.[10] Those long-term tests can extend over several years and can cost on the order of $1,000,000.

1.1.3. The budget constraint problem

Looking over all of these possible ways of purchasing additional information on a particular chemical, the government agency may notice that each test provides imperfect information of a particular kind, and at some specified cost. The costs for the most expensive tests are high enough to rule out exhaustive testing of all chemicals; here as elsewhere choices must be made.[11] But those choices can be made rationally only if we know what we expect to learn from the purchased information and what we will do with that information.

---

[8]The chemical literature is, to say the least, voluminous; it is almost certainly true that the quality of the best published results has improved radically, along with instrumentation and quality control methods, over the past two decades or so.

[9]For a relatively nontechnical discussion of the variety of short-term tests currently available, see Raymond Devoret (1979); the original article is Ames, McCann, and Yamasaki (1975).

[10]For a description of one of the largest and most ambitious programs of long-term testing, see, for example, the summary reports of the Carcinogenesis Bioassay Program, conducted by the National Cancer Institute of the United States National Institutes of Health.

[11]In the years 1977 through 1980, the testing budget of the Office of Pesticides and Toxic Substances probably was on the order of $20 million. But this is not a very useful guide to the resources currently expended on testing in the United States, because considerable resources are available to both other government agencies like the National Institutes of Health and to the private sector for these purposes. Further, many other countries, and some major international organizations like the United Nations, have programs of their own. Assembling information on these programs in some useful form would be helpful not only to us, but clearly to anyone involved in the forward planning of testing and regulatory programs. While this exercise probably would not be very difficult, it might well be tedious.

What do we learn from this listing of characteristics? One thing we learn is that it is too early to pick one or another institutional setting of framework as the way in which to analyze the toxics management problem. We know too little about the "distribution" of existing information, and about the costs of generating and transferring new information, to be dogmatic. We believe that the difficulties encountered by the American Toxic Substances Control Act (TSCA) implementation effort reinforce both these points.[12]

1.2. Implementation Efforts: Lessons from the American Experience

Efforts to design and implement toxic chemicals management policies have led to somewhat different results in the many countries concerned with the problem. Nevertheless, because the underlying problem is the same, there are some similarities across national programs. For that reason, observations based narrowly upon the American experience may be more generally applicable.

The American Congress enacted the Toxic Substance Control Act (TSCA) in 1976.[13] The executive branch agency charged with implementing that legislative mandate, the Office of Pesticides and Toxic Substances (OPTS) of the United States Environmental Protection Agency (EPA), is now four years old. Though judgments of administrative performance in a problem as novel and as complicated as the toxic chemicals problem should be tempered with mercy, many observers have been less than

[12]Some of these issues are discussed in a monograph, in progress at Resources for the Future, on the TSCA implementation experience.

[13]Toxic Substance Control Act, Public Law 94-469, 94th Congress of the United States, October 11, 1976. Also see the legislative history of the Act.

merciful. The agency, and the program, are seen by many as having failed to identify, justify, pursue and defend an implementation strategy.[14]

Such judgments often rest upon some simple summary statistics of agency performance and upon projections of what continuation of present performance will mean. In 1979, the OPTS issued its listing, or "inventory," of existing chemicals. That listing was required by the logic of TSCA. Because TSCA treats existing and new chemicals asymmetrically, as any workable toxic chemicals management must, an inventory of existing chemicals is needed to mark the boundary between "existing" and "new" chemicals. The American inventory lists approximately 55,000 chemicals in commerce.

By comparison, about 100 existing chemicals have thus far been identified as candidates for scrutiny and possible restriction.[15] Comparison of these small numbers to the large number of entries in the inventory is inevitably simplistic, but may not be misleading. The perception that present methods for defining a management strategy for existing chemicals are hopelessly cumbersome may be accurate.

The new chemicals management problem is substantially different than the existing chemicals management problem. The practical alterna-

[14]The US General Accounting Office, the accounting and inspection arm of the United States Congress, has just completed an extensive and very critical analysis of TSCA implementation. A less detailed critique is to be found in the briefs filed by the Natural Resources Defense Council in Natural Resources Defense Council, Inc. vs. Douglas Costle, Administrator and United States Environmental Protection Agency, United States District Court, Southern District of New York, 79 Civ. 2411, July 11, 1979.

[15]One measure of the number of chemicals thus far selected for serious attention is the number of chemicals listed by the Interagency Testing Committee. Other plausible measures support same order of magnitude estimate, about 100.

tives for management of a new chemical are much broader than the alternatives for managing an existing chemical. A new chemical can, in principle, be tested extensively before introduction into commerce. Concern over the rate--between several hundred and one thousand per year--at which new chemicals were being introduced in use and commerce was in fact largely responsible for passage of TSCA by the American Congress.[16]

Nevertheless, OPTS performance in processing applications by prospective importers and manufacturers to introduce new chemicals has yet to be tested by serious cases. Of several hundred such applications received, more than eighty percent were for either polymers or intermediate chemicals: polymers are unlikely to pose any serious health or environmental risk, and intermediate chemicals are unlikely to be released to the environment in substantial quantities.[17] The OPTS application process and requirement fails to focus information generating effort on those chemicals which really are candidate problems, but instead diffuses it over many that almost certainly are not. In each year since the passage of TSCA, several hundred applications to introduce new chemicals have been submitted to OPTS. Less than 50 such applications have been processed, with about 10 chemicals barred from introduction. Again, many have found the OPTS performance in designing and implementing a strategy for new chemicals management wanting.

---

[16]See United States Congress (1976).

[17]It is widely believed that, prior to TSCA, new chemicals were being introduced into commerce at about the rate of 1,000 per year. In Fiscal Year 1980 422 applications were submitted to EPA; as of April 30, 1981, 217 additional applications had been received by EPA.

Why have several years of effort and several hundreds of millions of dollars seemingly produced so little? We suspect that the explanation is relatively simple. The problem was never adequately named and analyzed so that sensible implementation strategies were never articulated.

More important, what is to be done now. For a problem as complex as the toxic chemicals problem, there can be no simple answers. Nevertheless, this paper aims at providing answers of a kind. It does so by offering several definitions of the problem and then drawing the implications for implementation. The premise of the paper is that sharp deputation over simple, idealized, explicit implementation strategies can help in the design of the real thing.

1.3. Overview of the Paper

Sections 2, 3, and 4 are the heart of the paper. In each of those sections, one perspective on the problem of designing strategies and institutions for managing toxic chemicals is presented. We have presented three different complementary ways of looking at the problem because we are far--in understanding and perhaps in time--from any simpler view of the problem.

Each of the three perspectives corresponds to a particular view of the technical and institutional barriers to effective toxic chemicals management. Perspective 1 has been given the name "centralized decision." The name assigned to Perspective 2, "decentralized information gathering and generation," is essentially self-descriptive. The same might be said for what we call Perspective 3, "team-theoretic approaches to allocating testing resources to firms."

Finally, a word on the organization of the paper. Each of sections 2, 3, and 4 is divided into two subsections; in each case, the first is literary and the second mildly technical. The nonmathematical reader should be able to get the gist of the paper by reading only the literary sections.

2. PERSPECTIVE 1: CENTRALIZED INFORMATION GATHERING, GENERATION, AND DECISION

2.1. Literary Recapitulation

We have noted that several very different kinds of information are useful in toxic chemicals hazards management, and that both possession and understanding of much of that information is "highly decentralized". Assembling all that information in one place with relatively little distortion and with the contextual information typically essential to accurate interpretation will be extremely costly if not impossible.

Nevertheless, there are compelling arguments for attempting and paying for, some measure of information centralization.[18] To cite one of

[18] The informational and incentive issues here are somewhat subtle. At one level they can be avoided by remembering that a great deal of information about chemicals is already in the public domain, so that a government agency hoping to use its information-generating resources efficiently must face up to the issue of how to use that information in allocating its resources. Something like the formal development in the text of the paper will be a necessary first step for the agency.

But the subtler issues are both important and intriguing. A reviewer of this paper has noted that both confidentiality and equity have proven troublesome in TSCA implementation, and so they have. Let us try to say why, and think about what can be done.

Confidentiality is an issue because firms understandably want to appropriate the full benefits of information they have generated privately. But from a broader social perspective, private property in information--"confidentiality"--makes sense only if the gains from the induced innovation outweigh the losses from restricting access to the information.

Almost all serious students of the innovation process are convinced that reducing the ability of firms to appropriate the benefits of innovation will reduce innovation. But it does not follow that all information generated by a firm in researching and developing a new product should be protected from other claimants. For purposes of argument, suppose that two distinct and separable kinds of information are generated. Information *I* bears on the production process and on those characteristics of the new chemical which account for its market value. Information *II* bears on the potential health and environmental hazards of the new chemical, and perhaps on the health and environmental hazards posed by related chemicals.

Then it is clear that while there is some justification for protecting confidentiality of

the most compelling, accurate information developed by one firm can, in principle, be shared with other firms at relatively low cost. Thus, if two or more enterprises perform the same test or tests on the same chemical, resources allocated to duplicative testing are essentially wasted. In part, for this reason, efforts to bring data potentially relevant to toxic chemicals management into once centralized bank are already fairly advanced.[19] Suppose, then, that this centralized data base contained all the information we ever would be able to use for designing control and regulatory policies for new and existing chemicals. How would we proceed?

Let us develop an answer to this question, an answer that is specified more mathematically in the following subsection. The process might proceed through several steps or stages. First would come choice of a management objective or objectives, then would come characterization of the entries in the data base in terms amenable to decision analysis. Finally would come an optimization exercise, leaving us with a designated

---

Information *I*, there is little or not such justification for protecting confidentiality of Information *II*. In fact the best arrangement would be full protection of Information *I* and full disclosure of Information *II*. (In passing we note that the question of why firms should be assigned the role of, or expected to, generate Information *II* has at least two good answers: they probably have some comparative advantage in producing Information *II*, and the two kinds of information are very likely joint products.) Thus the real issue is the extent to which Information *I* and Information *II* really are distinct and separable. I know of no serious examination of this question.

Finally, let us turn to the problem of equity in allocating the costs of testing among producers of a chemical which, the regulatory agency decides, must be tested. Life is full of problems, and this happens to be a particular case of one that is solved every day elsewhere. The provision of every public good inevitably involves a redistributive transfer; the provision of information on a particular chemical is precisely the provision of a public good. I have nothing to add here to the literature on sharing the costs of a public good. I do repeat one suggestion that has become a commonplace: let firms producing a chemical negotiate among themselves the arrangement for sharing the cost of the required testing.

[19]Several large integrated chemical information systems are in the advanced stages of development; like all such systems, they are being built in a modular way, and the least developed modules often are those that are most important for our purposes. For an example, in the Chemical Information System the Ames test data base module is still under development; this is hardly surprising, since Ames' original paper was published in 1975.

strategy: a list of instructions, including a testing program and a control strategy. In what follows, we mean by a "testing program" something like a list of which chemicals to test, and in which order. By a control program we mean a list of which chemicals to subject to which regulatory restrictions or controls, and in what order. Clearly, running a toxic chemicals management program in blind obedience to the optimization exercise would probably be disastrous. Equally clearly, exercises of this kind are invaluable guides in designing a real program. We proceed to justify that claim, and begin by retracing, in more detail, the steps of the optimization exercise.

The first step requires choice of an objective or objectives: we suggest choice of expected net benefits deriving from chemical use as a sensible initial benefit measure. In the simple, tractable but highly unrealistic case in which no two chemicals are substitutable for one another in any use, net benefits from a set of chemicals are a simple sum of net benefits from individual chemicals. Thus, our problem is reduced to identification and estimation of the latter.[20]

We do that in two stages. First, we estimate the net internal economic benefits associated with use of the single chemical. Here net internal benefit implies deduction of all internal costs, but not the external costs, of the particular single chemical. Those internal costs include all costs incurred in generating the information upon which the ultimate expected benefit rides.

[20]This question, and many other similar practical questions, are taken up in work in progress at Resources for the Future.

Next is the really hard step, estimation of the externality costs associated with use of the single chemical. To do no more than touch upon the difficulties involved in making that estimate, suppose that for the particular chemical in question the only such costs are the health-related costs of occupational and general population exposure to that chemical. To do the calculation we need three numbers: a shadow price at which to value health risk,[21] an estimate of the exposure of individuals to the particular chemical, and an estimate of the potency of the chemical in producing the particular health effect. By definition, potency is a summary measure of the relationship between dose and response.

Assume that there is unanimity on the appropriate shadow price of health risk. That assumption lets us focus upon the unique difficulties of the toxic chemicals case. It is a sad fact of life that the remaining two numbers required for estimation of the health cost of the particular single chemical--exposure and potency--are subject to serious uncertainties for many chemicals. This is true even for those chemicals which are both ubiquitous and suspect carcinogens, mutagens, or teratogens.

We might proceed by simply making point estimates of both of those quantities, but proceeding that way will foreclose the option of using our approach in just those ways that make it potentially valuable. For however little we know now about the particular chemical under investigation, we may know more in the future--as we subject this, and closely related chemicals, to tests of various kinds. Clearly what is needed is a sys-

[21] The "shadow price of health risk" means the monetary value of the resources society is willing to forego in order to reduce either some particular, or several, health risks faced by individuals at the margin. This may be equal to individual willingness to pay for health risk reduction at the margin, but the two values could differ for many reasons.

tematic way of integrating new information with existing information on the chemical under scrutiny.

A systematic procedure for that integration is available[22]. To use it, we need do only two things. First, our current information on the chemical's potency and on the prevailing level of human exposures to that chemical must be expressed probabilistically, thereby summarizing in this uniquely convenient form our uncertainties about each of those numbers. Second--the next stage of our centralized decision procedure--we need an appropriate and consistent characterization of the various kinds of information available from chemical testing and from other sources, such as exposure measurements.

Suppose we agree to summarize our initial information on our particular single chemical by a probability distribution on that chemical's potency.[23] The standard procedure for revising, or updating, the probability distribution in the light of new information is a procedure called Bayes' Rule. To apply that rule, we need an expression for the joint probability distribution of that potency variable and the random variable we

[22]Naively, Bayes' theorem tells us how, but almost everyone realizes that Bayes' theorem is a mechanical procedure which is far from capturing what goes on in even the everyday induction in terms of which much workaday science runs. To put the argument in Bayesian terms, much of the work is finding the "right" likelihood function, and we know of no efficient way to routinize that process.

[23]Here and elsewhere in the paper we use the term potency and (logarithmic exposure-adjusted) potency interchangeably, and perhaps confusingly. The variable k_i always means the latter: the exposure-adjusted logarithmic potency of chemical i. Let us be somewhat more fussy and precise, at least in this. First, exposure-adjusted means: take the toxiocological potency and multiply by the number of individuals affected (the population at risk can be stratified by dose levels). We use the logarithm of this number for two reasons. Because of the huge range in potencies and in exposed populations--each ranges over several powers of ten--logarithms are more convenient. And by taking the logarithm, we transform a variable which must range over the positive real axis to one that ranges over the whole real axis. This is convenient if we want to take a normal distribution to represent the distribution of k_i, and thus a lognormal distribution to present the (original, before taking logarithms) exposure-adjusted potency.

are observing--in whatever kind of testing or information-gathering, we undertake. That joint distribution must be grounded in some systematic understanding of the relationship between the logarithmic potency variable and the random variable being observed.

Such theories--plausible ones--are hard to come by. Just how hard will become clear as soon as we list the various, and diverse, information sources from which we can learn more about the chemical under scrutiny.[24] A short list--shortened by simplification and aggregation--must include:

1. searches of the chemical and biological literature;
2. structure-activity correlation;
3. short-term or bacterial testing; and
4. long-term or animal bioassay testing.

Now we can be explicit about why we are still far from the systematic understanding we need. Begin with literature searches: it seems plausible that toxic chemicals to which more people have been exposed have been more extensively studied. Forms for the joint distribution of logarithmic potency and measures of the amount of attention given a chemical by toxicologists and epidemiologists follow from the tentative assumption that more serious problem chemicals have attracted more and more serious attention.[25] Random variables characterizing the literature, such as the number of articles in a given period, should be correlated

[24]Keep in mind the limitation of any such formal scheme, and remember the reservations expressed in footnote 22.

[25]Clearly, this is tentative, and could even prove disastrous if taken too seriously: surprises do happen, and especially in chemical toxicology.

with logarithmic potency. This is conjecture. We do not know, because we have not yet tried to find out, if there are strong relationships between variables characterizing searches of the literature and the logarithmic potency variable.

Structure-activity correlation[26] (SAC) presents a simpler case, because the kind of information generated by structure-activity correlation is in exactly the form required for Bayes' Rule revision of probability distributions on the logarithmic potency of the particular chemical. Because SAC is in its infancy, the joint distributions inferred will typically not be very "tight." As SAC evolves, it can be expected to become a more discerning, and for that reason, more valuable tool. That evolution will be reflected in successively "tighter" joint distributions.

Short-term,[27] or bacterial, tests have multiplied in number and ingenuity since publication of the initial test procedures by Bruce Ames. Short-term testing can be brought within our framework--Bayesian revision of probability distributions defined on logarithmic potencies--in several seemingly plausible ways; two will highlight the range of choice. At one extreme, negative (respectively positive) Ames test results might be interpreted as evidence that the chemical in question is not (respectively is) carcinogenic. At another extreme, short-term test results can be interpreted as telling us something much more detailed about the carcinogenicity of the particular chemical.

[26]For a technical survey of the state of the art in structure-activity correlation, see A. Stuper, et al (1979). This is a field in rapid development. The relationship with pattern recognition and intelligent data bases, two rapidly developing subfields of artificial intelligency, is apparent. Structure-activity correlation has already proven its worth in the design of chemical syntheses; it may evolve into a far more reliable tool in support of toxicology.

[27]Again, see Devoret (1979) and Ames, et al. (1975).

At present, it is far from clear which of these views--or which intermediate position--makes sense. Supporting the first view of what can be inferred from short-term test results is the observation that not all known human chemical carcinogens test positive in short-term tests.[28] Buttressing the latter view is the empirical observation that the rate of growth of revertent bacteria in short-term tests is strongly correlated with the carcinogenic potency, at least for chemicals for which measurements of the latter exist.[29]

Under either of these assumptions, a joint probability distribution for Bayes' Rule can be written down. But the choice between the two assumptions is critical. Because short-term tests are relatively inexpensive, they are an important potential source of valuable information on chemical toxicity--if we believe, and act as if, the information generated by short-term tests is relatively reliable. If, for example, we act as if short-term tests can distinguish only imperfectly between noncarcinogens and carcinogens, those tests can serve only as a coarse screen. If, on the other hand, we act as if revertent bacteria growth rates convey valuable information about carcinogenic potency, we may in many cases be spared the expense and delay of long-term testing.

---

[28]At least some of the reasons for this are understood. The Ames test is a test for chemical mutagenicity; that is, for how effective the tested chemical is in forcing revisions in the genetic material. Some such revisions are associated with the breakdown in the cell growth process called cancer, but there are other ways in which that control system can break down. That the cell growth control system is complex, with many levels between the ultimate "hardware" genetic level and the cellular level is certain. The clearest evidence comes from studies of skin cancer. See, for example, John Cairns (1978).

[29]The Ames test uses bacteria lacking in the gene critical to synthesis of an enzyme essential to growth. When placed in an environment deficient in that enzyme, only bacteria with "revertent" mutations--mutations which restore the gene critical to that synthesis--can grow. If a strong carcinogen is placed in the medium, the rate of reversion, and thus the growth rate of bacteria, is accelerated.

For if anything is uncontroversial about long-term testing, it is the formidable cost, time, and quality control requirements for such testing.[30] Depending upon the particular test design, adequately controlled bioassays with enough test animals to give statistically significant results can take from one to three years and cost from several hundred thousand to one million dollars. Because those tests are aimed at establishing causal relationships between exposures (or doses) and health effects (or responses) of various kinds, great care must be taken to insure that even trace amounts of other potentially carcinogenic contaminants are not present, a requirement which imposes additional quality control costs. For all these reasons, the checking and validating of bioassay test results by replication doubles or triples the cost. To date, there has been little such replication.

If proposals to perform long-term tests are to be systematically evaluated, and if the results of long-term tests are to be used systematically, these results must be expressed in a form suitable for revision of pre-long-term testing probability distributions. Specifically, we must know something about the joint probability distribution of "true" carcinogenic potencies and of some random variable (or variables) characterizing the results of long-term tests. Such a joint distribution can be developed from our existing stock of long-term testing results, and used in the revision step: given a long-term test on a previously untested chemical, it tells us how much to weight that result.

[30] Recently, much attention in the United States has been focused on the quality control procedures underlying the animal testing done in the National Cancer Institute bioassay program and elsewhere. The difficulties in controlling the quality of these procedures is easy to understand: trace impurities of any carcinogen not under test can invalidate the rest results.

Thus far we have suggested only that the results of each of the four principal ways of generating information (on logarithmic potencies) can be summarized in a particularly convenient way. In that summary form, they can be used to systematically and consistently revise our estimates of the adjusted (logarithmic) carcinogenic potency of a suspect chemical.

But we have thus far said nothing about which, if any, tests on which chemicals should be done. We have remarked that exhaustive testing--running all tests on chemicals--is so expensive as to be essentially infeasible. Beyond that trivial observation, we have said nothing about how we would allocate a given limited testing budget among chemicals and tests. That is, after all, the practical question. A more or less ready-made approach to this question is available, and is provided by the branch of mathematical statistics called statistical decision theory.

2.2. A Mildly Technical Recapitulation

Let us see, in very brief outline, how the problem of allocating a limited testing budget among both chemicals and possible tests can be cast as a fairly standard problem in Bayesian statistical decision theory.[31] We also want to see "how bad" the large number problem really is: we want to know about how rapidly the computational burden of the optimization exercise proposed below for sequencing tests grows with problem size.

[31]See, for example, Blackwell and Girschick (1954), Ferguson (1967), or de Groot (1970). Remember that these are expositions of the theory of Bayesian statistical decision theory. The computational implementations of that theory in large-number problems raises additional, and somewhat novel, problems.

Start with our first task: casting the problem as a statistical decision problem. Figure 1 is an illustration of the way in which our four ways of gathering information on a particular chemical might be deployed against a single chemical. Where we have to resort to all four information-generating opportunities, we might successively improve our estimates of the (exposure-adjusted logarithmic)[32] potency k_i of that i^{th} chemical; the four, presumably successively improved, estimates are $k_i^{(1)}$, $k_i^{(2)}$, $k_i^{(3)}$, $k_i^{(4)}$ in Figure 1, and are obtained at costs C_L, C_{SAX}, C_{STT}, C_{LTT}.

The k_i's in that diagram are of course heuristic, for we begin with imperfect knowledge of k_i, and hopefully improve our estimate as we spend more on information on that chemical. But at each step we have only a more or less narrow probability distribution defined on k_i. What follows below is the standard Bayesian calculus for sequential revision of an initial, or prior probability distribution $f_0(k_i)$ on the (exposure-adjusted logarithmic) potency of a single chemical. We write down that calculus as if there were only one chemical to be tested and as if the four tests were to be made in the sequence indicated on Figure 1. We do so because half the art of applied Bayesian analysis lies in choosing a good probabilistic characterization of the kinds of information one has available;[33] literature search and biological test results do not naturally come in the form of joint probability distributions, and the usefulness of such information depends crucially upon choice of an appropriate form. The

[32] See note 23 above.

[33] Recall the difference between mechanical invocation of Bayes' Theorem and real-work induction; see note 22 above.

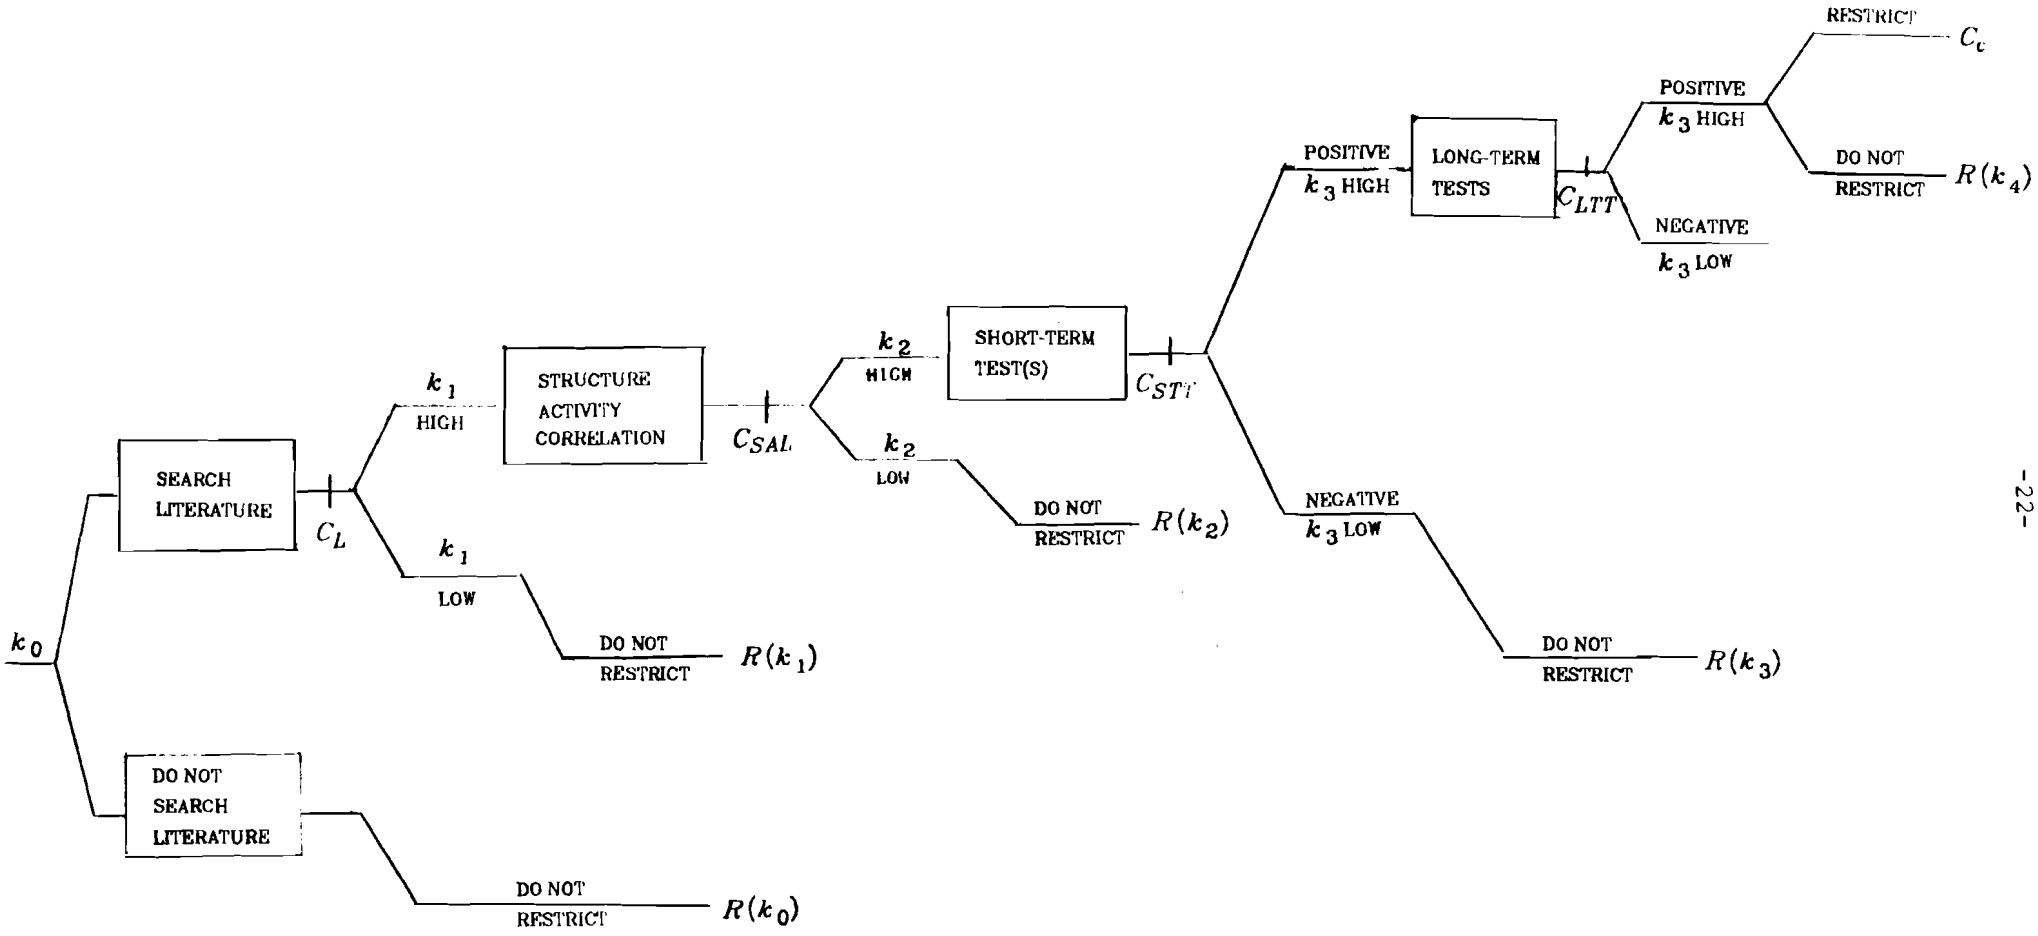

Figure 1. Decision Tree for Information-Gathering on Carcinogenic Potency of a Single Chemical

reader is warned that, in the general multichemical sequential case, not all chemicals will be subjected to all tests.[34] Of course, that would be ruled out in any event since the testing budget is constrained. But the distributions below are the essential building blocks of that general sequential case, and for that reason we have taken care in defining and specifying them.

Introduce notation as follows:

k_i	Prior probability distribution on the (exposure-adjusted logarithmic) potency k_i
1_1(DATALIT[i])	Joint distribution of DATALIT[i], k_i
$f_1(k_i \mid \text{DATALIT[i]})$	Post literature search distribution of k_i
1_2(STRUCTURE[i], k_i)	Joint distribution of STRUCTURE[i], k_i
$f_2(k_i \mid \text{STRUCTURE[i]})$	Post structure-activity correlation distribution of k_i
1_3(AMES[i], k_i)	Joint distribution of AMES[i], k_i
$f_3(k_i \mid \text{AMES[i]})$	Post short-term testing distribution of k_i
1_4(BIOASSAY[i], k_i)	Joint distribution of BIOASSAY[i] and k_i
$f_4(k_i \mid \text{BIOASSAY[i]})$	Post-bioassay distribution of k_i

[34] Here we are sloughing over many subtleties and many potential problems. The complexity result will depend upon how the problem is cast; there is no best way. The worst-case results typically of complexity theory may not be particularly helpful as guides to the computation problem for real data in this area. In any event, this is work in progress and work to be done.

The successive distributions f_0, f_1, f_2, f_3 of the variable k_i are related by the usual Bayes' Rule revision formulas:

$$f_1(k_i \mid \text{DATALIT[i]}) = \frac{1_1(\text{DATALIT[i]} \mid k_i) f_0(k_i)}{1_1(\text{DATALIT[i]})}$$

$$f_2(k_i \mid \text{STRUCTURE[i]}) = \frac{1_2(\text{STRUCTURE[i]} \mid k_i) f_1(k_i)}{1_2(\text{STRUCTURE[i]})}$$

$$f_3(k_i \mid \text{AMES[i]}) = \frac{1_3(\text{AMES[i]} \mid k_i) f_2(k_i)}{1_2(\text{AMES[i]})}$$

$$f_4(k_i \mid \text{BIOASSAY[i]}) = \frac{1_4(\text{BIOASSAY[i]} \mid k_i) f_2(k_i)}{1_4(\text{BIOASSAY[i]})}$$

In each successive equation, we have simplified notation by suppressing some of the previous stage conditioning values: thus $f_1(k_i)$ in the second equation stands for $1_1(k_i \mid \text{DATALIT[i]})$, and so on.

Nothing is easier than writing down formalisms; much more difficult is the prior conceptual work guiding the choice of formalization. What, then, can we say about the appropriate forms of the functions $f_1, f_2, f_3, 1_1, 1_2, 1_3$ which we have so blithely written down above? Elsewhere we have written on this question; here we content ourselves with a few words on the logic of those recommended initial choices, since the real work of implementation will require substantial refinement of those initial choices.

Each such choice reflects a commitment to a theory of, or at least a view of, the process by which the information to be exploited came into existence. It may be plausible to suppose that chemicals to which more

individuals are exposed and which are more toxic have drawn more attention from toxicologists and epidemiologists:[35] that supposition guides one form of the joint distribution 1_1. It may be plausible to suppose that structure-activity correlation provides good relative, but poor absolute, information on the k_i variables. Again, that supposition leads immediately to a particular function form for the joint distribution 1_2. Similarly, for short-term or bacterial testing, the relevant supposition is that such tests discriminate powerfully between noncarcinogens and carcinogens, but only poorly between carcinogens differing, even by a few orders of magnitude, in carcinogenic potency. For long-term or bacterial tests, the relevant supposition is that such tests give good information on k_i, if at relatively high cost. These latter two suppositions, like the first two, lead naturally to formalizations of the corresponding joint distributions, here 1_3 and 1_4.

Now let us remember that our real problem involves a decision about which tests we will apply to which chemicals and in which order. Because of the "large numbers problem", this is naturally posed as a sequential decision problem, but only can be practical if the computational burdens imposed by the large numbers problem are not overwhelming.[36] So let us pose, more or less rigorously, the sequential decision problem we face, and then let us see how rapidly the computational burden grows with the "problem size." The obvious measure of problem size here is, of course, the number of chemicals N_C.

[35]Again, we take note of the importance of surprises in toxicology; see note 25. The real question remains: how to characterize the existing literature as an information resources, and how to use it efficiently.

[36]See for example, Aho, et al. (1974) or Garey and Johnson (1979).

This decision problem, like any other, must be driven by an objective function describing just what we are trying to accomplish with a toxic chemicals testing program. Here is one such objective function; others are possible and may even be better, but one will do for illustrative purposes.[37] The testing program optimization problem is taken as

$$\max \sum_{i \epsilon N_C} \left[b - E \left[k_i \mid f_{n(i)}(k_i) \right] \right] \left[g(i) \right]$$

Here we have drawn on our assumption that the benefits associated with individual chemicals are independent and additive; b is the benefit per chemical, net of (internal) production costs, but gross of possible externality costs arising from introduction of that chemical into commerce. The subscript $n(i)$ is an ordered subset of the integers, 0, 1, 2, 3, 4 and indicates those tests which have been run, and the order of which they were run, in the optimum program, on chemical i. If none have been run, it consists of the single value 0. The probability distribution $f_{n(i)}(k_i)$ is the result of Bayesian revision in the order in which tests are performed. The multiplicative coefficient $g(i)$ is 0 or 1, as the chemical is banned from or allowed into commerce. Thus, this objective is nothing but the expected net benefits of chemicals remaining in commerce.

Given this (or any other plausible) objective function, we can turn to the problem of constructing the implied optimum program. The theoretical problem was settled long ago by the work of Wald, Blackwell and Girschick,[38] and others.

[37] Again, this is work in progress on TSCA implementation at Resources for the Future.
[38] See Wald (1947) and Blackwell and Girschick (1954).

Here is a very brief summary of what that line of work tells us. Suppose we are given a loss function for a decision problem. That loss is defined on $A \times S$, with A the space of actions and S the set of states of nature. We do not know which state of nature prevails, but we can, at cost c_j, make an observation on a random variable r_j for which the joint distribution $f(r_j, s)$ is known. Then Wald and Blackwell and Girschick tell us how to choose a sequence of observations, how to decide when to stop, and which action from A to take when we do stop.

Our practical problem is easily seen to be similar: the states of nature are the $[k_i]_{i \varepsilon N_C}$, our actions are $[g(i)]_{i \varepsilon N_C}$, and our four kinds of tests allow us observation--at some cost--on variables whose joint distributions with the k_i's we think we know something. The novel feature of our problem is the large numbers problem: how reasonable are the Wald-Blackwell-Girschick rules when the number of chemicals N_C becomes large, say 1,000 or even 10,000? If the time to compute a good testing program is bounded by some fairly low-order polynomial in $|N_C|$, say $|N_C|^3$, things may be tolerable. If the dependence is exponential, say $\exp(|N_C|)$, the scheme described above is obviously of no practical importance. It is easy to show that the bound is polynomial; for the two-test case, it is exactly $|N_C|^{3.39}$ We mention here that the application to the chemicals case of the Wald-Blackwell-Girschick apparatus is not exactly straightforward, in part because the tests we have described give information on many of the k_i's simultaneously.

[39] See note 36.

3. PERSPECTIVE 2: DECENTRALIZED INFORMATION-GATHERING AND GENERATION: LIABILITY-BASED INCENTIVE SCHEMES

3.1. Literary Recapitulation

We have suggested that information relevant to toxic chemicals management problems is widely held and both expensive and bothersome to transfer without distortion. A similar remark applies to candidate strategies for generating additional information. A little reflection on the degree of informational centralization implied by the "centralized decision" model highlights the need for much chemical testing and exploration of alternative control policies to be performed in a decentralized fashion--by individual firms, and even in individual plants.

That much is easy to say. It is harder to insure that an appropriate amount of (decentralized) effort is being expended, and that relevant and reasonably accurate summaries of that information are being transmitted to some "center" for use in those decisions which cannot, or should not, be decentralized. Looked at this way, the problem becomes one of incentive system design: what incentive system accomplishes those ends?

Here we identify, in an informal way, the principal issues of the incentive design problem.[40] An idealized version of one such incentive system would require that firms generate or assemble and then transfer

[40]There are other places where such problems arise, though they are inevitably somewhat different in character. One example is the insurance industry: information about individual subscribers is valuable to individual firms, and is to some extent transferred among and shared by firms. Neglect of these arrangements may have led some theorists to conclusions about the stability and efficiency of the insurance industry that are, to say the least, counter factual; see Rothschild and Stiglitz (1976). But they are hardly to be blamed. Work on the question of the value of alternative information gathering and sharing schemes is both difficult and in its infancy.

to the designated regulatory agency or governmental body, specified information on suspect existing, or new, chemicals intended for introduction into commerce. The designated information might include estimates of occupational exposure, environmental releases at several hypothetical production levels, the associated general population exposures, and carcinogenic potency. Those are simple summary measures of the kinds of information typically required for hazard assessment--or assembled in testing. They may or may not be the best summary measures; the necessity of choosing a few such summary measures for transfer is clear.

Equally apparent are the incentives that firms and individuals may have to either engage in too little testing, or to transmit strategically distorted summary measures.[41] Those incentives are particularly strong for chemicals which may affect large populations at low levels of exposure, and for those chemicals which may be implicated in health effects which cannot be unambiguously traced to any single chemical. Submission by firms of "strategically" low estimates of the carcinogenic potency of a proposed new chemical, for example, can be expected to push a government decision on that new chemical toward the outcome favored by the firm.

Against those incentives to insufficient testing and strategic misrepresentation, at least two contrivances can be deployed. Scientific norms requiring the submission of sincere, or honest, estimates exist and

---

[41]Implying that someone, or some institution may be tempted to "lie"--the euphemism is strategic misrepresentation-- raises hackles. No such implication is intended here. We suggest only that additional incentive to tell the truth can in Dr. Johnson's phrase, "concentrate the mind." On quite another issue, some insist that existing *ex post* liability schemes can do the job. For a characteristically brilliant and rather devastating attack on this position, see Judge David Bazelon (1980).

matter, but may alone be inadequate. Those norms can be reinforced by a system of complementary economic incentives explicitly penalizing insincere, or strategic, transmission of information on an existing or new chemical.[42] Properly designed, those incentive schemes should both prompt the right amount of testing by firms and insure the transmission of sincere information by firms to the government since such liability-based incentive schemes may, and should, be a part of many national toxics programs, the design problem is a practical and timely one. A formal attack on part of that problem is presented below in the next subsection. Here we content ourselves with an informal description of those results.

The purpose of a liability scheme is to encourage individual decision-making units--typically firms--to act, in generating, using and transmitting information on toxic chemicals, in the broader social interest represented by the government agency responsible for toxic chemicals management. Firms and government agencies typically will have conflicting objectives; liability-based incentive schemes aim at reconciliation.

First let us focus on part of the design problem. Assume that firms know the truth about their prospective chemicals, but that they transmit "strategic," as opposed to sincere, hazard assessments: that is, they transmit hazard assessments contrived to induce the government to act as the firm wants it to act. In particular, the firms may have both its own

---

[42]There are obvious resonances with several strands of the existing literature, and even with some of the venerable literature of probability theory. For the older literature, see de Finetti's (1972) well-known scheme for forcing a risk-neutral individual to reveal his subjective probabilities. The exercise presented here extends this approach to the case of a principal-agent problem where neither principal nor agent are risk-neutral, and where their risk preferences may differ. for an early statement of the principal-agent problem, see Ross (1974). For a survey of results in the closely related area of incentive-compatibility, see Green and Laffont (1979).

best (or "sincere") estimate of the likelihood that the chemical will later be proved carcinogenic, but may transmit to the government "strategic" estimates of those probabilities which differ from the sincere, or "best," estimate.

Two sources of the temptation to report "insincere" estimates may be particularly important. First, even if the chemical in question is in fact carcinogenic, detection and correct attribution of the effect may be impossible over the latency period of the chemical.[43] Even at post-exposure times greater than the latency period, detection and attribution may be nearly impossible because of the difficulty of reconstructing individual exposure histories. Second, if health effects are produced synergistically, the situation becomes even more problematic. Identification of a single agent as causally responsible may be logically impossible: in this case health effects are the joint products of exposures to individual chemicals and responsibility for those effects cannot be unambiguously allocated to individual chemicals.[44]

Taken together, these technical and physiological facts strongly suggest that the firm's strategic estimates of hazard probabilities will be significantly smaller than sincere estimates. The incentive design problem can now be spelled out: the government should set its liability schedule to compensate for the firm's temptation to submit strategic estimates of

[43]Much of the evidence bearing on the chronic--as opposed to the acute--health effects of environmental pollution has come from aggregate environmental epidemiology. For a methodological critique of this, both of work and for quantitative estimates of the residual uncertainties in our estimates of health impacts, see the author's forthcoming monograph.

[44]The problem referred to is related to the problem referred to, in the economics literature, as the joint-cost allocation problem; here the joint product is an undesirable environmental contaminant produced synergistically by two effluents. "In principle" if only the synergistic product is hazardous, only one firm should be "discouraged" by the liability scheme from polluting.

hazard probabilities. Roughly speaking, the condition determining the best liability rule is: the firm's strategic hazard probability estimates, when calculated in full knowledge of the liability or indemnity conditions imposed by the government, should coincide with the "best," or "sincere" probability estimates it would transmit to the government in the absence of any temptation to strategic misrepresentation.

Some technical difficulties arise in implementing this condition. Nevertheless it should be apparent even prior to detailed formulation and calculation that the implied liability or indemnity provisions for misrepresentation are quite large. An example may help: say the strategic hazard assessment is a factor of 10 smaller than the sincere estimate of carcinogenic potency. Then the imposed penalty for misrepresentation might have to be as large as 10 times the expected loss to redress the balance. If, for example, $1,000,000 in losses result from a single chemical, the corresponding liability for misrepresentation might have to be as high as $10,000,000.

Before continuing, we remind the reader that we have treated only part of the incentive design problem, and in doing so have deliberately set aside the design requirement that the incentive scheme prompt the right amount of information generation by firms.[45]

[45]We separated the two requirements by assuming that firms know the truth about their chemicals. Of course, they typically do not. Examination of the case in which they do not leads to some technically difficult and intriguing problems in economic theory: here we can do no more than provide a clean statement of those problems, which is often half the battle.

First suppose that the incentive scheme has successfully transferred to the firm the government's attitude towards risk, summarized by the utility function u. Suppose further that the firm can generate information only by repetition of one kind of test, each repetition costing e dollars. If the firm begins with a prior probability distribution $t_F^{(0)}(s)$, then successive tests lead to successive Bayesian revisions of the probability distribution which we call $t_F^{(n)}(s)$. The government's purposes will have been served if the firm tests n^* times, where n^* solves the maximization problem

$$n \max E\left[u(c(s)-ne)\right]$$

3.2. A Mildly Technical Recapitulation

Usually when we think of incentive and liability arrangements for toxic chemicals, we have in mind the liability firms should bear for damages their products may inflict on individuals. If those liability schemes are well designed, they provide incentives to firms to avoid inflicting such damages. But here we have something very different in mind: incentive schemes aimed at insuring the accuracy and veracity of information on new chemicals submitted by firms to the government regulatory agency[46]. For new chemicals, such information is necessarily an important basis for government decision making. Willfully transmitted misinformation can seriously distort that decision process. How can liability for such misinformation save us from that fate?

Here is a very simple setup for the design of a liability scheme. Introduce the following notations:

u_G	Government (or social) utility
u_F	Firm (or private) utility
b	Social benefits of single chemical use
S	Set of states of nature; s ε S

The expectation is computed with respect to the distribution $t_F^{(n)}x$.

As stated this problem is already respectably complex: it is a nonlinear integer programming problem. But now remember that, superimposed upon this problem, we have the principal agent problem: in general the firm utility function u_F and the government utility function u_G differ, and the incentive scheme must force the results they give to coincide. We have a novel combination of two staples of the economics of uncertainty: the principal-agent problem and the search problem.

[46] Again, for the relevant strands in the economics literature, see note 42.

D	Detection/nondetection states; d = 1 (detection), d = 0 (nondetection)
$c(s)$	Social cost of single chemical use
$h(s)c(s)$	Civil liability schedule
$h(s,d)$	Firm's perceived civil liability schedule
$\pi(s)$	True distribution
$m_F(s)$	Firm's optimal revealed distribution (or message) to government
$t_F(s)$	Firm's best-estimate distribution
γb	Profit accruing to firm
$P_F(d)$	Firm's detection/attribution estimate
$S \times D$	Set of slates of nature

In this setup we have introduced two utility functions u_G and u_F, representing the risk attitudes of the government and the firm. What u_G should be is a matter of public policy; presumably only empirical work can tell us what u_F is.[47] Both firm and government agree that net benefits associated with introduction of the new chemical into commerce are equal to b, and the firm knows that a fraction γ of those benefits will accrue to it as profit, for a total profit of γb.

But what is uncertain are the social costs imposed by the chemical, and the government's prospects for correctly identifying the chemical as an offender in the event that it does prove troublesome. The former class

[47]Under the extreme assumptions of the Arrow-Debreu (1959) model, firms are risk-neutral and maximize, expected profits; when complete futures markets exist, all risk is borne by "consumers" or by the "government". But, because complete futures markets do not exist and for other reasons, real-world firms are not risk averse. Just what the relevant function u_F is must therefore be established empirically. Similarly, choosing u_G is a question of public policy. The issue is how risk-averse do we want to be about toxic chemicals-related risks?

of uncertainties is captured by the set S of possible toxicity/exposure states s; the latter class is captured by D, the set of possible detection/attribution states. Taken together, the set $S \times D$ is the set of states of nature, and element of that set (s,d) identifies a particular toxicity and detection state.

In the toxicity states s, the chemical in question imposes social costs $c(s)$. Neither firm nor government, of course, know which s is the true state; the whole point of the communication process about to be described is to encourage the firm to submit the appropriate information.

The information actually submitted to the government by the firm we write as $m_F(s)$, with m for "message:" that information is in the form of a probability distribution on the combined toxicity/exposure states of the chemical. This may or may not be identical with the firm's "best," or sincere, estimate of the chemical's hazards. The latter we write as $t_F(s,d)$, indicating that the firm has a joint distribution on the states of nature $S \times D$. In some cases that joint distribution may factorize conveniently. Below, for purposes of exposition, we will assume such factorization, so that:

$$t_F(s,d) = t_F(s)\, p_F(d)\ ,$$

with $p_F(d)$ the firm's marginal distribution on D.

We are aiming toward formalization of this situation: the firm transmits its message $m_F(s)$ to the government, and the government decides, on the basis of that message, to either bar introduction of the chemical into commerce, or to allow such introduction. Suppose the firm knows the government's decision rule; that is, the firm knows how the

government maps the received message into a decision. Then it is apparent that the temptation to misrepresentation exists In principle, the firm can decide what it wants the government to do about the chemical in question--presumably allow introduction, since otherwise no application for such introduction will be made to the government--and then send a message which will lead the government, "by the nose," to that conclusion.

Against that temptation, the government can array only a liability schedule $h(s)$, a fee to be assessed by the government, against the firm, in the event that subsequent events suggest that the original message was in fact a misrepresentation.[48] The principle upon which the liability schedule $h(s)$ should be designed seems clear enough: it should guide the firm to reporting of its best-estimate probabilities. Message $m_F(s)$ and sincere belief $\pi_F(s)$ should coincide, or at least be close enough so that the government agency is not led to the wrong decision.

Let us see if this can be arranged "in general"--that is, without imposing special assumptions on the state spaces and functions. This can be done by examining a particularly simple case which nonetheless captures most of the difficulties of the general case. In that simple case, the space of states of nature $S \times D$ is as simple as can be: $S=[0,1]$, with 0 for nontoxicity and 1 for toxicity, and $D=[0,1]$, with 0 for nondetection and 1 for detection.

Figure 2 illustrates the expected utility maximization calculations which firm and government must perform in this particular case. The

[48] Again, see Judge David Bazelon's address, cited in note 41.

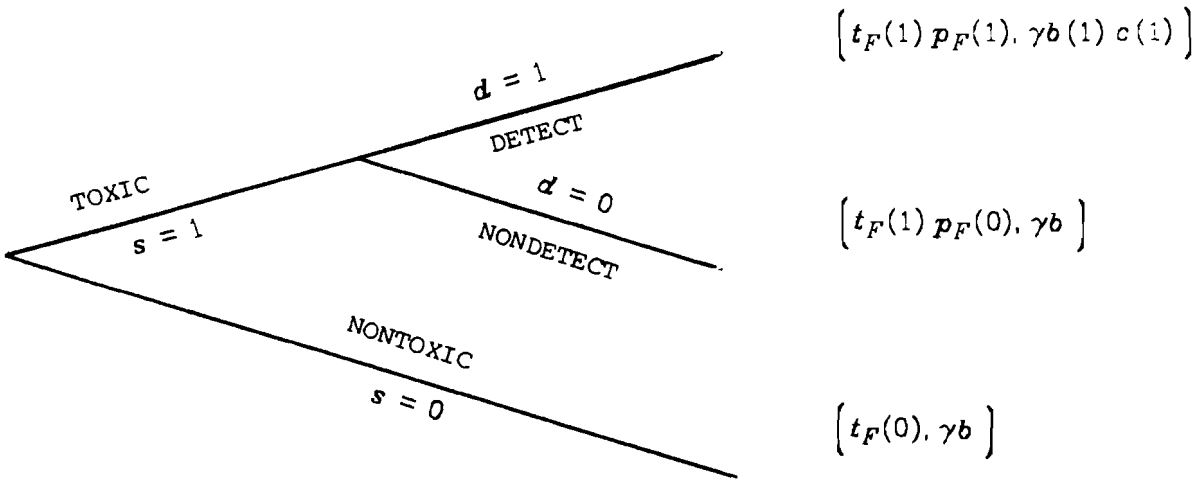

Figure 2a. Decision Problem of the Firm

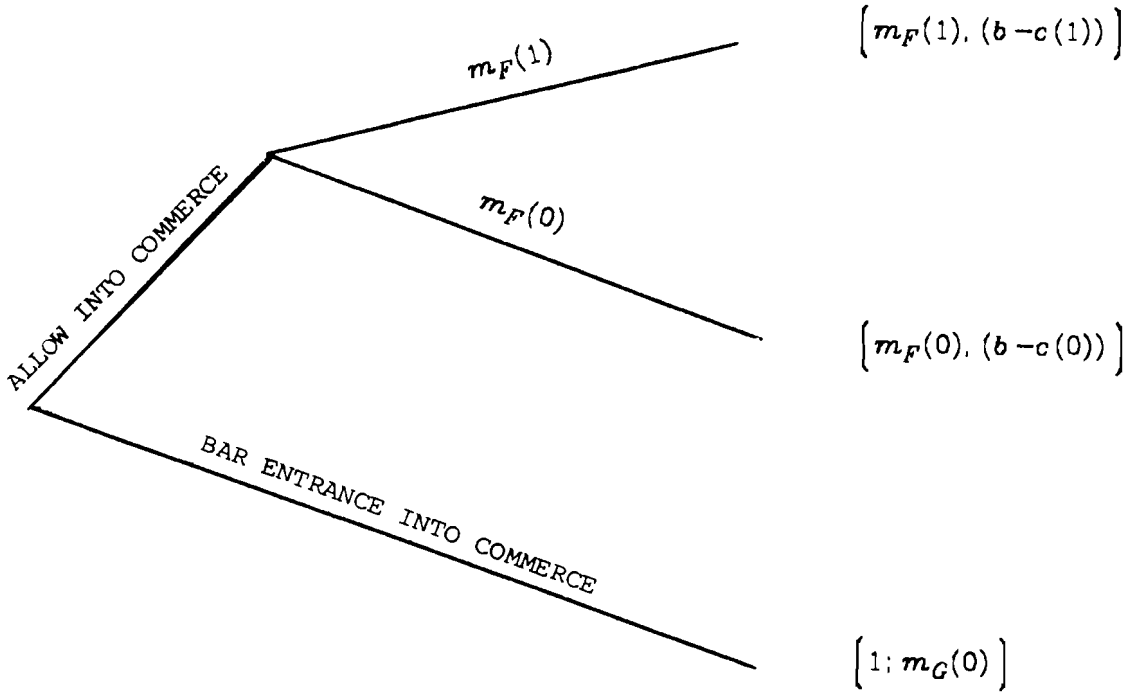

Figure 2b. Decision Problem of the Government

firm (Figure 2a) chooses a message $m_F(0)$, $m_F(1)$; the government (Figure 2b) takes that message as given and then computes expected social utility from the two alternatives "bar from entrance into commerce" and "allow into commerce." If the chemical is barred, social utility is $u_G(0)$, since no benefits or losses accrue; 0 does reflect the opportunity costs of benefits b foregone by restriction. If the chemical is not barred, the government proceeds to calculate expected social utility with the firm's message probabilities taken at face value. For example, if the chemical proves toxic, social utility is $u_G(b-c(1))$. Thus, we can write the government's decision rule $g_G[m]$ as a functional on the firm's message m. First compute

$$g_G[m] = m_F(0)\, u_G\,(b-c(0)) + m_F(1)\, u_G\,(b-c(1)) - u_G(0)$$

and then decide according to: if $g_G[m]$ is positive, allow into commerce; otherwise do not allow into commerce.

We return now to the firm's problem. The firm is armed with sincere estimates $t_F(s)$ (t for "truth") of the probabilities that the chemical is toxic, and with an estimate $p_F(d)$ of the probability of detection/attribution in the event that the chemical does prove toxic. Its decision rule is therefore based upon a computation of the functional $f_F[t,p,h]$ of those probability distributions and of the civil liability schedule $h(s)$. That functional is given by:

$$f_F[t,p,h] = t_F(1)\, p_F(1)\, u_F\Big[\gamma b - h(1) - c(1)\Big] + \Big[t_F(1)p_F(0) + t_F(0)\Big]\, u_F(\gamma b) - u_F(0).$$

The issue of liability schedule design can now be posed in a more or less precise way. Can the government choose $h(s)$, the civil liability schedule, so that the firm is led to submit a message $m_F(s)$ to the government which never, or hardly ever, leads to conflict between government and firm decisions? In other words, can $h(s)$ be devised so that the firm never has an incentive to submit information to the government which will lead the government to decide in the firm's interest, rather than in the more general social interests? If so, we say that $h(s)$ supports consistent decisions. But consistency alone is not enough--there is no particular virtue in consistent, but wrong decisions. We want something more; we want those consistent decisions to be accurate, in the sense that the firm's strategic message m_F is identical with its sincere message t_F. This can, under some additional assumptions, be accomplished by adjusting the dependence of the liability schedule $h(s)$ on the firm's message; for example, the liability payment $h(1)$ can be taken to be proportional to $(1-m_f(1))$. The constant of proportionality is chosen to guarantee, again under additional assumptions (including the assumption that the firm chooses an $m_F(1)$ maximizing the expected utility of profit), both consistency and accuracy.

4. PERSPECTIVE 3: CENTRALIZED ALLOCATION OF TESTING RESOURCES TO FIRMS: TEAM-THEORETIC MODELS

4.1. Literary Recapitulation

In discussing the design of a liability schedule for misrepresentation of the hazards posed by a new chemical, we have assumed that firm and government objectives are divergent. Firms maximize the expected

utility of profits; the government imposes upon firms an incentive scheme guiding those firms toward socially desirable action--the transmission of sincere chemical hazard assessments.

This section builds upon a radically different set of assumptions about the incentives of firms and the government. Here we suppose that there is no divergence between what firms and the government want with respect to toxic chemicals management policy: both are intent upon arriving at toxic chemicals management decisions consistent with social objectives. But here we depart from the informational assumptions underlying Perspective 1. In Perspective 1, all information was assumed costlessly accessible to the government agency charged with toxic chemicals management; in this section we make the (somewhat more realistic) assumption that such information is widely dispersed, and that full informational centralization is impossible, for reasons that are technical, economic, or both.

Under these new assumptions, there may still remain a critical role for the government agency: that role is essentially one of coordination. Coordination can insure that decentralized decisions to generate information on potentially hazardous chemicals are not wastefully duplicative, and that resources expended in generating additional information are allocated efficiently. In this context, efficiency means that more attention is paid to the more serious potential hazards. We want to derive some rules of thumb for the government to use in allocating total chemical testing resources among firms, and in choosing along "communication systems" between firms and the government. Use of the phrase "allocating final testing resources and effort among firms" should not be taken to

imply that all such resources and effort are under the control of the government or "center". Everything that follows is entirely consistent with a situation in which all chemical testing and evaluation funds are internally generated by firms. In both cases, the real issue is the best allocation of total testing resources and effort.

The issues raised by the need to choose among communication systems are somewhat subtle. By "communication system," we mean no more than the kinds of information transferred among firms, and between firms and the government. Some such system will necessarily be chosen, or will evolve, for even though decisions on particular chemicals will inevitably require particular kinds of communication and attention, an enormous number of chemicals will, over the coming decades, be processed by government agencies charged with toxic chemicals management. Some standardization and routinization of information transfer and decision procedures in inevitable.

Clearly, requiring that all information generated be transferred to all parties is out of the question: all parties would then drown in information. The hard-copy documents generated during the course of a single bioassay on a particular chemical, for example, will typically fill a small truck. The practical issue in communications-system design is thus what kinds of information transfer should be required, both among firms and between firms and the government?[49] We proceed here by examining a particular problem, one which suggests that a natural framework for thinking about these questions must combine team-theoretic notions with

[49]For the existing system, and some discussion of its evolution, see the author's forthcoming monograph on TSCA implementation.

more recent results in the area of computational complexity. Below, in the next subsection, the whole exercise is repeated somewhat more formally.

Figure 3 is schematic of the particular problem of interest. A "center"--the agency regulating toxic chemicals communicates with firms, makes observations on the external environment (does experiments). Firms--too do experiments, and in general communicate among themselves and with the center. (For simplicity two firms have been drawn in Figure 3, but any number might have been drawn.) In general, both the center and individual firms can take actions.

Figure 3 is so general as to be vacuous; it can be made more specific in a great number of ways. In fact, there will be one well-specified version of Figure 3 corresponding to each way of assigning observation, communication and decision responsibilities to firms and to the center. Because the number of possible well-specified versions of Figure 3 is so large, in practice we can do little more than write down, and compare, a few such well-specified versions. And to have a sensible standard for comparison we must have an objective function for the center.

Here, in words, is the promised well-specified example (formal treatment is postponed to the next subsection). There are n firms, each of which produces some amount of the same numeraire consumption good with chemical production-process inputs. Each firm production imposes both occupational and general population exposures to its chemical inputs, and each firm is characterized by a production/transformation frontier linking chemical inputs and numeraire good output. Each firm moreover has developed a subjective joint probability distribution

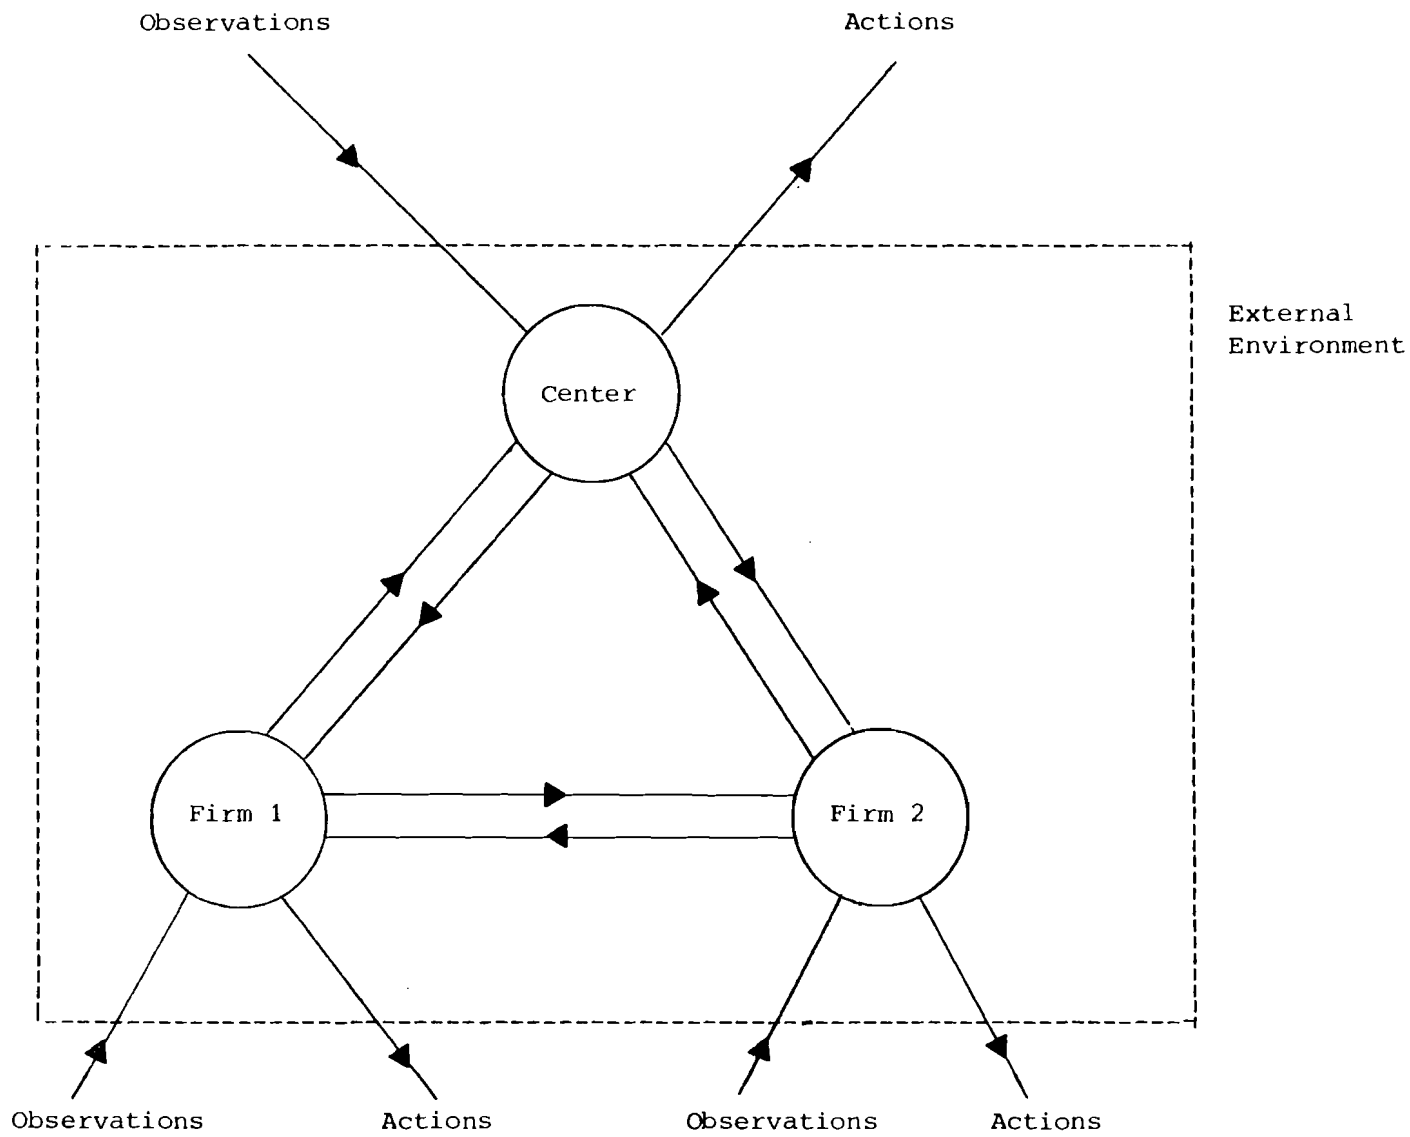

Figure 3. Communication Between Firms and the "Center".

function defined on chemical potencies, and on occupational and general population exposures. That distribution has been built up from experience and previous experimentation. Additional information is available to the firm on that distribution, but at a price: after the fashion of statistical decision theory we assume that the firm can learn more by drawing upon another joint distribution of potency and exposure variables.

The center is endowed with a testing budget (expressed in terms of the numeraire good) and is directed to act in accordance with specified shadow prices for the health risks of occupational and general population exposures. The center can learn more about particular chemicals by drawing on its own joint distribution of potency and exposure, incurring some unit cost with each draw: neither that unit cost, nor that distribution need be identical with those describing the firms. Finally, there are explicit costs attached to communication between firms and the center, and to communication among firms.

Now consider two proposed alternative assignments of observation, communication and decision responsibilities. In the first, all testing is done by firms, and the testing budget is allocated by the center based upon received summary measures of firms' joint distributions. In the second, all testing is done by the center, again based upon received summary measures of the firms' joint distributions. In both cases, final abatement decisions--decisions to trade off numeraire-good output against occupational and general population exposures--are taken by firms subject to the mandated social shadow price of health risk.

Which alternative is superior? The novel feature of the framework we have introduced is the possibility of answering this question. In principle, the answer is simple: the alternative offering the higher expected net benefit. The relevant benefit is numeraire good output net of health risk and testing costs. The latter are calculated at the relevant social shadow prices.

4.2. A Mildly Technical Recapitulation

Now let us repeat the example of the previous subsection somewhat more formally, and then tie that more formal version to the extensions of team theory that we believe are demanded by the issues at hand. Let us be brave, and introduce the following notation and variables:

I	Set of firms; $i \varepsilon I$
k_i	Potency of i^{th} chemical
q_i	Numeraire-good output of i^{th} firm
x_i	Occupational exposures to i^{th} chemical
u_i	Random variable for experiments on i^{th} chemical
$n^{(i)}$	Number of tests on i^{th} chemical by i^{th} firm
$n^{(i,g)}$	Number of tests on i^{th} chemical by the government
c	Unit cost of testing
C_T	Total testing budget constraint
$t_{(0)}^{(i)}(k_i, q_i, x_i)$	Firm i prior (joint) distribution
$t_{(1)}^{(i)}(k_i, q_i, x_i)$	Firm i posterior (joint) distribution
$e^{(i)}(k_i, q_i, x_i, u_i)$	Joint distribution describing experiments available to firm i

$t_{(o)}^{(i,g)}(k_i, q_i, x_i)$	Government prior (joint) distribution on chemical i variables
$e^{(i,g)}(k_i, q_i, x_i, u_i)$	Joint distribution describing experiments (on chemical i) available to the government
$\bar{t}_{(1)}^{(i)}(k_i, q_i, x_i)$	Summary measure of firm i posterior joint distribution
$\bar{t}_{(1)}^{(i,g)}(k_i, q_i, x_i)$	Government posterior (joint) distribution on chemical i variables
p	Consumption-goods shadow price of health risk

The setup we are after should describe the following situation. Each distinct firm uses one chemical to produce some quantity of the numeraire good q_i, and in so doing also produces occupation exposures x_i to that chemical. Because the firm's knowledge of the technological relationship between output and occupational exposure is imperfect, the firm's initial knowledge of that relationship is described by a joint probability distribution $t_{(o)}^{(i)}$. Because the firm is uncertain about the chemical's potency k_i, that variable also occurs as an argument of $t_{(o)}^{(i)}$, even though sensible experiments to improve information about the value of that variable will be very different from those that will improve information about the other arguments of $t_{(o)}^{(i)}$.

Now the firm can obtain better information on the relationship between the variables k_i, q_i, x_i, but only at a price: by paying the unit cost of testing c the firm can make an observation on the random variable u_i, which is tied to the variables of interest by the distribution $e^{(i)}(k_i, q_i, x_i, u_i)$. That latter distribution is known to the firm, and is a standard way of describing experimental possibilities. If the firm does buy information by testing chemical i, we assume that the purchased

information is used to revise the firm's prior distribution $t\binom{i}{0}$. We assume that the revision method is Baye's Rule, and the resulting (posterior) distribution is labeled $t\binom{i}{1}$.

The information available to the firm is embodied in $t\binom{i}{1}$, and if the information is very detailed that distribution may be thought of as parametric, but with many parameters. In general, transmission of more detailed information can be more costly, so that only some summary measures of any given distribution will be transmitted to the government. We use $t_{(1)}^{-(i)}$ to denote the summary measures of the distribution actually transmitted to the government. In principle any number of such summary measures, from zero to a full set of summary measures completely specifying the distribution, can be communicated. Any particular realization of Figure 3 will, of course, require a commitment to one such set, and that set will then describe the messages sent by firms to the government.

Now let us turn to the government. The government has resources C_T to be committed to testing, and will allocate those resources based upon its own information and the information obtained from individual firms. In particular, the government has prior distribution $t\binom{i,g}{0}(ki, q_i, x_i)$ on the i^{th} chemical variables, and has in its possession joint distributions $e^{(i,g)}(ki, q_i, x_i, ui)$ describing its opportunities for learning more about each of these chemicals. It has been instructed to use a (social) numeraire-good shadow price of health risk p in making its decisions. The government has also been saddled with the unenviable responsibility of organizing the framework in which testing and regulation will proceed: that entails, at the very least, deciding who should do what

testing, deciding how the results should be communicated among firms and the government, and choosing decision rules for chemicals.

Remember that it is specific realizations of the general scheme of Figure 3 that we wish to rank, and that, in order to do that ranking, we need an objective function. Here is one such particular realization, and one such function. In the realization, both firms and the government do testing, in accordance with a centrally-calculated testing program $(n^{(i)}, n^{(g,i)})$: by definition that program dictates the intensity of testing of each chemical by government and firms. Subsequent to execution of the mandated tests by firms, firms communicate to the government summary measures of their posterior distribution: that is, it is the $t_{(1)}^{(i)}$ that are transmitted to the government. Given this information, and given the results of the government testing program, the government revises its own prior distributions $t_{(0)}^{(i,g)}$ arriving at the posterior distributions $t_{(1)}^{(i,g)}$. The testing program is designed by computing:

$$\max \left\{ \sum_{i \varepsilon I} q_i - p \sum_{i \varepsilon I} k_i x_i - c \sum_{i \varepsilon I} \left[n^{(i)} + n^{(i,g)} \right] \right\}$$

subject to

$$c \left[\sum_{i \varepsilon I} n^{(i)} + n^{(i,g)} \right] \leq C_T$$

and with the expectation calculated with respect to the distribution $t_{(1)}^{(i)}$. That distribution recognizes all possible outcomes of experiments described by the distributions $e^{(i)}$ and $e^{(i,g)}$. Final calculation of production levels is done only in the light of the particular results obtained by firms and by the government in their testing programs. Call those latter distributions $t_{(1)}^{(i,g)}(k_i, q_i, x_i)$: then output variables q_i are assigned by

solving

$$\max_{(q_i)_{i\in I}} \left\{ \sum_{i\in I} q_i - p \sum_{i\in I} k_i x_i \right\}$$

The maximization is unconstrained.

That completes our description of one realization of the schematic of Figure 3. The reader can easily see that many others are feasible, and may even be superior according to the (first) objective function. An example: the center might communicate only the shadow price p to individual firms, asking them to estimate the social optimal number of tests for their own chemical. Those estimates, returned to the center, can be compared with the testing budget constraint: if, taken together, they exceed the constraint, an iterative process can be used to ration testing resources.

Many other particular realizations can be described, and may be of interest. But that is work to be done: our purpose here is simply to suggest that a combination of team-theoretic and statistical-decision theoretic notions can sustain some rigorous work on the institutional design problem, and may even enrich team theory in a way that only practical applications can.

To persuade the reader that such enrichment is a possibility, let us indicate the formal relationship between team theory[50] (as formulated by Roy Radner) and the kinds of specific institutional design problems described above. The reader with no interest in the relationship to Radner's work can, and perhaps should, skip the following pages.

[50] See Marschak and Radner (1972), Radner (1972), and Groves and Radner (1972).

Radner's description of a team[51] is sufficiently abstract that it can be summarized with a few pages of notation and several definitions. Finally, we need the following spaces and functions characterizing information access and task assignment in a team:

$N \cup \{o\}$	All agents ("team members") and nature (o)
$T[o,i]$	Observation set of $i \varepsilon N$; $t_{oi}(s) \varepsilon T[o,i]$; $t_{oi}(s) = \left(a_{oi}(s), e_{oi}(s)\right)$ signal and noise
$T[i,j]$	Message space (communication set), i to j_i, $i, j \varepsilon N$
$T[i,o]$	"Action set" of $i \varepsilon N$; $T[i,o] \stackrel{?}{\equiv} A_i$
$T[i]$	Internal output message set of $i \varepsilon N$; $T[i] = \prod T[i,j]$ $j \varepsilon N$ $j \neq i$
$R[i]$	Internal input message set of $i \varepsilon N$; $R[i] = \prod T[i,j]$ $j \varepsilon N$ $j \neq i$
$\tau(i)$	Task function of $i \varepsilon N$; $\tau(i): R[i] \rightarrow T[i]$
N	Network, or formal organization; see below:

$$N = \left\{ N \cup \{o\}; \left[T[i], R[i], \tau(i) \mid i \varepsilon N \cup \{o\} \right] \right\}$$

Figure 4 may help the reader to organize these notions in his or her mind; essentially what we have is a designated set of agents capable of communicating with one another and with the external environment. Among the latter form of communication, we include any activity bringing information within the organizational boundary: literature searches,

[51] See Marschak and Radner (1972).

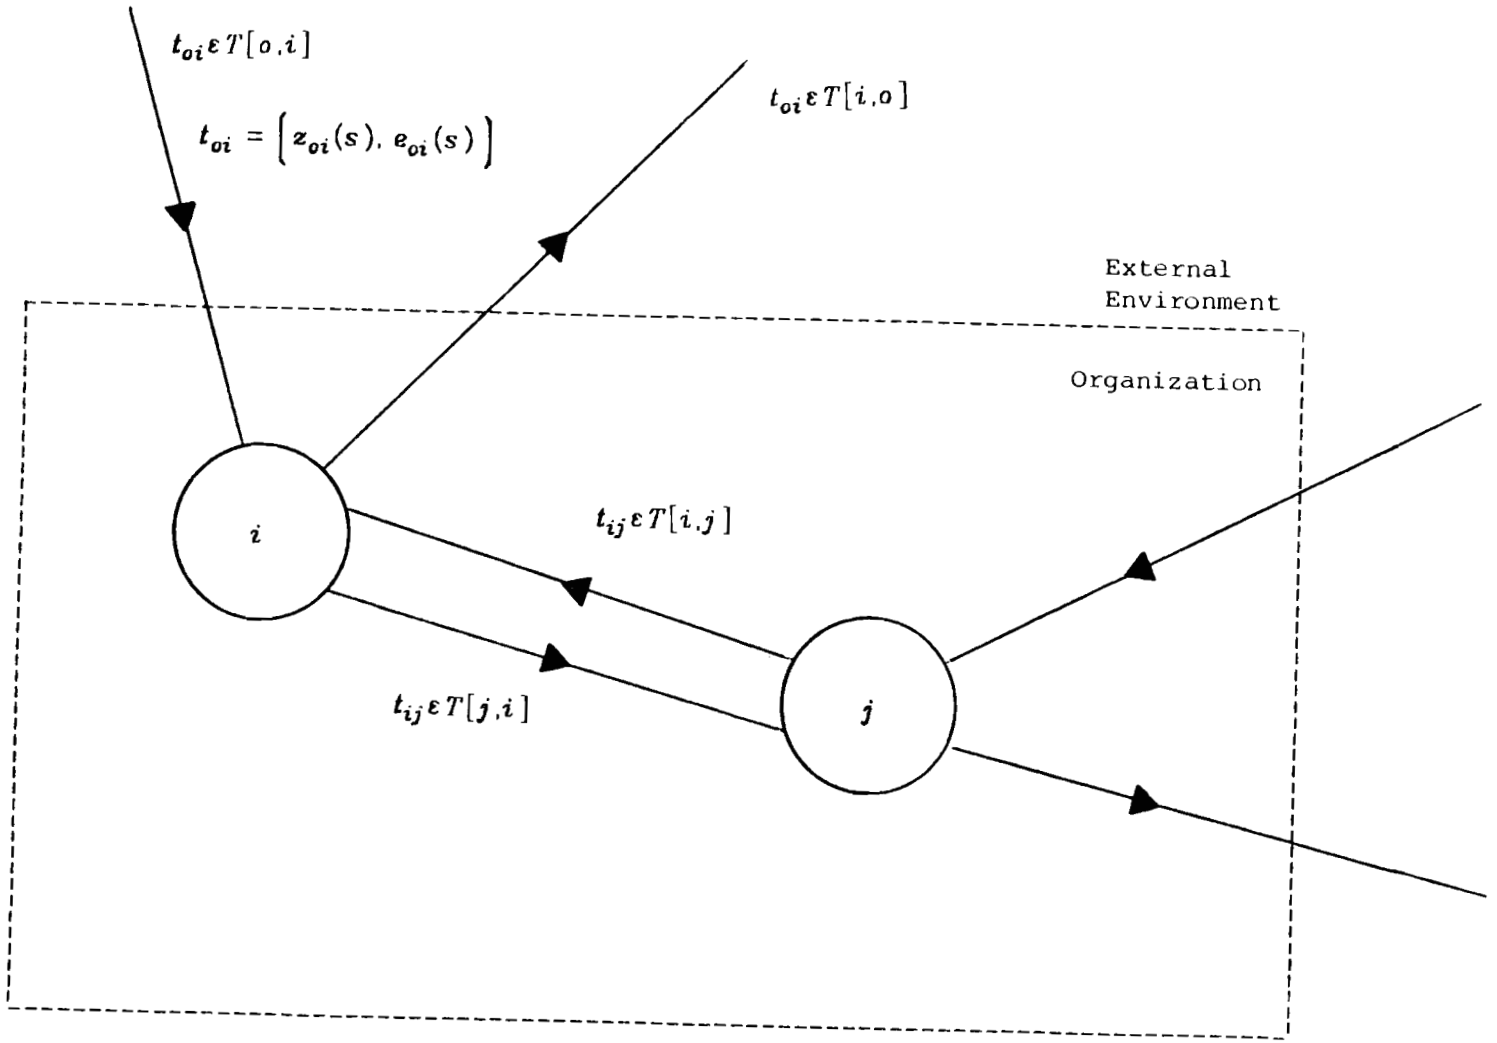

Figure 4. The Message-Task-Signal Network Description of a Team.

discussions with other firms, and laboratory experiments all count here.

Two alternative summary characterizations of teams can be based upon the above concepts; they will differ in their usefulness in guiding empirical investigation. Marschak and Radner (1972), in their seminal book, introduce two such characterizations.

A team in payoff-function form (PFF) is characterized by:

$$T = \{ S, M, A; W_G, C_I \}$$

Remember that the gross payoff and net payoff mappings are real-valued maps:

$$W_G: A \times M \times S \rightarrow R^1$$
$$C_I: A \times M \times S \rightarrow R^1$$

The related expected-value mappings are defined in the obvious way:

$$EW_G: A \times M \rightarrow R^1$$
$$EC_I: A \times M \rightarrow R^1$$

A team in network computation form (NCF) is characterized by:

$$T = \{ S; N; W_G, C_I \}$$

The PFF and NCF characterizations of a team differ as follows. In PFF, the team message sets and team action sets are assumed known; that is, all observation, communication, and decision assignments for all team members are assumed known. The information cost function C_I is likewise assumed known. In NCF, the team message set M and action set A are constructed by the explicit assignment of decision and communication responsibilities to team members. Similarly, the information cost function C_I is built up from explicit assumptions about the unit costs of

the informational functions as performed by specific team agents.

We have presented this summary of team theory only to make the following points: that the relevance and usefulness of the abstract theory depend upon the possibility of computing realistic objective functions for alternative institutional arrangements of practical interest. Moreover those possibilities cannot be even guessed at short of very particular realizations of the abstract constructs of team theory. Consider, for example, the tradeoff between communication costs and the benefits associated with better-informed decisions. In our example, firms communicate summary measures of distributions describing what they know to the government, and governments communicate testing programs to firms. We suspect that particular assumptions about the costs of communicating summary measures of the relevant distributions will be decisive: in any event this can only be examined by computation in some particular examples. But team theory, thus specified, does give us a basis for serious discussions of the merits of alternative arrangements in situations where information and communication costs are nontrivials, as they almost certainly are in the toxic chemicals testing case. Moreover we have a basis for testing the sensitivity of a computed ranking of alternative institutional arrangements to assumptions about the quality of the prior information held by firms and by the government, and to assumptions about the comparative advantages of firms and the government in various kinds of testing.

REFERENCES

Aho, Alfred V., et al. 1974. *The Design and Analysis of Computer Algorithms*. Reading, Mass.: Addison-Wesley Publishing Company.

Ames, Bruch N., et al. 1975. "Methods for Detecting Carcinogens and Mutagens with the Salmonella/Mammalian-Microsome Mutagenicity Test," *Mutation Research*, vol. 31, no. 6.

Balch, Michael, et al., eds. 1974. *Essays on Economic Behavior Under Uncertainty*. Amsterdam: North-Holland Publishing Company.

Bazelon, David L. 1980. "The Judiciary: What Role in Health Improvement?" Paper presented to the Annual Meeting of the Institute of Medicine, October 15.

Blackwell, David, and M.A. Girschick. 1954. *Theory of Games and Statistical Decisions*. New York: Dover Publications.

Cairns, John. 1978. *Cancer: Science and Society*. San Francisco: W.H. Freeman and Company.

Debreu, Gerard. 1959. *Theory of Value*. New York: John Wiley and Sons.

de Finetti, Bruno. 1972. *Probability, Induction, and Statistics*. New York: John Wiley and Sons.

de Groot, Morris H. 1970. *Optimal Statistical Decisions*. New York: McGraw-Hill Book Company.

Devoret, Raymond. 1979. "Bacterial Tests for Potential Carcinogens," *Scientific American*, vol. 241, no.2.

Doll, Sir Richard. 1977. "Introduction," in H.H. Hiatt, J.D. Watson, and J.A. Winsten, eds. *Origins of Human Cancer*. Cold Spring Harbor, L.I.: Cold Spring Harbor Laboratory.

Ferguson, Thomas S. 1967. *Mathematical Statistics: A Decision-Theoretic Approach*. New York: Academic Press.

Gary, Michael R., and David S. Johnson. 1979. *Computers and Intractability*. San Francisco: Freeman and Company.

Green, Jerry R., and Jean-Jacques Laffont. 1979. *Incentives in Public Decision- Making*. Amsterdam: North-Holland Publishing Company.

Groves, Theodore and Roy Radner. 1972. "Allocation of Resources in a Team," *Journal of Economic Theory*, vol. 4, pp.415-441.

Harris, R.H., et al. 1977. "Carcinogenic Hazards of Organic Chemicals in Drinking Water," in H.H. Hiatt, J.D. Watson, and J.A. Winsten, eds., *Origins of Human Cancer*. Cold Springs Harbor, L.I.: Cold Spring Harbor Laboratory.

Hiatt, H.H., J.D. Watson, and J.A. Winsten, eds. 1977. *Origins of Human Cancer*. Cold Spring Harbor, L.I.: Cold Spring Harbor Laboratory.

Interagency Testing Committee, U.S. Council on Environmental Quality. 1977. "Initial Report to the Administrator," *Federal Register*, October 12.

Marschak, Jacob, and Roy Radner. 1972. *Theory of Teams*. New Haven: Yale University Press.

McGuire, C.B., and Roy Radner. 1972. *Decision and Organization*. Amsterdam: North-Holland Publishing Company.

Radner, Roy. 1972. "Teams," in C.B. McGuire and R. Radner, eds., *Decision and Organization*. Amsterdam: North-Holland Publishing Company.

Ross, Stephen A. 1974. "On the Economic Theory of Agency and the Principles of Similarity," in M.Balch, et al., eds., *Essays on Economic Behavior Under Uncertainty*. Amsterdam: North-Holland Publishing Company.

Rothschild, Michael, and Joseph Stiglitz. 1976. "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information," *Quarterly Journal of Economics*, vol. 90, pp.629-650.

Stuper, Andrew J., et al. 1979. *Computer Assisted Studies of Chemical Structure and Biological Function*. New York: John Wiley and Sons.

U.S. Congress, 94th. 1976. *Legislative History of the Toxic Substances Control Act*, Public Law 94-469. Washington, D.C.: GPO.

U.S. Environmental Protection Agency. 1979. *Toxic Substances Control Act Chemical Substances Inventory*. Washington, D.C.: Office of Pesticides and Toxic Substances, GPO.

Wald, Abraham. 1947. *Sequential Analysis*. New York: Dover Publications.