UNITED STATES OF AMERICA
NUCLEAR REGULATORY COMMISSION

DOCKETED
USNRC

'83 FEB -2 A11:06

ATOMIC SAFETY AND LICENSING BOARD
Before Administrative Judges:
James P. Gleason, Chairman
Dr. Oscar H. Paris
Frederick J. Shon

In the Matter of CONSOLIDATED EDISON COMPANY OF NEW YORK (Indian Point Unit 2) POWER AUTHORITY OF THE STATE OF NEW YORK (Indian Point Unit 3)	Docket Nos. 50-247-SP 50-286-SP January 31, 1983

DIRECT TESTIMONY

of

ISAAC LEVI, Ph.D.

On Behalf Of

FRIENDS OF THE EARTH, INC.

and

NEW YORK CITY AUDUBON SOCIETY

D503

8302030210 830131
PDR ADOCK 05000247
T PDR

CONTENTS OF TESTIMONY OF ISAAC LEVI

QUALIFICATIONS OF ISAAC LEVI

I am currently Professor of Philosophy at Columbia University in New York City, where I teach graduate level courses in probability and induction, and the philosophy of science. I was a Guggenheim Fellow and Fulbright Research Scholar in the United Kingdom in 1966 and 1967, and a visiting scholar at Cambridge University in 1973. From 1973 to 1976 I was Chairman of the Department of Philosophy at Columbia University. My curriculum vitae is annexed to my testimony.

For over twenty years I have been engaged in extensive inquiry in the logic of probabilistic judgements and statistical inference, decision theory, and the application of these fields to questions of scientific knowledge and practical action. As can be seen from the annexed bibliography, I have numerous publications in these fields, among them, my recent book, The Enterprise of Knowledge, MIT Press, 1980. The appendix to The Enterprise of Knowledge presents an application of the approach developed in the book to the problem of assessing risk in the nuclear industry. That appendix, in a somewhat less technical version, appeared in Social Research in 1981 as: "Assessing Accident Risks in U.S. Commercial Nuclear Power Plants: Scientific Method and the Rasmussen Report". The discussion in my testimony parallels that of the article, which critiques the Reactor Safety Study (WASH 1400).

SUMMARY OF TESTIMONY OF ISAAC LEVI

1. Summary of Section One. The proponents of nuclear power are often ready to acknowledge the possibility of dire consequences resulting from serious accidents at nuclear power plants; but they hasten to assure us that the probabilities of major accidents or long term damage to the environment are so trivially small that the risk is acceptable. Defenders of reliance on nuclear power plants thus appeal to the principle of "maximizing expected utility". In other words, they take into account the possibility of dire consequences (costs) and beneficial consequences (benefits) and weight each possibility with an estimate of the probability of its occurrence. If, however, probability assessments cannot be made numerically definite, several "risk curves" may become permissible to use in computing "expected utilities", and it may not be possible to choose between options (e.g., whether or not to allow a nuclear plant to operate) on the basis of expected utilities.

Thus, the critical question to address is whether assessments of probability supported by the available data are sufficiently definite to justify a verdict one way or another in terms of "expected utility". If assessments of probability warranted by the available data are not sufficiently definite to justify a verdict in expected utility terms one way or another, it is entirely rational

to appeal to minimax criteria or worst possible case analysis. To do otherwise would be to indulge in an unreasonable form of wishful thinking.

It is widely agreed by both sides of the nuclear debate that the worst possible consequences of permitting nuclear power plants such as Indian Point to run are worse that the worst possible consequences of prohibiting their use. Hence, minimax considerations would argue in favor of refusing to permit the continued operation of the Indian Point plants.

With regard to probabilistic risk assessment at Indian Point, my conclusion is that the methodology described in the Indian Point Probabilistic Safety Study (IPPSS) for making probabilistic assessments is not adequate to justify our reposing condidence in the assessments made in that study.

2. Summary of Section Two. In principle, I am sympathetic to the use of event and fault tree analyses, such as in IPPSS, provided they are carefully and thoroughly implemented with scrupulous respect for possible incompleteness due to oversight. They must also be based on adequate estimates of the probabilities of failure rates for components of the complex system.

This discussion does not address the question of the completeness of IPPSS. However, IPPSS goes seriously astray with regard to estimating probabilities. Although the authors of IPPSS go to some considerable effort

to leave a different impression with their so called "method two" or "probability of frequency" framework and their "level two definition of risk", they have in effect responded to a major criticism by the Lewis Committee of the Reactor Safety Study (WASH 1400) by stonewalling--i.e., refusing to modify their method of reporting subjective or credal probabilities so as to give a range of distributions rather than a single one. Since IPPSS has no basis for making non-arbitrary assignment of definite distributions, its failure to offer a range of distributions represents a serious defect in the study.

3. Summary of Section Three. The preceding defect would perhaps not be so serious if the authors of IPPSS could claim that they have evidence and background information sufficient to justify assigning a definite credal or subjective probability distribution over the possible values of objective failure rates. If, on the basis of the available information, there is no good reason to pick one definite probability distribution rather than another, the choice of any one such distribution rather than another is arbitrary. Rather than make such arbitrary judgements, one ought to regard as permissible all distributions which are not ruled out by the available evidence. In this sense, the judgements of priors made by the authors of IPPSS are arbitrary, just like the judgements of priors made in the Reactor Safety Study (WASH 1400). As a consequence, the IPPSS

verdicts about posteriors cannot be trusted except, perhaps, in those cases where the plant specific data are so abundant as to swamp the impact of an arbitratily chosen prior.

What is required is a recomputation of IPPSS data against a broad spectrum of priors utilizing Bayes' Theorem. Pending such reassessment, one should proceed as if none of the posteriors can be trusted.

4. Summary of Section Four. The claim made both by the Lewis Committee and the reports of Sandia National Laboratories on IPPSS and the Zion Probabilistic Safety Study, that the use of log normal priors has negligible effect on the outcome of probabilistic assessments, cannot be trusted. They assessed the import of one arbitrarily chosen prior (log normal) by comparing it with another arbitrarily chosen prior. Such comparisons are deceptive and of doubtful value.

5. Summary of Section Five. IPPSS assumes (as did WASH 1400) that failure rates are distributed log normally, except that it assigns 60% probability to the "assessed range" rather than 90% probability (as in WASH 1400). This enables the authors to ascertain new 5% and 95% confidence limits. There are two components to this line of reasoning: (1) it is assumed that there is as much probability of a failure rate below the lower limit of the assessed range as there is above the upper limit; and (2) a log normal distribution is used to assess confidence intervals. If we suppose that the prior credal probability

is indeterminate, the use of a log normal distribution is arbitrary and, by the same token, it is far from clear why symmetry should be assumed in extending upper and lower bounds of assessed ranges.

Given the foregoing substantial and serious defects in IPPSS, a procedure for making sensible probability judgements is suggested. Tables are presented for the purpose of illustrating the difference between the arbitrary approach of IPPSS, and an approach which takes seriously the complaint that we lack a basis for making precise estimates of probability and should make assessments in terms of ranges of probability distributions. It should be apparent that data from plant experience will not begin to give definitive judgements until the data become very ample.

Pending a complete reassessment of the data on Indian Point in a manner which respects the importance of reporting the indeterminacies in probability judgement, one should not rely on the analysis contained in IPPSS.

6. Summary of Conclusion. The conclusion reached on the basis of these considerations is that we have thus far no rational grounds for supposing that the probabilities of serious accidents at Indian Point Units 2 and 3 are sufficiently low to justify operating these plants. Minimax considerations, i.e., worst possible case analysis, argue in favor of refusing to permit these plants to continue to operate unless and until adequate reassessment of Indian Point data results in a contrary conclusion.

SECTION ONE

RATIONAL CHOICE AND NUCLEAR POWER PLANTS

Assessing risks and taking decisions on such assessments in determining whether to permit the operation of nuclear power plants raise important issues concerning human knowledge and rational decision making which transcend the particular points which may be under dispute in the formulation of an energy policy. Precisely the same kinds of issues arise in other areas of public policy making such as in the control of various forms of industrial pollution, the licensing of the use of drugs by the FDA, or in the use of military intelligence information in the design of an arms policy. And, of course, they also appear in deliberation about investment policy formation, the design of investigations in pure science, and in other contexts of deliberate decision making.

Opponents of the operation of nuclear power plants focus attention upon the worst possible consequences of reliance on nuclear power. They invite us to consider the consequences of a major accident and of the possible long term damage to the environment due to such an accident or to the storage of radioactive waste.

Proponents are often ready to acknowledge the possibility of such dire consequences; but they hasten to assure us that the probabilities of major accidents or long term damage to the environment are sufficiently small

to warrant the risk.

Defenders of reliance on nuclear power plants thus appeal to the principle of "maximizing expected utility". In other words, they not only take into account the possiblity of dire consequences (costs) but the possibility of beneficial consequences (benefits) as well. They weight each possibility with an estimate of the probability of its occurence. Such weighting yields an assessment of the expected benefits and costs of reliance on nuclear power. The expected value of allowing nuclear power plants to run is claimed to be superior to the expected value of alternatives and, therefore, they conclude that operation of nuclear power plants is to be recommended.

From the point of view of expected utility maximizers, those who focus on the worst possible case without taking into account the probability of its occurrence are irrational. The latter are invoking the so called "minimax" criterion for rational decision making which favors choosing that option from a range of alternative policies for which the worst possible consequence (or "security level") is no worse than the worst possible consequence of any alternative. They argue that the worst possible consequences of operating such plants, such as the consequences that would flow from a serious reactor accident, are much more threatening to individuals and society than would be the worst possible consequences of

some feasible alternative.

Even though the minimax criterion for decision making under uncertainty was explored and taken seriously by very eminent statisticians[1] (such as A. Wald), the various versions of minimax theory (minimax loss, minimax risk, minimax regret) have come in for criticism. The chief line of such criticism is that minimax theory counsels us to adopt an unreasonable paranoia or pessimism concerning what "Nature" is likely to do. What is wrong with this view, so claim the critics of minimax, is that it leads us to assume that what Nature is likely to do depends upon our values--what we judge to be losses and gains. If we change our values, what Nature is likely to do will change as well. If wishful thinking is to be avoided, say the critics of minimax, so should the negative of wishful thinking.

But minimaxers should not be bullied by maximizers of expected utility into conceding their own irrationality. There are situations where minimaxing makes very good sense. The blanket and indiscriminate sniping at minimaxing in the name of Reason can only serve to give rationality a bad name.

The principle of maximizing expected utility is entirely cogent as a criterion of rational judgement when one can identify a precise "risk curve"(using the jargon of the Indian Point Probabilistic Safety Study (IPPSS))

---

1 - The locus classicus is A. Wald, On the Principles of Statistical Inference, Notre Dame, Indiana, University of Notre Dame Press, 1942.

which pairs with each possible "scenario" a probability of occurrence and an assessment of benefit or loss. To do this, however, one must be in a position to make numerically definite probability and utility assessments or, at least, to make assessments sufficiently definite so that the ranking of the policy options with respect to expected utilities is clear.

If, however, assessments cannot be made numerically definite, several risk curves may become permissible to use in computing expected utilities. In that event, it can happen that one option is optimal in expected utility according to one permissible assessment whereas an entirely different one is optimal according to another. In that event, considerations of expected utility become impotent for the purpose of deciding between the options.

Thus, if it should turn out to be the case that the probabilistic assessments made by various safety studies are too indeterminate to allow us to render a verdict as to the merits of operating or refusing to operate nuclear plants on the basis of an appeal to considerations of expected utility, we can no longer appeal to the argument that the probability of a serious accident is sufficiently low to warrant permitting nuclear generating plants to operate. The point is not that we can claim that the probabilities are not sufficiently low. We may not be in a position to decide whether the probabilities are low enough or not.

If that should turn out to be the case, it makes eminently good sense to appeal to another criterion which makes no assumptions about expected utility or probability of occurrences of diverse scenarios to render a verdict.

When neither permitting nor prohibiting the use of nuclear power is ruled out by appeal to considerations of expected utility, it is perfectly reasonable to focus on the worst possible consequences of the several alternatives and pick the option for which the worst possible consequences of that option is better than that of any alternative. Minimaxing comes into its own when maximizing expected utility fails to render a verdict.[2]

When a minimax criterion is used as a secondary principle in those instances in which maximizing expected utility fails to render a verdict, the objection that minimax is a paranoid practice carries no weight. That objection is grounded on the assumption that someone who uses a minimax solution is proceeding as if his probability judgements were those which made that solution optimal in expected utility. That is not the way the criterion is employed when it is used as a secondary principle. In this latter use, the criterion is employed as a substitute principle when no judgement about probabilities is sufficiently definite to allow us to make a judgement

---

2 - This idea is elaborated in some detail in my book, The Enterprise of Knowledge, Cambridge, Mass., MIT Press, 1980, especially chapters 4 and 7.

on expected utility. Appeal to worst possible case analysis (or to the consideration of "security levels") is intended to introduce a consideration distinct from expected utility in order to render a verdict when expected utility cannot do so. No assumptions (not even "as if" assumptions) about probabilities are made.

Objection may, perhaps, be made that minimax criteria are ambiguous in their application. Thus, Wald favored using what is often called a minimax risk criterion which differs from a minimax loss criterion and also from a minimax regret criterion. If it is claimed that there must be some firm decision made between these criteria as secondary criteria as a matter of rational principle, I quite agree that the difficulties become considerable. However, on the view outlined here, how the agent assesses security levels for the various options is a matter of ethical, political, economic, aesthetic or other value judgement just as is his assessment of the utility or cost of each possible "scenario". The choice of a method for fixing security levels of feasible options ought not to be settled in advance by criteria for rational choice but is itself part of the value commitments of the decision makers or the communities whose interests they serve.

I assume that it is widely agreed that the worst possible consequences of permitting nuclear plants such as Indian Point to run are worse than the worst possible

consequences of prohibiting their use. Perhaps this assumption is mistaken. But it does not appear to have been disputed by the advocates of nuclear power. And critics of the use of nuclear power seem to think the assumption sound.

The main thrust of advocates of nuclear power has been to belittle the relevance of pointing to the worst possible consequences of using nuclear power plants. They do so by insisting that everything depends on which policy bears maximum expected utility. We are urged to look at the probabilities that accidents and other dire consequences will occur and the probabilities that they will not occur, to calculate the expected benefits and losses of operating nuclear power plants accordingly, and to compare these expectations with those associated with other options.

Opponents of nuclear power insist on the relevance of focusing on worst possible consequences. They do not seek to show that the prohibition of nuclear power bears greater expected value than promotion. If they could establish this, they would prove their case. But such an argument is not necessary for the defense of their view. If assessments of probability warranted by the available data are not sufficiently definite to justify a verdict in expected utility terms one way or another, it is entirely rational to look at security levels. Since it seems to be agreed that security levels are higher if one does not

operate nuclear power plants than if one does, it is entirely rational to invoke a minimax argument and to recommend prohibition of the operation of nuclear plants.

Thus, the critical question to address is whether assessments of probability supported by the available data are sufficiently definite to justify a verdict one way or another in expected utility terms.

My conslusion is that the methodology described in the IPPSS for making probabilistic assessments is not adequate to justify our reposing confidence in the probabilistic assessments made in that report. Without a thorough reassessment of the data from the IPPSS and acquisition of additional data, we should conclude that considerations of expected utility do not warrant deciding one way or the other whether Indian Point Units 2 or 3 ought to be permitted to continue to operate. To conclude otherwise is to indulge in an unreasonable form of wishful thinking.

Without such reassessment and the determination of sufficiently definite probabilities by appropriate methods, rational decision makers ought to maximize security by refusing to permit the continued operation of Indian Point Units 2 and 3.

SECTION TWO

ON THE PROBABILITY OF FAILURE

I shall not comment on the use of event and fault tree analyses by the Indian Point Probabilistic Safety Study (IPPSS) in deriving probability distributions for the responses of Indian Point Units 2 and 3 to diverse initiating events. In principle, I am sympathetic to the use of such approaches provided they are carefully and thoroughly implemented with scrupulous respect for possible incompleteness due to oversight. But I lack the engineering knowledge needed to assess the extent to which the authors of IPPSS did, indeed, apply event and fault tree analyses in an effective and complete manner.

The advantage of using such analyses is that they enable us to exploit judgements as to the probability of failures of components to operate as intended in order to estimate probabilities of serious consequences--such as core meltdowns and failures of containment systems. Data concerning the performance of the overall complex system may be too sparse while information concerning the performances of components may be better and might, given an adequate model of the function of the system, be used to derive the desired overall probabilities.

But if this advantage is to be exploited, not only must our analyses of the complex systems be adequate so that we have a good understanding of how probabilities

of initiating events are propagated, we must also have adequate estimates of the probabilities of failure rates for components of the complex system. IPPSS provides an account of the methodology employed by it in making these estimates. It purports to explain methodological improvements over the Reactor Safety Study (WASH 1400) developed in response to the Lewis Committee Report[3]. In what follows, I explain by means of an illustration why their improvements are illusory.

Consider then the problem of ascertaining a probability of a given pump failing to perform as intended in the interval from time t to time t + dt for small dt. If we assume that the pump is subjected to a regular schedule of maintenance and inspection, we may assume that there is an objective statistical probability or, as I shall call it, <u>chance</u> p of the pump failing to perform in an arbitrarily selected interval of time of small length dt. Suppose we then can think of the interval from time t to t + dt as selected "at random" from all such intervals in the history of the pump. We are then justified in adopting as our subjective or credal probability that the pump will fail in that particular interval the value p (that is, the chance of the pump failing to perform in an arbitrarily selected interval of time of length dt).

Under the assumptions specified, the chance p identified is taken to be approximately equal to λdt where

3 - NUREG/CR-0400, Risk Assessment Review Group Report to the U.S. Nuclear Regulatory Commission, 1978.

λ(lambda) is a parametric value characterizing the failure rate of the pump. If the failure rate λ of the pump is known to be equal to λ*, then the chance is known and from that we can justify a given credal probability judgement.

Contrast the aforementioned predicament 1 where the agent knows that $\lambda = \lambda^*$, with predicament 2, in which the agent does not know the true value of λ but knows that it falls in a given range of values. What subjective or credal probability should the agent assign to the hypothesis that the pump will fail in the interval from t to t + dt?

Let me distinguish three cases:

Case 2a. Let the agent have credal probability judgements concerning the probabilities of various values of λ. That system of judgements is characterized by a probability distribution and, if the distribution is continuous, will be representable by a density function $f(\lambda)$, or if discrete, by a probability function $M(\lambda)$. Given such a distribution, we can then compute an expected value for λ. That will be $\int \lambda f(\lambda) d\lambda$ or $\sum \lambda M(\lambda)$, as the case may be.

Relative to that distribution and the expected value of lambda computed therefrom, the agent is then committed to a way of assigning a credal probability to the hypothesis that the pump will fail in the interval from t to t + dt. It is equal to the expected value of lambda.

In case 2a, I suppose that this expected value is equal to λ^* of predicament 1.

Thus, the agent in predicament 2a makes the same credal probability judgement concerning what will happen in the interval from t to t + dt as does the agent in predicament 1. The important difference between these two predicaments concerns the grounds on which they make their judgements. In predicament 1, the agent knows the failure rate and, hence, the objective chances of a failure in an interval of length dt. He knows no such thing in case 2a but has, instead, a subjective opinion expressed by a credal probability distribution over possible failure rates.

Case 2b resembles case 2a except that the credal probability distribution does not yield the expected value for lambda equal to λ^* but, say, to some other substantially different value λ^{**}. And that is used to determine the credal probability of failure in the interval from t to t + dt.

Case 2c also is a case where the agent does not know the true value of λ. However, in this case, the agent does not represent his uncertainty by a single credal distribution over the possible values of λ. Instead, it is represented by a set of permissible credal distributions (which is required to be convex). These are the distributions the agent has not ruled out for use in making probabilistic assessments and computing expectations.

And from these distributions, he obtains an interval of expected values of λ. This interval might, for example, be the interval λ^* to λ^{**}.

Observe that the set of values in the interval in case 2c is not a set of unknown values of the objective failure rate. Rather, it represents the agent's indeterminate state of subjective or credal probability judgement concerning failure in the interval from t to t + dt. Instead of having a numerically definite credal judgement as in cases 1, 2a and 2b, he has a numerically indeterminate assessment.

Suppose that the agent is offered a gamble where he wins S - P utiles (or units of value) worth of prize if the pump fails in the interval from t to t + dt and receives -P utiles otherwise.

According to Bayesian doctrine, the agent should accept the gamble (rather than get nothing for sure) if S is positive and P/S is less than λ^* in cases 1 and 2a, and also if S is negative and P/S is greater than λ^* in these cases.

That is to say, there is not the slightest difference in the way the agent should make decisions in these two cases on his credal probability judgements concerning failure. The fact that his probability judgements are grounded on different information in the two cases should not make the slightest difference in cases 1 and 2a.

The analysis of 2b is like 1 and 2a except that λ^{**} is substituted for λ^*.

However, Bayesians can offer no analysis of case 2c; for they have nothing to recommend when probability judgement goes indeterminate. To be sure, if S is positive and P/S is less than λ^*, the expected utility of the gamble is positive no matter what permissible distribution is used; and if S is negative and P/S is greater than λ^{**}, the same is true. But when P/S is in the interval from λ^* to λ^{**}, the gamble bears positive expected utility according to some distributions and negative according to others.

This analysis recommends refusal to gamble in such cases on the basis of the fact that refusal is the minimax solution. This recommendation coincides with the recommendations of C.A.B. Smith according to the theory of upper and lower "pignic" probabilities.[4]

I rehearsed the similarities and differences between cases 1 and 2a, 2b and 2c in order to emphasize a critical ambiguity in the way an important observation made by the Lewis Committee Report may be understood. According to the Lewis Committee Report "it is preferable not to try to come up with a point estimate - a single number - for a failure probability, but rather to content oneself with bounds".[5]

---

4 - C.A.B. Smith, "Consistency in Statistical Inference and Decision", J. Royal Statistical Soc. Ser. B, vol. 23 (1961), pp. 1-25.

5 - NUREG/CR-0400, pp. 8-9.

This recommendation may be construed as advice not to fix on a definite value of λ as the true value in deciding on a definite numerical judgement of credal probability as in case 1, as to failure in the interval from t to t + dt.

Alternatively, the recommendation may be construed as advice not to fix on a definite numerical judgement of credal probability for failure in the interval from t to t + dt. If one heeded this advice, one would adopt a posture rather like that in case 2c where one not only avoids making any definite judgement of the true value of λ but also avoids adopting a definite credal distribution over the values of λ, resting content with a family of values.

A fair reading of the Lewis Committee report suggests that they did intend the recommendation in the second sense. However, the authors of IPPSS, in responding to the criticisms made by the the Lewis Committee report of the Reactor Safety Study (WASH 1400), apparently have adopted the first interpretation. Hence, the authors of IPPSS have claimed that their "method two" or "probability of frequency" approach as described in sections 0.4.5 and 0.4.6 of IPPSS should be used to represent risk according to the "level two" definition of risk of sections 0.5.2 to 0.5.4 of IPPSS.

In a footnote to the opening paragraph of section 0.5, the authors of IPPSS claim that the valid point lurking behind "the major criticism by the Lewis Committee"

concerns the need to resort to the level two definition of risk to express uncertainty about risk assessments and failure rates.

The illustrations and discussion in IPPSS make it clear that the authors intend discussions of failure rates to follow the patterns illustrated by my cases 2a or 2b. According to IPPSS, we are to desist from attempts to come up with a single number as our estimate of the true value of the objective failure rate for the pump (or whatever other component is being considered) but rather assign a credal probability to the various hypotheses possibly true concerning the true value of the unknown λ.

The upshot is that even though IPPSS does desist from making a commitment as to the true value of some objective statistical probability or chance (or frequency, as the authors of IPPSS misleadingly call it), it is still committed to a definite numerical assignment to the expected value of the failure rate λ, and this expected value then becomes (or should become according to Bayesian doctrine to which the authors of IPPSS are apparently committed) the credal or subjective probability assigned to the hypothesis that the pump will fail in the interval from t to t + dt.

As we have seen when comparing case 1 and case 2a, when it comes to deciding whether to take risks, it makes little difference whether one takes the posture in case 1 or case 2a.

Furthermore, it appears that the method two favored by IPPSS is already incorporated in its essentials in the Reactor Safety Study (WASH 1400), at least when assessing failure rates of individual components.

Thus, although the authors of IPPSS go to some considerable effort to leave a different impression with their so called "method two" or "probability of frequency framework" and their "level two definition of risk", they have in effect responded to a major criticism by the Lewis Committee of the Reactor Safety Study (WASH 1400) by stonewalling--i.e., refusing to modify their method of reporting subjective or credal probabilities so as to give a range of distributions rather than a single one. Since IPPSS has no basis for making non-arbitrary assignment of definite distributions, its failure to offer a range of distributions represents a serious defect in the study, regardless of the correct reading of the intent of the Lewis Committee report.

SECTION THREE

BAYES' THEOREM

The observation made in Section Two would, perhaps not be so serious--although the failure to be straightforward would still be annoying--if the authors of IPPSS could claim that they have evidence and background information sufficient to justify assigning a definite credal probability distribution over the possible values of the objective failure rate λ. IPPSS does describe the methodology it uses to do just that. It is based on the use of Bayes' theorem and is summarized in section 0.14 of IPPSS.

According to section 0.14.1 of IPPSS, there are three types of information available for the "frequency of elemental events"--i.e., objective chances or, in the case of our example of the pump, the objective failure rate λ. They include:

- E_1 General engineering knowledge of the design and manufacture of the equipment in question.
- E_2 The historical performance in other plants similar to the one in question.
- E_3 The past experience in the specific plant being studied.

E_1 and E_2 is called "generic" information and E_3 "plant specific" or "item specific" information.

The task is to determine the probability $p(\lambda/E_1E_2E_3)$ for all possible values of the failure rate given the total generic and specific information. According to Bayes' theorem, this can be characterized by the following

equation:

$$p(\lambda_j/E_1E_2E_3) = \frac{p(\lambda j/E_1E_2)p(E_3/\lambda jE_1E_2)}{\sum_j p(\lambda_j/E_1E_2)p(E_3/\lambda_jE_1E_2)} \qquad (0.14\text{-}3)$$

As IPPSS explains on page 0-92, the posterior probability distribution in the light of evidence $E_1E_2E_3$ is then seen to be a function of the prior probability for λ relative to the generic information E_1E_2 and the likelihood of the truth of λ_i on the plant specific data E_3 relative to generic information E_1E_2.

IPPSS supposes that in many cases the likelihood function can be specified with good precision and gives illustrations (such as equations (0.14-5) and (0.14-6)) of what such likelihood functions would be in typical cases. Although one might want to examine the assumptions about likelihoods more closely, that is not the concern of this discussion. I shall suppose that precise likelihoods are available.

A problem, however, still remains. To obtain a definite posterior distribution for failure rates, one needs to have a definite prior distribution relative to the generic information. If that information goes indeterminate--i.e., if there are many permissible prior distributions--then the posterior probability judgement will also be indeterminate.

It is, to be sure, the case that, if the plant specific data E_3 is sufficiently ample, the range of permissible probability distributions will be narrowed down considerably and, in the limit, will converge on a precise

posterior. At least this will be the case when the likelihood functions are of the sort considered in equations (0.14-5) and (0.14-6).

However, the rate of convergence will depend upon the extent of indeterminacy in the prior family of distributions for failure rates. The greater the indeterminacy, the slower the rate of convergence to precision in the posterior, and the greater the amount of data which will be required to "swamp" all permissible priors.

Furthermore, the extent to which precision is achieved will also depend upon the amount of plant specific data available in E_3, which may not be very much at all. I do not know myself the extent to which one can claim that the plant specific data yield likelihoods which swamp all permissible priors. I will want to return to some aspects of this issue shortly. Nonetheless, there may be many cases where plant specific data are sparse and where it is not at all plausible to suppose that likelihoods swamp all permissible priors. In that case, posterior credal probability judgements for the failure rate may be very indeterminate indeed.

Such indeterminacy ought to be propagated throughout the entire causal analysis in order to make a judgement of probability of serious accidents such as core meltdowns.

It is clear from the explicitly stated methodology of IPPSS that they did not do this. They explicitly state procedures for adopting priors for failure rates and these procedures invariably involve the adoption of a definite prior distribution.

Not only does this procedure violate what appears to have been the intent of the Lewis Committee report--but it entails the making of arbitrary decisions about probability.

Let it be noted here that I am not complaining that the probabilistic assessments are grounded on subjective judgements. I quite agree that such subjectivity will be unavoidable. The verdicts reached are going to be judgements of individuals and in that sense will be subjective. And, as the authors of IPPSS note, subjectivity is not the same as arbitrariness.

If one has information on the basis of which every reasonable person ought to make the same definite credal or subjective probability judgement, then one's endorsement of that probability judgement is not arbitrary.

However, if, on the basis of the available information, there is no good reason to pick one numerically definite probability distribution rather than another, the choice of any one such distribution rather than another is arbitrary. Rather than make such arbitrary judgements, one ought to regard as permissible all distributions which are not ruled out by the available evidence.

In this sense, the judgements of priors made by the authors of IPPSS are arbitrary, just like the judgements of priors made in the Reactor Safety Study. As a consequence, the IPPSS verdicts about posteriors cannot be trusted except, perhaps, in those cases where the plant

specific data are so abundant as to swamp the impact of an arbitrarily chosen prior.

As I have already indicated, I am not in a position to determine the extent to which the plant specific data for Indian Point are sufficient to swamp the impact of arbitrarily chosen priors. What is required is a recomputation of IPPSS data against a broad spectrum of priors utilizing Bayes' Theorem. The authors of IPPSS have not done this. Consequently, we do not know the extent to which this can be done and should proceed, pending such reassessment, as if none of the posteriors reported in IPPSS can be trusted.

SECTION FOUR

LOG NORMAL PRIORS

The standard procedure followed in the Reactor Safety Study (WASH 1400) in selecting prior distributions over failure rates was to identify an "assessed range" of failure rates for the component under scrutiny, grounded on the "generic information" alluded to in IPPSS, and then to assume that the subjective distribution of probabilities over the possible failure rates is a log normal distribution with 90% of the probability assigned to values in the assessed range and where 5% probability is assigned to each of the tails outside of the range in a symmetrical manner.

The Lewis Committee report complained about the arbitrariness of this procedure and quite rightly so. However, it did not think that the use of log normal priors of this sort would distort posterior porbability judgements by more than a factor of 2 or 3, so that posterior judgements based on such priors would not be distorted excessively.[6]

The Lewis Committee did not explain the basis for its passing this verdict of the relatively neglibible effect of using log normals. One can only surmise that they examined some of the generic data (which for the most part apparently are not sufficient to do reliable goodness of fit tests) and noticed what they thought were distributions

---

6 - NUREG/CR-0400, p. 9.

which better fitted the relatively sparse data. They then compared the impact of using such distributions as compared with the log normals to reach their conclusions.

If that is indeed the approach used by the Lewis Committee, it is surely open to question. If the data are sparse, there is not the slightest reason to suppose that the prior credal distribution should be represented by the curve best fitting the generic data. Even curves which deviate wildly from such best fitting curves might yield such data with significant probability--especially if the data are not excessive.

Thus, to assess the import of one arbitrarily chosen prior by considering another arbitrarily chosen prior is deceptive.

What one should do is to look at the prior distributions which give the worst possible fits of the generic data and which still could with significant probabilities yield that data and where the worst possible fits are in all directions. In that case, the deviations from the log normal in some directions might have been quite significant indeed.

I am not suggesting that the prior distributions giving worst possible fits of the generic data should be adopted as uniquely permissible prior credal distributions. My point is rather that one should begin with a set of distributions bounded by such worst possible fit distributions.

It is just as arbitrary to pick one particular distribution from the set as it is any other. The reasoning of the Lewis Committee report appears to make this mistake.

The import of this mistake is even more apparent in a discussion of the same issue in the report on the Zion Probabilistic Safety Study prepared by Sandia National Laboratories[7] and in the similar treatment by Sandia Labs of IPPSS[8].

The practice of IPPSS and the Zion Probabilistic Safety Study (ZPSS) differed from the Reactor Safety Study (WASH 1400) in that they fitted only 60% of the total probability symmetrically within the assessed range according to a log normal, rather than 90%, as in the Reactor Safety Study.

The Sandia reports on IPPSS and ZPSS ask the question as to the effect of these prior distributions on their estimates of failure rates.

Sandia proceeds as follows: They look at the plant specific data E_3 and compute maximum likelihood estimates for the failure rate and variance of the failure rate from this data and compare them with the posterior expected failure rate computed from the log normal prior and the variance of that distribution.

---

7 - Sandia National Laboratories, "Review and Evaluation of the Zion Probabilistic Safety Study", Letter Report, March 5, 1982.

8 - Sandia National Laboratories, "Review and Evaluation of the Indian Point Probabilistic Safety Study", NUREG/CR-2934, SAND 82-2929, December, 1982.

In effect, as Sandia points out in section 2.6 of both reports, their procedure was to use a flat or uniform prior over the assessed range (or, strictly speaking, the logarithmic transformation of that range). As a result, Sandia concluded that for the most part, the choice of a log normal prior would not have a marked effect.

But once more, we have a comparison of one arbitrarily chosen prior with another, and such comparisons are of doubtful value.

Thus, the claim made both in the Lewis Committee report and the Sandia reports on IPPSS and ZPSS that the use of log normal priors has negligible effect on the outcome should not be trusted.

SECTION FIVE

ASSESSED INTERVALS AND WORST PERMISSIBLE PRIORS

As already noted, IPPSS follows ZPSS in broadening the assessed ranges of failure rates as compared with those used in the Reactor Safety Study (WASH 1400).

They appeal to experimental data suggesting that even experts often tend to adopt subjective probability distributions that are "too tight". They appeal, in particular, to a study by Slovic, Fischhoff, and Lichtenstein which claims that when individuals give interval estimates with 98% probability, somewhere between 20 and 50 percent of true values fall outside the limits of these estimates. Since much of the so called generic information relative to which prior distributions for failure rates are based seems to derive from the subjective testimony of experts, it seems quite sensible to follow the practice of IPPSS at least to this extent--to wit, in broadening the interval within which 90% of the probability is to be assigned.

But there are some serious difficulties involved in determining how this should be done.

When IPPSS takes generic information from the Reactor Safety Study concerning the failure rate of something like a pump, it takes the assessed range given by the Reactor Safety Study, assumes that with 60% probability the failure rate is in that interval, with 20% probability it is beyond the lower end of the interval, and with 20% it is beyond the upper end.

Then IPPSS assumes, as did the Reactor Safety Study, that failure rates are distributed log normally compatible with the restrictions thus made. This enables them to ascertain new 5% and 95% confidence limits.

There are two components to this line of reasoning: (i) it is assumed that there is as much probability of a failure rate below the lower limit of the assessed range as there is above the upper limit; and (ii) a log normal distribution is used to assess confidence intervals.

If we suppose that prior credal probability is indeterminate, the use of a log normal is arbitrary and, by the same token, it is far from clear why the symmetry should be assumed in extending upper and lower bounds of assessed ranges.

Granted that all of this is so, how should one proceed in order to make sensible probability judgements?

To begin with, let us bracket the problem of determining 95% confidence intervals. Let us suppose that we can claim that a certain interval is with probability 1 the interval within which the failure rate is to be found. My own judgement is that it is possible to specify in many cases, a large finite interval for which such confidence can be expressed. But very little in the conclusions to be drawn depends on it.

Even given such an interval, the problem remains of determining a prior credal distribution for failure rates in that interval.

If we adopt a prior credal state which is maximally indeterminate (so that it is the "convex hull" of all distributions which assign probability 1 to some failure

rate in the range and probability 0 to the rest), no amount of plant specific experience will enable likelihood functions to suppress the differences between the competing priors in the credal state so that with enough experience, large indeterminacy can be reduced. It would be as unreasonable to prevent large enough experience, even in principle, from leading us to relatively determinate judgements as it would to arbitrarily pick a prior single distribution out of a hat as IPPSS, ZPSS, and the Reactor Safety Study do. Both extremes should be avoided.

The idea is to adopt a prior credal state consisting of a broad band of distributions and, hence, considerable indeterminacy in prior probability judgement while still allowing for the possibility that, if data obtained from plant specific experience becomes ample, posterior distributions will be derivable that are confined to a very narrow band. How broad the band of prior distributions should be depends on how stringent our demands are on the amount of information that will be required to reach fairly sharp estimates of failure rate probabilities on the posterior data.

There are many different ways one might proceed in attempting to implement this approach. The arbitrariness involved here, however, can be rendered innocuous by making prior probability judgement sufficiently indeterminate, in which case there will not be much difference between the recommendations of the different approaches.

For purposes of illustration, let me suggest the following approach.

Let us fix, to begin with, on a suitable sharp distribution to serve as a benchmark. A sensible distribution for the failure rates would be a conjugate prior distribution for the parameter of the Poisson distribution --which would be a gamma distribution. In particular, however, one might begin with the misleadingly called "ignorance" prior which has the density $1/\lambda$ normalized to the assessed interval for λ or is the uniform or "flat" distribution for $\log\lambda$ over the corresponding interval for $\log\lambda$.

The proposal is to fix a band around such an ignorance prior such that all permissible prior distributions fall in that band. Then stable estimation theorems guarantee that with sufficient experience, posterior distributions will fall in a narrow band focusing on the maximum likelihood estimator. However, the rate at which posterior data will swamp all priors in the prior band will be far less rapid than it would be if we started with the ignorance prior as our sole distribution. In this way, we guarantee that maximum likelihood estimates are not taken too seriously until we have very ample data--which is surely an appropriate attitude to take given the importance of the issues which are to be settled.[9]

If our benchmark distribution is uniform over some

9 - for further discussion see, The Enterprise of Knowledge, chapter 13.

finite interval, we can view it as uniform over the unit interval from 0 to 1 and, when desired, alter the scale by whatever factor is appropriate.

We can then envisage the band as characterized by two lines: one above and parallel to the uniform density of 1 and one below and parallel to that density. Let the upper line be α (alpha) times higher than the benchmark (α is greater than 1) and the lower line be β (beta) times the benchmark (β being positive and less than 1).

To simplify, I shall suppose that these two numbers are so chosen that it is possible to assign an event of Lebesgue measure x for some positive number x less than 1 a probability of .5 according to a density function which assigns density equal to β to points in x and density equal to α to points in the complement of Lebesgue measure 1-x. This should be so no matter what events having these Lebesgue measures are selected. This procedure guarantees that $\beta x = \alpha(1-x) = .5$. Given the initially specified value of x, we can then determine the values of α and β.

Alternatively and, perhaps, more perspicaciously, we may fix in advance a factor k (greater than or equal to 1) such that $\alpha = k\beta$. Assuming that the bounds are so constrained that 50% of the probability can be assigned to events of positive Lebesgue measure x, all of whose points receive density β, and the remainder of the points have

density α, we can then determine the values of these two parameters and the value of x. x is equal to $k/(k + 1)$ so that $\beta = .5(k + 1)/k$ and $\alpha = .5(k + 1)$.

When $k = 1$, of course, the upper and lower lines collapse into the uniform density and we have the classic case of an ignorance prior over the unit interval. As k increases, the band within which permissible priors are permitted to lie becomes wider. As k goes to infinity, we may think of a case where 50% of the probability concentrated in a set of Lebesgue measure 0 (a finite or countably infinite set of points, perhaps) and 50% in the complement.

Thus, values of k index the extent to which we intend to be cautious in our initial assessments of prior probability in the sense that we will require a certain amount of data before we will concede that likelihoods swamp our prior ignorance.

Assuming that we have fixed on a value of k, we can then determine a <u>worst permissible prior</u>. In our case, it would consist of a distribution assigning 50% of the probability with density β to values of failure rates at the low end of the spectrum and density α to the rest.

Focusing on a worst permissible prior (wpp) does not mean that we adopt that prior as our prior belief. It is not any more representative of prior belief than our best possible prior. But we should ascertain whether

the posteriors calculated with worst permissible priors are swamped or not in such a way as to eventually lead to a calculation of expected utility favorable to the operation of the nuclear power plant. If it does, a strong case has been made for this conclusion. Otherwise expected utility has failed to render a verdict and minimax becomes operative and recommends forbidding such operation.

I have spoken, for the sake of simplicity as if the derivation of the worst permissible prior was based on the unit interval as the assessed interval. Of course, in most cases to be considered here where assessed intervals are for failure rates, neither those assessed intervals or the corresponding intervals for logarithms of failure rates are unit intervals. So the "median value--which is x for the unit interval--must be multiplied by the appropriate scale factor.

In table 1 below, I adopt as an assessed range used in the Reactor Safety Study (RSS) for "failure of air operated valves to remain open", as reported in IPPSS on page 0-95. Assuming that with probability 1 all failure rates fall in the assessed range from 2.8×10^{-8} to 2.8×10^{-7}, I take the logarithms of these end points to obtain a range for values for the logarithm of the failure rate. Relative to that interval, I use a uniform distribution for the benchmark distribution and obtain a worst possible probability distribution which I then transform

back into a distribution over the failure rates. I report the median and mean values obtained in this way by RSS using the log normal distribution and using my approach for the uniform prior (over logarithms of failure rates)--i.e., when k = 1 and also when k = 3, k = 19, and k = 99.

Table 1

	Median	Mean
RSS	8.9×10^{-8}	1.1×10^{-7}
k = 1	8.9×10^{-8}	1.1×10^{-7}
k = 3	1.6×10^{-7}	1.4×10^{-7}
k = 19	2.5×10^{-7}	1.8×10^{-7}
k = 99	2.7×10^{-7}	1.9×10^{-7}

Table 2 gives results for the same problem where the assessed range is extended as in IPPSS.

Table 2

	Median	Mean
IPPSS	8.9×10^{-8}	2.3×10^{-7}
k = 1	8.9×10^{-8}	2.3×10^{-7}
k = 3	2.8×10^{-7}	3.2×10^{-7}
k = 19	6.7×10^{-7}	4.1×10^{-7}
k = 99	8.0×10^{-7}	4.9×10^{-7}

The purpose of supplying these tables is to illustrate the difference between the approach of IPPSS to constructing priors based on generic experience and an approach which takes seriously the complaint that we

lack a basis for making precise estimates of probability and should make assessments in terms of ranges of probability distributions.

It should be apparent, however, that if we use worst permissible distributions as priors to calculate posteriors, sample data from plant experience will not begin to give definitive judgements until the data become very ample. If the data are scarce, it is to be expected that we shall not be able to calculate definite probability distributions for such events as core melts or containment failures, but will have to rest content with interval probability assessments.

The critical issue is whether it will be possible to obtain interval probability assessments sufficiently definite so that the upper probabilities calculated from the worst permissible priors are still below the required levels.

Clearly that will depend upon how cautious we are in adopting indices of caution k. My own view is that given the gravity of the issues involved, we should set k rather high. The main criterion, however, ought to be that when there is controversy concerning the degree of caution to be exercised, we should choose high values of k rather than low; for only then do we proceed in a manner which can reach conclusions acceptable to all parties to the dispute.

The main point to emphasize, in any case, is that the authors of IPPSS, like the authors of ZPSS and the Reactor Safety Study (WASH 1400), report sharp distribu-

tions grounded on the arbitrary selection of log normal priors. Even if there were no fault to be found with the causal analysis of the various systems studied and the way failure probabilities were propagated, issues about which I do not feel qualified to comment, the analysis offered by IPPSS remains faulty.

Furthermore, attempts to make comparisons of the probabilities of major accidents at Indian Point Units 2 and 3 with such probabilities for other plants will have to contend with the deficiencies of probabilistic safety studies for other plants of the same kind--such as the ZPSS.

Consequently, pending a complete reassessment of the data on Indian Point Units 2 and 3 (as well as for other plants) in a manner which repects the importance of reporting the indeterminacies in probability judgement, one should not rely on the analysis contained in IPPSS.

CONCLUSION

The argument of this discussion rests on two main theses:

(1) One should rely on calculations of expected costs and benefits in making decisions only when the probabilities available are sufficiently definite to render an unequivocal verdict. When such calculations cannot be used to decide because of indeterminacy in the assessments of probability (and, hence, of expected costs and benefits or in the distribution of risks), appeal should be made to minimax or worst possible case analysis.

(2) The calculations of probabilities contained in IPPSS gain precision by the arbitrary selection of prior probabilities for the failure rates of components and this spurious precision is propagated up through the entire causal analysis.

The conclusion which may be reached on the basis of these condierations is that we have thus far no rational grounds for supposing that the probabilities of serious accidents at Indian Point Units 2 and 3 are sufficiently low to justify operating these plants. Minimax considerations argue in favor of refusing to permit these plants to continue to operate unless and until adequate reassessment of Indian Point data results in a contrary conclusion.

June 1982

Isaac Levi

Born: New York, New York. June 30, 1930.

Married. Two Children.

College and University Education:

B.A.	N.Y.U.	1951	Philosophy; Minors: Greek and Mathematics
M.A.	Columbia	1953	Philosophy
Ph.D.	Columbia	1957	Philosophy

Dissertation Supervisor: Ernest Nagel

Honors, Awards, Grants:

1951	Phi Beta Kappa
1951	Pi Mu Epsilon
1955-56	W.T. Bush Fellow, Columbia University
1962-63	N.S.F. Grant
1965-66	N.S.F. Grant
1966-67	Guggenheim Fellow and Fulbright Research Scholar in the United Kingdom
1968-69	N.S.F. Grant
1971-72	N.S.F. Grant
1973	N.S.F. Grant Visiting Scholar, Corpus Christi College, Cambridge U.

Employment:

1954-56 Rutgers U. (part time)
1956-57 The City College (lecturer)
1956-57 Columbia University (lecturer)
1957-58 Western Reserve U. (Instructor)
1959-62 Western Reserve U. (Assistant Professor)
1961 (Summer) Columbia U. (Visiting Associate Professor)
1962-64 The City College (Assistant Professor)
1964-67 Western Reserve U. (Associate Professor)
1967- Case Western Reserve U. (Full Professor)
1968-70 Case Western Reserve U. (Chairman, Department of Philosophy)
1970- Columbia University (Full Professor)
1973-76 Chairman, Department of Philosophy, Columbia University

Professional Organizations:

A.A.U.P., A.P.A., Assoc. of the Philosophy of Science,
Board of Governors of the Philosophy of Science Association, 1975-76.

Scholarly Activities:

In addition to publications appended to this vita, I have read papers at several professional meetings including the A.P.A. Eastern and Western Divisions, The Br. Association for the Philosophy of Science, and special conferences held at Salzburg, Philadelphia, and London, Ontario. I have participated in several other conferences and read invited papers at many Universities including Oxford, Cambridge, London School of Economics, The University of Helsinki, Columbia, Chicago, Rochester, Michigan, Harvard, Boston University, Virginia, and others.

Summary of activities from 1980.

Isaac Levi

1980 Visiting fellow at Darwin College, Cambridge.

Offered a course of lectures at Cambridge.

Gave two talks at Oxford, one at LSE, one at U. of Bradford and one at U. of Warwick.

Gave a series of lectures at the University of Helsinki and received a medal from the University

1981 Read papers at Rutgers U., Lehigh U. and at the Seminar on Bayesian Inference at Carnegie-Mellon and N.J. Philosophy Assoc. Also at conference at Georgia Tech. on Information Theory.

1982 Read papers at the Universities of Chicago and Colorado and Washington University. Will read paper as part of symposium at Philosophy of Science Association meetings in the fall.

1983 I am an invited speaker on the Foundations of Probability and Induction at the 1983 Congress of the Division of Logic, Methodology and Philosophy of Science, International Union of History and Philosophy of Science. This congress will meet in Salzburg.

I have also been informed that my name is "on the list of scientists and scholars eligible for fellowships at the Center for Advanced Study in the Behavioral Sciences". I have also been told that my "name has been placed on our tentative roster for" 1983-1984. But otherwise I have heard nothing. But there is a remote possibility that I shall request a leave to visit the Center at Stanford for that period.

BIBLIOGRAPHY OF ISAAC LEVI

1957

Doctoral Dissertation, *The Epistemology of Moritz Schlick* under the supervision of Ernest Nagel, Columbia University (microfilm).

1958

Review of Roy Harrod, *Foundations of Inductive Logic*, *Journal of Philosophy* 55, 209–212.

1959

(a) 'Putnam's Three Truth Values', *Philosophical Studies* 10, 65–69.
(b) Translation of R. Carnap, 'The Old and the New Logic', in *Logical Positivism* (ed. by A. J. Ayer) (Free Press).

1960

(a) 'Must the Scientist Make Value Judgments?', *Journal of Philosophy* 57, 345–57.
I suggest that although scientists may make value judgements the values may be characteristic of the scientific enterprise and, therefore, different from moral, economic, political, prudential or other types of values.
(b) Translation of A. Meinong, 'The Theory of Objects', in R. M. Chisholm (ed.), *Realism and the Background of Phenomenology* (Free Press) (with D. B. Terrell and R. M. Chisholm).

1961

(a) 'Decision Theory and Confirmation', *Journal of Philosophy* 58, 614–625.

Bogdan, R. J. (ed.), 'Henry E. Kyburg, Jr. & Isaac Levi', 309–318.

Copyright © 1981 by D. Reidel Publishing Company.

(b) Review of *Minnesota Studies in Philosophy of Science, Vol. II*, *Journal of Philosophy* **58**, 241–248.
(c) Review of Danto and Morgenbesser (eds.), *Philosophy of Science* and E. H. Madden (ed.), *The Structure of Scientific Thought*, *Journal of Philosophy* **58**, 387–390.
(d) Review of A. Rapaport, *Fights, Games and Debates*, *Harvard Educational Review* **31**, 477–479.

1962

'On the Seriousness of Mistakes', *Philosophy of Science* **29**, 47–65.
This was my first effort to construct decision theoretic models of scientific inference utilizing cognitive or epistemic utilities along the lines suggested by 1960a. Both this paper and 1960a were written in ignorance of Hempel's writings on epistemic utility. The vagaries of publication resulted in a discussion of Hempel's work (1961a) appearing in print before this paper. In point of fact, 1961a was written after 1963a in which I explored the possibilities of construing Popper's corroboration measures as measures of expected epistemic utility and proposed using measures of relevance as measures of expected epistemic utility.

1963

(a) 'Corroboration and Rules of Acceptance', *British Journal for the Philosophy of Science* **13**, 307–313.
(b) Review of H. Leblanc, *Statistical and Inductive Probabilities*, *Journal of Philosophy* **59**, 21–25.
(c) Contribution to *Harper's Encyclopedia of Science*.

1964

(a) 'Belief and Action', *The Monist* **48**, 306–316.
(b) 'Belief and Disposition', *American Philosophical Quarterly* **1**, 221–232 (with Sidney Morgenbesser).
This paper, written in collaboration with Sidney Morgenbesser, first proposes the view of disposition predicates which I subsequently elaborated upon in 1967a and 1977b and which was adapted to furnish an account of statistical probability or chance of the sort proposed in 1967a, 1973b, 1977a, 1977b and 1980a.
(c) 'Utility and Acceptance of Hypotheses', *Voice of America Forum Lectures, Philosophy of Science Series*, No. 2.

1965

(a) 'Deductive Cogency in Inductive Inference', *Journal of Philosophy* **62**, 68–77.
(b) 'Hacking Salmon on Induction', *Journal of Philosophy* **63**, 481–487.

1966

(a) 'On Potential Surprise', *Ratio* **8**, 107–129.
My first essay on Shackle's theory of potential surprise. Further discussion is found in 1967a, 1972a, 1979e and 1980c.
(b) 'Recent Work in Probability and Induction' (reviews of books by I. J. Good, I. Hacking, R. C. Jeffrey and H. Törnebohm), *Synthese* **16**, 234–244.

1967

(a) *Gambling with Truth* (A. Knopf, New York) (reissued in paperback without revision in 1973 by MIT Press).
On pp. 240–241 of this book, I wrote, "Individuals and institutions strive to attain many objectives. At times, these ends conflict; at other times, they complement one another. Philosophers legitimately ask questions about the relative importance of different ends, including the cognitive objectives of scientific inquiry. But disparagement of cognitive ends (even when there are grounds for it) ought not to disguise itself by reducing these ends to practical ones. Truth, information, explanation, simplicity are desiderata that are different from wealth, love, security, health, peace, etc. They ought to be recognized as such. Such recognition is enhanced by showing how the ends of inquiry control the legitimacy of inferences."
In support of this view, I proposed an account of cognitive decision making designed to accommodate a limited range of problems. This account was the culmination of the exploration of accounts of epistemic utility I had begun in 1960a.
(b) 'Probability Kinematics', *British Journal for the Philosophy of Science* **18**, 197–209.
(c) 'Information and Inference', *Synthese* **17**, 369–91.
This paper was written while I was reading proof on 1967a. It responded to proposals of Hintikka and Pietarinen. The most important technical development in it, however, is the modification of the account of epistemic utility I proposed in 1967a. In particular, I no longer required every element of an ultimate partition to be as informative as every other and, given this modification of the model, was prepared to entertain extending the scope of the applicability of the model to the question of reaching conclusions concerning theoretical hypotheses. 1967c ought to be read along with 1967a.

1968

(a) Review of J. Hintikka and P. Suppes (eds.), *Aspects of Inductive Logic*, *British Journal for the Philosophy of Science* **19**.
(b) Review of W. Salmon, *The Foundations of Scientific Inference*, *British Journal for the Philosophy of Science* **19**, 259–61.

1969

(a) 'Confirmation, Linguistic Invariance and Conceptual Innovation', *Synthese* **20**, 48–55.

(b) 'If Jones Only Knew More', *British Journal for the Philosophy of Science* **20**, 153–159.
(c) 'Induction and the Aims of Inquiry', *Philosophy, Science and Method, Essays in Honor of Ernest Nagel*, ed. by S. Morgenbesser, P. Suppes, and M. White (St. Martin's Press), pp. 92–111.
(d) Review of *The Problem of Inductive Logic*, ed. by I. Lakatos, *Synthese* **20**, 143–148.
(e) 'Are Statistical Hypotheses Covering Laws?', *Synthese* **20**, 297–307.
This paper questions the widely held view that statements of chance or statistical probability can serve as lawlike generalizations in explanation or that so called 'inductive statistical explanation' can be cogently regarded to be a species of covering law explanation.

1970

'Probability and Evidence', *Induction, Acceptance and Rational Belief*, ed. by M. Swain (Reidel, Dordrecht), pp. 134–156.

1971

(a) 'Certainty, Probability and Correction of Evidence', *Nous* **5**, 299–312.
(b) 'Truth, Content and Ties', *Journal of Philosophy* **68**, 865–876.

1972

(a) 'Potential Surprise in the Context of Inquiry', in *Uncertainty and Expectations in Economics: Essays in Honor of G. L. S. Shackle*, ed. by C. F. Carter and J. L. Ford (Blackwell, Oxford), pp. 213–236.
(b) Invited Comments on Churchman (pp. 87–94) and on Braithwaite (pp. 56–61). *Science, Decision and Value*, ed. by J. Leach, R. Butts and G. Peirce (Reidel, Dordrecht).

1973

(a) 'But Fair to Chance', *Journal of Philosophy* **70**, 52–55.
(b) Review of D. H. Mellor, *The Matter of Chance*, *Philosophical Review* **82**, 524–530.

1974

'On Indeterminate Probabilities', *Journal of Philosophy* **71**, 391–418.
This is my first statement of an approach to probability judgment and decision theory involving a substantial departure from the Bayesian decision theory on which I had reluctantly relied in 1967a.

1975

'Newcomb's Many Problems', *Theory and Decision* **6**, 161–175.

To my knowledge, this paper contains the first discussion of the idea that the conflicting approaches to the Newcomb problem be viewed as invoking different principles of expected utility maximization. The idea is usually attributed to Gibbard and Harper although they themselves acknowledge that I suggested, as an alternative to the Bayesian approach, calculating expected utilities using the unconditional distribution over states when states are causally independent of options. Gibbard and Harper appear to think their own proposal to be different from this. However, it is demonstrably equivalent given the assumptions about probabilities of conditionals they adopt. To be sure, there is an important difference. I showed that in the case where the demon is perfectly infallible the non Bayesian prescription recommends the two box solution – which, in my view, is clearly absurd. Gibbard and Harper accept the absurd implication with equanimity. The paper also points out that if one calculates probabilities using conditional probabilities of states on acts, there is no definite solution to the Newcomb problem as stated by Nozick and others. Nozick thought otherwise because he confused the probabilities of options given states with the probabilities of states given options.

1976

(a) 'Acceptance Revisited', in *Local Induction*, ed. by R. Bogdan (Reidel, Dordrecht), pp. 1–71.
(b) 'A Paradox for the Birds', in *Essays in Memory of Imre Lakatos*, ed. by R. S. Cohen *et al.* (Reidel, Dordrecht), pp. 371–378.

1977

(a) 'Direct Inference', *The Journal of Philosophy* **74**, 5–29.
This paper discusses the important issue of direct inference from knowledge of chances or objective, statistical probabilities to judgements of credal probability. Aside from the pioneering discussions of Reichenbach and Fisher, the most thorough study of direct inference had been Kyburg's. I advocated an approach alternative to Kyburg's and explored some of the issues involved in the dispute. The topic is further discussed in 1978c and in considerable detail in 1980a.
(b) 'Subjunctives, Dispositions and Chances', *Synthese* **34**, 423–455.
(c) 'Four Types of Ignorance', *Social Research* **44**, 745–756.
(d) 'Epistemic Utility and the Evaluation of Experiments', *Philosophy of Science* **44**, 368–386.

1978

(a) 'Irrelevance', *Foundations and Applications of Decision Theory*, ed. by Hooker, Leach, and McLennen, Vol. I (Reidel, Dordrecht), pp. 263–275.
(b) 'Coherence, Regularity and Conditional Probability', *Theory and Decision* **9**, 1–15.

(c) 'Confirmational Conditionalization', *Journal of Philosophy* **75**, 730–737.
(d) Reprint of (1964b) and of (1977b) in *Dispositions*, ed. by R. Tuomela (Reidel, Dordrecht).

1979

(a) Translation of (1967a) into Japanese (Kinokuniya Book Store, Tokyo).
(b) 'Inductive Appraisal', *Current Research in Philosophy of Science*, ed. by P. D. Asquith and H. E. Kyburg (PSA, East Lansing, Mich.), pp. 339–351.
(c) 'Serious Possibility', *Essays in Honour of Jaakko Hintikka* (Reidel, Dordrecht), pp. 219–236.
(d) 'Abduction and Demands for Information', *The Logic and Epistemology of Scientific Change* ed. by I. Niiniluoto and R. Tuomela (North Holland for Societas Philosophica Fennica, Amsterdam), pp. 405–429.

Many authors have complained about the relativity of appraisals of inductive inferences according to my theory to the choice of an ultimate partition of maximally consistent potential answers and to other contextual factors like degree of caution or boldness and appraisals of potential answers with respect to informational value. I have repeatedly insisted that this complaint points to a virtue of my approach rather than to a defect. However, until I had developed the apparatus presented in (1980a), I was not able to give a systematic account of how conflicts in the choice of ultimate partition or in the values of other contextual parameters are to be handled. This paper offers such an account.

(e) 'Support and Surprise: L. J. Cohen's View of Inductive Probability', *British Journal for the Philosophy of Science* **30**, 279–292.

1980

(a) *The Enterprise of Knowledge: An Essay on Knowledge, Credal Probability and Chance* (MIT, Cambridge, Mass.)

This book has three objectives: to outline a view of the structure and aims of inquiry without seeking foundations for knowledge or the naturalization of epistemology, to elaborate a novel account of probability judgement, utility judgement and rational choice as part of this approach to an account of inquiry, and finally to exploit the proposals made in a critical review of diverse responses to some central problems of statistical theory.

Chapters 1–3 explain in some detail the epistemological outlook which motivated (1967a) as modified by the explicit endorsement of epistemological infallibilism first advanced in (198?). These chapters contain comments on the views of Quine, Kuhn, Feyerabend, Peirce and Popper which are further developed in Chapter 18 and in the appendix. Chapter 4 contains an overview of my approach to probabilistic reasoning. This is elaborated upon in considerable detail in Chapters 5–10 and Chapter 13. The central feature is the relaxation of the strict Bayesian requirement that states of probability judgement (credal states) be representable by unique probability measures and states of valuation be representable by utility measures unique up to positive linear transformations. Convex sets of probability measures and of utility measures are employed instead. This idea, together with an account of rational choice grounded on a lexicographically ordered series of criteria for the admissibility of feasible options is elaborated upon in Chapters 6–9. It responds to the boast that Bayesian theory offers the most general approach available to decision making by offering a theory of far greater scope allowing strict Bayesian theory to be a special limiting case. Chapter 5 discusses some narrowly technical questions concerning 0 probability and countable additivity. Chapter 10 discusses conditionalization and irrelevance from the point of view explained in (1978a). Chapters 11 and 12 contain the fullest account I have offered to date of my views on dispositionality, ability and statistical probability or chance and of the relations between these 'objective' modalities and corresponding 'subjective' modalities. Chapter 12 discusses, in particular, the important topic of direct inference. Chapter 13 provides an account of how to revise or to select prior states of credal probability judgement. Chapter 14 discusses ways to rationalize fiducial inference from a Bayesian point of view. Chapter 15 explores Hacking's specific version of a Bayesian rationalization exploiting the notion of likelihood. Seidenfeld's proof that Hacking's effort and all other efforts to convert fiducial inference to Bayesian inference are inconsistent is exploited to show that Hacking's law of likelihood cannot be used as a principle of inductive logic as Hacking thought. Chapter 16 discusses the ideas of Fisher on the connection between fiducial inference and direct inference and Kyburg's original and ingenious efforts to elaborate an improved Fisherian theory. An alternative developed by A. P. Dempster is also considered. Chapter 17 discusses the approach to statistical theory developed by Neyman, Pearson and Wald. The deficiencies claimed to be found in these theories are invoked in support of the approach to probabilistic reasoning advanced in Chapter 13. Chapter 18 returns to some general themes about objectivity and context dependence applicable to the views about the revision of knowledge developed in Chapters 1–3 and the views about revision of probability judgement developed in the light of the discussion of Chapters 4–17. The appendix on the Rasmussen Report illustrates some of the central themes of the book by reference to the methodological approach used in the assessments of the reliability of backup systems in nuclear plants according to the Rasmussen Report.

(b) 'Induction as Self Correcting According to Peirce', *Science, Belief and Behaviour, Essays in honor of R. B. Braithwaite*, ed. by D. H. Mellor (Cambridge), pp. 127–140.
(c) 'Potential Surprise: Its Role in Inference and Decision-Making', *Applications of Inductive Logic*, ed. by L. J. Cohen and M. Hesse (Clarendon, Oxford), pp. 1–27. Also replies to comments by P. Teller, H. E. Kyburg, R. G. Swinburne and L. J. Cohen and comments on papers by R. Giere, J. Dorling, J. E. Adler and R. Bogdan.
(d) 'Incognisables', *Synthese* **45**, 413–427.

This paper together with (1980b) presents a view of Peirce based on a more accurate reading of Peirce's ideas on probability and statistical inference than is usually found in the philosophical literature. In particular, (1980b) shows that Peirce had, as early as 1878, elaborated the Neyman–Pearson technique of confidence interval estimation for the binomial case and had an accurate appreciation of the conditions of its applicability. Moreover, textual evidence is adduced in support of the view (which Reichenbach himself shared) that Peirce's view of induction as self correcting was not an anticipation of Reichenbach's own view and does not rely on any dubious vindication arguments.

1981

(a) 'On Assessing Accident Risks in U.S. Commercial Nuclear Power Plants', *Social Research* **48**.
(b) 'Should Bayesians Sometimes Neglect Base Rates?', *The Behavioral and Brain Sciences*.

1982

(a) 'Dissonance and Consistency According to Shackle and Shafer', PSA (1978), II.
(b) *Review of Theory and Evidence* by C. Glymour, *Philosophical Review*.
(c) Review of Vol. 2 of *Studies in Inductive Logic and Probability*, ed. by R. C. Jeffrey, *Philosophical Review*.
(d) 'Liberty and Welfare', to appear in *Beyond Utilitarianism*, ed. by A. K. Sen and B. Williams.
(e) 'Escape from Boredom: Edification According to Rorty', *Canadian Journal of Philosophy*.
(f) Review of *Chance, Reason and Cause* by A. Burks.
(g) 'Direct Inference and Confirmational Conditionalization', *Philosophy of Science*.

198?

'Truth, Fallibility and the Growth of Knowledge', which is to appear in the Boston Colloquium volume 'Language, Logic and Method' with a discussion by I. Scheffler and A. Margalit. The editor has promised that it would appear within a year since 1975 when it was initially submitted to him.
The paper itself was first written and delivered orally in 1971.

This paper is the first expression of an important modification of the epistemological outlook I had advanced in (1967a) and had elaborated upon in (1970) and (1971a). I draw a distinction between the corrigibility of knowledge (which I accept) and its fallibility (which I deny). I explain why Peirce and Popper must reject the distinction as I understand it because of their shared outlook concerning the ultimate aims of scientific inquiries. Elements of the point of view of this paper are presented in (1976a) and (1977b). (1980a) elaborates upon it in detail.

In process: Papers on decision making under unresolved conflict and collective decision making.

Critical Literature

Reviews of *Gambling with Truth* (1967): I. Hacking, *Synthese* **17** (1967), 444–448; R. C. Jeffrey, *J. Phil.* **65** (1968), 313–322; K. Lehrer, *Noûs* **3** (1969), 255–297; J. Mackie, *BJPS* **19** (1968), 261; D. Miller, *JSL* **36** (1971), 318–320.

Reviews of *The Enterprise of Knowledge* (1980): D. V. Lindley, *Nature* (Nov. 6, 1980); R. Swinburne, *Times Higher Education Supplement* (Oct. 3, 1980); A. Margalit, *Times Literary Supplement* (Feb. 27, 1981), 237.

Discussions

Adler, J. E., 'The Evaluation of Rival Inductive Logics', in L. J. Cohen and Hesse (eds.), *Applications of Inductive Logic* (Clarendon Press, Oxford, 1980), pp. 383–384, 386n, 387–390, 393n, 403n; and 'Comments', *op. cit.*, pp. 419–420.

Cardwell, C., 'Gambling for Content', *J. Phil.* **68** (1971), 860–864.

Cohen, L. J., 'How Empirical is Contemporary Logical Empiricism', *Philosophia* **5** (1975), p. 359ff; *The Probable and the Provable* (Oxford, 1977), pp. 66, 124, 316ff; Comments on Levi's 'Potential Surprise: Its Role in Inference and Decision-Making', in Cohen and Hesse (eds.), *op. cit.*, pp. 64–66; 'What Has Inductive Logic to Do With Causality', Cohen and Hesse (eds.), *op. cit.*, p. 151, 173.

Gaa, J. C., 'Moral Autonomy and the Rationality of Science', *Phil. Sci.* **44** (1977), 513–541.

Gibbard, A. and Harper, W. L., 'Counterfactuals and Two Kinds of Expected Utility', in *Foundations and Applications of Decision Theory*, ed. by C. A. Hooker, J. J. Leach and E. F. McLennen (Reidel, Dordrecht; 1978), Notes 11 and 12, p. 161.

Giere, R., 'Foundations of Probability and Statistical Inference', *Current Research in Philosophy of Science*, ed. by P. D. Asquith and H. E. Kyburg, *PSA* (1979), pp. 518–519.

Goosens, K., 'A Critique of Epistemic Utilities', *Local Induction*, ed. by R. J. Bogdan (Reidel, Dordrecht, 1976), pp. 93–114.

Hacking, I., 'The Theory of Probable Inference: Neyman, Peirce and Braithwaite', *Science, Belief and Behaviour: Essays in honour of R. B. Braithwaite*, ed. by D. H. Mellor (Cambridge U. Press, Cambridge, 1980), p. 153.

Hesse, M., *The Structure of Scientific Inference* (U. of Cal. Press, Berkeley, 1974), p. 112 and 122n.

Hilpinen, R., *Rules of Acceptance and Inductive Logic* (North-Holland, Amsterdam, 1968), pp. 85, 93, 94–104.

Hintikka, J. and Pietarinen, J., 'Semantic Information and Inductive Logic', in *Aspects of Inductive Logic*, ed. by J. Hintikka and P. Suppes (North-Holland, Amsterdam, 1966), pp. 96, 107–108.

Hutchison, T. W., *'Positive' Economics and Policy Objectives* (Allan and Unwin, London, 1964), p. 103n.

Jeffrey, R. C., 'Dracula Meets Wolfman: Acceptance vs Partial Belief', in Induction, *Acceptance and Rational Belief*, ed. by M. Swain, (Reidel, Dordrecht, 1970), pp. 157–185.

Kyburg, H. E., 'Recent Work in Inductive Logic', *APQ* **1** (1964), p. 9; 'Conjunctivitis', in M. Swain, *op. cit.*, pp. 191–215; 'Local and Global Induction', in *Local Induction*, ed. by Bogdan, pp. 191–215; 'Chance', *J. of Philosophical Logic* **5** (1976), pp. 363, 365, 366, 371–376, 389–392; 'Propensities and Probabilities', in *Dispositions*, ed. by R. Tuomela (Reidel, Dordrecht, 1978), pp. 277, 284–285, 289, 295–299; Randomness and the Right Reference Class', *J. Phil.* **74** (1977), 501–521; 'Conditionalization', *J. Phil.* **77** (1980), 98–104, 113–114.

Leach, James, 'Explanation and Value Neutrality', *BJPS* **19** (1968), 93–108.

Lehrer, K., 'Induction, Consensus, Catastrophe', in R. Bogdan *op. cit.*, pp. 131–143; 'Truth, Evidence and Error: Comments on Miller', in Cohen and Hesse (eds.), *op. cit.*, pp. 133, 140; 'Justification, Explanation and Induction', in M. Swain, *op. cit.*, pp. 109, 120.

Mellor, D. H., *The Matter of Chance* (Cambridge U. Press, Cambridge, 1971), pp. 1, 79; 'In Defense of Dispositions', *Phil. Review* **83** (1974), 172, 174, 181.

Continuation of bibliography for Isaac Levi

1981

(a) 'On Assessing Accident Risks in U.S. Commercial Nuclear Power Plants', Social Research 48, pp.395-408.

(b) 'Should Bayesians sometimes neglect base rates?' The Behavioral and Brain Sciences 4, 342-343.

(c) Comment on 'On some statistical paradoxes and nonconglomerability' by Bruce Hill, Trabajos de Investigacion Operativa y Estadistica 32, 135-141.

(d) 'Direct Inference and Confirmational Conditionalization,' Philosophy of Science 48, 532-552.

(e) 'Escape from Boredom: Edification according to Rorty,' Canadian Journal of Philosophy 11, pp.589-601

1982

(a) Review of Theory and Evidence by Clark Glymour, Philosophical Review 91, pp. 124-128.

(b) Profiles of Henry E. Kyburg, Jr. & Isaac Levi ed. by R. Bogdan, Dordrecht: D. Reidel. xi,322 pp.

This volume consists of intellectual autobiographies by Kyburg and by me together with four essays, two devoted to Kyburg's work and two to mine, together with our replies and bibliographies of our work. The Profiles series is a series of such volumes designed to review the work of contemporary philosophers in mid career.

(c) 'Dissonance and Consistency according to Shackle and Shafer,' PSA 1978, 2.

(d) 'Conflict and Social Agency,' Journal of Philosophy 79,

(e) 'A Note on Newcombmania,' Journal of Philosophy 79,

(f) 'Liberty and Welfare,' in Beyond Utilitarianism ed. by A.K. Sen and B. Williams, Cambridge: Cambridge U. Press

(g) 'Truth, Fallibility and the Growth of Knowledge,' which is to appear this year in Language, Logic and Method ed. by R.Cohen and M. Wartofsky Dordrecht: Reidel.

(h) 'Ignorance, Probability and Rational Choice,' Synthese.