NUCLEAR REGULATORY COMMISSION

ADVISORY COMMITTEE ON REACTOR SAFEGUARDS

IN THE MATTER OF:

SUBCOMMITTEE MEETING

on

RELIABILITY AND PROBABILISTIC ASSESSMENT

Place - Washington, D. C.

Date - Wednesday, 5 December 1979 Pages 1 - 278

Telephone:
(202) 347-3700

ACE - FEDERAL REPORTERS, INC.

Official Reporters

444 North Capitol Street
Washington, D.C. 20001

CR8493

PUBLIC NOTICE BY THE

UNITED STATES NUCLEAR REGULATORY COMMISSION'S

ADVISORY COMMITTEE ON REACTOR SAFEGUARDS

Wednesday, 5 December 1979

The contents of this stenographic transcript of the proceedings of the United States Nuclear Regulatory Commission's Advisory Committee on Reactor Safeguards (ACRS), as reported herein, is an uncorrected record of the discussions recorded at the meeting held on the above date.

No member of the ACRS Staff and no participant at this meeting accepts any responsibility for errors or inaccuracies of statement or data contained in this transcript.

1573 002

CR8493

UNITED STATES OF AMERICA

NUCLEAR REGULATORY COMMISSION

ADVISORY COMMITTEE ON REACTOR SAFEGUARDS

SUBCOMMITTEE MEETING

on

RELIABILITY AND PROBABILISTIC ASSESSMENT

Room 1046
1717 H Street, N.W.
Washington, D. C.

Wednesday, 5 December 1979

The ACRS Subcommittee on Reliability and Probabilistic Assessment met, pursuant to notice, at 9:35 a.m., Dr. David Okrent, Chairman of the Subcommittee, presiding.

PRESENT:

DR. DAVID OKRENT, Chairman of the Subcommittee

MR. HAROLD ETHERINGTON, Member

DR. J. CARSON MARK, Member

MR. JEREMIAH J. RAY, Member

1573 003

PROCEEDINGS

DR. OKRENT: Good morning. The meeting will now come to order. This is a meeting of the Advisory Committee on Reactor Safeguards Subcommittee on Reliability and Probabilistic Assessment. My name is David Okrent. I am the Subcommittee chairman. The other ACRS members present at this time are Mr. Carson Mark and Mr. Harold Etherington.

Also in attendance -- and Mr. Jerry Ray, Jeremiah Ray. Thank you.

Also in attendance are several ACRS consultants, Mr. Lave, Mr. Lowrance, Mr. Shinozuka. I think we will have one or more other ACRS members and consultants coming in later.

Also in attendance are two ACRS fellows, Mr. Michael Griesmeyer and David Johnson, at the table.

The purpose of this meeting is to discuss the use of risk assessment methods and the establishment of risk criteria by government agencies and other groups. We would also like to get suggestions on appropriate goals and criteria which might be used for nuclear power reactors. The meeting is being conducted in accordance with provisions of the Federal Advisory Committee Act and the Government in the Sunshine Act.

Mr. Gary Quittschreiber, on my right, is the designated federal employee for the meeting. The rules for participation in today's meeting have been announced as part of the notice of this meeting, previously published in the Federal Register, November 20, 1979.

A transcript of the meeting is being kept and will be made available as stated in the Federal Register notice. And it is requested that each speaker first identify himself and speak with sufficient clarity and volume so that he can be readily heard.

Those around the table have nametags, so the recorder will know who you are, but if others would please identify themselves in at least enough time so that the recorder knows, that will be helpful.

We have received no written statements or requests for time to make oral statements from members of the public; however, I plan to run this meeting a little more like a panel discussion with reasonable participation from the audience, so if a member of the audience would like to make a point, please get the attention of the subcommittee chairman.

We will now proceed with the meeting. Let me just, by way of brief introduction, for the benefit of those who may not be aware of it — although the subject of quantitative safety goals is not a new one in this society and it is not a new one in the question of nuclear reactors, it has not been one that has been addressed to the point that the Atomic Energy Commission before, or the Nuclear Regulatory Commission now, have identified quantitative risk acceptance cri[illegible]a.

In May — more specifically, on May 16, 1979, the Advisory Committee on Reactor Safeguards wrote to the Nuclear Regulatory Commission, recommending that consideration be given by the Commission to the establishment of quantitative safety goals for overall safety of nuclear power reactors, and gave some reasons for why they thought this would be a useful thing to try to do. The next step chronologically was that one of the commissioners wrote to the ACRS asking for further comments in this regard.

And in August the ACRS said it would try to see if it could develop some possible proposed criteria that might be considered for use by the Nuclear Regulatory Commission, that could be published for comment, or this sort of thing.

I might note an additional item that has transpired in this regard — is that in NUREG-0585, which is entitled "TMI-2 Lessons Learned Task Force Final Report," this task force made one recommendation, namely, Number 11, which is entitled "Safety Goals for Reactor Regulation," and in this they recommended that the Commission develop definitive policy guidance or articulation of a basic safety goal for nuclear power plant regulation. They didn't specifically urge that this be strictly a quantitative goal, but they suggested that the goal would be supplemented, where possible, with quantitative risk criteria.

Then, more recently, on November 9, 1979, in a letter from the chairman of the Nuclear Regulatory Commission, Dr. Hendrie, to Dr. Frank Press, director of the Office of Science and Technology Policy, in which the Nuclear Regulatory Commission commented on the President's Commission report with regard to the accident at Three Mile Island. The Nuclear Regulatory Commission acknowledged the previous recommendations of the Advisory Committee on Reactor Safeguards, that the Lessons Learned Task force -- and stated that a safety goal of nuclear power plant regulation in terms of clear subjective criteria is needed and should be supplemented, where possible, by quantified reliability criteria.

So, in any event, there is a small recent history of this topic within the last several months, and what the ACRS has tried to do in the last few months is to see whether, via this subcommittee, one could develop both background information and possibly one or more alternative proposals for risk acceptance criteria, quantitative if possible.

But we shall see. In any event, the purpose, then, of this meeting, was to try to learn a little bit about what was being done in this regard in some of the other government agencies, and also to see what suggestions one might get for goals in quantitative risk criteria that might be used for nuclear power reactors.

So that is by way of brief introduction. And since I am a minute ahead of the agenda, I will try to continue that way, since it is probably the last time today that will be the case.

Our first speaker is Dr. Vincent Covello of the National Science Foundation, who I believe will discuss some NSF sponsored programs on risk assessment. I might note there are a few copies of the agenda on the table over there if anybody doesn't have that and would like to have one. So, Dr. Covello.

(Pause.)

DR. COVELLO: My name is Vincent Covello. I am the Program Manager for the Risk Analysis Program at the National Science Foundation. I have been asked today to describe very briefly to you the program in risk analysis at the Foundation, as well as to answer any questions that you might have about that program.

First of all, the program is very new. It is approximately five months old. It is located in the Division of Policy, Research and Analysis in the

Foundation. It is also part of the Technology Assessment and Risk Analysis Program as a whole. It is a subcomponent of that particular program.

Because it is located within the Division of Policy, Research and Analysis, all studies that are sponsored through this program have a policy focus, and I will come back to that later.

The program was created in response to several requests, specifically one from the House Committee on Science and Technology, which asked the Foundation specifically to develop a long-term program of research that looks at methods of risk assessment. We have been asked to sponsor research that would develop improved methods for risk analysis and risk assessment. We also received several requests from agencies including the Nuclear Regulatory Commission to also look at questions such as risk acceptability and general questions relating to issues and risks.

In response to these various requests the Foundation, approximately seven months ago, created this particular program. It was formally announced five months ago, in August. We are engaged in two activities, in response to the request.

First we are engaged in planning activities. The planning activities are threefold. First, we have asked the

National Academy of Sciences to assist us in planning a full-scale program in risk analysis. We hope that as a result of the Academy's work they will produce an agenda of research that we can use as part of our own planning for risk assessment. We expect that a program announcement will probably be announced and available to the public this coming summer.

Second, we have commissioned several papers on risk assessment on the state of the art, surveying where we are with regard to risk assessment and where we should be going.

And third, we have also established a liaison activity with various agencies, trying to assess their own needs and what types of activities at the Foundation would support their particular programs.

In addition to the planning activities for the full-scale program once it is in place, we are also currently engaged in encouraging the submission of unsolicited proposals dealing with a wide range of questions related to risk assessment.

What we have done recently, if everyone has the handout, we have sent out a letter to the research community, to approximately 4000 people, announcing the creation of the risk assessment program, and specifying the types of questions that we are interested in and the type of

kapHEE

proposals we would be interested in supporting. I think you will see in the questions listed there they cover a wide range of activities including how to determine how safe is safe enough, a question dealing with the perception of risk, risk acceptability, institutional/organizational constraints on risk decision-making -- these are the types of questions we are addressing. And they specify a very broad range of activities and we expect, at a later date, to be more specific in the types of proposals we are interested in.

At the present stage, though, since we are primarily interested in obtaining some of the best information and knowledge as to the state of the art, these are the broad outlines of the program.

The program has several constraints on it, though. In addition to proposals that deal with those questions, we are also asking, first of all, that the proposals submitted to the Foundation deal with risk on a generic level as opposed to dealing with specific applications. As I mentioned before, our task is to develop improved methods of risk assessment.

As a result, if specific applications -- for example, to nuclear power or to other energy systems -- are suggested, we ask that those particular studies be considered as case studies addressing a more general, broad issue in assessment.

kapHEE

The second constraint on the program is that all proposals have to have a policy focus. In other words, all the proposals and research responses should somehow or another address the questions that are being asked by decision-makers, and they should be addressing issues currently on the agenda, or expected to be on the agenda for risk analysis in the coming years.

The third constraint on the program is: we primarily deal with technological risk and not with natural hazards or with entrepreneurial risk. Technological risks such as energy, toxic substances, the whole range of risk that relate to science and technology — we cannot deal, given our mandate, with risks that don't fall within that category.

At the present stage, because we are involved in our planning activities, the type of research we are supporting range from between projects of $150,000 to $200,000, and a year to a year-and-a-half worth of effort. We have received as a result of the addendum to our program announcement approximately 150 letters of inquiry and preliminary proposals. We are also now already reviewing several proposals that have come out as a result of that letter.

We have also supported several studies that are ongoing. These particular projects were not supported directly under the risk analysis program. They were taken into the program after the creation of the program. The two are as follows: the one project, by Robert Cates and Roger Casperson at Clark University. The title of their proposal is "Methods for Improving Public Policy for Technological Hazard Management." The focus of that particular study is to develop a taxonomy of risk and to use that taxonomy of risk in developing approved methods of technological hazard management.

The second major project we're supporting in present times is one by David Okrent, the chairman of the subcommittee at the meeting. The title of that project is "Alternative Risk Managment Policies for State and Local Governments." The focus there is looking at types of risk that have to be dealt with at the local level and developing improved methods for management of those types of risks.

On both of those projects, there is a subcontract in Eugene, Oregon, which is looking at perception of risk and feeding that material into the main projects.

At the present time, the program has a budget of approximately $1.5 million. We expect within the next five years to increase the size of that budget to something larger than that, although it hasn't been exactly determined how large that budget will be. It will primarily depend upon the response of the research community to the types of questions we've asked. And as a result of that response we will determine how large the budget will be.

That is the general outlines of the program. I think it's probably easier for me to address questions than to go into too much detail about any of these particular studies we are addressing.

MR. LOWRANCE: Bill Lowrance, consultant. I have been concerned for a long time with how the NSF works with the other mission agencies in supporting research of this sort. I think this is a good example, where NSF could do generic research that could serve the needs of anybody from the FDA to EPA to NRC. And I wonder, mechanically, how you interact and whether you are soliciting those mission agencies for ideas, either research needs or people who might be interested.

DR. COVELLO: Well, first of all, one of the activities we engage in is direct liaison with the mission agencies. For example, the creation of the National Academy of Sciences Committee on Risk and Decision-Making was partially in response to a request from the Nuclear Regulatory Commission, from the office of Saul Bean, to look at questions of acceptability of risk. We discussed with the NRC the various needs, concerns, and used that as the basis for our own putting together of the Academy committee.

In addition we have liaison activities with each major agency that has risk assessment responsibilities. We share with them proposals for both review of our proposals to them -- we also review proposals on their side and that way we maintain a liaison of knowing what the other groups are doing and what their programs are.

A third things we are doing — and this is much more specific — is that we have asked the Academy to meet periodically with an advisory group we are at the present time establishing, of government regulatory agencies dealing with the question of risk. This will be approximately 14 individuals representing each of the mission agencies that has risk assessment responsibilities.

And we will have this advisory committee of government agencies meet with the Academy committee to indicate to them what the needs are of the particular agencies, and what types of questions they would like to see addressed from their perspective.

We hope through this interaction between the Academy and the advisory committee that the Academy will be more responsive to the needs of the mission agencies. We also hope to use that advisory committee as an information resource for our own program on the needs of the particular agencies.

DR. MARK: You use the word "technology" and distinguish between [illegible] nkebite, for instance. How broad an area is cover[illegible] [illegible]nology? Is it all industrial? Or scientifi[illegible] [illegible]ivity that affects the population?

DR. COVELLO: The actual outlines of it are still not clear. As you know, the definition of "technology" varies from person to person and place to place.

DR. MARK: I mean, the release of carbon dioxide by burning wood in Vermont isn't very technological.

DR. COVELLO: Right. But what we are looking for in the proposals is connection to science and technology in some direct, as opposed to an indirect or implicit way. So carbon dioxide would not be, but if we looked at the technologies that produce carbon dioxide or carbon monoxide, or other sort of types of toxic substances, we are looking at the actual production of technology that produces it, as opposed to the product of that technology.

There is the focus. It has to be tied back to the particular technology itself, as opposed to the product of that technology.

DR. MARK: It seems to me that there is something perhaps not caught up here that policy as well as technology may result in carbon dioxide. And is that within the scope of the matters of concern?

DR. COVELLO: I'm not sure I understand.

DR. MARK: Well, I referred to the burning of wood in Vermont. People are saving enormous amounts on their gas bills by doing that, and if everybody in the country should do that, this is not so much a matter of technology as a matter of policy. And is that going to be caught up by your net?

DR. COVELLO: I will probably have to leave that as a open issue, because we have not yet tried to strictly define the boundaries of what would be considered within the program, to find out what is technology and what is not. We are hoping as a result of our own internal planning activities and review to come to a more specific definition of what science and technology related risks are, and to use that as a boundary condition for the program.

Again, it is such an open area and subject to such controversy about what would be considered a technology and what is not a technology — there is no specification on what would be determining that, how that would be determined.

DR. OKRENT: Is there some reason why this is a new program in NSF? The question of risk and society and what constitutes acceptable risk, it seems to me, has been staring us in the face for some time. I would have expected NSF would have been ahead of the Congress, instead of behind it, in this regard.

DR. COVELLO: Well, the NAE study was sponsored by the National Science Foundation for technology assesment and the RANN program, the one on risk benefit, is that the one you are referring to?

DR. OKRENT: In 1972, but there has been no program.

DR. COVELLO: No, there has not been a specific program in risk analysis or risk assessment at the Foundation. Most of the activities have dealt with specific technologies. There has been a program in chemical threats to the environment. There is a program on earthquake hazard mitigation. There are programs that deal with different types of science and technology and risks associated with them.

There hasn't been a program at the Foundation which has focused on risk assessment per se. Questions related to risk have been dealt with as part of those other programs. What I believe has happened now is that as a result of both the Congress and internal planning as well is that it is decided that the Foundation should have a more visible and centralized location for risk analysis studies. And the importance of the question, the generic issues, that crosscut various types of agencies and various problems are being recognized -- and it is seen as beneficial to have a centralized location for risk analysis programs that will coordinate with each of the separate programs within the Foundation.

DR. OKRENT: Because, in fact, there was no home within NSF that one could find for such a proposal a year or year-and-a-half ago -- I know from my own experience. Dr. Shinozuka.

DR. SHINOZUKA: You specifically mentioned that your program excludes consideration for a natural hazard.

DR. COVELLO: Right.

DR. SHINOZUKA: It seems to me that a natural hazard obviously is an important component when we assess, say, safety of structures -- how we design structures against such natural hazards -- result in different levels of risk. I would like to know why, then, you do not entertain the proposal that might involve natural hazard consideration.

DR. COVELLO: At the present time most of the proposals dealing with natural hazards are dealt with by a program is the Foundation's earthquake hazards mitigation program, which primarily deals with earthquakes and natural hazards as well, including generic questions relating to natural hazards. It is possible that at some later date we will coordinate closely with that program. One of the activities relating to the planning activities of the Foundation is to determine whether or not those activities should be joined together. At the present time there is a functioning program dealing with natural hazards in the foundation, and it wasn't seen to be —

DR. SHINOZUKA: I am aware of that program that deals with earthquake hazard mitigation in the National Science Foundation, but I am not quite sure if they really focus on, let's say, the policymaking based upon risk analysis. That is my primary concern.

DR. COVELLO: There is one -- I brought along the program announcement for that particular — the earthquake hazard mitigation program, and they do have a policy research section. And I will read from it:

"Policy research: This particular element of the program strengthens and facilitates the adoption and implementation of technological, social, and management practices that minimize earthquake damage and facilitate

493 02 02

pv HEE

rapid recovery from earthquake disasters."

They have, in other words, within their mandate to deal with policy questions relating to natural hazards.

DR. OKRENT: Well, thank you, Dr. Covello.

The next speaker is Dr. J. L. Von Thun.

MR. VON THUN: It is Larry Von Thun, and there is no doctor.

DR. OKRENT: He is from the Bureau of Reclamation, and he's going to give us two presentations. The first discusses the mission progress of the seismic risk analysis interagency committee on seismic safety. And then he will discuss the Jackson Lake risk analysis. And if we are fortunate, he may also provide some suggestions on what we should do for nuclear power.

MR. VON THUN: I am with the Bureau of Reclamationn. I don't know how many of you know what the Bureau of Reclamation is.

DR. OKRENT: Why don't you tell us in 30 seconds.

MR. VON THUN: They just changed their name. I will have to tell you that, as well. The Bureau of Reclamation is a water resources agency dealing with the 17 western states. Most of the large dams in the western U.S. were built by the Bureau of Reclamation between 1902 and the present: Glen Canyon, Hoover, Shasta, most of those dams. As such, as an agency which is producing a product,

1573 021

493 02 03

pv HEE

we are more on the firing line in front of, say, states, even other federal agencies, with regard to risk assessment, people asking us, "What are you doing about risk assessment," rather than, say, being devoted to research.

So, what I will be discussing today is how we are trying to apply our risk assessment methods with the information that we currently have.

DR. MARK: Could I ask. You said you have 17 western states?

MR. VON THUN: Yes.

DR. MARK: And the other states, you are not active; but there, the Corps of Engineers might be doing some things?

MR. VON THUN: That's right. The Corps of Engineers is really applicable to the entire U.S. The Bureau of Reclamation, beginning with the Reclamation Act of 1902, only dealt with reclaiming or providing water, essentially, for the arid west. That is the differentiation. Whereas the Corps of Engineers works with navigation, and that can be within any of the U.S.

The first presentation is with regards to a committee on seismic safety in federal construction. After the San Fernando earthquake in 1971, a lot more attention began being paid to the seismic safety of structures, both dams, buildings, even nuclear power plants, although that

pv HEE

was well along the way as far as seismic risk assessment. And with the failure of a number of dams throughout the U.S. and generally more an dmore attention being paid to seismic safety, the President asked the Office of Science and Technology Policy -- I believe that was where it was initiated -- to form an interagency committee of all federal agencies concerned with building structures to come together and decide now the seismic safety could be improved.

And that committee has a number of subcommittees. They have committees on seismology, on site hazards, on critical facilities. And one of their committees is on risk analysis. And the question to this subcommittee is: how should federal agencies be using risk analysis?

My personal background with regard to how I got on this subcommittee stems largely from two events. One is the Auburn Dam in California -- which is not a dam as yet, it is just a proposed dam. It has been the subject of a great amount of public interest with regard to seismic safety. California has a Seismic Safety Commission which asked these questions of the Bureau and of any other agency which is planning on building a structure as to: can you verify its seismic safety? And this gets into the question of what is the true risk involved?

Carl Steinburge, who is the chairman of that Commission, asked me the question before a meeting similar

493 02 05

pv HEE

to this with regard to acceptable risk to actually come up with a better definition for "acceptable risk," because, as he felt, "acceptable risk" is neither acceptable now nor accepted by anyone, it is just a term we all use. And in grilling before the Seismic Safety Commission, the Bureau began to think a lot about what and how to apply risk assessment in that work.

Then, following the Teton Dam failure, we were revived by numerous agencies, and, without exception, the advisement was given to develop risk assessments in our work, both with existing dams and with ongoing dams. So, some of the work that we have done has been in that regard in trying to work with what is the seismic safety of our existing dams and how do we determine what the probability of risk of their failure is.

And it was through those types of efforts that I was involved in this committee or asked to be a part of this committee.

As far as the sccomplishments of our Subcommittee 5 on risk analysis, there have not been a great deal to date. I have provided you a three-page handout here which should not be considered the output of the committee, but rather my own personal input to the committee, because there have not been any decisions reached by the committee as yet as to how they are going to approach the problem of recommendations to other federal agencies on the use of risk assessment. But there are a few things which fit right into what you are asking for as far as the purpose of this committee: what is being done and what might be done.

The first thing that I would discuss in this regard is a general framework for risk assessment. We have three elements, as I see it, three elements with regard to seismic safety: one is what is the probability of occurrence of seismic events; the second is what is the probability of failure of a facility given the seismic event occurs; and the third is the estimation of the consequences of failure of the facility.

In reading the literature, the terms "risk, hazard, hazard potential, and exposure" are used interchangeably; there is no consistency. And the first thing I was recommending to this committee is that we develop a framework where we consistently refer to these different terms with regard to their application in decision analysis by risk-based methods.

In this vein, then, I suggested the term "exposure" is used in reference to the action or influence to which a facility is subjected. In other words, what is the exposure of a particular facility to an earthquake either shaking or faulting.

The second term, "hazard," I suggest that that be used in reference to the damage potential. In other words, when we create a dam, a nuclear power plant, or a building, we have created a potential hazard that can fail and cause some damage. Without the facility there, the hazard doesn't exist.

The third term, "risk," has been used in all of the elements we talked about. The risk of the earthquake, we talked about the risk of failure of the facility, and we talked about the risk of loss of life to individuals. I would suggest that, rather than use it in all those veins, that we use the term "risk" to refer only to the probability of failure occurring as the result of some exposure to some action such as an earthquake and resulting in some damage.

The total framework which we deal with I called "risk-based decision analysis." With all the work that we do at the Bureau, that is the way we refer to it. So, this is the framework that I would suggest using.

As far as residual risk goes, the definition is supplied in your handout, and this is how I would look at residual risk and think that it should be advanced to the public and within all the different agencies.

"Residual risk" is the probability of occurrence of a facility failure, the chance that a facility will fail, and that loss of life or economic loss which remains in existence. Some probability of that remains in existence.

After the loading conditions, the earthquake loading conditions or whatever loading conditions you have, after those conditions have been specified, appropriate designs to take care of those conditions have been prepared and additional safety precautions are established, what is left after you do all of those things — after you have assigned the facility, after you have specified what the earthquake is for design or the other loading provisions for design, after you have designed the facility to take that loading and after you have established safety precautions — that is "residual risk."

And when somebody accepts that, when the owner of a facility or the U.S. Government or somebody says, "Okay, there is still this much chance that there could be a fouiup or a problem," that becomes "accepted risk." And to date, that has not occurred, and that is really what I see you are talking about doing through this committee, is defining some number that becomes the "accepted risk."

We have talked about "acceptable risk," but I think in very few cases, if any, has an agency or an owner decided this is how much risk that I am willing to take.

With regard to what recommendations our subcommittee might make, the next two pages provide my sort of device or input to our subcommittee with regard to recommendations. If you look at the middle of the second page of the handout, it says: "The question then is: considering the current level of knowledge, experience, and practice in risk analysis, what are the possible types of requirements of federal designers and regulators that could achieve the above results" -- meaning lowering the chance of seismic hazards.

These requirements and regulations would occur in three different areas: One, the determination of seismic loading. Under the determination of seismic loading, then, there are three possible requirements that I have suggested.

One is to specify the technique to determine the level of loading to be used; the second is to establish a requirement that to determine the risk level of loading selected and compare it to other risks -- in other words, a person designing a building or person designing a dam or a nuclear reactor or whatever would be required under our suggestions to determine what is the level, what is the risk of that particular level of loading. Then item 3 would be just as -- or item C would be just as item B except we would specify what the methodology is that would require -- that would allow that determination of risk.

If we got to the point of saying a risk level of loading at, say, 1 x 10 to the 4th was an acceptable or was the acceptable risk of loading, loading only, then one approach would be to say, to actually calculate the loading, given that an earthquake has an occurrence interval of a certain amount, that has a probability of not being ceded in so many years and that would specify that would meet that particular requirement of loading.

As of right now, my recommendation would be that all we can require of federal agencies at this date is to ask that in each analysis the risk level of loading be selected, be determined, and that it be compared to the risks of other loadings so that we can get a handle on what is the likelihood of all of the loadings and people actually go through the process of seeing what the risk is. We are not at a stage — I don't feel we are at a stage where we can tell people building buildings or building dams that this is the level of earthquake loading that should be used and have that applied across the board.

Okay, the second area where we would require some regulation is in determining the merit of design provisions in the reduction of risk. This is to do with the question: if we add such and such, a reinforcement, to our design, or we add some additional safety precaution to our design, how does that in fact lower the risk of failure?

And the requirements that we could impose there are: specify that the alternative designs to reflect the different degrees of resistance to seismic loads be identified; require that the reduction in risk as a result of various design measures be estimated along with the cost of the measure; and require an estimate in terms of the savings in potential damages be computed thus allowing a cost-benefit relation to be drawn.

This alternative A here is really nothing more than what was suggested by Professor Whitman back in about 1971. He made a classical study of earthquake hazards in Boston and went through an analysis showing a cost-benefit for designing the buildings to various earthquake levels. Right now I don't see that we are much farther along than that.

And a possible requirement is that we just ask people designing buildings or any other structures to go through this process of looking to see what would we do to lower the risk of failure of this structure if we added this much reinforcement or if we made this particular change, look at the alternatives, and see what the benefits are.

In effect, this is very similar to the passage of the NEPA Act. There was no requirement other than -- or the basic requirement was that people look at the alternatives. And this is what we're talking about here: look at what the alternatives are, see and calculate what the reduction of risk is.

A second possibility is that we would specify that the risk that is computed be no greater than a certain amount. And the third one -- which is in error here -- is that require that that requirement be determined by a specific methodology; in other words, we would say that the risk of failure of a facility can be no greater than a certain amount and we will tell you how you're going to calculate that certain amount.

Those would be the possible recommendations with regard to determining the benefit of alternative designs.

And the third item there is with regard to the final product of determining what the risk of a loss of life or economic loss is once the structure in in place and, given that the event takes place, we could require that the residual risk of various potential damage levels be determined for certain structures. And then the second part of that is that we would specify how much that residual risk would be.

And as I say, these are not adopted by the committee. In fact, these have only been sent to them. But the timing of the meetings was such that they were going to meet yesterday and they didn't meet. They will be meeting next week. In fact, we have one member here from that subcommittee.

These are just suggestions to give you some idea of what is being thought about with regard to imposing some

493 02 13

pv HEE

sort of requirements or regulations or at least suggesting these for use by other federal agencies.

There is a member of the Nuclear Regulatory Commission on this subcommittee, by the way.

So, I am open to questions on what our subcommittee is doing.

DR. OKRENT: Dr. Lowrance.

DR. LOWRANCE: Bill Lowrance, from the seismically active area of Palo Alto, California.

How would item 2-B be pursued? This says "specify that the risk of failure given various seismic loadings associated with the structure be determined and that it be no greater than a given amount; that is, a risk ceiling be prescribed." Where do you suggest that come from?

MR. VON THUN: That's why I am saying that is somewhere maybe in the future. Right now we are not to the point that that can be done. When the committee was formulated, that was one of the things that was sort of thought about, that this might occur. In fact, in the directive we will get a little more to where that can come from later.

But basically, the only place that I see that it can come from is for people to begin applying the rules or begin applying the risk assessments, see what types of norms or standards can be developed with regard to what is a good

493 02 14

pv HEE

value to use as a risk level. If we don't make the studies of actual applications, we won't have an idea whether 1 x 10 to the 6th or 1 x 10 to the 8th or to the -8th is a valid number. So right now the number isn't available.

DR. OKRENT: Dr. Lave.

DR. LAVE: What is the uncertainty level associated with either estimating the seismic risks or with estimating the probability of the failure of a structure in relation to a particular type of risk?

MR. VON THUN: It is very great. I would say that the first one is quite a bit less. What the loading is is less than that can be estimated better than can what the probability of a failure is given that loading occurs.

The reason for that is we have had as few earthquakes as we have had. We have had a lot more earthquakes than we have had failures due to earthquakes, so we don't have a very good statistical base on which to decide how a dam or a structure is going to react to the earthquake. Buildings may be a little different than dams, but we don't have — we have, for instance, in the United States only a couple of embankment dams that have suffered damage due to earthquakes among the thousands and thousands of embankment dams that we have.

DR. LAVE: Let me pursue this a little bit more. How is the probability of a seismic event estimated? Is it simply historical frequency?

MR. VON THUN: No. The probability of a seismic event has undergone -- the estimation has undergone a great deal of work, and it is continuing to undergo a great deal of work.

But is based upon what is the geologic capability of a particular area, and that is referred to as the "maximum credible earthquake," or in the nuclear reactor field as the "safe shutdown earthquake." This is the largest earthquake that can be postulated to occur for a particular seismic region. There we're talking about the probability of an event that could vary from once in 100,000 years or once in 500,000 years down to where it could occur once every three or four hundred years. For smaller-sized earthquakes usually the historic frequency is the method that is used.

e2

mgcHEE

DR. LAVE: What steps have been taken or can be taken to try an quantify the level of uncertainty that is associated with each of the two events, either the seismic loading or the structure failure in case of a particular seismic loading?

MR. VON THUN: There have been a few studies. There was recently a conference in Pasadena on earthquake engineering, and at that conference there were some suggestions made by Woodward, Clyde & Associates of how to provide some constraints on the uncertainty. If one allows the uncertainty to go unchecked, then at very low risk levels it becomes a tremendously predominant factor.

So they were suggesting that certain modifications be applied to that uncertainty function as it approached the low risk levels. So other than a few attempts to look at this uncertainty question, there hasn't been a lot done.

DR. LAVE: Has there been any attempt at all to take a look at the uncertainty level with respect to the mean probability or the most frequent probability?

MR. VON THUN: Yes, there is.

DR. LAVE: And what kind of sub-numbers do you come out with?

MR. VON THUN: I don't have them offhand. Do you mean as far as, say, the mean level of acceleration is such and such, and what is the bounds on that?

DR. LAVE: No. I was thinking about — you have here probabilities of seismic events of, for example, .003, and obviously there's a large uncertainty associated with that. Is the uncertainty equal to the mean or ten times the mean or a thousand times the mean?

MR. VON THUN: I think one of the best examples of what that uncertainty is is in a paper that Dr. Okrent had a lot to do with where the experts in seismology were asked for a number of different nuclear reactors, nuclear reactor sites, as to what the low level risk would be for this maximum credible earthquake. Each of these investigators postulated from available information what the risk of a certain intensity earthquake was, and that gives an idea of what the range is. And it was quite a range.

As far as a detailed assessment and application of uncertainty, I am certainly aware of it, but I don't think it has been done. I am aware of its need to be done, but it really has not been done to the degreee which you are asking the question about.

DR. LAVE: Just one final one here. What could be done to lower uncertainty levels?

MR. VON THUN: Well, with regard to the seismic event, there isn't really anything that can be done except do a good job of analysis of your site. We don't have any othe information than the geologic record which we might be

mgcHEE

lucky and excavate through a fault and find that it has had so many displacements in so many thousand years and be able to estimate what the earthquake is on the basis of those displacements.

At the Auburn site in California, we went through a tremendous study and really came up empty-handed as far as being able to definitively say what the recurrence interval of earthquakes in that area was. So it isn't a foolproof thing, but it is something that should be done.

We don't have any more historic records than we have, so we can't do any better on that. But my feeling is that the uncertainty in that is not so great.

With regard to performance under seismic loading, we can make additional tests. We can do better studies. We can develop dynamic analyses methods to try and predict the response, but without prototype failures we really will always have a certain amount of uncertainty. The approach that we're taking at the Bureau is to try to make that calculation on as many structures as possible in order to assess the relative risks of these failures, and in that way, get rid of the uncertainty.

If nobody makes any analyses, if we always keep it in the research mode, then we really will always have that uncertainty sitting there as far as the actual structures. I think once we start the studies on actual structures, we

mgcHEE

will reduce some of the uncertainty.

DR. OKRENT: Dr. Wilson?

DR. WILSON: I was just wondering -- I'm talking about the uncertainties -- one thing that worries me is one tends to talk about the expected value of an earthquake or a hazard and the uncertainty around that, whereas quite often that is not what the decision-maker needs to know. He wants to know the probability of exceeding a certain bound, and hopefully he is going to set that bound so the probability is very low. And since the distribution in which the uncertainty is described and the width of that distribution is certainly not very simple, are you addressing that specifically when you are discussing the uncertainty?

MR. VON THUN: Yes. Generally what is done is an upper bound is taken, rather than the mean value, either in all the steps or certain of the steps to assess the loading. We don't take the upper bound at the first stage, the upper bound at the second stage, and the upper bound at the third stage. We may take the upper bound in certain regards, but that problem is being addressed. We do not typically take the expected value in each case, and we look at a number of possibilities.

In a later study, here you will see how that is done.

DR. OKRENT: Dr. Castenberg?

DR. CASTENBERG: Bill Castenberg from the ACRS. I'm curious about your definition of "hazard" and "risk." If I understand correctly, risk is probability, and so you could have a dam near a large city and a dam out in the country that had the same risk where the probability failure is the same, but they might have different hazards because of population. And if I understand that correctly, most of your work is geared toward risk reduction rather than hazard reduction. Is that correct?

MR. VON THUN: Well, the work is with regard to risk-based decision analysis. Now we're not to the point where those types of decisions really have been made on the basis of risk analysis. But I think fundamentally the question is right. We would look at what the risk is at this site versus the risk at another site, and that is what we want to deduce.

DR. CASTENBERG: Do you mean, in the sense of this definition, really probability?

MR. VON THUN: Yes.

DR. CASTENBERG: Or do you mean risk in a more general sense where you include hazard as well as probability?

MR. VON THUN: Hazard is always included. When you talk about the resulting certain level of damage, that would always be included, so that the risk of, say, 40

HEE

lives would be the same, no matter where you were working, but the damage level, the risk of a certain damage level occurring, that is what we're talking about.

DR. OKRENT: Could I ask the following question? Auburn dam which you mentioned was estimated by Mr. Seidegren to have the potential for causing up to three quarters of a million fatalities were it to fail suddenly, which is indeed a large number. Is there, in your way of thinking in regard to the evaluation of accepted risk or acceptable risk or whatever way we do it, is there some limitation that enters from the magnitude of the event, independent of the probability, or should we always factor in a probability times the maximum hazard and look only at the product? Or how do you propose one deals with this question?

MR. VON THUN: My proposal -- the way I would look at that is, each case has got to be considered differently with regard to what are the defined levels of damage in the case of any dam, and Auburn is a good example. You would look at what is the probability of damage level of $2 billion for property damage, say, wiping out Sacramento, doing damage to Sacramento, and the probability of so many lives lost. And when one deals with that number, that residual risk that was associated with that loss, he would refer to that specific loss, so that it wouldn't be lost in front of the eyes of the public or in front of the eyes of the Commission or anybody. He would say, this is the risk of that loss. And the decision might be made that no matter hat the risk was, even if it was — as long as it was non-zero — that the hazard of putting the structure there is so great that we won't accept any risk, and therefore the structure should not be sited there.

That would be on way to look at the problem. At another site where you had very little potential damage downstream, than a much higher risk would be maybe accepted.

I don't know whether that answers specifically what your question is.

DR. OKRENT: I was just wondering if your Interagency Committee was planning to come up with possible numbers to put in such a thing or how the Bureau of Reclamatior approaches the same question. They have a lot of large dams which, let's say, don't have perhaps the three quarter million potential, but they certainly have hundreds of thousands, I would guess.

MR. VON THUN: The Interagency Committee I don't believe will come up with the numbers. We don't really have the background on the Committee or even the diversity to allow coming up with those numbers, nor have we [illegible] e the studies. But I think that the report of the overall Committee is going to be out within a few months. It is supposed to be out within a few months, and there is just no way that those numbers could be developed and generated. And I don't think it is right at this point to do it because there isn't enough background to establish what those numbers should be.

As far as what the Bureau's approach is, which by the way, their new name is the Water and Power Resources Service -- the Bureau's approach, which is really my approach or one that I think should be used, is that we can't establish the numbers until we have looked at a number of different sites and started getting some information.

The example, which I'll show you in a moment, on Jackson Lake is currently the only real example where we have some idea -- and maybe not even all that accurate an idea -- but some idea of what the risk is on one structure. Decision-makers need to have a whole background of what these risks are in existing facilities and on other types of risks, not only for seismic risks but for flood, for normal resevoir loading, for any of the loadings. They have to get an idea of what risk are we exposing the public to involuntarily compared to what risk the public is exposing them to voluntarily in order to come up with that number.

So to just pick the number right now is not valid. What we feel is that we should go through this, look at all of our existing dams or a large number of them, find

mgcHEE

out what the actual absolute risk of a failure is to a number of events, and then we will be in a position to pick that number.

DR. OKRENT: Dr. Shinozuka?

DR. SHINOZUKA: I am seriously concerned about the uncertainty involved in the estimation of probability of failure of a structure. This point has been raised, but since I am more or less working in the structural analysis and design area, I feel that the uncertainty involved in the estimation of probability of failure of the structure would be something we really have to address ourselves to.

I feel at this time that the methodology has not really been established. If you look at some of the papers dealing with this problem, the best you can see is the application of, shall we say, first order statistics involving just expansion of the first two times, evaluating variances, and apply certain subject judgments in estimating the uncertainty and crank out numbers.

It is even difficult to apply this to a relatively simple structure consisting of a number of elements. So my point here is that I think at all levels, federal levels and also university community research for example, I think we should take up this problem more seriously and look into it very carefully so that these risk assessment methodologies can become reliable.

MR. VON THUN: I absolutely concur with you. I think that the weakest link in what I will be presenting here and the weakest link that faces us is in predicting how the structure is going to fail, if it is going to fail, and what the likelihood of its failure is. That is the weakest link, and it is not in risk assessment methodology. The methodology for the risk assessment there. It is in understanding the failure, being able to decide what its response is, and what its likelihood of failure is.

Once somebody can tell you that, then the risk assessment can go ahead and be performed. And even maybe an add-on to that is the foundation, like on a dam structure. We might know quite a bit about the materials that go into the dam and quite a bit about how the dam is going to respond itself, but the uncertainty in the foundation is even greater.

DR. SHINOZUKA: My difficulty is that although you mentioned that methodologies have been established, but if these methodologies require the information that we may never be able to get, then we are in trouble.

MR. VON THUN: That's right. And our approach really is not to say that this is going to work. It is to try and see if it works and see where the loopholes are. That is our approach, rather than feeling it is an absolute foregone conclusion that it will work.

DR. OKRENT: Maybe it would be good to go on to your next part, and then we can discuss both parts of your presentation.

(Slide.)

MR. VON THUN: What I would like to talk about is the second handout. The top got cut off. It should say, "Decision Analysis Model to Determine Resevoir Restriction Level at Jackson Lake", and the situation at Jackson Lake which happens to be in the Teton area -- the Grand Teton, Jackson Hole country.

This was a dam that was built back when a method called hydraulic fill was used where the dam was actually placed by putting in wet material and letting the wet material run out, so it was a very loosely placed dam. And a dam placed like this is subject to what is called liquefaction. When an earthquake takes place, the soil shakes. The core pressure or water pressure within the material builds up, and the material can actually flow.

Van Norman Dam in California, in response to the San Fernando earthquake, had partial liquefaction, and Sheffield Dam in 1926 had a complete failure due to liquefaction.

So as part of the reanalysis of our existing dams, which amounts to about 300, this was one of the dams that we took a look at, and it was recognized that this site was near to where there were faults, and there could be strong shaking at this site, and it was subject to liquefaction.

Professor Seed and Professor Lee from the Univeristy of California took a look at this site and made the recommendation. So we had a resevoir there that was quite important as far as the public was concerned. A number of people go there for recreation benefits. The lake also supplies irrigation benefits in the lower Snake Valley below this dam, so as far as the Bureau was concerned, it was important to maintain the resevoir as high as practical but also to reduce the risk.

So what I'm going to go through here is the decision-based risk analysis that we used in order to decide what the level of the lake should be kept at. To do this, the first point was to look at what the hazard was as a function of resevoir level for the potential failure modes, and we had a failure mode where the entire dam went out at its base, height, and one way it went out at the top of the zone where it had been placed by hydraulic fill. So we looked at two different modes of failure, and we showed that as the resevoir load lowered what the potential damage downstream would be, given that a failure occurred.

I'm going to go through these in a little more depth. I just want to go through the steps first.

The second thing was to estimate the probability of various levels of earthquakes causing liquefaction at the dam. We had several source areas that could produce earthquakes, and they could produce earthquakes of different sizes with a different probability. Given that that probability occurred or that that earthquake occurred, there was some probability based again on our evidence from around the world with regard to liquefaction that an earthquake of this magnitude might and it might not produce liquefaction in the dam.

A stronger earthquake would have much more likelihood of producing liquefaction, so we had a probability that the earthquake would occur, a probability that the dam would liquefy from this earthquake given a probability that liquefaction would occur at the dam site. If liquefaction occurred at the dam site and the dam had several probable failure modes, it might completely level out. It might level out only a few feet or something in between.

So we had to estimate the probability of damage as a function of resevoir level, because if we said that it could go through these different steps of failure and we put the resevoir at different levels, then there were different probabilities that the dam would fail for each of these failure modes.

If there are any questions as I go through this, you can ask them.

MR. ETHERINGTON: Yes. Why does a magnitude seven have different liquefaction probabilities in A, B, and C?

MR. VON THUN: Because of distance.

MR. ETHERINGTON: Oh, I see. It is distance.

MR. VON THUN: One source may be closer to the site than another. I will have one up there in a little more detail that shows that.

DR. CASTENBERG: On the column, "Earthquake Probability", is that probability per year? How do we interpret that?

MR. VON THUN: That is annual probability. Those are hypothetical numbers.

DR. CASTENBERG: Right.

MR. VON THUN: But this is how we did it with regard to annual probability, and then over the lifetime of the structure. We typically work either with annual

e3

probability or probability in 100 years -- one of those two modes.

(Slide.)

The next thing to do is to compute the probability of overtopping as a function of reservoir restriction level for each failure mode. So, we would take the probability that the earthquake could occur times the probability that it would cause liquefaction times the probability that there would be a certain failure mode on the dam, and then look at the reservoir level and see whether that caused overtopping or not.

And that gave us an absolute probability of overtopping as a function of the reservoir level, and we were able to work out this probability for the mode 1 type failure, which was in the upper part of the dam, and the mode 2 type failure. You see the mode 2 type failure was not -- in this particular case does not show much sensitivity to a drop in reservoir level.

Then the fifth step is to examine the likelihood of overtopping and downstream hazard potential for the current criteria versus any proposed revised criteria and then evaluate the total reduction in the risk of each mode, the acceptability of risk for each mode, and consider the benefits of reservoir elevation, operating procedures for maximum benefits within the range of acceptable risks.

Now I will go through those steps in a little more detail.

(Slide.)

pv HEE

This is — what I am showing now is sort of a fundamental approach that we would use with regard to seismic probability consideration in any of the studies that we do. We have a site here, and we have several sources that could affect this site with regard to earthquakes. Here is a fairly large — this is where the Hebgen Lake event occurred. We said a maximum earthquake of 7.5 could occur at this site. This is quite some distance away. Anything less than the seven to 7-1/2 range would not produce a problem at the site.

We had -- there is another source here, the intermountain seismic belt. It could produce an earthquake of 7.5 which might affect the site. And then there are two sources near the site that could produce earthquakes from 6.0 to 7.25.

(Slide.)

Now, each of those sites could produce earthquakes with this probability. The question was asked earlier about how we make the determination with regard to the probability of the event. This is based on historic record. There is no earthquake of 7-1/2. There is an extrapolation to get to 7-1/2, to get the probability of that magnitude of event. There happens to be geologic evidence in the area that shows us how much offset has occurred over the last 10,000 years, and putting that into magnitude 7-1/4 earthquakes down to six earthquakes to account for that displacement, we are able to come up with a recurrence relation based upon geologic evidence, and then that is correlated with the historic evidence of seismicity in that area. We have a fairly good correlation.

So this type of relationship was used to get the annual probability of earthquakes of a certain range, and we have to talk about a range -- seven to 7-1/2, six to 6-1/2, like that -- in order to encompass the total probability of earthquakes.

And the other question about uncertainty that was asked -- this is how we take care of -- the gentleman asked about the range of expected value and so forth -- this is how that is taken care of: by looking at the bigger earthquakes. They have a higher probability of causing damage, but they have a lower probability of occurring. But by taking them into groups like that, you can look at that fairly realistically.

(Slide.)

The next thing to do is to say, "Okay, now, from all of these earthquakes we've looked at" -- I think I got that out of order, but the idea is the same: we need to look at which earthquakes can cause liquefaction and which cannot cause liquefaction. These, we say that any earthquake within any of the zones at a certain distance and

pv HEE

a certain magnitude that falls in this range, we will say that the probability of that causing liquefaction is zero and it no longer needs to be considered in the risk analysis.

This has to be low enough so that we encompass even a minor amount of liquefaction. To get that, we use the worldwide data and the advice of our consulants on what might cause liquefaction.

(Slide.)

These Xs are examples of where liquefaction was caused at various distances, and you can see from this that where we have historical examples of liquefaction, we are up in here with fairly large magnitude earthquakes, but there have been some cases where we had small magnitude earthquakes. This lower-bound curve then encompasses all possible liquefaction-producing events. But if we did have an earthquake up in this level, the chance of there being liquefaction has got to be considered greater than if we had an earthquake at this level, say, at this distance. And that is one of the reasons for those probabilities varying.

DR. OKRENT: The Sheffield Dam you mentioned, which failed and liquefaction occurred, what was the magnitude of the event or the intensity at the site?

MR. VON THUN: I can't remember offhand. I think it was on the order of a magnitude seven, but I don't

493 04 05

pv HEE

remember offhand.

(Slide.)

This is just a reproduction of the little box that we saw earlier except these are the actual numbers that were used in the assessment. The A-1 source and the A-2 source each had the same probability of earthquake, but because the A-1 source was closer it had a higher probability of liquifying the dam. We only considered earthquakes from the other two sources in the range 7.25 to 7.5. The one at a greater distance had a lesser probability of liquifying the site.

Then, by summing all of these, all areas combined -- because we can't just look at one of the areas, we have to take the total probability of liquefaction and combining all of the areas -- this said that there was essentially one chance, annually, one chance in 100 of there being some liquefaction at the site.

This was a lot greater than I would have guessed, than I guess even right now, but this is the way, when we went through the analysis, these are the subjective numbers that were applied. I think we would make actually -- this probably should be maybe more on the order of one in 1000 because the dam has been there for 50 years thus far and there hasn't been any indication of liquefaction. But if one had to make it an assessment of putting a dam there had

493 04 06

pv HEE

there never been one, I think there isn't anything basically wrong with these estimates from the data that we have worldwide and from the estimates that were made on the condition of the structure.

But this is exactly the problem that you raised. It is our estimate of what the structure would do is — whether there would be liquefaction of the structure or not is pretty tough, but we tend to be conservative in all our estimates.

(Slide.)

Engineers in general tend to take the conservative approach.

Now, the part that is not here in any detail is the next phase, where we say -- but it is in the report which was handed out -- which again is saying what the structure will do given that liquefaction occurs.

Now, when Sheffield Dam failed, it failed down to 20 percent of its height. When Van Norman Dam failed, it only failed about to 70 percent of its height; in fact, the reservoir was not lost when Van Norman failed. Those are the only two examples that we have.

And what we did was: we said there are three possible failure modes — 90 percent of height -- I can't remember the others -- 10 percent of height, and maybe 50 percent of height. It is in the report. We said there are three different possible failure modes of the structure, and three different possible earthquake levels: the strong earthquake, the medium earthquake, and the small earthquake. The small earthquake had some probability of causing a full failure down to 90 percent, but it had a lesser probability.

So, by taking each of those conditions of the structure and each of the probabilities of earthquakes at different levels, we made an estimate of what the performance of the dam would be under the strong shaking. And this is by far the weakest part of the analysis. In any case, making that analysis, we were able to look at what the risk level would be due to lowering the reservoir elevation. And we found that if we lowered the reservoir elevation to 6756 we were cutting the risk by 50 percent, what we consider to be our absolute risk.

It wouldn't make a lot of difference that these numbers are probably not all that accurate, but we feel that they are relatively accurate. This might be a couple of other zeros on here or maybe one other zero. But the shape of the curve, we feel, is fairly accurate.

So, from this we saw that we could reduce the risk from this mode 1 type failure which we thought was the most likely mode. We thought we could reduce the risk considerably.

The other element after looking at how the risk of overtopping could be reduced, the next question is: as we lower the reservoir, the more we lower the reservoir the less damage even if we had overtopping.

(Slide.)

So, those two things work for you: you keep reducing your damage as you go down, and you reduce your likelihood of damage. So, this relationship that showed the damage cost, right here, versus reservoir elevation, we can then look at what is a good combination between lowering the probability of there being overtopping and lowering the damage with a new reservoir level.

And it turned out that if we went to any flow greater than 50,000 second-feet, then we felt that we were getting to fairly high damages: $500 million in property damage and associated loss of life, risk of loss of life. It turned out that any flood less than 30,000 cfs would stay within levees provided downstream by the Corps of Engineers. So, it turned out that that happened to coincide quite well with this risk-reduction level of 6756. Actually, 6756-1/2, which hit about right here, would keep us within 25,000 second-feet as being the maximum flood that could be produced downstream.

So, the decision was fairly obvious: to stay below this level and reduce the risk of overtopping by about

493 04 09

pv HEE

half of what it was if we did not change the criteria.

We then proposed the -- the way our agency works is we have regional offices that deal with the public and deal with actually making the designs. Our office in Denver does all of the analyses. So, we proposed this to the people who are in charge of operating the plant. They said, "Let us propose an alternative rather than just leaving the reservoir at one fixed level all year. Let's look at how we might operate the reservoir so that we can get maximum utilization for irrigation and for recreation and keep the risk within a level that you've supplied."

So, they proposed a different operating criteria, which in fact kept the risk level just as we had specified, although they did raise the height of the dam during one month of the year.

(Slide.)

So, the hazard during one month of the year is a little higher.

I didn't mention in this total probability formulation we also took into condition the operating criteria which showed the reservoir up and down during the year. The probability that the reservoir was at a certain height was taken into account in the total risk assessment.

That is the overview of that study.

DR. OKRENT: What did you do about the mode 2 failure? That was the more serious one.

MR. VON THUN: The mode 2 failure, we essentially said that its risk was lower. This was in an area that was essentially not as subject to liquefaction. We said that the mode 2 failure had a much less likelihood of occurring, and so its total risk was less. But we did absolutely nothing about changing it.

The only way it is affected is that when we lower the reservoir to this level, its chance of failure if it did occur would be less flow out of the reservoir. But there wasn't anything specifically done or any decision based on the mode 2 failure.

The same results occurred as far as lowering the reservoir, but the chance of a mode 2 failure wasn't really a factor.

DR. OKRENT: I agree. But the mode 2 failure probability is only a factor of 10 smaller on your figure.

MR. VON THUN: That's right.

DR. OKRENT: I gather the mode 2 failure leads to larger amounts of water.

MR. VON THUN: Yes.

DR. OKRENT: Greater damage, greater loss of life. I don't know. You didn't mention what kinds of loss of life could be associated with a mode 2 failure. What would it be in the summer?

493 04 11

pv HEE

MR. VON THUN: In the summer there is recreation downstream where there are boaters. The chance of the mode 2 failure actually producing overtopping -- I mean, producing a worse flood wave isn't really known. It just means that the failure would occur lower in the structure. Actually, the chance of overtopping is similar.

DR. OKRENT: So you have no basis for assuming that a mode 2 failure means more water or greater flooding?

MR. VON THUN: That's right. The difference being -- well, the total risk of its flooding is the same as the mode 1 failure.

DR. OKRENT: I guess I don't understand. I would have assumed, if you are failing down to a lower level, unless you assume once overtopping occurs you lose the whole dam anyway -- what is your assumption?

MR. VON THUN: We made several different assumptions about how that would fail. We had a 200-foot breach, a 400-voot breach, and an 800-foot breach. So, it was the same under both conditions.

But really, the answer is we essentially ignored the mode 2 failure as far as making any decisions on what to do about the reservoir and whatever risks are remaining, whatever residual risk is there for mode 2. Since it was not essentially impacted by lowering the reservoir, we did not do anything. We did not make a decision on that basis.

Now, this is not — just for information, this is not the final answer. At Jackson these were interim criteria while we decide what to do about the structure in terms of reinforcing it or rehabilitating it. But in the meantime we wanted to come up with a restriction level that would be meaningful. And the analysis actually said we can't do much about the mode 2 failure.

DR. OKRENT: I understand that. But what I am getting at is the probability of overtopping from a mode 2 failure, if I read the graph, roughly is about 2 x 10 to the -4 per year.

MR. VON THUN: That's right.

DR. OKRENT: And the only thing you can do is take the water out of the dam to some level where there would be no flooding, in order to avoid it, I guess is what you are saying.

MR. VON THUN: That's right.

DR. OKRENT: So you are accepting some such risk here.

MR. VON THUN: That's right.

DR. OKRENT: That doesn't surprise me that it is this sort of magnitude. I think it is probably larger at various other dams.

MR. VON THUN: And that is what we intend to find out. The approach that we're doing is that this is the first; we actually made a decision on this basis; there were other analyses that were companion to this, but the decision was made all the way to the commissioners' offices on the basis of this. But it is the first example where the decisionmakers have actually seen any numbers like this on which they are making a decision.

Now, the program that we have right now is a reevaluation of existing structures, and we are taking two dams -- one in a concrete dam and one in an earthen embankment -- and trying to go through the total risk situation -- this is just seismic risk -- and get a comparison of what is the risk under just normal reservoir loading, what is the risk under earthquake loading, what is the risk from overtopping due to a flood, that we have allowed some probability, what is the risk due to a landslide -- all of these different risks -- and put them all in one package.

The University of Utah is doing exactly the same thing for us at another dam, and there are several other universities that are making these sort of studies, trying to move into the more practical application.

I think MIT has a grant -- I am not sure -- I think it is an NSF grant to make a study of total risk assessment.

DR. OKRENT: They do have one, I believe, some

pv HEE

sort of risk-benefit methodology.

MR. VON THUN: My feeling is that that is where we need to go in any of these things, is to try and generate these numbers. If we find out that we have serious deficiencies -- which I believe we will -- then that is the area that we ought to look into as far as risk assessment.

I think that we have spent plenty of time to date in looking at some of the minor ramifications, like the statistical analysis of peak acceleration; we have done all kinds of refinements on peak acceleration. It isn't even a particularly good parameter to use in judging the performance of the structure, and yet we have had study after study that looks at that particular parameter.

We need to go and find out the areas where we don't know very much and see if anything can be done. Another example would be, say, looking at a concrete dam. We do make hundreds of tests on the cylinders that go into the concrete that goes into a concrete dam, and we can make a good statistical evaluation on the chance of that being less than the 4000 psi which we planned for it to be in there.

But there isn't a concrete dam around that's going to fail compression due to that mode of failure. A foundation failure which we have very little assessment on is the way that it would fail, and that is the type thing where we have to do this total risk assessment and see where we really have the weak areas and at that point decide whether the method is usable or not usable.

DR. GRIESMEYER: You said you had an option of raising the level for one month so that you can use it for recreation or better irrigation. And presumably, you kept the expected risk of overtopping constant; you lowered it a little bit during your low time and raised it a little bit.

MR. VON THUN: That's right.

DR. GRIESMEYER: And this is good if the expected value of risk is a good thing to limit. Now, if the uncertainties are large, it may be really during that one month you have an unacceptable risk.

MR. VON THUN: That's right.

DR. GRIESMEYER: Even though the expected value over time is constant.

MR. VON THUN: That's right. And if it was regarded — if we had regarded that it was unacceptable, then that's what we would have done. We said, "No, you can't do it because we still have a large enough hazard here that we will not accept the chance, the one in 12 annual chance, of there being an earthquake during that period of time."

DR. GRIESMEYER: And then you also have uncertainties in these estimates. If you're lucky, you've got a couple more zeros there in front of it; if you're unlucky, you have a few less zeros and then it becomes a more serious event.

MR. VON THUN: And we did not, in this study, go into uncertainties. We intend to try and put in uncertainties in a meaningful manner on the work that we're doing now. We have a team composed of some people who are expert in risk analysis; the rest of the team is composed of people who work in dam safety and design. We feel that it is good to have those people in as a coordinated group.

And we intend to have the whole process reviewed by university people working in risk analysis, when we are through, to try and get an accurate assessment on an accurate use of the uncertainty. If we allowed uncertainty to go unchecked, I am sure that it would completely dominate our studies because there is so much uncertainty.

What we are to try and band — the question was asked about how do we deal with structure performance — to try and get some control on this number over here, we took all of the data developed by the Corps of Engineers under the National Dam Safety Act, when they went and got ifonmation from the 49,000 dams that are in the U.S. We have information on all dam failures. We are going to — we have a program where we have all of that information on a data base — we are going to put dams in categories of height, type, location in the country, so that we know what their exposure is and develop an exposure function and then look at what their performance has been.

And that will give us some idea here, as an empirical value, than, on the other hand, we're going strictly on what I call the "calculated approach," where we're trying to have engineers actually compute, given that this is the loading, whether it be reservoir loading, maximum reservoir loading, or earthquake or flood, what is the probability that that dam will fail. And then we will have that calculated number to compare against the performance number and see whether there is any correlation at all.

&4

DR. OKRENT: Dr. Shinozuka?

DR. SHINOZUKA: I would like to know your assumption about liquefaction. When you say liquefaction has occurred, does that mean liquefaction has occurred throughout the dam or at a certain location of the dam?

MR. VON THUN: We assume that it occurred at a certain location of the dam.

DR. SHINOZUKA: And you have a procedure from which you can then evaluate probability of the dam failure?

MR. VON THUN: Yes, and also probability of damages. Each of these types of cases, as I mentioned earlier -- every one has to be considered on a case by case basis. There was a certain zone, many of the dams that are built are built in parts. First one agency builds this part, and then the farmers add this part, and somebody else adds another part. And so, at this particular site we had a certain area that was subject, more subject to liquefaction than other areas. And so that was taken into consideration.

DR. SHINOZUKA: Another question. This probability of dam liquefaction under certain adverse conditions, these probabilities will be given by some experts, or there are ways in which you can compute these probabilities?

MR. VON THUN: There are ways in which you can compute it. This particular analysis was not done, you will note when you read the paper. Did you make adequate copies?

There are dynamic analysis methods that can predict the performance of the structure. These are plagued by not having enough real life examples to compare with. But they can give you an estimate of what the performance has been on the basis of what we've seen in those analyses, plus what we saw at Sheffield Dam.

That is how the analysis was made. All you have here is the summary report. The total report has all of the appendices, the comments of Professors Seed and Lee, the studies on the hazard that the USGS made and others. I just brought the summary report along, and I don't know whether it has the Sheffield earthquake magnitude in it or not.

If there are no more questinos on Jackson Lake -- are there?

DR. OKRENT: I think Dr. Wilson has a question on something.

DR. WILSON: Yes. I would like to ask a general question. I have a colleague who has been connected with dams for some years, Arthur Cassagrande, with whom I've discussed these matters. And I know he has always been -- he maintains, of course, that a properly designed dam will just not fail.

And there, of course, the question is on the adjective "properly." And when we try and get around to that, it turns out that there's a difference of opinion of himself and some other people in the business. And of course, all the dams he designed he feels are properly designed, and so on.

Is there any way of taking account of what appear to be or what certainly, on some of the older people in the business, are major differences of opinion in how one should go about some of these dams?

MR. VON THUN: On how they should be designed?

DR. WILSON: Yes. As far as I can make out, if perhaps they don't persist in the younger people who are designing dams now, know about it; but they persist certainly in the people of Arthur Cassagrande's age, the difference of opinion as to whether some dams are well-designed or not well-designed.

And if you picked two consultants working on this from one group of opinion, you might completely not get the proper spread of the uncertainty.

MR. VON THUN: I don't know. There are certainly ways to crank that in. When we are dealing with remote problems, we start with, say, an earthquake that has a maximum credible earthquake, I mentioned earlier, might have a likelihood of one times 10^{-5}, the likelihood that the dam failed if it is designed well. If you have someone like Arthur Cassagrande and he would say, gosh, I'm sure my dam isn't going to fail.

And then you're talking and you say: Well, how

mte 4

sure are you? Is there one chance in a hundred that it will fail? And so you then talk about one chance in a hundred times one chance in 10,000 that there would be a failure event. And then you talk about, well, how much failure is failure. And so a total failure -- there might only be one chance in 100 that, given that there is some failure, there might be a total failure.

So now you're talking about something like one times 10^{-9}. To crank in a difference of opinion on how one person thinks of a design versus how another thinks of a design may not be all that meaningful in generating that total number. But there are certainly ways to do it, because you could then say: Well, if he says this is the way, and he says this is the way, and you really think there is serious concern, then you would just lower that factor in of one to 100 which the person said could occur maybe to one in 10, to account for the fact that there is a dispute over how it should really be designed.

So I think that that is about as responsive as I can be to that question. We have tried in the Corps of Engineers data, which records whether or not the structure is engineered or not engineered, because a number of the 45,000 dams in the U.S. were not actually engineered at all -- and also, some of them are inspected regularly and some are not inspected regularly.

Ace-Federal Reporters, Inc.

1573 069

mte 5

In our function of deciding which ones are more likely to fail or not to fail, we are going to crank in that sort of information to make that judgment.

Okay, the very last sheet of your handout on the Jackson Lake paper doesn't have anything to do with the Jackson Lake paper, but it was just included in this handout. This is from another study that we made, again only to do with seismicity. This was more of a study to decide what number we would assign as being an acceptable number for seismic risk.

And our problem here is that in certain areas of the country, rather than having a specific fault to deal with, we might have earthquakes at random location. This is more prevalent in the areas such as Nebraska, Wyoming, North and South Dakota, Kansas, and those areas. Here an earthquake could occur right under the dam site, with some probability. Or it could occur at a certain distance from the dam site with some probability.

The likelihood of occurring right under the dam site is extremely remote. And the question here was: What type of probability should be assigned in deciding the distance from the site that one should assume the design earthquake? And in making this study which was handed out, which I believe Mr. Quittschreiber delivered to the members of the Committee, I took a look at what some of the probability factors that

mte 6

were used for seismic design throughout the country. And this table is a presentation of what those are, at least the ones that I found.

Just for reference, the 100 year flood shows a probability of annual occurrence at .01. A number of structures throughout the country are designed to handle the 100 year flood. But the probability of exceeding the 100 year flood is actually quite great over a 100 year period. And the probability of a 200 year flood is even greater.

The California legislature developed a criteria that said that any fault that has shown movement within the last 10,000 years is to be considered an active fault for purposes of locating residential structures.. In other words, if a fault had moved, theoretically, 10,001 years ago, then you could site your house on top of that fault. So the accepted probability there is an annual occurrence of .0001, and in a 100 year period it's probability of, say, reactivation, is .01.

The probability of not reactivation or not exceeding this in a 100 year period is .99.

The seismic risk map, which is now, I believe -- has to be considered a misnomer. We don't talk about the seismic risk, because we really should talk about risk as being the total picture. I would rather refer to that as the seismic exposure.

In any case, that shows a probability or gives the probability of a particular level of acceleration occurring without being exceeded, a 90 percent probability of not being exceeded in 50 years, which boils down to a once in 475 year event. And that probability of nonexceedence in a 100 year period is .81.

There was a study done, reported at, I think I said earlier, the conference that I was referring to. I think that report was in Pasadena. Actually, it was at Stanford. This report, the case history, is for an MCE; was reported two years ago at Pasadena. And this was an example used by an investigator, and his probability for the MCE was .00004, and that gives a probability of occurrence in a 100 year period of .0004, and nonexceedence in a 100 year period of .996.

The NRC criteria of an earthquake, an active fault, now called a capable fault, of one movement in 35,000 years, or multiple movements in 500,000 years, I have interpreted those to mean -- just for this illustration, that would mean that if a fault had occurred -- a fault had last moved 35,001 years ago, then it would be considered inactive. So anything that is an accepted probability for seismic loading.

And the occurrence of multiple movements, I said more than two in 500,000 years. That boils down to .000006 of that type of event being exceeded, and so that would give a nonexceedence in a 100 year period of .997 and .9994.

Those are just some examples of what could be interpreted as quantitative numbers for just the seismic risk -- I mean the seismic exposure.

DR. OKRENT: With regard to the NRC criteria, I think if you were to look at the return period for what is called the safe shutdown earthquake, you would have a higher probability of occurrence than the 100 year period by quite a bit. In other words, I think the numbers you are extracting from the criteria used for an inactive fault, a lower probability of occurrence than one gets for the design against seismic shaking.

MR. VON THUN: This is the only quantitative numbers that I had.

DR. OKRENT: I realize that. I'm just mentioning this in passing.

MR. VON THUN: You're saying that the number would be like less?

DR. OKRENT: The exposure to shaking is larger than this by quite a bit.

MR. VON THUN: But there isn't anything quantitative.

DR. OKRENT: Not in the criteria. You have to evaluate it on a site by site basis.

MR. VON THUN: Yes. It would absolutely have to be, because there isn't anyplace, I think, that you can put it aside where you would not have some shaking. But as far as

mte 9

deciding whether the fault is inactive or active-- but the Bureau criteria which we used for the Auburn site was movement once in 100,000 years was active. Anything greater than that was inactive. And that gives you the number .999.

So what we said in this study was that, as far as an order of magnitude estimate, that it was certainly conservative enough. And so, in making a determination, where what really was involved here is that, here is our site, there are random earthquakes occurring, and we had to come up with a distance from the site where we would say a design earthquake would occur or a maximum credible earthquake would occur.

And so, if this distance was based on the probability of that number right there, if there was a .999 chance that there would be no earthquake within this zone during a 100 year period, then that distance was specified. And in this case, it happened to be something like, I think, 22 miles. There were very few earthquakes around the area, but that gave a quantitative way of assessing where we would place the design earthquake.

On a completely arbitrary basis, we would have had to say, say if we had ignored the risk assessment, we would have had to say that the earthquake could occur right under the dam. So this is the only attempt that we have made to do a quantitative number like what you are talking about. And I really think that it is great that people are struggling with putting quantitative numbers on just to find out whether it can or cannot be done.

It would be nice to do it and nice to be able to tell the public, this is what we're thinking of in terms of a total risk. Whether we can actually do it or not is another question.

MR. RUBINSTEIN: Was there any physical basis associated with that distance?

DR. OKRENT: Would you please give your name?

MR. RUBINSTEIN: David Rubinstein, NRC.

Was there any physical basis?

MR. VON THUN: The actual faults are handled as actual faults. In other words, we have more than one earthquake for which we would design if we had an actual fault, say, located here, that we knew the distance to, then that would be used in the analysis, as well as the random earthquake.

But in this case there was no structure. This distance was not based on any structure. It was just based on a seismotectonic zone.

MR. HARBER: Gerry Harber, NRC.

Is that not based on a continuation of the relationship?

MR. VON THUN: No, not at this point. A continuation relationship would then -- would occur here. Once you decided

mte 11

where the earthquake was, then you would say: All right, we will consider a magnitude six or seven or whatever the earthquake is at 22 miles. And there would be a certain amount of attenuation to the site.

But the differential in acceleration was not considered. We tend to specify the earthquake distance, focal depth, and then look at what its effect on the dam is. In most critical facilities, I think that is the way that most of the people that I am familiar with -- the geologic assessment is made first of when and where the earthquakes will occur, and then we take that earthquake and specifically work with it, rather than developing what might be considered isoacceleration maps and working with just an acceleration.

MR. HARBER: But your circle there is based on the seismicity. How did you get the value of the radius of the circle? Based on the seismicity of that area around the earthquake within that area?

MR. VON THUN: Yes. In other words, you take a large area like this, and here's your site here, and throughout this large area you say that there is a random distribution of earthquakes, with some distribution, like I showed earlier, for a certain probability.

If you take one little spot on that within this large area, that little spot has a certain probability of there being an earthquake of certain magnitude in that area

mte 12

out of the total area. So if you think about now, what we're really saying is that within the site area there are a number of these little unit areas that have a certain probability of producing an earthquake of a certain level. When you sum those and they equal or exceed the .999 probability of nonexceedence, then you define the area of limitation, where you say that there is .999 probability that there won't be an earthquake within this zone.

MR. HARBER: But what happens if the isoseismal area of maximum intensity is larger in diameter than your circle?

MR. VON THUN: I would say if that were the situation, it would have to be handled in a different manner. That would move you clear out, and you would say there wouldn't be any earthquake considered at all. I don't think that can happen.

DR. OKRENT: I'm going to have to interrupt this discussion, because we are about 11 minutes behind the agenda. It's not your fault.

Thank you very much. It was a very interesting presentation.

Why don't we take ten minutes and then resume.

e-5 (Brief recess.)

MC GEE

DR. OKRENT: Our next speaker will be Michael McGee, who I am advised is an environmental scientist with the Office of Environmental Quality at HUD. He has a background in geophysics and geology as well as in public health, and worked with the Corps of Engineers before joining HUD. Mr. McGee?

MR. MC GEE: Thank you. My name is Mike McGee. As was said, I am with the Department of Housing and Urban Development. I am also a member of the subcommittee that Larry has referred to on the earthquake hazard assessment. I would heartily concur with many of the observations which he has made on his specialty, and in particular emphasize the need for recognizing the distinctness of the terms "hazard" versus the term "risk," and how many people interchange the two.

I think the question most often asked of me or of people in my office is, namely, what the Sam Hill Housing and Urban Development is doing in the process of risk assessment, hazard analysis, because very few people think of the Department as having any interest at all in such. I have been in this particular specialty area now for about 2-1/2 years, and the Department for the past five years. We have been concerned about the aspect of hazardous materials from the standpoint of community safety, the safety of HUD-supported projects, whether they be housing projects or

kapHEE

community development projects, community growth.

Our basic issue is land use and the incompatibility thereof with regards to hazardous materials operations of facilities versus dense popluation centers and how community growth is being tailored. We have found, from our experience over the past, let's say five years, that the definition of risk assessment or namely answering the question of what is an acceptable risk, is really a problem in human nature. And it seems to be directly proportional to the public's awareness of that particular hazard.

The simplest case where the public is largely ignorant of a danger, they will accept almost any definition of what an acceptable risk is. Where the public's awareness of that hazard may be, let's say, equal to the actual extent of the hazard as known by scientists, engineers, given the state of the art limitations, the public is amenable, reasonably amenable, to some sort of reasonable definition of acceptable risk.

The most profound situation seems to be where the public's awareness is, shall we say, magnified of what the actual hazard or danger seems to be, and here the definition of an acceptable risk becomes a very sensitive issue, and very site-specific for that portion of the public's awareness.

We have been tending toward an education type approach, or an answer to try and address this problem, the idea being to make the public aware of all pertinent facts of the problem, trying to assure them that whoever the decision-maker is defining that acceptable risk, he has all essential elements well in hand.

The case history, if you will, which made us aware that the problem existed, was in South Carolina. Columbia, South Carolina, to be specific. It was a bulk storage facility of a utility company which contained 40 60,000-gallon propane tanks, arranged in two rows. These tanks were of such a construction that they were horizontal, much like an elongated frankfurter, in their design, and such tanks, when they are involved in an incident, tend to rupture quite violently, with the ends proceeding almost like a missile and in roughly parallel alignment to the long axis of the tank.

In this case, the proposed housing project was located right directly near the ends of those tanks, the nearest house being 82 feet away from a 40,000-gallon -- or a 60,000-gallon container of propane. And the community had absolutely no idea why any of the engineers in the Department were somewhat concerned about the residents of the project which might move into that housing.

Our engineering analysis, which we had to proceed upon — because, after all, the Department hadn't been aware of the problem, didn't have any criteria to apply, was that an explosion involving one tank would have wiped out 60 percent of the residents of that total project, and we were talking about a project involving 1100 units of public housing.

The significance of Columbia, South Carolina, then, began to get us thinking as to how to define the problem, what sort of criteria we might need to prevent this type of thing from occurring again. And we began to question it: well, don't the local communities have codes or safety conditions? Because, after all, in recent years there has been a public trend away from federal regulation, an over excess amount of red tape, forms, et cetera, just for the sake of regulation.

We found that there wasn't, in fact, any, that the local communities didn't have safety standards, weren't aware of the problem as we had not been aware of the problem, or did not enforce whatever codes they did have. So we then recognized a need for safety provisions at a community level regarding hazardous materials facilities.

The principle involved two fundamental concepts which we defined. One is a safety separation distance. This is no 100 percent guarantee that -- it would be defined by an environmental safety standard giving an acceptable degree of risk from some adverse effect. In crude English, putting enough distance between that hazardous material container and, let's say, a public housing site or some occupied community facility, so that an incident at the hazardous material site would not adversely impact or seriously impact the residents or occupants of that facility.

The other fundamental point was to define a recognized danger zone. This danger zone would be that area physically impacted by that hazardous material, when an incident occurred with whatever conditions were operative at the time the incident occurred. And our principle was that safety separation distance, which we would prescribe by some means, would at least equal the worst case danger zone from that hazardous material container or installation or facility.

We then funded research, awarded a contract, and I have the several results garnered from that contract. One was a departmental guidebook with procedures to analyze hazards, namely, hazards from fire, hazards from blast or explosion and hazards from toxic substances, and some environmental safety standards. These environmental safety standards addressed thermal radiation and blast overpressure.

The safety standards which resulted were, one, designed to achieve safety for buildings, occupied buildings. The precise safety standard was 10,000 BTUs per square foot per hour of time interval, and this was to prevent self-ignition of the structure which might be in close proximity to an intense source of thermal radiation for a period of time.

The thermal radiation standard for the protection of people was 450 BTUs per square foot over an hour time interval, and it was designed to prevent the occurrence of a second degree skin burn on a person who might be in close proximity to that intense source of thermal radiation.

The blast overpressure standard was 0.5 psi and it was designed for the safety of the building structure to prevent the failure of that structure or the failure of a major component thereof, which might adversely impact anybody inside that building, from the blast energy.

Now, a little bit about the nature of the exposure. Now, note, I have been talking in terms of risk, hazard, but I haven't really talked about exposure. With respect to duration, we were saying the thermal radiation mgiht be expected to last for as long as five minutes in that particular building. The thermal radiation might be expected to last as long as two minutes, to a person caught in an exposed area at the time of an incident. And of course, with respect to blast overpressure, the exposure would be instantaneous.

The assumptions were predicated that the local fire department could respond within a certain set period of time and act effectively to reduce the amount of thermal radiation incident upon that building's surface. The ability of a person to react and take protective cover predicated the definition in two minutes.

Now, a lot of people would say that if somebody was threatened by a fire, it would certainly take drastic action long before two minutes occurred. Well, we as engineers, we were looking for a safety margin, and with respect to the statutory dictate of housing and urban development, we were looking for rather large safety margins in our result, for this reason, and that is the nature of the exposed population at whatever time that incident might occur.

HUD had to be concerned about low and moderate income families. These families, it was felt, would have a high probability of large numbers of small children. A small child in a crib would not be able to take the same action as, say, one of your average human beings in the sample population, your normal statistical average, i.e. the child would not be able to run away effectively.

Point number two: HUD was concerned with elderly citizens, people with restricted mobility. So once again, a 10 or 15-second reaction time to a disaster might not be appropriate, especially if the project was designed for elderly citizens.

And point number three: handicapped citizens might actually necessitate a rescue or physical removal by someone else in order to ensure that they reached safety from such an incident.

So these were the concerns that we had when we built in our safety margins, and asked the research consultant to come up with environmental standards. The present status of our regulation now, or our policy now, if you will, is that we have produced an advance notice of proposed rulemaking which was published on the 10th of September in the Federal Register.

We are presently finalizing a departmental regulation addressing the hazards of thermal radiation and blast overpressure, and we are taking the position of attempting to accurately recognize hazard scenarios in the worst case and attempting to come as close as we can in practical terms to achieving a zero risk. By that, I mean if somebody defines a potential danger, no matter how safely that tank might be designed, no matter how safe you might call for certain housing design elements or blast wall protection or some shielding, there is always the human element involved, and we are talking about more sensitive portions, if you will, of the total national population --

kapHEE

namely, elderly citizens and handicapped people and so on.

The duration of exposure and the exposed population caused us to get into risk assessment, which we have emphasized analyzing the hazard to achieve a low level assessment of risk. As a case in point, fire analysis, for us, we found was basically asking the question of: what is the worst fire hazard for a facility?

Some of the facilities we have come in contact with, chemical plants, large refineries, various sources which have had a large array of chemicals and/or fuels in a large number of various designed containers. Our answer, which we found, was you basically had to ask: what is the most flammable material present on that installation? What is the largest container containing a dangerously flammable material on that installation, and finally, what is the nearest container to your proposed housing site?

The largest safety separation distance which would be calculated with resepct to that specific facility ultimately came from one of those three which had previously asked the questions, and this was the one which would be employed with respect to the whole facility, for that site or that project which was being evaluated as to its suitability for housing or community development or what have you.

Our search through records for data and

1573 086

information pertinent to aid us in this -- we have found that there is almost no information which is directly applicable. We had to take into account a worst fire hazard, and nobody really asked people the question of where they were standing whenever they received, like, a second degree skin burn from an incident, or how many buildings were involved, radius-wise, from the source of the fire.

So we were immediately faced with a limitation, a state of the art limitation on the amount of data that we had.

We found we also had to take into account additive factors, namely, what conditions, particularly with regard to weather or just the human situation, might be additive, might make a disasterous incident even worse yet -- environmentally speaking, a worst case condition. Such additive factors with regard to fire would be high wind conditions, prolonged drought conditions and a delayed response by the community or the local fire department to enact safety measures and appropriate protective measures.

Interestingly enough, we did identify one actual occurrence of such an event. It occurred in Chelsea, Massachusetts, where approximately one-quarter to one-third of the entire urban community burned down and the same hazard scenario, interestingly enough, repeated itself within a 70-year time span, namely, an industrial source, actually a rag shop, caught fire during high gusting wind conditions. The community, used to fires in the industrial sector, attempted a delayed response. The water conditions or the fire-fighting resources were at a low ebb because of prolonged drought conditions and the same script repeated itself, I think in 1908 and again in 1973.

The third incident which hit the community occurred one year later, in 1974. So what we have found is, in attempting risk assessment or hazard analysis on these type of disasters, you have to be confronted with a very large absence of pertinent data that you might want. And we have had to go the route of research to answer a lot of the questions which we have had pressing us.

We have had very little tangible that has come forward and we have tried to expose our problem and our need for information in as many different directions as we could go. This has been educational not only in the aspect of us to the public, but in many cases of ourselves in attempting to refine hazard analysis.

With respect to explosion analysis, our approach was the same. What is the worst case hazard? Are there any additive factors which might make that a worst worst condition? And I think the most spectacular explosion, short of a nuclear weapon, is the BLEVE, the Boiling Liquid Expanding Vapor Exposion. Most people think of the incident which occurred in Crescent City, Illinois. If you are familiar with the National Safety Transportation Board, they have some spectacular photos of a fireball which is coming close to engulfing a small midwestern community, and it involved a relatively small container, 30,000-gallon liquid propane tanker.

The spectacular point about it was it is one of the only photos available of its kind. The photographer happened to be right there at the right instant, and he said that he had terrific problems trying to hold down the light exposure in order that he wouldn't lose the photograph.

These two hazards, thermal radiation and blast overpressure, we treat in our proposed regulation. Now, we say blast overpressure because we do not get into fragmentation or missile effects, which we have become aware are associated with catastrophic explosions of fuel or chemical facilities.

These, the statistical probabilities that would be associated with trying to pin down acceptable risk and an appropriate safety separation distance, are literally too large to try and encompass with information as it exists right now or as we are aware of it right now. I can give you one example later of such an improbable incident.

The third hazard which we have chosen not to treat at this time, but to subject to future research, is that of toxic substances. Many of these have hit the headlines from time to time in the newspapers and usually involve mobile sources.

Now, once again, this has brought us to an interesting dilemma. By the nature of our responsibilities, we are treating stationary sources only. This is as opposed to most of what you hear, involving a railroad tank car or a truck tanker. These are the purview of the Department of Transportation, and as such we are limited in what standards we can apply to railroad rolling stock or truck traffic.

The policy limitation which would hit us were we to treat mobile sources would be that we would have to restrict or apply safety separation distances and restrict housing from major highways and the rail lines, which have a large occurrence of chemical tanker traffic.

Toxic suostar.es. These we have found are associated with a large number of elements of uncertainty: wind direction, wind speed, surface wind flow, like vallye wind flow, atmospheric stability, and just the question of an acute toxi limit of threshold brings into play all sorts of questions about exotic research regarding human health effects, what would be a safe dosage for a sudden, intense burst of a chlorine gas cloud or an ammonia gas cloud. That in itself would be a fertile field for research.

We have found two limits or two constraints, even attempting to refine ourselves to the more concrete aspects of the problem, basically fire and blast. One is the flashback, where you have a flammable vapor cloud; and the other is combustion products. Here what I'm talking about, if you had a chemical container, like perhaps vinyl chloride, which became involved in a fire, one of the considerations would be that vinyl chloride is considered a hazardous vapor. And yet, a combustion product, if you submit vinyl chloride to fire, is phosgene, which is a military poison gas and something altogether different, as opposed to defining a safe exposure level or an acceptable risk.

As I say, we have only had about five years experience at this. We are right willing to admit we are still learning the trade, if you will. The more we compare notes with other people, the more we attempt to educate the public as to different aspects of the hazard, the more we become aware of certain factors that we have to consider and certain elements which have to be incorporated in any risk assessment technique.

Our other areas of interest have included the formation of an environmental hazards task force at the secretarial level, which is attempting to address these issues at a policy level for our department; chemical landfills as a separate issue, and this being necessitated by the results coming out of the Love Canal incident; radiation, namely indoor exposure to radiation over reclaimed phosphate areas, and radiation exposure due to either mining tailings or radioactive impurities in construction materials.

All these have been hazardous aspects that we have had to address or we have found a need to address in terms of community safety and public housing. The constraints that we now have have either surfaced as a result of people calling to attention various cases and asking for an analysis from our field offices or from our assembling of a case history file of incidents.

And one that struck me as intriguing relates to the question about why we are treating blast overpressure and yet ignoring considerations about fragmentation or missile effects. This involved an incident of a refueling facility in Port Newark, New Jersey, and one of the results was that a section of a propane tank was propelled into the air as it ruptured from a series of explosions at the facility, and upon returning to earth it penetrated the ground and ruptured the underground water main which was supplying water for the fire fighters who were fighting that particular fire.

So an assumption that we might have made as to the duration of such a fire at a certain facility would have been shot to pieces by such an unusual occurrence such as this.

As I say, we've only been at it five years and we're still learning the trade ourselves.

Thank you.

DR. OKRENT: Dr. Lave?

DR. LAVE: I'm curious about two things. First of all, how is it that you protect severely handicapped people in the event of a fire, especially if it's a multi-story building and you're not on the bottom floor?

MR. MC GEE: We have had a case occurrence in Pennsylvania where they asked for an addition to an existing elderly facility. The hazard was two propane tanks which were located about 50 feet away. The community was totally in support of the project. And our answer, once somebody arrived on the scene and negotiated with the officials concerned, we took advantage of natural shielding of the land and we specifically asked that a blast wall be put into the interim space between the wing which was to be added and the propane tanks.

The thermal radiation portion of the threat to the elderly residents, we asked for optical shielding. They were going to put porches on the one side of the building. We haven't got the design parameters, because our limitation there is we're in the Office of Environmental Quality, so we have to interact with the architectural and design standards in housing within our own establishment.

We are in the process of negotiating such considerations.

DR. LAVE: Is it fair to interpret your answer that your answer is you don't expose them to the risk, and that's your way of protecting them?

MR. MC GEE: We try not to, yes, sir.

MR. LAVE: A second question is why is it that you're using the maximum credible accident, somehow defined, as a design criteria? Here I would have thought that there is no justification for doing that, that you would want to design for the expected event and not the maximum credible event.

MR. MC GEE: This came out of our research work with the consultant firm. For me to go beyond that, I would be speculating.

DR. LAVE: Okay. Then let me just say that I think there is absolutely not justification for using that criteria, rather than using the entire distribution and trying to take a look at some sort of average risk, particularly where you're talking about cases of a small number of people being killed as a result of the maximum credible accident. I just don't see any reason whatsoever to use that. And the short-term consequences of using that are to overdesign structures, so that you have less adequate, less available public housing for people who need it, and therefore they suffer from having not as much housing as they need.

MR. MC GEE: Well, we have also tried to come to grips with the problem that the facility or the installation may grow at some future point in time. If I may, there seems to be two or three practical limitations which we have to consider, simply because of the current need for housing in the country, and the current need for more and more fuels being conveniently stored for energy reasons, and more and more chemicals being utilized by a fairly large number of industries.

If we were to assume this is the boundary between, let's say, an industrially zoned property with some need for hazardous materials, and our proposed project site, when we review the project we can't make things retroactive to things in the past, nor can we treat any change which might occur in the future.

So if some facility had a large container here, then our assessment, our distance which might be calculated -- we realize that the developers are prone to walk the fine line. If you tell them 999 feet away you may build structures and you can build those structures, ten structures to an acre, they will go 999 feet and build 10 or 10.1 to an acre, if they can do so.

The question was asked of us, instead of fuel A, some time in the future fuel B is put into this container, the implication being that fuel B carries a much larger separation distance than fuel A; what to do then?

Another point was, suppose a container at some point in the future is put in on the industrial property, carrying the same chemical or fuel in this container as was in this present container, the same separation distance. The population here would now be in jeopardy. Once again, we have no control in practical terms over such occurrences.

And it was also pointed out that community facilities being as they are, even though the separation distance might be here, the children of that community might just elect to play right here, in which case the safety design measure that we were trying to achieve would be totally aborted if an incident were to occur.

So we recognized some elements that, however well intended the policy might be, were just simply out of purview.

DR. LAVE: But I don't understand what conclusion you draw from that. You know, it's sort of like saying, if my grandfather -- my grandmother had wheels, she would be a trolleycar. If somebody were dumb enough to put a facility, a storage tank, right next to where your property was before, then there's no design criteria you could have used which would have protected you.

What else do you say after having said that?

MR. MC GEE: Well, all I'm trying to do there was to reconcile the maximum credible, as opposed to an average, occurrence, if you will.

MR. RICHARDSON: Jim Richardson, NRC.

There are certainly some societal benefits to be gained occasionally by taking a certain amount of risk. Has the HUD arrived at an agency policy on acceptable risk, what level of risk is acceptable to the public?

MR. MC GEE: Not at this time.

MR. RICHARDSON: Is there a goal to establish any such standard?

MR. MC GEE: We are looking to interacting with other agencies to try to define such, or to achieve something consistent with what other agencies may be arriving at on the basis of their research. But not at this time; there is nothing in existence. And we have this not as a specific goal, but as a general principle.

DR. OKRENT: Dr. Castenberg?

DR. CASTENBERG: I thought you said at the beginning of your talk that you were surprised that local communities, local regions, had not developed criteria of their own?

MR. MC GEE: We found no uniformity.

DR. CASTENBERG: Have you thought at all about withholding a project until that local government or local agency came up with some acceptable standards that they developed as a community, rather than you developing the criteria here in Washington?

MR. MC GEE: This again would have been before I joined the Department, and I can't really say what concerns were existing at the time of the research contract. Our experience started in hazard analysis and risk assessment after the research contract was turned back to us by the consultant, which was five years ago. So this would have been approximately seven or eight years ago.

I don't know, to be precise.

DR. OKRENT: With regard to the question by Mr. Richardson, a housing development poses other kinds of risks to its inhabitants besides those that arise from the storage of chemicals nearby. There can be fires that arise inside the building, crime inside the building. Earthquakes can cause damage to the building. And I'm sure we can think of one or two more.

Is there some effort to look to see whether there are important, let's say, gains in safety that could be made in many of these areas; how the risk from such things would compare to the kinds of things you have been talking about; whether there is a more effective way to spend the money at some point; and also, what constitutes an acceptable level of risk in regard to things of this sort?

I have to assume if you spent more money in a housing development, you could improve its fire resistance against internally caused fires, and possibly reduce the incidence of crime which threatens one's personal life, not only his property.

MR. MC GEE: We are assessing other environmental hazards, things like what is an acceptable level of air pollution, what is the effect of noise on a community or the wellbeing of the community, what are the effects of ground water contamination on a community, particularly communities out West where ground water as a source may be important for their livelihood.

We are addressing socioeconomic factors in appropriate lighting, whether people may get hurt from the lack of such; traffic accessibility. There is an effort being made to address all environmental hazards, whether they be physical environmental hazards or socioeconomic hazards, by the Department before each and every project is funded. We had some involvement in the Love Canal incident. Some properties were inherited back by the Department. Fortunately, someone told me that the person looking over the project application of Love Canal remembered that somewhere there had been a chemical facility, and they had disposed of some sort of chemicals way back when. And so he wrote in a precautionary statement in the projects being applied for that the Department should not consider anything except those homes which were sited on original land, i.e., no landfill type of sites.

And that thing in itself, I expect, saved the Department a good deal of embarrassment when Love Canal did surface when it did. And this was well before there was any such thing as environmental assessment or many of the terms that we have today.

We don't have risk assessment in precise quantitative terms. In the absence thereof, we tend to approach it as close as we can to zero risk, the standpoint being, if there is a potential danger there, then it becomes a question of when this hazard actually will occur, as opposed to if it will occur.

DR. OKRENT: Dr. Shinozuka?

DR. SHINOZUKA: The fact that you intend to provide the people with zero risk, isn't that the basis for your choosing maximum credible situations?

MR. MC GEE: In part, sir, in very large part.

DR. OKRENT: But isn't it a fiction that you are providing zero risk, since there are these other risks which in fact may be fairly substantial? You're trying to make it, I will call it, zero for purposes of discussion with regard to one specific area, let's say the storage of hazardous chemicals in the vicinity. But it is not a general goal that you are trying to achieve for all aspects.

MR. MC GEE: One constraint that we realized with regard to blast overpressure is in the standard as specified by the research consultant for what money that was available at that time. There is some probability that someone could, after all, be standing in front of a window, and although the building itself might be secure from the blast wave incident upon it, that person could just happen to be standing there looking out of the window when such a blast would occur, would most assuredly some adverse health effects.

Trying to pin precise numbers on what is the possibility of a person being impacted by shattered glass from a explosion which is, according to the industry, fairly negligible to start with, is very time-consuming to consultant and research firms, and very expensive and very iffy to pin down; and kind of like the example of what is the probability that a large section of propane tank will take out the water main which is being utilized to fight the fire.

There are too many "if's" in the problem, I guess, is what I'm saying.

DR. OKRENT: Well, we are a few minutes behind the agenda. Thank you very much for an interesting discussion. It is my first knowledge of some of the cases you mentioned, although I was aware that HUD had some kind of program of this sort.

I'm going to propose that we break for lunch and reconvene, according to the agenda, at 1:30.

e-7

(Whereupon, at 12:35 p.m., the meeting was recessed, to reconvene at 1:30 p.m. the same day.)

1573 102

CR 8493
HEER
t-8 mte 1

AFTERNOON SESSION

(1:35 p.m.)

DR. OKRENT: We are going to have a shift in the printed agenda and the next speaker will be Mr. Snyder from American Nuclear Insurers.

MR. SNYDER: Thank you, Mr. Chairman.

My name is Phil Snyder. I am here -- I don't know if I can say representing the insurance industry, because that is a rather broad term to say anyone is representing. But by my experience in that area is why I'm here. I am a professional engineer, started life as a chemist, spent eight years with Reynold's Aluminum doing chemical and metallurgical research, and then out of that, for some strange reason that I've never quite figured out, I ended up in the insurance field and have been working there for the last ten years, doing exactly what we are discussing here today, and that is, assess risks, translate these things to underwriters so they can make financial decisions, and work with the reduction of risks.

Now, the topic, of course, is use of risk assessment methodology, goals, and the criteria used by the insurance company. The initial draft that my boss wouldn't allow, I just said, well, we don't use any, and said that was simple, which isn't quite true. And he pointed out to me all the shortcomings in my statement.

We do not as an industry sit down and perform

mte 2

esoteric mathematical studies of the risk, such as some of the things we have heard -- dam failure, aircraft crash, specific seismic probabilities at a specific location. But we do use historical facts concerning many events, and we use a branch of science called actuarial science, which is a little bit of mathematics and a little bit of waving a magic wand over a black pot, and put this result into our business.

Now, before we get too far along, I need to give you some brief idea of what insurance is. Insurance is not gambling. We don't sit down and play a poker game and shoot craps. Insurance is the transference of risk. The risk has to be known, identifiable, and already existing, before we in the insurance business can, for a financial fee, accept this risk, to take it off of someone else's shoulders.

In a poker game, two people sit down and actually create risk out of nothing, and that is a difference. And if you have something that is speculative or is in fact gambling, you'll find you cannot buy insurance on the outcome.

Now, back to this risk assessment. The simple way, of course, is for the people who are like Travellers', Mutual, Liberty, somebody that's in the homeowner's insurance, because they have got just the ideal situation. They are writing single family dwellings all across the United States, and you find that these are very easily categorized into similar things.

mte 3

Almost all of them, without doubt, have bathrooms, kitchens, bedrooms, living areas. A certain percentage have garages. There are only certain broad categories of construction. You have frame, you have masonry, you have masonry veneer, and so forth.

Values also tend to lump together. You will find that probably 99-1/2 percent of these will fit between $30,000 and $200,000 in value.

So you have a tremendous, broad, and relatively homogeneous population, and the actuaries can sit down with these numbers and they can predict, at least on a nationwide basis, right down to the seventh and eighth decimal place what the loss experience in residences will be for the coming year.

And the statistical base is so large that, say, something like the conflagration in San Francisco after the earthquake, which destroyed a very large segment of the city, that doesn't affect the statistics at all. That is just way down in the sixth and seventh decimal place.

Now, that is fine and dandy for us. Unfortunately, in the business I'm in, American Nuclear Insurers, policy-writers on such places as Three Mile Island, we are not working with a broad base of small-valued homogeneous items. We have specific large-value nuclear power plants, fuel fabricating facilities, research institutions, and so forth around the country.

And for those of you that do play with statistics on occasion, there is nothing worse than the guy come running in and say: Look at here, we've got a population of three and an experience of X; now what's the frequency going to be over the next ten years? Then he doesn't like the answer he gets.

So what do we do at American Nuclear Insurers, at other insurance companies in this type business? I mean, let's face it, General Motors has insurance, Ford Motor Company, say Celanese Chemical, large papermills. They all carry insurance. They're industries. Say U.S. Steel; you don't have a large, homogeneous population of steel mills, either.

And the things you do statistically just don't apply. Yet we can profitably assess these risks and provide insurance for this type facility. And the way we do it is to sit down and hire a good engineering staff, which operates at many levels.

You start out with actually having engineers in the field. Usually these are fairly young people. Most of them are either fresh out of college or -- often we find that they are just either retired or have come off a hitch of, say, Navy duty. In mechanical areas, a machinist's mate makes a good starting point for certain types of engineering inspections. These report to supervisors with more experience and training, and right on up the line.

And pretty soon in any industry, be it automotive, steel, heavy chemical, primary metals production, you have a large body of experience of looking in detail at very small, discrete things, which funnels in like tributaries into a river. Finally, somewhere at the top, you can make industry-wide assessments for broad rating purposes, and you can make, to some degree, pretty good individual facility assessments that this place is going to blow up or it's not, or it is a good risk or it's not.

Let's get back to some of this. And I will do a little horn-tooting specifically in nuclear power plants. There is no experience. The first one didn't go into commercial operation until either the late 50s or early 60s. I forget when Dresden started up. We haven't had, thank goodness, a large history of big accidents to play with. So statistically we just have a broad range of confidence.

But what we in the insurance area find is two things: One, in any industry, in any area of insurable loss, even the homeowners, it is not the big total 100 percent loss of something that we look at, that costs us the most money and so forth. It is the incremental loss. It is when you burn out your kitchen because the grease caught the curtains on fire. That is not the total value of the dwelling. There are very few of those in the country each year. But you have an entirely different situation when you look at how many kitchen fires were there. And these, of course, are less than the policy value.

We carry the same thing in nuclear power plants. You find you have turbines, valves, motors, pipes, structures, electrical circuitry, control boards, pressure vessels, piping. And we, at least, say to some extent that the pipe carrying steam is a pipe carrying steam. An instrument cable is an instrument cable, whether it happens to be sensing the temperature of molten metal in a steel mill or whether it happens to be carrying a signal for a containment pressure sensor.

And we look then at this. It is an instrument cable. In the case of fire protection, which is my area of greatest experience, we know that cables burn. And the insurance industry knows that cables burn. In fact, in 1956 -- not '66 or '76, but in '56 -- the Factory Insurance Asociation, which is now called the Industrial Risk Insurers, published a book which covered cable fires, switch gear installation and proper installation and protection to prevent this.

Now, in the nuclear industry we sometimes tend to reinvent the wheel and things that may happen in a steel mill or a paper mill obviously can happen to us, and we sit down and write new regulations in a vaccum and they get enforced. Bright young design engineers sit down and design things out of existence.

At this point, as I usually do some place in every one of my talks, I will throw out Snyder's law again. Number one, it is that everything that can burn, will. Number two is, ignition sources are free. And, number three, Murphy was an optimist.

And the insurance organizations around the world that participate in writing nuclear facilities -- and this is a great, tightly knit fraternity that brings all of the insurance money available together, so that the plants in the United States, the plants in Germany, the plants in England, Switzerland, Sweden, are afforded the maximum insurance coverage that they can get.

We participate in all of them in the free world. But they look--the insurance business' greatest experience in large facilities is with fire insurance. They look at the nuclear industry worldwide. They determined that the risk from a large fire was too great. So, starting in 1972, there were several international meetings which resulted in the publication of international guidelines for fire protection in nuclear power plants.

This was issued prior to the Brown's Ferry fire, and we like to toot our horn again and say, if Brown's Ferry had been constructed and operated in accordance with those guidelines, they would not have had their disastrous fire.

We didn't insure Brown's Ferry and we didn't expect

Brown's Ferry. But I will go so far as to say that if we had, we would have prevented it, because you just never know. Again, Murphy was an optimist.

There are lots of other cases where our practices in the insurance industry for reducing risk, at least to our dollars and the policyholder dollars, just doesn't quite track with the regulatory atmosphere we see today. Piping is a good example.

We know that you have pipes in paper mills, chemical plants, fossil-fired coal plants. You name it, everything has got pipes. You pipe water, steam, what have you. We also know that these things are subject to failure. There are quite a few piping failures throughout the country each year. And our way to ensure the least risk to us is to insist on full compliance with the ASME Boiler and Pressure Vessel Code, and that goes for the design, the construction, and the maintenance and testing.

And if you closely read the federal regulations in this area in the nuclear business, you will find only equivalency is required. And we sit back and say: Well, we don't know what equivalent is. We don't have the staffs available to decree or study that something is equivalent. So we just fall back and say: Fine, you meet the ASME Code. When you have the certified stamp by all the authorized inspectors and the operating certificates are in line, we will insure it.

Of course, next year, when it's time to test it, you have to keep that up, too.

Now, back to the actual, I guess more in-depth of risk analysis. As I said, we don't really do a rigorous mathematical analysis of our risks, except in those areas that are subject to good statistical work, such as homeowners insurance, life insurance, where the risk is of early death and you have a tremendous population to calculate for.

When you write a billion dollar chemical facility and you are writing a policy that covers not only the actual physical value of property at that location, but also the continued operation of it, which we call business interruption insurance -- and say, a large refinery could purchase a policy that, in the event of loss by fire, storm, explosion, lightning, whatever, it is covered; that the insurance company will pay for their loss of profits until that production is restored.

And if you look at the production values of some of the refineries and chemical plants these days, you definitely get the attention of the underwriters, who immediately come over to the engineers and say: Hey, what's going on? Can I write this? Is it going to blow up tomorrow? And we always very carefully weasel it around and say: No, it is not likely to.

But he is the guy that signs the policy, but he always has a memo from the engineering side that told him it was okay.

But we divide our loss studies in this high-risk, high-severity, low-frequency situation, into sort of three levels. We have what is called a probable loss. This is the case where, back to Snyder's law, where it says ignition sources are free. We assume some accident occurs. But then we assume a reasonably adequate performance by the plant personnel, by designed-in, built-in safety features and so forth. These are your normal, nonconsequential industrial accidents: a gas pipe fitting, a furnace breaks, the low pressure sensor senses this and the valve goes shut. That is an accident. Everything works like it should. And away we go.

We basically don't insure these. These are handled by policy deductibles. We sort of consider them just normal occurrences.

The second level we look at we call a maximum probable loss, which goes a little further and assumes some adverse performance, either in the case of the operators or on installed safety equipment or other such things. The Brown's Ferry fire and the TMI incident we put in this category of maximum probable loss.

And since we are in high limit, infrequent but high severity incidents, we have a third category which we call maximum catastrophic loss, which anticipates a serious failure of equipment and adverse performance of personnel.

And we also anticipate that the loss will progress beyond those areas where you would normally expect it to bound itself by some passive means.

To get these concepts across, we usually use the example of a turbine. In the first case, that is a probable loss, we would assume that the turbine froze a blade, and this is held within the casing. The vibration sensors which they have sense the imbalance and the machine is safely shut down. As I said, we have deductibles and exclusions for that type thing. It is a normal operating occurrence. There is no problem and no insurance collection by the client.

At the second level -- in fact, we actually study things like this to get estimates of the dollars involved -- we would make the assumption that this blade or several blades actually penetrate the casing, do it in the worst possible area, sever a lube oil line, and a fire starts. But to fit in this category, we would then assume that the operators took the correct actions to begin bringing the turbine down to standstill, lubrication was maintained, and the fire protection systems did activate and control the fire. These are the type things that start getting headlines to be written up in the paper. If it is a nuclear plant, it will make national news.

The third case, though, we also look at -- and very few people seem to go that far -- as we assume, without attaching any probabilities yet, that everything that we have talked about before happens. The blade goes up. It ruptures the oil line and the fire starts But in this case the fire, we would say, might progress such as it burns out the wiring which powers the lubrication oil pump for the turbine.

We heard about the propane tank breaking the water main. We could make some assumption that the fire protection system was impaired for one reason or another. The turbine is sitting there burning, attempting to coast down without lubrication. We seize a shaft, break the unit in two or three pieces. The missiles progress out in sort of predictable paths, hit a transformer in the switchyard, and that causes an electrical failure here, there, and yonder.

And these things do happen, fortunately not very frequently, and we do look at them. And a fair amount of our engineering work is in trying to prevent this, such as relocation of equipment, redundant lubrication systems, protective control wiring, this type of thing.

We use, as I previously mentioned, a large staff -- "we" means the industry -- a well-trained and experienced inspectors and engineers, and we inspect. And we consider there's a big difference between inspect and audit. We inspect all of the facilities which we insure on a rather detailed basis.

In just the mechanical equipment area, our inspectors will average about 45 days per year per unit in a nuclear power plant. In doing this, he is physically checking motors, valves, piping, hangers, you name it, for their degree, how well are they maintained. He spends very little time, relatively speaking, reviewing compliance with maintenance criteria. Instead, he goes down and looks at the pump, or he looks at the motor.

And it is surprising how many times we will get a report in that the motor is fine, the pump is fine, but it's got no lubrication oil in it. And when we check, we find that the documentation is sufficient to show that this is properly maintained, but the end result doesn't agree with that.

e-8

1573 115

We use these engineers, and we use them subjectively and objectively. The subjective area is, in my opinion, the most important, and this is their personal opinion of the plant management, and these come from various different types of inspectors.

We have inspectors that are experienced in the fire Protection areas. One of the first things we ask of these inspectors is, what is your opinion of the management? Are they cooperative, and on and on and on? We send in mechanical equipment inspectors. We ask them the same question: What is your opinion of the management? Is it good?

We will send in nuclear specialists, and this is where we, of course, get into most of the auditing procedures. But still one of the first questions we ask them is: What is your, Mr. Inspector, your personal opinion of management? And this subjective opinion, we feel, is the most important item in assessing the risk at a particular facility. And although it may be extremely difficult and impossible to quantify, it is our feeling that if you have poor management, you have a definite increase in risk.

The next areas we look at, of course, get into the objective area where you can define and quantify such as we look at specific plant protection. Does it mean the appropriate building codes? Does it comply with the

mgcHEE

appropriate ASME codes? How many fire protection systems are there? Of what quality are they? Are they build in accordance with the applicable codes? What are your actual demonstrated results of maintenance procedures? What type of protection do you have against flood, tornado, earthquake, lightning, and so forth? These, we look at and can verify and do use in actually developing insurance rates and so forth.

A third gross overall area that we look at is in that area of specific hazards, such as rad waste handling systems. You've got liquid, solid, gaseous, and so forth. There are all sorts of ways to compress, compact, concentrate, transport — you name it. We insure the results of this, either if the equipment breaks or the resultant contamination, so we have specialists who look at these specific hazards.

For instance, if you are concentrating waste and using a medium that has known hazards of combustibility, we look at how is this combustible medium handled? Is it safely stored? Are the purchase orders written in such a manner that if you think your order of methyl-ethyl ketone, you don't get methyl-ethyl ketone peroxide. That's not too likely to happen anywhere, but it is just an example.

For instance, we ask, because of the known hazards of plastics, that you use certain grades of plastic film for

mgcHEE

your covering protector and so forth. At this point, we have a specialist who's checking to see that they are, in fact, ordering and are, in fact receiving the appropriate material.

In this area, we issue bulletins. You will find — I think we copy the NRC as well, but we have bulletins ranging from everything from how we want your cooling towers constructed to proper design for off-gas systems at BWRs. And we are most of the time proud of the fact that we don't have to follow these rather slow federal regulatory processes and can do things in a hurry.

Of course, doing them that way, we're not always right, but at least we do something. As a for instance -- and we can go back to Brown's Ferry again — at the time of the Brown's Ferry fire, the nuclear insurance industry had identified that penetration seals, those cables through far areas, were a problem and that the combustibility of cables was a problem area.

We had started testing in several areas before the fire. In fact, our first cable fire tests were run several months before Brown's Ferry. And after the fire, we were able to get an area of penetration seals, a full-blown approval and testing program underway, and by fall -- that is between the fire being in March and either late September or early October — was the first full-scale industry test.

mgcHEE

We had our program down, documented the approval procedures on the way, and industry testing underway. We continued testing cables, and we issued our own guidelines for cable construction, cable installation, cable separation, even though we had a lot of argument with the IEEE and a few other organizations. And we did that, had our own guidelines out by early 1976 in a final form, and unless I have missed something drastic, the federal government still does not have definitive regulations in this area. They are still in draft. There are still hearings going on. And there are only beginning to be national concensus standards coming out in this area.

So we are proud of the fact that we can identify areas of hazard and areas of risk, and at least for our purposes of protecting our dollars, move promptly to do something about it.

Now, the purpose of this Subcommittee, as I understand it -- and I may be wrong -- is to review the NRC program on quantitative risk criteria and to possibly see if there are methods, et cetera, that could be used to reduce the risk to the public. And I say, fine. If we are looking at a broad population, a statistical probability of earthquake, that's the only way you can do it -- take a large population. You know where the faults are. You make assumptions, and so forth.

mgcHEE

But we are dealing with your licensing. We are insuring specific nuclear power plants. We don't write an insurance policy on a nuclear power generating industry of the United States and have this big, broad, wonderful average to work with. We write specific facilities, and the NRC licenses specific facilities, and it doesn't do us or you too much good to know that some event has an entire population probability of occurring only once every 10,000 years because no one stops to address the fact that that is fine, but does it happen to be once in ten years at this particular location, and once in a million at all the rest?

That's what we have to deal with in the insurance industry — is which one is the one that's going to get us, not what is the industry-wide average. It is sort of a case of figures don't lie, but liars figure.

If I can offer a suggestion — and I doubt very seriously if it can be implemented under the regulator systems -- but that suggestion is that to reduce individual risks which will have the net result of reducing overall risks, you concentrate on the management of the facilities.

And the example I will give of that is the Dupont Corporation. They have, or they are well recognized at least in my business of being Mr. Safety — they don't blow up plants. They don't burn them down. They don't injure employees — anything like the rest of American industry.

mgcHEE

Their personnel safety, fire and explosion safety, transportation safety records are so good that they honestly cannot afford to purchase insurance because the increased cost to the corporation would just be prohibitive, because no insurance company could afford to write those great of values, even knowing the low risk, for a premium that anywhere approaches their in-house cost for maintaining and operating a safety program.

And the key to the Dupont program and quite a few other chemical people in the chemical industry is management responsibility from the Chairman of the Board to the President of the company to the individual plant manager. There is a real, deep belief and dedication to doing it safely.

When a Dupont plant puts up a sign that says, "Safety first", they really mean it. They don't do it because the National Safety Council sent them a box of them. In fact, they take it so seriously that a Dupont plant manager who has too many accidents at his plant is in jeopardy of losing his job, even if he happens to be setting a productivity record for the whole corporation. They will not tolerate managers who don't run a safe plant.

And by having this personal responsibility, knowing it is my job, we get something done. They believe in it. The foreman on the line who is supervising the people, he knows that he will get fired when the plant manager gets fired. So he is interested in being safe.

I think one of the best examples of that in the federal field and something that the Commission should definitely look at is the outstanding lack of success of the OSHA program. This was forced onto American industry as the great, wonderful thing that will make the workplace safe for the worker, and if you look at the statistics, there has been no improvement in accidents whatsoever.

The plants that were safe operations before OSHA are still safe operations. The plants that weren't safe operations before OSHA still aren't safe operations. And there again is that management commitment.

OSHA requires compliance with regulations, and there are two ways to comply with regulations. You can go out and do only what you absolutely have to as a minimum to meet the letter of the regulations and the auditor checks you off, or you can really want a safe operation -- find out what the regulation means, where it came from, and go out and do it on purpose, make it better in some cases.

In our own area, you can see this when we interview and discuss operators. You have operators in these plants. He's got a license. He's a licensed operator; he's been through the training program. And there is one guy -- he wanted to be a plant operator. That was

mgcHEE

one of his big goals. He wants to be a good one. He studies extra time; he goes further than he has to; he learns what makes a plant tick. He spends time talking to other people in other departments -- what are you doing -- and he's a darned good operator, and yet he may have had the same score on the examination as the guy who only did what he absolutely had to.

And I think from your own experience, you will know that that second guy is a much poorer operator than the first one I discussed.

And with that, I will answer any questions that I can.

DR. LAVE: You made this interesting statement to begin with that there were small eventualities that wound up costing money. Is that only because of experience to date, or do you expect that that will be true in the future? If so, why do we need Price-Anderson?

MR. SNYDER: Well, I heard a similar question earlier today. We need Price-Anderson because the insurance industry does not have sufficient assets to provide all the coverage necessary to the public.

DR. LAVE: Come on, now. What kind of accident are you talking about?

MR. SNYDER: Well, what was the release at Three Mile Island? I mean, offsite it is a minor accident, yet

1573 123

mgcHEE

look at it from the insurance company's standpoint. For the next 20 years, you're going to be battling lawsuits -- possibly if the court climate is such, paying exorbitant claims. There is no limit to what the courts may allow in a very minor accident.

So we think — and we may be wrong, but the Congress somewhere made the decision that we do need nuclear power. Someone made the decision that the available private insurance wasn't enough to give the public a good feeling, so they added some on top.

DR. LAVE: Without getting to that, could I get back to the first part of my question? I understood you to say, it is the little things that have cost you money so far. And let me just ask again, do you expect that will continue to be true in the future, absent some really colossal occurrence?

MR. SNYDER: Okay. You may have misunderstood that. You're right. Let's take a hypothetical plant of any type that is worth $100 million, and we write an insurance policy for $100 million. The probability of paying that $100 million is just vanishingly remote. What we pay are the $10,0000; $30,000; $50,000 every month, two months, three months, type losses. And these add up over the entire history of an insurance organization to be a much greater sum than the individual, infrequent, what we call

mgcHEE

"shock loss."

DR. LAVE: Would that also be true, for example, for supertankers carrying oil?

MR. SNYDER: I don't know. I'm not in the marine insurance business. But depending upon how the policies were written, which might cover contamination to cargo, late delivery, this type thing, it very possibly could. If there was only a catastrophe policy written that we're only going to pay you when the thing finally sinks, then, no, that wouldn't be true.

DR. LAVE: That must be just a lovely kind of a case for an insurance company, if your claims are relatively small, much less than the face value of the policy -- lots of little tiny claims instead of a couple of really big ones.

MR. SNYDER: Well, no, it is not, because they won't buy the insurance if our premiums are exorbitant, so we have to sort of balance in there, if we are writing a $100 million facility, they aren't going to pay $50 million a year for insurance.

DR. WILSON: But I don't understand that one, because the premiums that a utility company pays now are such a small fraction of what the charges appear on my electricity bill, I mean that it's after all -- the reason they buy insurance is because they have to. If you multiply

mgcHEE

it by ten, it still won't be on my electricity bill.

MR. SNYDER: That's true.

DR. WILSON: So that answer, I think, is just wrong. It just can't be right.

MR. SNYDER: Let's back up a little. We are in the free market situation that, with the exception of what is required by Price-Anderson, the utilities can elect to buy or not to buy our insurance for damage to their property.

Now because Price-Anderson is mandated by law, and we are a monopoly, we feel somewhat bound to keep the rates as low as possible. In fact, if you look at the structure after ten years, we return all unexpended premiums to the policyholder. So what the utility paid in ten years ago, which wasn't used to pay for losses that resulted from operation during that year, is returned to him this year.

DR. OKRENT: Can you say whether you have differentials in rates for insuring different nuclear power plant properties because of the management?

MR. SNYDER: No, not for management.

DR. OKRENT: So all of the managements have turned out to be the same in your eyes from a dollar protection point of view?

MR. SNYDER: From a legal rating point of view, by being able to look at lots of things, if we have a plant that we do feel the management is extremely poor, we can find enough other things to complain about which are the result of that poor management to actually get an increase in rate. But it is not and cannot legally be the cause of our judgment of their management capabilities. It has to be based on physically verifiable, you-guys-don't-maintain-this-and-it-doesn't-work type situations.

So it is the same result.

DR. OKRENT: Dr. Castenberg?

DR. CASTENBERG: Early in your talk in regard to fires, you used the phrase: "The risk was too great, and therefore we went and looked for some guidelines." When you used the phrase "the risk was too great", was that a qualitative judgment or a quantitative judgment made -- that the risk was too great? How did you arrive at that phrase, or how did you arrive at the decision that the risk was too great?

MR. SNYDER: It is almost like a multiplex. We talked about rating. Obviously, if we could charge the full value of the plant, we wouldn't care what the risk was because we could never pay out more than that for the rates that we have to charge, and I'm talking general industry now. They pay maybe an average of six cents per year per hundred dollars of value of insurance, which is generally about one tenth of the average homeowner's insurance. So when we look at a facility and then we find historically here and there something new is cropping up like concentrated cable fires, we transpose that to whatever we're working with, and we say, worldwide we have experienced a certain number of fires of values, dollar values at a certain level, can we afford at our present rates to accept that risk of occurrence in our plants?

And when we find that we can't, because we are in competition with other organizations who aren't going to raise their rates because we do, we have to improve the protection and the design.

DR. CASTENBERG: So you have some threshhold limit in dollars paid out which is a criterion?

MR. SNYDER: Yes, but it is not fixed.

DR. CASTENBERG: I see.

MR. SNYDER: It floats.

DR. OKRENT: Dr. Wilson?

DR. WILSON: The thing I'm not clear on, in some other industries, certainly in for example gasoline tank trucks and things of that sort, the insurance industry sets standards. What is it? The Fire Protection Agency in Boston actually does it, and in fact they have got other standards for some of the industry even beyond that, and premiums for insurance are judged by the use of those standards.

So in a certain sense, the insurance industry did a fair amount of the regulation of the industry through this practice. Now I don't know of any example in the nuclear industry. Is there any, for example, example where you would set higher premiums for Three Mile Island-2 than Three Mile Island-1 or anything whatsoever in the nuclear industry where you were setting a criteria and force it on the industry by adjustment of the premiums? Or is there any contemplated?

MR. SNYDER: Yes. We do that for each individual facility. I am not an expert in the rating of the public liability insurance portion, but I know they have a published schedule that is public knowledge that they use to assess plants which cover such things as surrounding population density. And all of these factors go in, and different plants have different rates.

On the property side, the direct physical damage to the utility company property that I work with, we have a Rating Bureau which is an independent organization, Nuclear Insurance Rating Bureau, that sits, I guess, quarterly in New York City. We have to report to them a certain list of facts about the individual plants. What is the size? What is the value? What type of plant is it? Who built it? Where is it? What kind of public fire department do they have nearby, and on and on and on? Many reports. And they establish, based on a lot of things including general published insurance rates for specific hazards, and they come out with a rate. And that is what is charged to the plant.

1573 130

DR. WILSON: The second part of the question, is there any feedback, and if I take the Pilgrim plant near Boston, which happens to be a local one, it is a little closer to Boston than, for example, Maine Yankee up in the corner of Maine, and might, by your population density, therefore satisfy some different criteria. Is there any known case in a siting decision to a utility company, say, do they bill Pilgrim-2 or do they bill Maine Yankee 2? Is there a choice of these things influencing what the industry actually does?

MR. SNYDER: There should be, but the only time I've ever seen this effect occur is when we negotiated with the utility and said, what you propose is just not insurable. There will be no insurance available, period. Then the listen. Otherwise, as you previously mentioned, the cost of insurance is such a small part that it is generally ignored.

DR. OKRENT: Thank you, Mr. Snyder. This was an interesting talk. I appreciate your coming here. I understand that Mr. Ellett and Mr. Richardson are here now. So we can revert back to the planned timing, albeit a little delayed. Let me ask: would it intolerably inconvenience Dr. Page, for example, if he is 40 minutes behind where he is currently scheduled? Because we have slipped, is that tolerable?

DR. PAGE: That is okay.

DR. OKRENT: Then I would propose that we hear from either or both Mssrs. Ellett and Richardson from the Environmental Protection Agency.

MR. ELLETT: My understanding, you're going to have other people from the EPA this afternoon -- I think you will see similarities and differences between how each of these offices factor risk analysis into their decision-making process.

The Office of Radiation Programs uses a variety of considerations in setting radiation protection standards and guides. Risk is, of course, only one factor in the decision-making process leading to regulation. I believe you have copies of our policy statement of March 1975. It has been published in the Federal Register and several other places.

I would like to call your attention to the last paragraph in that statement, where it says the linear hypothesis by itself precludes the development of acceptable levels of risk based solely on health considerations. Therefore, in establishing radiation protection positions, the agency will weigh not only the health impact but also social, economic and other considerations associated with the activities addressed.

As is also stated in this policy statement, the

Office of Radiation Programs believes it has an obligation to make available to the public numerical estimates of the risks to human health from radioactive contaminants. The agency recognizes that these estimates are not precise and that they are truly estimates and not predictions of environmental impacts.

While such estimates are an important component of the standard-setting process, they are not a substitute for the current federal radiation protection guides, and I would like to recall the first federal Radiation Protection Council publication, where injunctions are given that any radiation exposure must be both useful and necessary in the first place, before there is any consideration of whether an allowable risk is all right, an exposure must have an expected benefit compared to the health risk imposed, and finally, that all exposures be as low as practicable.

Moreover, to quote the Federal Radiation Council directly, they say there can be no single permissible or acceptable level of exposure without regard to the reason for permitting the exposure. I think our office has backed away from the idea that there is some kind of magic number applicable to all situations involving radiation exposure. And I believe this is reflected in the standards we've set.

The Office of Radiation Programs has authored several radiation protection standards and guides. None of these have identical risk estimates, although in some cases the estimated risks are not too different.

However, the basis for selecting the limits varied quite strongly from one case to another. Now, one reason for this is the legislative mandate for various types of environmental regulations differ greatly depending upon what is said in the enabling legislation. For example, the Safe Drinking Water Act of 1974 specifies that the national interim drinking water regulation shall protect health to the extent feasible using technology and techniques that are generally available, taking cost into consideration.

In contrast, under the Atomic Energy Act, the administrator has the authority to establish such standards as he may deem necessary or desirable to protect health or minimize danger to life or property. There is no reference to cost.

I would like to outline how the agencies establish radiation standards under each of these differing authorities. First, I would like to talk about the Safe Drinking Water Act. Proposing limits for radioactive contaminants of drinking water, the agency was faced with two quite different situations. Man-made radioactivity usually enters the environment from controllable sources already subject to a high degree of governmental regulation. On the other hand, radioactivity is ubiquitous in ground water throughout the United States. Selecting the limits for man-made radioactivity correspond to a four milligram per year dose with an estimated risk of .4 to two deaths per year for each million persons exposed. That is, of course, exposed at that maximum limit, and we would expect most people not to be at the maximum limit.

This limit was not selected on the basis of a de minimus risk, but rather, to quote the Federal Register notice on these regulations, it was chosen on the basis of current levels of radioactivity in community water systems and to avoid undesirable future contamination of public water supplies. The limit for radium was set a at a concentration corresponding to an intake of 10 picacuries per day. Estimated risk at this level of radioactivity is .7 to three deaths per year for each million persons exposed.

Selection of this limit was based on the cost of obtaining lower concentrations and agency's uncertainty on the number of community water systems that would be impacted. It was not selected on the basis that this was an acceptable or de minimus risk.

Indeed, the risk of the drinking water maximum contaminant limits are probably higher than allowed for any other drinking water contaminants.

Turning to the uranium fuel cycle standard, which was established under the Atomic Energy Act, a somewhat different approach was used. The agency examined the cost of risk prevention for different parts of the uranium fuel cycle and strove for some comparability in terms of cost per cancer death averted.

Mr. Richardson was very instrumental in developing the uranium fuel cycle standard and will talk a little bit more about the detail of how we developed our uranium fuel cycle limits.

MR. RICHARDSON: I would like to apologize for not having any prepared statement for you. The rationale that was used for the uranium fuel cycle standards -- I think it's well known to most of the people here, and it is documented to the extent that we were able to write it down in the environmental impact statement that went along with that standard when it was promulgated a couple of years ago. And I am sure there are copies of that available. We're about out, so I didn't bring 20 for everybody.

As Dr. Ellett pointed out, the standards for the uranium fuel cycle had a heavy base in cost as well as risk, because we are permitted to do that by the Atomic Energy Act, which doesn't really place any restrictions on the methodology that you use for setting standards.

We looked at, I think for the first time, the question of long-term risks to large numbers of people

kapHEE

spread over large distances.

That is to say we calculated the collective dose commitments to the world's population in the case of isotopes like crypton, and those standards reflect the risks associated with those assessments as well as the risk associated with doses to people close to facilities. And that is probably the most significant difference between previous standards and those standards for the uranium fuel cycle.

The question of deciding at what point you had spent enough money to avoid risks was a central issue in the standards, and it wasn't handled very cleanly because you can't handle that kind of question very cleanly. What we did was rank all of the major control methods by cost and by the amount of risk that was avoided through their use. A cost effectiveness ranking -- and we found that it was convenient to divide all of those options, each of which had its own dose level attached to it, into three broad categories.

Quite arbitrarily, those that fell above a half million dollars or so per life of lethal cancer, those that fell less than $100,000 for each of those types of events, and then of course, a third class, which were those that fell in between. That ranking implies an evaluation of life that falls somewhere between $100,000 and a half

kapHEE

million dollars or so.

We purposely left it very vague. When we looked at the options that were available to us for setting standards based on that kind of a ranking, we discovered that there were very, very few levels of control that fell in the third or middle category and it became rather easy to say, Yes, you should do all those things that cost more than half a million dollars per life -- or you shouldn't do those things, excuse me -- and that you obviously should do those things that were down $100,000 or less, because those value judgments were, it appeared to us, generally acceptable areas where there wasn't much dispute about whether you should spend money at those levels.

There are always people who say you should spend infinite amounts of money, but they weren't very reliable. And there were very few people who said you shouldn't spend money when the cost of life is less than $100,000. We published our proposals and if you read them carefully enough -- we didn't make it terribly explicit, but if you read it carefully, it is all in here. And out of the very many comments that we received, not a single person, not one person questioned this implied judgment that life -- the cutoff between when you should spend money to save lives and when you shouldn't -- fell somewhere between $100,000 and half a million dollars.

DR. WILSON: Do you have that case where it is published?

MR. RICHARDSON: Yes, it's here. All of the comments are in Volume 2, and the discussion of the cost considerations is in Chapters 4 and 5 of Volume 1.

Well, that gave us some confidence that we were on the right track, and we went ahead and I think that is really all that I have to say about the question of criteria for establishing standards.

We used a number which was directly related to health effects, not to man rems or some other surrogate. We used a range of values rather than a single number, and nobody objected. I guess we are both open for questions.

DR. OKRENT: I will try one. Let me invent a hypothetical situation that, from the fuel cycle examination, you found the same kind of dispersion among fixes -- in other words, that there were some cases where it cost $100,000 or less to defer premature death, or it was over half a million dollars, but you estimated that the risk to the individual that resulted when you used the cost of life way of deciding what to implement -- where that risk was 10 to the minus three per year to the individual, what would you have done?

Would you have said it's okay, what we need to do is meet this value of life criterion? If that wasn't okay, if it wasn't -- would 10 to the minus four have been okay, would 10 to the minus five?

MR. ELLETT: It happened in the uranium fuel cycle standard, but this is a hypothetical question.

MR. RICHARDSON: It happens all the time. And I left out an important part of the discussion of those standards. You have to have two things to worry about when you establish standards, at least two things to protect people, and one of them is what is the collective impact and what is the maximum individual impact. And they are quite different and independent things.

There you are faced with -- well, it is another kind of a degree of arbitrariness. The range of evaluation for human life is certainly arbitrary, but you can go to society and say, Well, is it acceptable, what you've done? I guess the same thing is true of individual risk. In the case of the uranium fuel cycle, it was a fairly arbitrary decision.

We had situations at milling operations where some of the individual risks were fairly high, but the cost effectiveness of eliminating them did not fall in the range I just discussed. The situation for iodine emissions at reactors sometimes falls into the same category. And there, we got quite arbitrary. We looked at the level of risk that had been attained in the industry in most of the components

kapHEE

of the fuel cycle and we looked at the levels of risk that were attainable using cost effective methods, and they all fell at levels of dose which were below 25 millirems, and so we arbitrarily said, Well, for the sake of equity we will impose this kind of a limit.

Quite frankly, back in those days when we did this, we did not look to see what the lifetime risk was to an individual explicitly. It turns out now to be fairly high compared to some other levels of acceptability used in the agency for other standard-setting activities. 25 millirems per year is -- I've forgotten the number -- it's on the order of 10 to the minus three or four per lifetime of doses, certainly worse than 10 to the minus five or six, which is what the agency strives for.

And we have to depend upon arguments that say, Well, that 25 millirems is really not the limit at which reactors and other facilities operate. It is the standard under which the regulations for ALARA sit -- as the framework in average exposures, for even common exposures to individuals are generally even much smaller than that.

I think if we were setting the standard today we would not be able to set a standard that was expected to be approached often at that type of level. But to answer your question, how do you arrive at the acceptability of a number 10 to the minus four or five, six or seven, I have no

kapHEE

wisdom.

MR. LAVE: You made a statement that the agency strives for 10 to the minus five or 10 to the minus six per lifetime. Would you clarify that?

MR. RICHARDSON: Yes, the agency has recently published the cancer policy proposal.

DR. OKRENT: If I take that 10 to the minus five or 10 to the minus six per year in drinking water, for example, is that for everything in the drinking water, or is it for each item?

MR. ELLETT: Actually, it would be separate for man-made radioactivity and radium. That was just for administrative simplicity, we felt that anybody who was heavily impacted with man-made, for obvious reasons, surface water, so they would be subject to radium and vice versa. It isn't for each isotope; it is for all isotopes, sum.

DR. OKRENT: I thought that there was a goal that the agency had which was being applied more broadly than for radioactivity.

MR. ELLETT: That goal was articulated much later than the drinking water regulations. The drinking water regulation was the first regulation for radioactivity. All other maximum contaminant levels for drinking water were set on the basis of threshold that was obviously inapplicable to radioactivity. I think at the time they were established, the agency considered the risk limits rather high for radioactivity, but given the amount of radium contamination and fallout contamination, there wasn't much that anybody could do about it.

DR. OKRENT: No, but I am just trying to understand the answer to the question raised by Dr. Lave, which I thought said that there was a goal by EPA recently announced, that the lifetime risks from something -- I want to find out what the "something" is, should be. I don't know what you said. 10 to the minus five or 10 to the minus six.

MR. ELLETT: I think the EPA speakers that will come after me will be better qualified to speak on that than myself. I can only tell you what we have done in the Office of Radiation Programs.

The standards that I discuss with you were set before that policy was ever articulated by the agency. It doesn't become a retroactive truth.

DR. OKRENT: Somebody among the EPA speakers will be able to define the policy, and to what it applies? Is Dr. Page saying yes?

DR. PAGE: Dr. Anderson is here, too.

DR. ANDERSON: What was the question?

MR. ELLETT: I think it is a safe assumption.

gsnHEE

DR. OKRENT: Okay, I gather that it will automatically evolve.

DR. ELLETT: We will be here for further discussion.

DR. OKRENT: Are there other questions now for Drs. Ellett or Richardson?

(No response.)

DR. OKRENT: Thank you. I guess it will be useful to hear the succeeding speakers. I gather that Dr. Anderson is here, and she is also with the Environmental Protection Agency, and if I recall correctly, very interested in various aspects of toxic substances and so forth.

DR. ANDERSON: Are you ready for me?

DR. OKRENT: Yes, please.

(Pause.)

DR. ANDERSON: I apologize for walking in just at the last minute, so I don't know what I've missed. But I looked at the part outlined on the program for me and I represent the part of the PA that works in carcinogen risk assessment. And recently, we are enlarging the office to include the possibility of doing in informal ways risk assessments for other health effects.

Can you hear me, or should I put this on?

DR. OKRENT: I guess that it is preferred if you can.

DR. ANDERSON: I notice specifically you said I would discuss case histories in the use of risk assessment and risk criteria in this office which is being formed around the carcinogen assessment group.

And I would make suggestions in appropriate goals and risk criteria that might be used for nuclear power reactors.

The first order of business is something I feel comfortable with. The second order of business, I certainly will have to leave to this committee and to the Nuclear Regulatory Commission.

I think that some of our experiences might perhaps be helpful in forming the basis for further thinking on your part.

The risk assessment activities in EPA have largely been involved with carcinogens and radiation, at least in formal ways of doing risk assessment, and having formally laid out risk assessment documents, become a part of the regulatory process.

So it is this area that I will discuss and I hope as a policy matter can provide some basis for your particular problems.

In 1976, the EPA did announce a policy for assessing risks associated with carcinogens and you can probably recall that the EPA was the first agency to do this. There are four major regulatory agencies involved in carcinogen regulation, including OSHA, Consumer Products Safety Commission, and the Food and Drug Administration.

But I think it is important to put risk assessment in an appropriate social context when we think about the use of risk assessment because it is within this context that we have found that risk assessment serves a particularly useful purpose in some case and in other cases, has really been able to provide very little for the regulatory decision-maker.

So I would like to talk about three areas that have certainly affected recent risk assessment within EPA. The first is what I call social policy, which clearly is laid out by the Congress.

And I'm not going to get into details of regulatory authorities, but only to suggest the differences in measure where risk assessment may be plugged into regulating carcinogens.

The second is the approach that EPA has taken in a scientific way to lay out these risk assessments so that we have some uniform procedure agency-wide for assessing risk and presenting the information.

And finally, I'm going to give you on this overhead projector some examples of cases where the agency has used risk assessment as a part of the process and has actually taken an action which has either reached the final stages of regulation or is formally proposed.

In this area of social policy, suppose Congress decides right off the bat what is going to be regulated, how it shall be regulated, and who the regulator shall be?

Now there obviously are things that are not regulated, so risk assessment in those areas becomes an academic exercise, at least as far as regulators are concerned. Within the establishment that I spoke of, the Occupational Safety and Health Administration has an act which requires that it protect worker health to take feasibility into account.

So feasibility becomes the primary focal point for setting worker standards.

In the Food and Drug Administration, the Delaney clause, of course, is the clause that comes to mind. That is the only clause in the Federal Government that is an absolute ban clause. And clearly, the degree of risk assessment that is appropriate here or useful here is very much less than a risk benefit balancing situation, although there are other provisions of the Food and Drug Act which certainly do use risk benefit balancing approach such as in the drug provisions.

Also, in food contamination risk assessment has begun to play a role because of the importance of setting tolerance levels and looking at risks associated with inadvertent contaminants.

gshHEE

Within EPA, we have almost the full range of legal requirements which permit us to use risk assessment in a widely varying number of instances and in many different ways.

And it also leads the scientist in the agency, and that certainly involves my group, to do risk assessments in more or less depth, depending upon what the practical framework is.

So EPA has seven major regulatory authorities under which we regulate carcinogens. And I'm using carcinogens, of course, as example because this is the one area where the agency has clearly stated a policy for using risk assessments and is doing risk assessments in a very consistent fashion.

These seven areas are covered under the Clean Air Act, the Water Act, the Drinking Water, Pesticides, Toxic Substances, Solid Waste and Radiation.

And you've already heard about radiation, so I don't have to talk about that. Nowhere in our provisions do we have an absolute ban clause, but we do have a number of different requirements in the FWPCA, or the Federal Water Pollution Control Act.

We have provisions or sections of the Act which are technology-based alone. Therefore, risk assessment can play very little role in actually regulating under that provision

gshHEE

where we have a provision which requires that national water quality standards based on health alone, risk assessment is almost the only compelling force in setting that standard.

And I'll show you some examples of that.

Under the Clean Air Act, we do have a provision which appears to be primarily health-based. So, again, risk assessment becomes a very important tool.

And pesticides and toxic substances require overt risk benefit balancing. And clearly, here is an opportunity for the full use of whatever risk assessment can provide in a risk benefit balancing approach.

So EPA now has this broad social authority or legal authority to regulate substances which are being subjected to risk assessment.

So the question then is how are we conducting these risk assessments so that we are able to present some coherent picture to the people making decisions under these various acts.

EPA has adopted a process for doing this and insofar as possible, the agency has decided that regulation will take place in a two-step process to the extent that the statutes permit this.

The first step will be a step that lays out fully, as fully as possible, the risk assessment part of the regulatory consideration.

gshHEE

The second step will lay out, insofar as is possible, the social and economic impacts. The two will be considered together as is appropriate under the Act.

Doing the first part, which is our part today, the risk assessment part, the agency decided that it would adopt a consistent approach. And I think, as with any risk, the idea is to define, first of all, how likely the risk is to occur.

With carcinogens, this is the qualitative part of the risk assessment and we have decided to take a weight of evidence approach, recognizing that only rarely do we know for certain that something is a human carcinogen.

We express this qualitative part of the assessment in terms of the certainty of evidence, the weight of evidence, or the strength of the signal, taking everything into account.

The second part of the risk assessment is quite a separate part of the process and that is on the assumption that the risk exist, just how bad is it in terms of public health.

Needless to say, in these two areas, in carcinogenesis, the tools aren't all that certain. We would like to have, for example, more human epidemiology data so that we can close the gap between the laboratory experience and predictions and what we are actually observing.

In the few rare cases where we haven't, we regard this as precious data in our possession.

It is the basis for doing the risk assessment in carcinogenesis and it certainly could stand strengthening.

Nevertheless, given the tools that we have now, we are going about this risk assessment work expressing the uncertainties, trying to clearly identify the gray areas and, in short, trying to describe these risk assessments as fully as we can to the regulator so that the uncertainties do come through loud and clear.

The carcinogen assessment group in EPA, with which I have been affiliated for 3-1/2 years, has been largely responsible for this work. We have done most of these risk assessments that I'm speaking of, and we are charged with insuring the adequacy of risk assessment activities agency-wide.

We have now to date looked at about 100 or so cases, something over 100. We certainly have found that our information is most often not what we have liked to see. There is data missing and information largely incomplete, and so forth.

But given all of the deficiencies in the tools and the lack of data, we still are able to use risk assessment to a large extent within the agency in making our decisions.

I have described the weight of evidence approach and the

gshHEE

idea of expressing public health impacts in terms of magnitude and the public health problem.

I think the one tough area that the agency has not addressed, and it is the reason I can't help you with the second part of my assignment today, and that is exactly what level of the agency would be comfortable with accepting as an acceptable risk or as a safe risk, or at some level that we're even going to set as a target for regulatory action.

I think there are a number of reasons why this is particularly difficult for the EPA, where there are different legal requirements, as I described in the beginning, where risk benefit comes into play.

Then, clearly, some assigned level of risk isn't going to always apply because you are never going to duplicate or very rarely duplicate exactly the balance between the risk and the benefits.

So it makes sense in those areas for the risk level to fluctuate, depending upon this balancing act.

In some areas, we have absolutely been driven to propose levels based on some target level of risk. And as I said, I will show you that example. But on the whole, I think these legal requirements cause particular problems for the EPA.

We don't want to lock the agency into one target level and then find where there are absolutely no benefits at all, that we are really stuck with at least that level. And there are many other reasons that I won't go into.

I have a feeling that you probably heard a good deal about that.

I think one thing that has complemented the process in EPA with regards to carcinogenesis is the inter-agency activity.

The IRLG risk assessment doctrine has been adopted by whole agencies and is titled, "The Basis for Risk Assessment."

It was published in the Journal of the National Cancer Institute in July of this year. I was on the panel that wropte this document and we found it dovetails nicely with the work that we've been doing inside the agency.

But I think of primarily importance to you would probably be these case examples, which I do want to show you. These are just some of many, many circumstances we have looked at in EPA where risk assessment has been used. I had to be careful in selecting these because I did not want to bring you in some regulatory program where we really haven't decided what to do yet.

So these are examples in the public forum already and where a decision has been made. And I have selected them to demonstrate the different uses that risk assessment has that have been put to within the agency.

(Slide.)

The Office of Air Programs is faced with a large number of possible carcinogens in the atmosphere to which people are being exposed both in high exposures around sub-population groups exposed to higher levels around point sources, and also nation-wide ambient exposure.

The office has a Section 112 that I mentioned that appears to be largely a health-based section of the Act and the question is: How can they implement a regulatory policy? They have decided to use risk assessments in the following ways.

They first have selected this list of organic chemicals. This list has been selected because they think they are the chemicals that they think that people are exposed to at the highest levels. They are very familiar organic solvents -- chloroform bromide ethylene, dichloride ethylene oxide nitrosamines, tri-chlora ethylene, tri-chlora ethylene fina-iodine chloride, carbon tetrachloride.

So the idea was to present these to the carcinogensis group for risk assessment treatment.

We went through -- we have been through many of these, not all of them, but we have been able to present the Office of Air Programs with, first of all, a qualitative statement on how likely these agents are to be human carcinogens.

So we have a weight of evidence, or qualitative part of the risk assessment which they can use.

In addition, they asked us to do something called a unit calculation, and that is to assume that the standard individual is exposed to one microgram per cubic meter for a lifetime, and give them a related health risk number so that they can use this number in setting priorities.

So I have a few examples to show you.

(Slide.)

They are using these unit calculations to make decisions about which agents appear to be the most potent carcinogens because they have to set some priority, they have to have some way of approaching their particular problems.

They are expressed in increased individual lifetime risks. You will see that they fall into categories of one chance in 10,000, one chance in 100,000, one chance in a million, and so forth.

The office couples these with some knowledge of exposure in order to set priorities based on human hazard.

DR. OKRENT: Those are lifetime risk, you said?

DR. ANDERSON: Right, based on exposure to one microgram per cubic meter.

(Slide.)

This particular office then is faced often with looking at residual risks and also looking at relative risks in order to decide both what category should be regulated, as well as whether or not they have done an adequate job in

gshHEE

regulating to some such standard.

This came up in the case of vinyl chloride. This table represents a comparison of risks associated with a variety of air pollutants.

In the case of vinyl chloride, Table 1 expresses the lifetime probability of cancer death due to maximum exposure and Table 2 expresses the lifetime probability of cancer due to average exposure.

The second column is the number of exposed individuals and the last column is the total cancer cases expected on an annual basis.

DR. WILSON: When you say maximum exposure, you mean maximum exposure now taking place? I mean, some of these clearly, there were fantastic exposures in the past.

DR. ANDERSON: Right. This is from EPA's point of view based on an assessment of who's being exposed now in the environment over which we have some regulatory options.

That is, people living in bands around these various plants where vinyl chloride is an environmental pollutant.

DR. WILSON: So that would not include the work place?

DR. ANDERSON: That's right, it does not include the work place. So we are looking at the possibility of regulatory exposure to those individuals living near these plants in the case of vinyl chloride.

gshHEE

So the idea was to look at some population groups at high risk and here you see before regulation, these individuals living closest to these plants had a fairly high increased lifetime risk due to the vinyl chloride exposures.

And there were certainly not an insignificant number of people being exposed at these high levels.

On the whole, we had then an average increased risk of 10 to the minus 4, which is still certainly a considerable risk, and 5 million people being exposed at this level.

The agency took regulatory action. The resulting exposures reduced these risks by an order of magnitude.

So then the question came up, in fact, it came up in an impressive fashion -- we were sued to lower the standard. And the question was, should we cash our chips here? Should we make the companies involved in this particular pollutant spend more money? Or did we think that we had done an adequate job?

Our office of general counsel used a relative risk concept in arguing that we had, for the time being, done an adequate job in the vinyl chloride case. And they, in fact, presented these two tables in their brief, demonstrating that we had yet to regulate arsenic, benzene, cadmium coke ovens where we had fairly substantial risk to a fairly large number of people, both in the maximum category as well as the average category with the total impacts, in this last column obviously being fairly considerable.

In the case of coke ovens, 149 expected cancer deaths per year.

So I think this is a good example of the use of risk assessment in looking at procedural risks as well as a very practical example of where risk assessment has been used in a relative fashion to make practical considerations and practical arguments for and against regulation.

By the way, these numbers are presented with decimal points, not because of the precision of the numbers, but we found if we take the decimal points away and round the numbers off, people simply can't duplicate our exercise.

And so they don't know where the numbers came from.

So I meant to emphasize those uncertainties. Now in one unique area in the agency, we have had to face the fact that the law, as it is written --

a11

DR. OKRENT: Excuse me. Did your counsel win his case based on the position that it would be well to work on benzene and arsenic and so forth?

DR. ANDERSON: The case was won. And I don't know to what extent in the final decision we agreed that we would immediately work on benzene or arsenic because it was won some time ago. And I think that we haven't really achieved all that much yet, but we are working on it.

(Slide.)

pv HEE

In this particular case -- I realize this is difficult to read -- but this is an area in one of our acts where we don't have the option of considering technology or feasibility or cost and benefits or social and economic impacts.

The Act clearly says that we have to set a nationwide water quality criteria number, and it has to be a number, and our general counsel's office told us that if the number was zero it wouldn't be a practical matter that these particular cases are tried at local district courts and that the judge sitting out in the middle of someplace in the midwest being presented with a zero level would simply throw the thing out of court and never believe it for one minute.

So, in other words, the case has come up in EPA where we had to set some number, and the way we have gone about it is to use risk assessment as the basis for doing it.

These are proposed water quality criteria that have been published in the Federal Register. Comment has been invited. And they have not been finalized. I think generally the comments have been reasonably favorable, at least the methodology or the use of risk assessment has not been -- has not really been attacked all that heavily. Some of the methodology, yes. And in some particular cases, the data base for specific carcinogens. But as I said, these are not finalized.

The iaea here was to take the list of water pollutants that we had agreed to consider in the court consent decree to look at the dose response slope, the incidence slope, choosing the most sensitive species, and to come up with an associated water quality concentration that is associated with an increased lifetime risk of 10-5.

Now, we assumed -- we had to make basic assumptions about exposure, so we assumed here in the footnote a lifetime daily consumption of two liters of water per day and consumption of fish. And we took bioaccumulation into account because that is one of the driving factors in setting water quality criteria. So, for these carcinogens, we came up with these numbers.

Then, our water office decided, "Well, why use just 10-5? Why not propose a range?" So, they proposed the number associated with 10-7, 10-6, and 10-5, and asked for comment. So, they haven't really set one level yet, but you can see what they are driving at.

DR. LOWRANCE: What does that 10-5 number mean?

DR. ANDERSON: It is the average increased lifetime risk of getting cancer from exposure to these particular chemicals based on an average daily intake of two liters of water contaminated at that level per day for a lifetime, over 70 years, and a consumption of .0187 kilograms of fish per day with the appropriate bioaccumulation factor.

DR. OKRENT: For each one, so if four of them or 10 of them were in your water, you would have 10 times the effect?

DR. ANDERSON: Right. And since these are nationwide -- this point has been raised, you know, how do you take into account multiple exposures from a single river which happens to be very polluted? And the answer has to be that, first of all, these are obviously very low concentrations, and we simply haven't been able to take co-carcinogenesis into account in setting these levels, because where we have the information we have taken it into account but we for the most part just don't have that kind of information.

And also, since this is a nationwide number, it is not a localized number. Localized circumstances can be taken into account when state implementation plans are approved by the agency. Since these are criteria levels that are target levels, they are not enforceable levels. And then a state presents a scheme for meeting these standards. If they have some exceptional circumstance, the agency can grant a variance. So if they have a particular circumstance that can be taken into account at that stage of the process.

pv HEE

DR. WILSON: I understand the numbers in the table, taking the two highest numbers -- the TTCD and the DBNA -- I presume those are different because BCME accumulates more; is that right?

DR. ANDERSON: Yes, that's right. And in this case, where there is bioaccumulation or substantial bioaccumulation, obviously the number was driven by that rather than the two liters of water per day.

This is a case where risk assessment has been used, and it is unique in the area, as far as I know, where it has been used as the only basis for proposing a number that is saying that protecting public health.

DR. OKRENT: So, if I understand correctly, this is a proposed goal and, in fact, it is proposed in terms of 10-5 and 10-6 and 10-7?

DR. ANDERSON: The level associated with that, making certain assumptions about exposure into account.

DR. OKRENT: And if it were adopted, it would still be a goal?

DR. ANDERSON: The way the Act is constructed, these water quality criteria are criteria to be aimed for in the state implementation plans, so these are not enforceable levels, although our lawyers tell me they do take on a fairly substantial significance.

DR. OKRENT: That is like the way you are fixing

pv HEE

up the smog in Los Angeles.

DR. ANDERSON: Congress did this. Don't blame us.

DR. OKRENT: I am not blaming you. I am only saying you can't always meet EPA's goals.

DR. ANDERSON: Yes. And the reason I started the discussion by mentioning the importance of the social constructs largely created by the Congress, the extent to which risk assessment is useful, at least in EPA, very much depends on the practical circumstance.

(Slide.)

Finally, I wanted to show some examples of where risk assessment has been used in reaching benefit balance decisions. These are --

DR. OKRENT: Could I come back to the last point. Does EPA have any basis for measuring the cumulative risks from drinking water supplies with regard to carcinogenic effects as they exist now and, let's say, including radiation and without radiation for whatever it is -- some very big city water supplies, not those fortunate enough to have the water they make beer out of?

DR. ANDERSON: I think that what you're driving at is something we would very much like to do. We haven't been able to do this, and, in fact, again, in a practical framework where we have, for example -- I know you're using

pv HEE

this as an example -- I guess what you really want to know is have we been able to take total body burden into account? And I don't think we have. It would be ideal, but we haven't been able to do it.

We would like to be able to say that we speak in terms of very low risk anyway and we hope that when they are all put together for an individual exposure we are doing reasonably well. But we really don't know. So, we aren't approaching it that way.

There are, in fact, some areas in EPA where risk assessment is being left out of the process because it is so hard to do. And in the case of drinking water there has been a regulation to put charcoal filters in drinking water supplies to simply take who knows what out, and they really haven't been able to say what it is they're taking out. Their regulation in 10 city public water supplies where there are a lot of people taking drinking water from there and where the water is dirty enough and they say, "We think this makes sense. We can't spell it out in risk-assessment terms."

DR. OKRENT: I vaguely recall seeing a paper by Dr. Page when he was at Cal Tech on that area.

DR. ANDERSON: Perhaps he can shed more light on this.

Now, this is an area where we certainly use risk assessment probably for the longest period of time. Pesticides certainly came into focus in the agency early on. And risk assessment, the risk assessment process, initially was written very much with the pesticide situation in mind.

In short, the amendments to the Federal Insecticide Act require a balancing of risk against benefits. This permits an opportunity to squeeze risk assessment for all its worth. And we would certainly like to have better tools. In each of these cases there was a part of the risk assessment which is not reflected on this chart, and that is the qualitative assessment. This is very important to EPA's process because the data can certainly give more or less certainty to the signal that something is in fact a human carcinogen. So, keep that in mind when we look at these.

In the case of chlorobenzilate, the qualitative data was not the strongest we'd ever seen. I think the response was in the male liver, which some toxicologists really don't regard as providing the strongest evidence at all. Nevertheless, we regarded this as some evidence that chlorobenzilate might pose a risk to humans. On the assumption that it was a human carcinogen, we did some qualitative risk assessment work. The primary use of chlorobenzilate was for citrus fruits. The exposure assessment provided to us assumed a nationwide exposure of 220 million people. The calculation provided these lifetime increased cancer risks due to the exposure on the citrus.

Now, the second is the lifetime risk to the applicators because of their obviously higher exposure. The first number, 10-6, is the lifetime risk to the people eating the fruit because of some residues. In short, the agency regarded this -- or the office of pesticide programs and ultimately the administrator -- as a low risk. But in the case of citrus, the benefits are exceedingly high. The registration was retained, but in some moderate uses where the benefits were regarded as very minimal -- such as I think there was a minor use to control spiders around boat docks, and that use was canceled -- it was decided that even a low risk outweighs a very low benefit.

DR. LOWRANCE: Point of clarification. I don't quite understand how these numbers work. I think I am being simpleminded and just missing something. But would you just run across, say, the top line and tell how you got to the right-hand column?

DR. ANDERSON: Yes. I think this is important, and I have left it out. In using quantitative risk assessment in EPA, we try to emphasize that a two-pronged approach is appropriate. First of all, with EPA's mission it is very important to look at the increased individual

pv HEE

risk because there can be pockets of high exposure; in this case, the applicators, for example. So, it is important to focus on the individual's increased lifetime risk.

In this case, I didn't, when I made up this slide, I didn't have the number of applicators, so I couldn't fill in the last column on total impact. But in other words, we have an increased individual's risk, and then we have a total nationwide impact. If we only look at the total nationwide impact, we could have very few total cancer cases nationwide, but we could have pockets of exposure where individuals are suffering very high risks. So that is why we emphasized both.

Alternatively, where we have, in this case, 220 million people exposed to very low levels with an increased risk of 10-6 with a nationwide number of expected cancer deaths per year of 1.5 or so, you can see that there could be circumstances where 220 million people would be exposed to a risk that would be considerable, and you would see that impact in those total numbers.

DR. WILSON: I think one thing that is troubling us: if I take 220 million and multiply it by 2 x 10-6, I get 440; I divide by seventy per life, and I get six instead of 1-1/2.

DR. ANDERSON: It should come out to 1-1/2.

DR. WILSON: I multiplied 220 by 2 x 10-6, and that's 440. I divide by 70, and that is six.

DR. ANDERSON: Okay, maybe I have made an error here.

DR. OKRENT: Well, in fact, if you look at the first line under "amitraz," where the numbers are similar, one gets eight, which is what Bill Lowrance said.

DR. ANDERSON: Yes. That evidently is an error. I apologize. You should get the answer by multiplying the first column by the second column and dividing by 70.

DR. LOWRANCE: I was just trying to be sure that I understood.

DR. OKRENT: It is probably 7.47, if you want me to guess.

DR. ANDERSON: Yes, that could be a typo. As long as you understand what I am getting at.

Nevertheless, that was regarded, whether it's one case or seven cases, as a low enough risk relative to the benefits and the registration for citrus use was retained.

In the case of amitraz, there was an application to the agency to consider this for registration. The qualitative data was very weak, indeed. We had a signal that it might cause cancer in one sex at one dose level of a tumor which was very high in the historical controls. Needless to say, blip of a signal was very uncertain, so we said this could be regarded as suggestive evidence only and it would be very desirable to have another study.

But on the assumption that amitraz is a carcinogen, we went ahead and looked at the dose response curve we could construct from the data we had and said if it is a carcinogen the impacts would look something like this for consumption of apples for which the application was to use amitraz on apples and pears. We looked at nationwide impact on apple consumption eaters, and we came up with an increased individual risk of around 10-6 and total impacts in the case of apples and pears of eight and six.

For the applicators, you will see that the risk is somewhat higher.So this permits the program to look specifically at special labeling requirements for the applicators.

In this case, the risk-benefit balancing decision came out -- I think it was pears -- that there were no good substitutes for, so they decided to grant a three-year temporary registration until more data could be generated. In the other case it tipped the other way because there were other pesticides available and the benefits were not thought to be as considerable, and that registration was not permitted.

So you can see how these things can teeter depending upon just how strong that benefit side of the case is.

DR. OKRENT: What would you think are the uncertainties in your estimate of lifetime probability of cancer deaths due to exposure? Is it larger than a factor of 10?

DR. ANDERSON: This is something that we dont't have the best information, and I think, as I said in the beginning, we have very few cases where we can actually close the gap between the predicted cancer cases and observed cancer cases. In other words, there is great debate on the shape of the dose response curve within species extrapolation and even greater debate when you cross species lines.

We are now doing something in the carcinogen assessment group which I think will help. This whole quantitative assessment business is based on essentially six cases which were presented in 1975 as part of an NAS pesticide report, where I think three of the cases came out of the nose. Two of them were off by an order of magnitude lower. And one was off I think it was -- yes, it was off by an order of two orders of magnitude.

In short, the basis for doing the qualitative assessment is not all that certain, but it is stronger than the quantitative and the qualitative. We have about 25 cases where we can compare the animal results with human results, and the correlation is pretty good. There is one case where it is not that good, where we don't have the animal data, on arsenic.

DR. OKRENT: Well, let me pose the question a different way. We have been talking about several chemicals here. Let me assume if we consider 100 of these, can one rule out that for one of these that the lifetime risk turned out to be 10-3 where you thought it was 10-6?

DR. ANDERSON: No, we can't rule that out. What we have done in EPA is taken, as best we can, consistent assumptions and the same extrapolation model. Now, if you use a different extrapolation model, you get answers all over the place. So, the way the agency has been using this information is very much in a relative sense and not as an absolute correct number.

The other part that makes these numbers uncertain aside from the uncertainties in the shape of the dose response curve and across inter-species extrapolation, certainly is the lack of good exposure information. Exposure assessment is a very weak area indeed. So, when I started out, I said these numbers are uncertain. I want to emphasize again they are uncertain. But it is the best we have.

DR. OKRENT: Now, I am not faulting you in any way.

DR. ANDERSON: No. I just wanted to be sure, if I hadn't emphasized that.

DR. OKRENT: We face in the nuclear safety area large uncertainties with regard to many predictions, and one of the thorns in our side is what to do when you talk about quantitative criteria in the face of these uncertainties. And I am trying to ascertain a little bit what it is one thinks could be the range of uncertainties here in the kinds of things you are looking at and how EPA thinks it is reasonable to, A, define the uncertainties and, B, to act in the face of them.

DR. ANDERSON: What we have done is we certainly know about the six cases that have been published. In addition, we say that we take the linear model because we want to -- we don't want to underestimate the risk because if we are underestimating the risk we could have a national disaster and not know it. But at the same time, if we are way overestimating the risks, it is not being particularly helpful to the process of setting regulatory actions. We hope we are not vastly overestimating, and certainly we don't want to underestimate.

What we are trying to do right now as an activity in the agency is to take the 30 or so possible cases where we have some human data and we have the animal data and we can compare the predicted human response with observed human response. Now, this doesn't mean we have the best epidemiology studies in all these cases, but we can at least set an upper bound to see how close we are coming. And I think in the near future we will be able to say a little more about this. It looks like thus far -- and we're still going through our calculations and we are adding more cases -- that in about 90 percent of the cases we are coming out within an order of magnitude, again taking very consistent assumptions and the very same extrapolation model in every case.

DR. WILSON: Arsenic would be way off; wouldn't it? I mean, there is no way you can predict the human taking the animal data from arsenic, you can get the human data.

DR. ANDERSON: Yes, in the case of arsenic we don't have the animal data.

DR. WILSON: Well, there is some animal data, but it doesn't give the right answer.

DR. ANDERSON: Yes, that is the one exception where we really don't have good animal backup data.

DR. WILSON: If there is another chemical out there like arsenic and all you have is animal data, then that might be the catastrophe.

DR. ANDERSON: Well, I am not sure I understand.

DR. WILSON: Well, we have some -- I know of about 60 different attempts to find carcinogenesis due to arsenic, and they basically either find very small carcinogenicity in animals or none at all. And then we know it does -- it is moderately potent for humans. In fact, it has been disputed over the years, but I think it is now generally agreed that it is.

But if only you had animal data and then you are trying to predict human data and make a plan, "Shall we use arsenic as a pesticide," the answer would clearly be "Yes." And you might well be wrong. And I am just worrying about that. These are the sort of things we are getting in the nuclear business, too.

DR. ANDERSON: But this is the qualitative side of the question. This is where we keep saying we are more certain than we are on the quantitative side, and that is: how likely are all of these animal bioassay tests to give us appropriate signals to indicate cancer risk where it exists? In the case of arsenic, that is the only really good exception in 26 cases.

DR. LAVE: Benzene?

DR. ANDERSON: Well, there are some indications if you pick the right model you might get leukemia. That is another good example.

DR. WILSON: It is also the numbers on benzene. On benzene, what numbers we do have, if you put an upper limit on animal data, it is not inconsistent with human data; whereas, the arsenic, the animal data is actually inconsistent. It is not that it doesn't exist; it is inconsistent.

DR. ANDERSON: Yes. In the case of aphlitoxin, of course, we only looked at responses in rats. We would have thought aphlitoxin was not a human carcinogenic. So, picking the right animal model certainly is important.

Just wandering through these quickly -- I think you're getting the drift here of what we are trying to do -- chlordane and heptachlor, which the agency canceled and suspended, had substantial nationwide impacts mainly because the persistence and bioaccumulation in the range of the hundreds. So, the agency did cancel and suspend those uses. It retained only some underground uses for termite control.

Finally, this is a case of kind of an after-the-fact risk assessment where the agency, before the guidelines were adopted for risk assessment, set an action level for kepone in fish in the James River. Most of you remember, I am sure, the James River incident. This was set largely on economic grounds: how much could we permit and still not kill the fishing industry down there.

When we did the risk assessment calculation after the fact, we found indeed the risk was somewhat high relative to other risks, where carcinogens had been

pv HEE

regulated within the agency.

I think the agency is not taking a risk benefit approach to setting these action levels, and there have been a series of --

DR. OKRENT: Excuse me. So there has been no decision?

DR. ANDERSON: This was not changed, because the fishing industry has largely -- they haven't had a lot to fish down there, recently. They haven't changed this but they have taken a different approach to setting action levels, one where it is considered in the process of setting the action level. This was before the agency really started to do risk assessment work, that we were forced to use safety factors and a crystal ball, and some knowledge of the economic circumstances, in setting these levels.

DR. OKRENT: So, consumption was permitted, or was not?

DR. ANDERSON: Yes, with these levels. A few more cases came to mind that I thought might be interesting for you all. We have some localized problems where it's very hard to get something across to the public. Several examples. One, we have a discharge of nitrosamines in our Region 5 office in Indiana. There was sort of local panic, you know, nitrosamines belonging to a family of well known carcinogens, and the citizens were quite concerned. There were two drinking water communities taking the water supplies downstream from this discharge, so the question certainly wasn't a qualitative question of, are nitrosamines cancerous or not, but rather, what kind of impacts are we experiencing?

We had some exposure levels for these local communities and we did this kind of risk extrapolation work and we found that the levels to the people, the individual increased lifetime risk levels to the individuals living in these communities, was on the order of 10 to the minus seven. It turned out that we emphasized that these numbers are uncertain, that this is a relatively low risk estimate and certainly doesn't indicate a public health emergency; nevertheless, precautions seemed to be -- it seemed to be reasonable to take precautions to reduce exposure. It seemed to serve the regional office well. They were able to get this across to the communities and the stir really did die down fairly rapidly.

We had another similar discharge from a point source into air in upper New York State. It was a trichloroethylene discharge, and again we had some exposure levels at monitored levels around the plant -- in this case, it looked like the public health risk was in the neighborhood of 10 to the minus five, so what we're doing with these numbers is we are comparing these now, but we have kind of set up a scale and we're using the same model with sort of a scale of where regulatory actions with regards to carcinogens have fallen.

And we can say this is somewhere on that scale, it is somewhat higher, or it is the upper end or the lower end of that scale. That is about as close as we have come to really setting some target levels. But again, this was a case where it seemed to emphasize that there was not an enormous public health problem there, certainly not one calling for immediate closure of the plant. EPA certainly does have that authority. I think we recognized we have to use it very carefully.

But if we thought that the circumstance was certainly bad enough, we could simply shut the plant down.

DR. OKRENT: A 10 to the minus five number would arise from continued exposure at that level?

DR. ANDERSON: Yes. Assuming the individual lives there and is exposed for a lifetime of 70 years to the level around the plant, yes, the lifetime average exposure.

DR. OKRENT: Okay.

DR. CASTENBERG: I noticed in some of the entries in the last table -- used 220 million, the total population of the U.S. If you had an extensive list of all the chemicals used in the environment, how many of them would the average person in the United States be exposed to? 1000 of them? 50 of them?

DR. ANDERSON: I don't know. That list of 42 that I started out with, which is just the Office of Air Programs list of likely organic solvents to which people are exposed, certainly gives us a good start. And that is only their first cut. I really don't know the answer to that. I think our Office of Toxic Substances could probably come closer to having information of that sort. I just don't have it.

DR. CASTENBERG: The reason I bring it up is I was also concerned about this 10 to the five number -- not concerned, but interested in the 10 to the minus five number. Then the numbers on your last page are 10 to the minus six numbers, for total U.S. population, the 220 million. The numbers on here are 10 to the minus six, and give you numbers like eight deaths, 10 deaths, six deaths, and, at 10 to the minus five, then those numbers would be 80 and 60 instead of eight and six -- and if I am exposed, or if you are exposing the population to 1000 of these things on the list, or 10,000 of these, those numbers start to get very large.

DR. ANDERSON: Yes, if in fact the exposures are to the same people at the same levels and so forth, although when we're talking about 10 to the minus five I think it is fair to keep in mind that everybody has a one-in-four chance of getting cancer, anyway, and a one-in-five chance of actually dying of it. So we're talking about a

kapHEE

one-in-100,000 or one-in-a-million chance here. So, you could be exposed to quite a few things before you would have an incremental increased risk. That would be really outstanding; in fact, we sometimes look at some of these risks as levels that are below levels that would be detected by epidemiology. No detectable cancer when you get down this low.

So these are pretty low levels. But needless to say, if an individual is exposed to enough different things at low levels, presumably there is an accumulation of risk. There is no question about it.

DR. LAVE: Could I ask you about the bioassay data? In the case of aphilitoxins, as you mentioned, you get very different numbers depending upon whether you're extrapolating from mice or rats. How do you choose which animal to extrapolate from?

DR. ANDERSON: We always pick the most sensitive —

DR. LAVE: I thought that was your answer, because that, then, means you don't have a consistent set of methods. In some of the tables you've got human epidemiology data; in some of these you have quite a number of different animal species that were tested; in others, a single animal species; and therefore the ones where there's a single animal species, unless it happened to be the most sensitive one -- which is not very likely -- then you may be off by a factor of 1000 in the table.

DR. ANDERSON: There's a little bit of intimation emerging that is somewhat helpful in this regard. Bruce Ames and Ken Hooper and some of his people are putting together -- they are computerizing all of the animal bioassay work that has been done and I recently heard Ken Hooper present a paper -- these results were not published. They were alluded to in a Science Magazine article.

What they are saying they are finding is where you have a positive response at all, and you threw out all of your negative responses, that you are coming out with a reasonable measure of potency; in other words, where you have a positive response, you don't get potency measurements all over the board. If your test model isn't going to pick up the cancer signal at all, then you would just get a negative response.

But this is certainly a problem. It is one of the problems. When we started doing cancer risk assessment work in EPA, we recognized it applies both to the qualitative work as well as the quantitative. Chemicals have been tested in very uneven fashions and not only that, some have been tested in 30 systems and some in only one, and not only that -- the quality and the design and conduct of the study varies all over the place.

DR. LAVE: But when I first asked somebody at the FDA that question, they said to me, Well, of course the animals you use are the ones that metabolize whatever it is that we are exposing it to, the same way that humans are exposed, and metabolize it the same way that humans do. You obviously don't use an animal that doesn't metabolize it the same way.

And that made a lot of sense to me, except then when I talked to some of the people over on the other part of FDA, they said, Well, of course we don't do it that way. We use the most sensitive species.

DR. ANDERSON: Well, it would be nice if we could do what you first said; that is, be able to follow metabolic pathways and say, This test model has the same metabolic pathway as the human.

But most of the time we really have no earthly idea what the metabolic pathway is, so we couldn't possibly do it.

DR. WILSON: Wasn't part of the answer to [illegible] that all of the animals one uses are, in fact, mammals, so in that sense there is a metabolism? We're not using insects and they're not using fish.

DR. ANDERSON: Well, of course the reason we're using rodents, again, goes back to the case of the 26 chemicals where we have some correlation information. If we had the same kind of data base for some other test model, I would think we could use it in the same way.

There is nothing particularly unique except for the amount of data, the data base.

DR. WILSON: Incidentally, we have a paper coming out which should be in press about now, comparing animal carcinogens in the different animals and it does have the results you described.

DR. ANDERSON: I think that's going to be very helpful.

DR. WILSON: We agree almost exactly with Bruce Ames' output. I sat down with him and we went over the computer output.

DR. ANDERSON: We have talked to Ken about coming in and presenting some of this, because it goes right along with the work you are trying to do, that is, the correlation project of predicted and observed for every case we applied.

DR. LAVE: Can I just see if I can paraphrase what you and Dick just agreed to, to see whether I understand it? That is that where you find that a substance is a carcinogen across a number of species, that the potency across species will be quite similar; does that paraphrase it?

DR. WILSON: Within a factor of 10, yes.

DR. LAVE: I would have thought some of the things

kapHEE

we have been talking about earlier contradict that, but they may be the rare exceptions.

DR. WILSON: I think there may be three or four exceptions. Arsenic is the big one.

DR. OKRENT: Would you give your name?

DR. CACHERINI: Part of the answer to that question, in regard to the one part of FDA and the other part of FDA -- it depends upon whether one is looking for an estimate of quantification of the risk or whether they are trying to determine whether there is or is not a risk.

In the first case, one would require a more appropriate model to the human than in the second case, where one's looking strictly to see if they can incite a risk in an animal model. So I think the appropriateness of the use for the actual number of quantification of the risk requires a model which is consistent with the human.

DR. OKRENT: Dr. Page?

DR. PAGE: There are two other points that go along with what is coming out of Bruce Ames' laboratory. One is that part of the exercise was to figure out whether or not tests that came out negative were compatible with tests that came out positive, in the sense that if you looked at the potency that was applied by the positive test, you could try to figure out what the power function was of the negative one, to see whether or not it was likely that it would show up negative when that was given.

A lot of these inconsistencies between species turns out to be an inconsistency of the test design, where if you took into account the power of the test then you would see that these differences of potency aren't nearly as big as what would first appear if you looked at the test just sort of at face value.

The second point is that --

DR. WILSON: Isn't that just a statement that they were seeing positive -- people were saying positive or negative before and as soon as they put it in a number, saying an upper limit, and saying, this is less than something or other, then the inconsistency appeared? If you just say positive or negative, you can get apparent inconsistencies. But then you have to include the test.

DR. PAGE: That's correct. If you looked at the test and figured out what the power function was, and you graphed the tests -- that is what Hooper's people are doing -- then it looks like -- it looks much more compatible.

DR. ANDERSON: There's no question about it. You do simply have some chemicals that the test model is inappropriate — for whatever reasons, they can't give enough dose to have the power to reach an adequate level.

I have just one other example I wanted to mention, and this is something that is plaguing EPA, and we are

kapHEE

having to sit back and think hard about what to do -- and that is the likely impacts of dieselization.

It is a case where risk assessment is undoubtedly going to play a large role, and I think the synfuels development also goes along the same pathway. The way we are looking at this is we need some answers fast, as usual, and so it is not possible to get all of the research information we'd like to have.

So the way we are approaching this, from a risk assessment point of view, is to attempt to get some information to compare relative potencies with other combustion products for which we have more experience, such as coke ovens, cigarette smoke, coal tars -- and then get some relative comparison with diesel exhaust, because the question is not whether or not diesel exhausts are carcinogenic -- we know they are, as are most combustion products. But the question becomes: how much of a public health hazard are we likely talking about as we go along this dieselization pathway?

We very reluctantly did some very preliminary risk estimates based on coke oven information. And no matter how many caveats we put in and how much we emphasize we didn't particularly want to do this, our Office of Air Programs said they needed some information and something was better than absolutely no idea, and that they were compelled to give at least some signal.

They wanted to know whether the dieselization was likely to impose a large public health problem, so we did do this, reluctantly, and we did find that it is conceivable that diesel exhausts are as potent as coke oven emissions, and if that data is at all applicable, and if their exposure estimates are at all accurate, that we could be -- it raises a warning signal. We could be dealing with a public health problem of some magnitude.

But we now have some research in progress where we are trying to get some answers. We have been trying to compare these potencies using skin paintings and intratracheal data, and we hope to have a biological basis for doing this work. But I think this is an example of trying to use the best tools we have in a very uncertain circumstance.

But nevertheless, having some rough idea of -- or ballpark estimate -- is better than having simply no information at all. So, in short, I think at EPA from the experience that I've been talking about, the experiences with carcinogens that I've been talking about, I have been able to show you examples where we use risk assessments to set priorities, to look at relative risk, when we are considering further regulatory action, to look at residual risk where we have taken an action and we want to know if

kapHEE

that is good enough, in risk/benefit balancing decisions and in one unique case where we have actually proposed these criteria based on increased individual risk level.

This concludes, I think, my contribution to this. And as I said, I simply can't help you with the second part of my assignment.

DR. OKRENT: Dr. Leachman?

DR. LEACHMAN: Yes, I am from the House Science and Technology Committee. I have a question with how much we have to know about -- the chemical carcinogens is something I know a little bit more about; namely, the radioactive or rather the radiation effects.

I think I understand correctly that the radiation effects we know vastly more than we do about the chemical carcinogens, and we also realize the difficulty that the chemical carcinogens act in different ways, whereas radiation generally acts the same way. But from a societal or even a legal point of view, we are faced with the realization that in the case of radiation we are now looking back in a somewhat retrospective way, and realize that we are very cautious, that most of the radiation we assumed in a very pessimistic way that there is a linear dose effect relationship.

We set up standards and we have the public very concerned. Now, in retrospect, we realize that was overly

kapHEE

stringent, and the concern was not properly placed. I wonder whether there is a similar apprehension coming along in the chemical work, in the sense that you do not -- I sense that you do not now know the dose effect relationships of small doses and maybe there is hardly any effect, and maybe this is being overly stringent and overly cautious, notwithstanding the fact that Congress has laid it on you.

DR. ANDERSON: A couple of things. In my ignorance, I guess I didn't realize that the dose response at very low doses of radiation had been all that thoroughly --

DR. LEACHMAN: Well, there are two types of radiation, the low linear energy transfer, low LET and the high LET. Most of the radiation that people encounter, except in buildings, is the low LET and there the dose effect relationship is generally believed to be much less at the very low levels, than previously expected. This is on the basis of current data.

DR. ANDERSON: From BEIR.

VOICE: You have only read a draft of that. I think that is a limited interpretation by one group.

DR. LEACHMAN: Let's just say the majority, then.

DR. ANDERSON: I thought there was some debate, and I don't want to talk about it except to say that in general circumstances for chemical carcinogenesis that is

kapHEE

less certain than the radiation area. That has a much more solid foundation. When we talk about low dose exposures, several of the models do approach the area of the low dose. We are using, again, the best tools we have, as we can close this gap I've talked about, between laboratory experiments and some human observation, and I think it is possible, we certainly have many things that are suspect carcinogens which are in wide use --

DR. LEACHMAN: Could I ask the fellow one question? Do you use the Ames test at all? Do you always have to use the animal test?

DR. ANDERSON: We have taken a position on that, which is an official agency position and that is that in the qualititative sense, we regard the Ames test and some other short-term in vivo and in vitro test data as suggestive of carcinogenesis in the present correlation, but we have not yet been willing to use that information alone to form the basis for saying that we really believe that the chemical is likely to be a human carcinogen.

DR. LEACHMAN: But I was asking in a different sense.

DR. ANDERSON: In the quantitative sense, no, we don't use the data from the Ames test to do quantitative extrapolations.

DR. LEACHMAN: That's not the sense I am asking.

I am asking: can you use the Ames test, which is so simple, cheap -- and you can get good statistics, to try to establish a dose effect relationship?

DR. ANDERSON: That is what I was getting to, and we don't. We are not doing it for carcinogenesis at all. We don't think it is appropriate. However, the agency is attempting to write some guidelines for another health effect, mutagenesis, and looking at what, if anything, would appear to be appropriate for a risk extrapolation for an endpoint of mutagenesis, not carcinogenesis.

And that point has not yet been answered, and the agency has not taken an official position on it.

DR. OKRENT: Well, thank you, Dr. Anderson. I found that very interesting and informative. And I look forward to having someone keep us better informed in the future than we have been in the past.

The next speaker is Dr. Page, also from EPA. But first, why don't we take a short break?

(Recess.)

DR. OKRENT: If we can reconvene, I think our next speaker is Dr. Toby Page, who is also with EPA, indeed I believe from Cal Tech.

(Pause.)

DR. PAGE: Let me just start with slightly more than the usual disclaimer. Steve Jelenick was the one who was invited to come here. He is the Assistant Administrator for Toxic Chemicals, and since I have been working on the concept of unreasonable risk, I was the one to fill his shoes. However, I don't speak for EPA. I'm really a guest scholar on a leave of absence from Cal Tech, so you can't hold EPA responsible for anything I say. So that is a strong disclaimer.

Now, having said that, I think what I would like to do is just sort of talk a little in continuation of the kinds of things that you were talking about with Betty Anderson, and bring that into the approach that I sort of see might be emerging from the toxic substances program and try to draw some parallels with the kinds of problems that you people may be having with safety programs for nuclear power.

And of course the last is what I know the least about, so these will be sort of shots in the dark, but maybe they'll be useful and maybe they won't be.

Okay, now. One of the things that emerged from

mgcHEE

the discussion with Betty Anderson is this number of 10 to the minus 6 or 10 to the minus 5, and basically this is looking at some sort of criterion for acceptable risk that is focused on what is entirely on the risk side.

This is just not really in the cards for the toxic chemicals control program, basically because the legislative history requires balancing of cost, benefits, and risks. That is quite clear. And so, what we are really directed to do by the legislation is to develop a concept of unreasonable risk that is up front about the balancing process.

I think sort of as an aside there always is a balancing process, even in the Delaney Amendment which ostensibly says it does not have a balancing process. The problem with the Delaney Amendment in my mind is not that it precludes cost-benefit analysis. It is that it forces cost-benefit analysis underground, so that we find people unwilling to start the machinery to do appropriate tests because they are scared of how they might come out. If they came out positive, then they'd be forced to do something that would lose control of the process. So what we find is that in 19 years, only three chemicals falling under the Delaney clause which seems a bit strange -- there are a lot of carcinogens in the food supply.

So this sort of gets to the end of what I wanted to say, which is basically, if we think of a cost-benefit analysis not just on the chemical at hand but on the whole structure of incentives and institutions, we may get much more power and use and guidance out of it that way.

And in a way what I'm saying is, if we are going to try and set -- if we're going to try and make decisions about how you're going to have street lamps, traffic lights, or red and green, or whether you're going to have policemen and stop signs or yield signs for a city that has got 40,000 intersections -- it probably does not make too much sense to do a thorough, year-long cost-benefit study on each one, each intersection.

You could do a cost-benefit analysis on the cost of information, and you may need to develop some rules of thumb that say we're going to be pretty crude, we're going to set a green light for 20 seconds and a red light for 20 seconds, and then if there are problems, we may have to readjust, but we want to get a lot of traffic lights out. And this may be a better technique than spending five years or one year or whatever on a particular intersection.

Now the parallel I'm trying to draw here is that there are two levels to think about it. One is the level of the specific chemical, the specific problem at hand, and the other is the level of the entire institutional structure. And the latter forces you to be more qualitative in the

mgcHEE

analysis, less quantitative. But it allows lots of quantitative thinking.

And now I'm really sort of getting to the end of what I wanted to say, but I may as well say it right now, that it allows thinking in terms of trying to think through the principal agent problem, the cost of information. There are all of these disciplines that have a quantitative punch to them in the sense that they are rigorous.

On the other hand, their primary value I think is in developing rules of thumb and guidance towards what kinds of institutions make sense. And to sort of just say it straight out, it appears to me that that is really what the Three Mile Island report is trying to say -- that we focused a great deal on the particular minor problem at hand, whether or not a valve is going to fail, and we have left out of consideration the institutional structure whereby we worry about such questions as how operators are trained, what kind of incentives are placed on the operators to learn their material, to know what the regulations are, what kinds of substitutes there are for the regulations, legal fee applicants versus something else more decentralized, how we look at design errors ex ante rather than ex post, which is really a game of incentives. So I think it is that side that needs stressing, and it is that side that comes out of the Three Mile Island report, and I hope it is that side that will sort of clarify the rest of the remarks I want to make.

So having said that, the Toxic Substances Control Act mentions the term "unreasonable risk" at least 43 times. It is hard to know how many more times it is in there because it's very hard to count them. Your eyes glaze over as you go through the Act. I am thankful to Harold Greene, a law professor at Georgetown University who made that comment before another session.

Now the term sometimes comes up "may present an unreasonable risk", and this is a trigger for testing, and sometimes it comes up with "presents an unreasonable risk", and this is a trigger for precautionary action. It may also come up in terms of "does not present an unreasonable risk", and this come out in terms of a responsibility to write something in the Federal Register to show why you're not regulating something in which you started a process. So whichever way you go, you have to worry about that.

There are some 43,000 chemicals on the list of inventory of inorganic chemicals that are not being used as pesticides. If they're used as pesticides and in commercial list, they may be on the list as well. Cosmetics are not on it. Tobacco is not on it. Things like that are exempted from the Act. So we still have a residual very large number of chemicals, being added to by about 400 per year, and you

mgcHEE

begin to see why I have this image of setting street ligh s in a city with 40,000 intersections.

If you decide you're going to regulate one at a time and it takes five years to regulate one, then you're talking about a large number of years. It sort of reminds you of the kind of years we talk about for half-lives for radioactive things that you have to worry about. There are cheap tests like the Ames test that I think cost something like $200 to do a simple version, $1000 to do a really good version, to do the long term bioassay, the kind that Betty Anderson would like to see before she comes out and says that it really is a carcinogen -- that may cost a half a million dollars per shot -- so you're not going to be able to require every chemical to be tested, every suspicious chemical to be tested under a bioassay condition in order to decide whether it is really worth regulating or not.

Some very crude guess, which I guess is just sort of drawn out of the air but has some sort of guidance to it, that perhaps one to five percent of the synthetic organic chemcials today that have been created since the Second World War may be carcinogens, mutagens, or teratogens and therefore subject to regulation.

One of the things I should mention, however, is I think that it may be sort of an interesting point, is that right now almost the entire focus of concern is on cancer.

On the other hand, it may turn out that the mutagens and teratogens pose greater harm socially than the carcinogens, partially because they destroy the gene pool for many generations as opposed to just kill one person today. This may also be a concern to be considered with radiation, because the same processes that cause cancer radioactively may be able also to cause mutations you don't see for two or three generations and then come up and have a larger impact.

Now the problem gets joined, I think, in a parallel way with the nuclear safeguards as it does with toxic chemicals, and that is that if we are trying to be remedial, the problem is sort of preidentified for us. We have an accident. Clean it up.

We find that asbestos insulation workers are — half of them are dying from lung cancer, so we know that asbestos is a bad agent. But our problem is to be precautionary. And then we have to ask the question, how much precaution is enough, and that becomes a very tough problem because the problems don't preidentify themselves when you're trying to be precautionary, like trying to look for a needle in a haystack, and until the needle announces itself, it is very hard to find.

Now our problem in some ways, I think, may be a bit more structured than yours, and that is because we do have a formal test that we can do on chemicals to see whether or not they are carcinogens, mutagens, or what not, and we know we have some testing protocols, and it is kind of stylized. And the parallel problem for you, I think, is how do you smoke out a design defect before it has actually gone into an accident, and there is no sort of nice methodology you can grind and come out with an answer that tells you that, yes, indeed, this was a design defect, although after the accident you very often can identify it.

So in our sort of stylized world, there are sort of three basic ingredients that go into a concept of unreasonable risk. One of them is the baseline information that you have on hand before you make a decision whether or not your're going to require testing. This involves what's in the literature, what is in the premanufacturing notice, what is know, what inferences you can make about it, this sort of thing.

Then the second characteristic, the second ingredient is the charateristics of the test. And for, I think, all people who have enough of a sense of statistical hypothesis testing to sort of realize that when you go into a test, you can get hurt two ways. You can have a guilty chemical come out innocent, or you can have an innocent chemical come out guilty. You have false positives. You can have false negatives. And you trade off the probability of one against the other by setting the critical region of the test.

It just happens that for the kinds of effects we are trying to find, this trade-off is quite unfavorable, and just a sort of quick example of this -- the most used bioassay test is sort of the NCI test that requires 50 animals. It requires a five percent significance level, which means a five percent chance of a false positive, and for a medium level carcinogen, it offers very low power.

And example of this is, suppose that we have a ten percent response rate, background response rate for the animal, realizing that this does not mean that the animal is super-sensitive because after all, we have a 25 percent background response rate, so these mice are less sensitive than we are, and suppose we're looking for something like benzene, that might double this background rate under some increase of dose below its maximally tolerated level. And so, in other words, the background rate would jump from a 10 percent response rate to a 15 percent response rate. What is the chance that the test is going to find this effect, this 50 percent effect?

Well, the chance is less than one percent. So here we are, we have an NCI type cancer test that has a five percent chance of a false positive and a 99 percent chance of a false negative, and in decision theory terms, this just does not make much sense. The reason that it doesn't make

mgcHEE

much sense is, for many precautionary actions that we can take, the cost of the precautionary action may be low compared to the cost of not taking the precautionary action is the adverse hypothesis is true and this chemical really is a carcinogen at the potency we're talking about.

So to say that again, if we have -- the third ingredient is really the ratio of the cost of a false negative to the cost of a false positive, and so that is where sort of the economics comes in, and when we begin to worry about what the costs of information are, that tells us how much up front work we need to do to try to figure out what this ratio might be before we begin to proceed in requiring the test or requiring precautionary action.

Let me put the matter sort of the other way. Suppose that the cost of information was zero, that we could do these cost-benefit analyses at no cost and no time, then it would make sense to do everything up front in the sense that before we require tests, we would figure out what regulation would be in order if the test came out positive, at what level, and what the benefits foregone of the chemical are once you regulate it, what the costs of the chemical are in terms of not regulating it, and living with the carcinogenic or mutagenic or teratogenic properties, to do all that work up front.

However, when we begin to worry about the cost of

mgcHEE

information in terms of both time and resource cost, because these are expensive kinds of things to do, then we are beginning to be in this traffic light setting problem where it may not make sense to do as much up front work on the cost-benefit side before we require the tests, as up front work on the baseline information side to see whether or not this hypothesis looks like there really is an adverse hypothesis.

So what I'm saying is that, instead of worrying about your cost-benefit analysis only in terms of this chemical, if we worry about the cost-benefit analysis and the level of the entire decision-making process, we may make much better use out of it.

So that is the point. Let me just try to sort of draw this parallel a little more. It seems to me that in the big accidents that have happened and the interesting near-accidents in nuclear power, we have a combination of three factors. We have the mechanical failure, the design error, and human error, and they sort of intermix. They come together.

And the problem is that one of these three is easy to study and the other two are real hard to study, and you can guess which one is the easy one to study. It is the mechanical failure. And that is what has gotten the lion's share of the attention.

mgcHEE

And this situation is a little bit like the old joke of the man who lost a coin, and he's looking under a streetlamp for it. A bystander comes along and says, "Can I help you find this coin." And the man says, "Oh, yes, help me find the coin." And the man says, "Where did you drop it?" Then he said, "Oh, across the street." And he says, "Well, why in God's name are you looking here?" And he says, "That's where the light is."

So the problem is to think more systematically in a cautionary way of smoking out some of these design areas and some of these incentive systems that would make human error less likely. And I think that may be where the big payoff is. I think that is really the point I want to make.

Let me make just one other point and try to get some reaction or discussion. I think when we take these three ingredients into account that involve the toxics problem, they also involve the nuclear safety problem with a vengeance, in a sense. And that is that when we are dealing with the ratio of cost to mistakes, the downside costs of being wrong and not finding some design error may be thousands and billions of tons bigger than the downside cost of being overly precautionary for some little tidbit.

So when we start looking at incremental cost of safety and the whole potential on the downside, it is very large. And so that sort of steers you to make all sorts of

mgcHEE

precautionary pile-ups, and the question is where do you stop.

DR. WILSON: On that, I'm not quite sure what you're saying there, because -- excuse me -- I'm stopping you there because it seems appropriate, because are you trying to imply that the costs of not finding an error and the actual cost of having to fix it up, or there is another cost which we find very clear from Three Mile Island and which is probably not the cost of fixing up Three Mile Island, but the cost of the public concern that causes — which is an indirect cost, which is very much greater than the actual direct cost?

DR. PAGE: Well, I mean both.

DR. WILSON: Are you meaning both?

DR. PAGE: Yes. Maybe it would be worth — I don't know, how do you want to do this? You want to open this up to discussion?

DR. OKRENT: Well, why don't you finish? We have time.

gshHEE

DR. PAGE: Let me just give you a little stylized example of this to sort of pinpoint what I mean by these asymmetries.

Suppose that you're feeling a little poorly and you walk into your doctor's office and you say, I'm feeling a little poorly and so he gives you his check up and he says, don't worry too much. I think you're okay.

On the other hand, I've examined a lot of people like you and I find that one out of 20 people like you has this rare tumor.

And you begin to start shaking. So you have this 5 percent background level, and that's what I mean by the baseline information, a 5 percent chance that the adverse hypothesis is true.

And you say, well, what should I do? And he says, well, I think we should give you the test and the test is pretty good. It has only -- if you're due to have a tumor, there's a 90 percent chance that it will show up positive. If you don't have the tumor, there's a 90 percent chance that you will show up negative. Come back in two weeks and take the test.

So you take the test and you come back in two weeks and he says, look, you'd better sit down. The test came out positive. So you sit down and he says, but don't worry too much. It is still unlikely you have the tumor. And you

gshHEE

say what? How can this be? I thought it was a good test.

And if you applied basis therom to the information I have given you, it turns out that you really do only have a 30 percent chance of having this tumor, even after the test is true.

So then he says, well, should I have the operation or not?

Let's stop right here for a second. One point is that I think most people, when they are dealing with low probability events and things that require taking into account badly formed baseline information and new information, they do a real poor job. There's a whole history on this.

Kuhneman and Turfsky have an article in Science that show that people handle low probabilities very poorly in updating the situations, which is part of the problem that I think that you face as well as we face in toxic chemicals.

So you say, well, should I have this operation? It was only a 30 percent chance of having the turmor. And the doctor said, I think you should because the cost of dying of cancer is at least 10-fold times higher than the cost of the operation, as nasty as it is.

And you say, well, if I believe in subjective probabilities, and that's what the meaning of 10-fold means in the betting situation, then I'm willing to reject the bet and take the precautionary operation.

So then you take the operation and you come back and find out that you were one of the lucky people that didn't have the tumor after all, that it was a precautionary operation and you sort of wonder how do you avoid malpractice suits?

And he says, well, basically, I give all of my information up front. People know this is precautionary to begin with. And you say, well, it is a mistake ratio of 2 to 1 to unnecessary operations, to unnecesary operation. And he says, no, that's not true at all. My mistake ratio was 19 to 1.

I have for every false, negative, I sustain 19 false positives.

Now if you grind through basis theorum again, you will find that is also true. And that's another place where people don't think very well about it.

So then the question becomes, when we look at toxic chemicals, we find lots of examples of false positives. You can quickly find 20 or 30 examples of false positives, but you can only find 2 or 3 examples of false negatives.

So the question is have we been fooling ourselves as to the right balance of our mistake ratio? Have we not been precautionary enough?

So that is a question of looking at expected values and seeing whether the numbers make sense or not.

DR. WILSON: I did not catch the sign of that statement. Have we been too cautious, was that?

DR. PAGE: No, I'm saying have we been fooling ourselves in the sense that we have paid a lot of attention to false positives and we have protected ourselves a lot against false positives. We have been unwilling to classify something as a carcinogen until there is overwhelming evidence that it is a carcinogen. And in doing so, we find lots of classifications that go from safe to dangerous, which is the sign that we have had a false negative -- I'm sorry, which is a sign that we've had a false positive.

I'm sorry, this is a sign that we have had a previous false negative when we take a chemical which we think is safe and then we find it's dangerous.

Now there are a couple of cases where we take a chemical which we classify first as dangerous and then we loosen up and say, ah, maybe it's a little safer. The question is which way do these loosening and tightening operations go and what is the ratio of them?

So that becomes a way of calibrating what you think as to how precautionary you should be. And also, it gives you a way of approaching the problem of trying to decide what people are really implicitly judging as to the underlying likelihoods of that, of the hypotheses of the ratios of cost.

Now when I said that you face this problem with a vengeance, what I meant was that two of these asymmetries can be summarized in terms of what used to be called the zero infinity dilemma, where the first half of the zero part is that the probability of the adverse hypothesis is considered to be many-fold times less than the probability of the benign hypothesis.

For chemicals, we may be dealing with the probability of the adverse hypothesis being sort of one percent, five percent. These are sort of the subjective estimations that people will give you as to -- well, for example, if you ask people — I hate to answer this people, but you ask them anyway. You say, what do you think the probability is that the cancer rates might double in the next 30 years due to all these synthetic organics we've been pumping into the environment in the next 20 years?

And they will hem and haw and you force them to have a median estimate from which they think it is just as likely to go more rather than less.

And I've done this a few times and it comes out somewhere from 1 to 5 percent, which is high considering the downside risks.

On the other hand, I have a feeling in the nuclear safety area you are willing to push these priors way down many orders of magnitude less than that.

gshHEE

So we're talking about what kind of prior that you are willing to live with. So that is what I mean. These asymmetries may be more striking in the nuclear power field than they are in the toxic chemical field.

Let me just stop there and engage in a discussion.

DR. OKRENT: I have a question. You said that you were working on unreasonable risk, I think that that was the term. And I was wondering if you were going to quantify unreasonable risk for me in some way.

DR. PAGE: I think where this line of analysis goes is, first of all, it says it will be unlikely that we will come out with -- it is possible that we will come out with a rule of thumb that says we will live with chemicals that have risks less than 1 out of a million. But that is a rule of thumb.

That might be like setting the traffic lights in the city. It is possible. It could come out that way. But I think it's unlikely that it's going to come out that way. I think it's much more likely that it's going to come out to be a statement that says that the standard of proof under which you required this much evidence rather than that much evidence in order to take a precautionary action is dependent upon some sort of critical ratios which have to do with your feeling on the ratio that the cost of a false negative to the cost of a false positive, your feeling of the priors that has to do with the level of suspicion that you have on the chemical before you go into the process of gathering more information on it and how sensitive the test is, how favorable your trade-off curve is between false positives and negatives.

And that has something to do with whether or not you're going to require a half million test or a million dollar test or a $2000 test.

So that I think what is going to happen is that there are going to be qualitative packages of what needs to be done in what situations that are not going to be terribly formalized. They are going to be judgment calls, but they are going to be guided judgment calls.

DR. WILSON: Coming back to something that you are saying, when Elizabeth Anderson was talking about this false positive and false negative question, a thing which everyone said where it is not found to be carcinogenic and people suggested that they haven't looked carefully enough -- if one looks through the data on this question of people just saying, it is carcinogenic, it isn't.

And it is only just recently that this quantitative thing has been answered at all to actually get in there. Is that actually getting into the level of the unreasonable risk question? This risk is 10 to the minus 10. Even leaving yourself a factor of 100 or 1000, it really isn't worth

gshHEE

bothering about. But this risk is 10 to the minus 1, and I have really got to move hand and earth to get it done.

Is that sort of thinking likely to come about?

DR. PAGE: I think so, sure, and even more so. I mean it may turn out that if you are looking at -- here's an example of what we're talking about. Suppose that you're looking at a medium level carcinogen, maybe benzene as an example, which has an extraordinarily large exposure.

Then for that kind of carcinogen, you may want to not live with a 5 percent significance level in the test. In other words, you do a test and you find that it is carcinogen of this potency with a P value of 10 percent, as opposed 5 percent.

That may be enough for a precautionary regulatory action because the exposure is so high and the power of the test is low for the potency that you are concerned about for regulatory purposes.

So in other words, this potential ratio of cost of wrong decisions either way helps guide the process of what kind of sensitivity makes sense in order to take the precautionary action.

DR. LAVE: That would say that you would not do a 50 rad test.

DR. PAGE: That's right. Let's say that you want to have a more expensive test with a more favorable

gshHEE

trade-off all along the line and then still adjust your fine-tuning the test between false positives and negatives.

So instead of coming out with a single one rule of thumb, I think that we're going to get a number of rules of thumb.

DR. OKRENT: It seems to me that what you've been discussing is what level of information and what level of uncertainty goes with what anticipated level of risk.

And I would argue in fact that in the nuclear reactor game, the same thought processes are gone through, except they are not so labelled.

There are some devices that you have to build and test and test and test to show that they are going to work. And there are others for which there exist general design criteria. And the designer says, yes, I designed it according to these criteria.

That is taken as probably good enough with a minor audit. And in a sense, these are related to how important you think the function is and what happens if the function fails.

So, what I'm still interested in is whether you see a chance of defining unreasonable risk, unacceptable risk, acceptable risk. I don't care which adjective you put in front of the word, but is there some way you see of quantifying one of these in the field in which you are working, and with an uncertainty band, if you so wish?

DR. PAGE: I think instead of having an absolute

gshHEE

number that says one out of a million, that's the sort of thing that you're asking about, whether or not things might tend toward a number like one out of a million is acceptable, one out of 100,000 is not acceptable.

Is that right?

DR. OKRENT: That is being done implicitly by the actions that are taken anyway. I mean, we looked at a list of things that might be in water and so forth. And if you are accepting this in water, whether you say it or not, your best information is that this is what is going on.

DR. PAGE: Well, I think that in EPA, especially in the radiation program, there is an effort to try and develop a rule of thumb that is sort of absolute like that.

But in the office of toxic chemicals and substances, this is less likely to happen because of a mandate to do upfront balancing of cost and benefits. That is written right into the legislative history.

So that the kinds of rules of thumb that are likely to emerge when you put all of this together and you worry about the cost of information I think are going to say something about what standard of evidence is appropriate in this situation and what kinds of mistakes we're willing to live with in the aggregate, that kind of thing. And how long, if you have to do a two-year cost and benefit analysis before you come out with a regulation, this may b too long. So

gshHEE

what else can you do?

I think what I am suggesting is that the path of this kind of approach which really is an application of worrying about incentive compatability and this sort of thing is to worry about what kinds of incentives need to be created so that the system will work better sort of on its own.

And it seems to me that this is very much the spirit of the Kemeny Report. It says that we spent much too long worrying about will this valve fail, sort of without humans around, and not enough time worrying about what kind of incentive structure do we need to have so that utilities are going to have trained operators, or so that they're going to have the valves and alarm systems and have the alarm systems as hierarchical so 100 alarms don't ring all at once, or so that you can see the meters, they are not all hidden from view, or all are not hidden from view.

DR. OKRENT: You see, I can find an equivalent analogy in the practice of medicine.

There was a time in the '20s when doctors were rather careless about prescribing radioactive materials as health improvements. And in the 1950s, they were rather careless about giving children radiation treatments.

And they now, I think, many of them are more cautious, but not necessarily all because there are a certain number that give X-rays to protect against malpractice suits, and so forth.

I have little doubt there is a change in people's thinking after they know more. I assume that in the practice of, for example, looking at toxic chemicals, things have evolved and there are some things that, as you said, were thought to be negative or false negatives, and so forth.

And you could, I think, look at the nuclear reactor business and say, they had a false negative because they thought that by, in fact, putting in lots of equipment and not relying on the operator, they had protected themselves against something.

You can see parallels, I think, in each technology. And I guess I tend to expect them. I guess you would like to have perfection, but I haven't seen it, as I say, in the practice of medicine, and I've been in the middle of some of the errors that the practice of medicine has made.

DR. LAVE: May I paraphrase something that Toby said to make sure it is clear?

I think that in part of the agreement, or at least part of what he was answering was that if the general perception is that nuclear power is totally unnecessary, then the acceptable risk level is going to be enormously different than if the general perception is that nuclear power is really necessary for our existence.

DR. OKRENT: I agree.

DR. LAVE: But then when you say what is the acceptable risk level, that is being postulated on some general view as to what contribution nuclear power will make toward our social well being.

Insofar as that changes rapidly, for example, and you can certainly expect the perception as to what the acceptable risk level will be will change rather rapidly.

DR. PAGE: That can be sort of added to, and I was just reading through the Three Mile Island report this morning. One of the things they say that struck me is that the accident itself, even though it didn't go its full path and people did not get heavily dosed, was clearly unacceptable from a social policy point of view. And another Three Mile Island accident in the next few years would also be unacceptable in the sense that it would have enormous political ramifications and what would happen to the industry.

So I guess part of what I'm saying is that a notion of acceptable risk has to do with sort of the individuals that may suffer because they are getting individually dosed with radiation because there has been some accident.

But another part of the notion of acceptable risk has to do with the entire energy program and what kinds of accidents are acceptable in a political sense of setting

gshHEE

down the industry or keeping it going.

And sort of the flavor that I got from the Three Mile Island report was that we have to look at the latter.

DR. LEACHMAN: I'm curious about what Mr. Lave is saying and what you, the speaker, are saying. Are you talking about the headlines here or are you talking about some realities when you're talking about whether nuclear energy is needed and how people perceive these things?

What are you going to do instead? What are the comparative risks? I don't know what your points are.

DR. PAGE: Well, the point is that if people believe that --

DR. LEACHMAN: What is "people." I don't understand what you mean. Are you talking about a public opinion poll? Are you talking about the Union of Concerned Scientists? Are you talking about what the knowledgeable scientists believe?

What is "the people"?

DR. PAGE: I'm really talking about the entire legislative process in the sense of whether or not a moratorium gets voted in Congress or not.

DR. LEACHMAN: Well, that is decided.

DR. PAGE: Well, it's not decided because there could be another Three Mile Island accident tomorrow.

So the question is what is the level of safety that we're

gshHEE

willing to live with in reference to the likelihoods of other partial accidents and what the likely public response to it is, as it is filtered through the political system?

I mean, those are real. The attitudes are real. And if people get very upset and they close down the industry, that is a real decision.

DR. LEACHMAN: There is a real indication a couple of years ago with a vote, and it was a legislative consensus 2 to 1 against shutdown.

Now I'm not sure what people you're talking about and what these opinions are that you're speaking about.

DR. PAGE: The impression that I got from reading this Three Mile Island report is that the repercussions of another Three Mile Island accident might reverse that vote.

So this becomes part of a calculus as to what is an acceptable risk. And especially if you believe that nuclear power is necessary because there are no good alternatives.

Then that puts a different flavor on what safety needs to be done, so that there won't be these political repercussions.

DR. LEACHMAN: I'm not sure what context you're sayning that in. Are you saying in comparative risks or something else?

DR. PAGE: What I'm saying is we tend to think in terms of sort of the real risk of people getting hurt and

gshHEE

how many cancers there are going to be and that sort of thing.

But there's also a political risk that has to do with Congress voting down the nuclear industry.

Five years ago that would be considered sort of inconceivable, but not it is conceivable.

DR. WILSON: That is the standard, in a sense. That distinction is not much different -- on the second type of risk is the sort of decision analysis that has been taught at Harvard Business School for some years.

I mean, that's the standard business risk. And the type of risk that we're talking about, risk to life and limb, is a different one.

DR. PAGE: I should keep those separate.

DR. WILSON: Yours is an older type. Businessmen are taught to face that type of risk.

DR. PAGE: Maybe I wasn't clear enough in separating the two.

DR. LAVE: Toby, you were being a little slippery, I think, talking about some of the asymmetries, particularly about the path or what is the probability of a chemical that was classified as a carcinogen now being exonerated versus a chemical that was not classified as a carcinogen now being classified.

And I think that your point is well taken. But one has

gshHEE

to be careful in applying the point.

In particular, I think it is extremely unlikely that a substance which was once classified as a carcinogen would be subsequently exonerated. Takes cyclamates, just for a couple of reasons. One, it is clearly an uphill battle to try and reverse it later on. That is the burden of proof is clearly on you to prove that it is not a carcinogen, which is probably an impossible task, as they are finding out with cyclamates.

And so, in general, there's no additional testing that happens.

Secondly, I think that there is always an natural path of some substance being not classified as a carcinogen. That is, if we have a new chemical come on the scene, of course it is not classified as a carcinogen until we get some evidence on it.

So that as evidence begins to accumulate, things will be classified as carcinogens.

I think in general your point is well taken about looking at the decision calculus. But the asymmetries are much more difficult to account for.

DR. PAGE: That is a good point. Basically, what you're saying is that not only do we have to worry about the discovery of these false positives and false negatives, but we also have to worry about the probability of discovering them and the means of discovering them.

And I think that this gets down to sort of what I wanted to talk about, which is that when you start asking these kinds of questions and you start asking how can I better discover these previous past mistakes, and how can I better anticipate them in the future. And then we get into the design question a little bit. I'm sort of looking at the design question as sort of equivalent to trying to find out whether a new chemical is a carcinogen ex ante, because it is just hard to tell whether a chemical is a carcinogen ex ante, just as it's hard to know in Three Mile Island or in Brown's Ferry the cables were lined and they should have been separated.

After the fact, it is easy to see. But before the fact, it's real hard.

DR. WILSON: Isn't the problem there both in the carcinogens and in the reactor questions a problem because you haven't specified properly your boundary line?

Now I would like to take as a naive principle that everything is carcinogenic as a starting point and that they may be, even drinking water is probably carcinogenic, but very low potency. And it's just a question of amount.

If you take that starting point and ask yourself, what is the number by which you have potency or some such thing, the slope which Elizabeth Anderson was talking about where you

gshHEE

had to reduce it to this particular discrepancy and it normally disappears.

e - 15

1573 224

CR 8493
HEER
t-16 mte 1

In reactor safety, the same thing comes up. If you are trying to say that something which appeared to be safe has now proved to be unsafe, do you always find that was the case? We would always find something starts by being safe, but it has never been in use. There's no experience with it. And then it automatically becomes unsafe the first time it fails.

So if you start again with the viewpoint, which I think in reactor safety we tend to start with, that any gadget is potentially unsafe, it is a question of the level of the failure level at which you begin to start worrying about it. And again, I think that particular problem disappears because it is -- I am essentially arguing for always putting upper limits on something when you don't know it, rather than just saying it is not there.

Does that sound right?

DR. PAGE: I think there is a lot of sense in what you are saying. I don't know -- I don't know whether -- I mean, the impression -- you see, I guess I'm just trying -- I'm sort of being swayed by this Three Mile Island report. The sense that I got is that people in the industry had the feeling that nuclear power plants were safe because there had been no accidents that involved people for a long time. And so there was a certain sort of complacency within the industry. And when an outsider said, we have to make these plants safer, then the reaction was, well, they're already safe and there's nothing that can be done.

But in reading the Three Mile Island report, it sounds like there are all kinds of ways they can be made safer. Some of these ways are at low cost, and they arise because we've been looking at just one aspect or we have been concentrating on one aspect of the problem and neglecting other aspects, which are sort of up for grabs.

There is a lot that can be done because it has been relatively unexploited, and that is in the design side and the training side. So I look upon that sort of as an optimistic, hopeful --

MR. LAVE: Although there is another point. TOSCA was an attempt by the Congress to switch the burden of proof of trying to correct bad situations that had occurred and trying to prevent situations that did occur, as was OSHA. And getting back into your basic framework, whether it makes sense to go for prevention rather than early detection depends upon what is the cost of prevention versus the cost of early detection.

And if it turns out that -- suppose it turns out that one chemical in 10,000 is a carcinogen. Then I think the judgment will have been wrong by the Congress and TOSCA would have been wrong from a social viewpoint. It would have been better to try and indulge in early detection, to count the few dead bodies that arise and say that's too bad, rather than to go through the elaborate steps you're going through. And I think that part of what the nuclear people keep on saying is that they have always felt under the burden of prevention rather than early detection. So it is clearly cheaper and easier to find faults when you are engaging in early detection, rather than the prevention phase.

And it may be that part of what the Three Mile Island report does is to show that the nuclear people were searching under the light rather than looking at these other areas that they couldn't quite deal with, and may have been unaware that they were searching under the light.

DR. OKRENT: I wonder if I could turn the discussion back to the kind of thing that Richard Wilson was bringing up a little bit ago. Is there, in your opinion, some level of risk that is small enough that if you know it with some degree of uncertainty to be stated, this is acceptable for society to impose upon somebody who receives no direct benefits, as may well be the case for various of the chemicals that you are dealing with?

DR. PAGE: If we adopted an acceptable risk of one out of a million and everybody is dosed by this chemical, that means 220 cancers per year. It certainly seems to me that there are many kinds of chemicals for which this would be an unacceptably high number of dead bodies in order to get the benefit.

And I can think of formaldehyde and toothpaste as an example. If that raised the risk to one out of a million for the entire population, it just wouldn't be worth it to get that nice, stingy taste in toothpaste. There are other ways you can do it.

I think there are times when one out of a million risk is far too high a risk. On the other hand, I think you can probably easily think of situations where we as a society have been willing to live and appear to be willing to live with a risk of one out of a million or more than that, because the benefits seek high in comparison.

DR. OKRENT: If I can paraphrase what you just said, the one in a million multiplied times 200 million --so this is where everyone is exposed -- imposes a large enough societal cost that there should be some benefit that justifies this. If there were a limited number of individuals exposed to the one in a million -- let me say a million instead of 200 million -- would that change your conclusion? Or 100,000? Or is there some -- would you be willing to say, well, 100,000 have this exposure, but the rest of society is much, much lower?

And then, would you still look strongly -- I'm just trying to see, is there some number that is low enough that, if the total societal burden isn't large, from an individual point of view you would think it was all right?

DR. PAGE: Well, I think the total number exposed is certainly an important consideration, and also whether or not -- if the exposure is voluntarily assumed. I think that is an important consideration. And if the market in which the risk is assumed voluntarily, if it is a good market and with well-informed people and people who are getting hazard pay and that sort of thing, then that -- and the market is really working and people aren't forced into it because they happen to live in Appalachia.

DR. OKRENT: But generally speaking, we're talking about places where people are ill-informed and it may not even be well known.

DR. PAGE: Well, I think these are all parts of what you mean as to whether something is acceptable or not. I think a risk that is imposed from one generation to another generation is different, because it has equity aspects to it, from a risk that is imposed by one person to himself as he drives his car down the highway.

Because of these things, I think it may be a mistake to get a simple number, and it may make more sense to, first of all, to try to figure out what the ingredients are and not just a shopping list, but how they actually interact, one with the other.

But I don't think that is where the real payoff is. I think where the real payoff is is in trying to make the incentive system work better, so that some of these choices aren't quite as imposed and nasty as they appear to be now. What I'm saying is that it appears that there is a lot of give in the system that could be exploited if we were to attend to the level of incentives and institutions, as opposed to accepting the menu as given and then deciding whether or not we're going to live with it, up or down.

DR. WILSON: The incentive system working better is something which I think most of us feel very strongly about. It is important. But there has been very little achievement of that aim. When you set a standard, it tends to be set in concrete, and then all incentive is just absolutely stifled. So the real question is, how can you achieve that? I mean, do you set a standard which can be moved up and down by some sort of procedure every few months -- or every few years, I mean -- if more data comes in, has a procedure for reviewing it?

And, associated with that, should you set a standard, an extremely tight one if you have no information at all and we just don't use the chemical; and then gradually relaxing as you get more and more information? I mean, at the moment -- well, it's not quite true. But in many cases we have no regulations at all if we have no information.

DR. PAGE: Well, I think these are the right kinds of questions to ask. And in the chemical field, for example, we have some of the same problem as you have in the nuclear field, where a lot of the regulation is done by industry and reported back to the central government, and the central government accepts, is in the position where it almost has to accept the word of the industry that these things have been done, whatever they are.

The situation in testing, the industry does its own testing of chemicals and the validation of tests. The incentives to do a good test rather than a bad test right now are backwards, in the sense that you have a better chance of getting a chemical accepted if you do a bad test.

Now, how are you going to turn this around? Well, double-blind testing, validation of testing, strict liability for the tests, and such, like the SEC has rules that hold the officers of a company accountable for bad information submitted in the financial data.

You can begin to think about what kinds of institutional mechanisms might lead to better tests and validation of the tests. And we might get out of this awful position of going and finding that enormous numbers of tests on pesticides are worthless.

DR. WILSON: Well, the ones you have described to us are one which are in a regression of a never-ending sequence of checks and checks, and not automatic feedback systems. I mean, I would like to have an automatic feedback so you could connect the loop once, and then it automatically stabilizes. At present the loop is completely unstable, and we have to put stop points here just to stop the whole thing from going to pieces.

And you are just -- the list of things are just a bunch of stop points and not a system to make it basically stable.

DR. PAGE: Although a third part of the testing scheme could have this feedback, where you end up paying people off or being bountyhunters, essentially.

DR. WILSON: That is why I like the idea of something being automatically treated dangerous right from the beginning, and there is some level of risk which you can accept, and then you can go away from that level by doing more and more tests. And then you climb up to here. And you might then say, well, no matter how safe it is, we can't do it.

DR. PAGE: Another example of this is in the design of the tests. The safe dose may be coupled to essentially, to the confidence level, the error bar in the result. So you get rewarded for a sharper test. And that was one of the motivations for the technique back in the beginning.

DR. OKRENT: Other questions for Dr. Page?

We have been more than generous in posing questions.

(No response.)

DR. OKRENT: Well, thank you. You are welcome to sit down around the table. We are going to have some general discussion, and you may join if you are so inclined.

The thing I would like for us to do at this point is to get comments from our consultants in particular on what would be fruitful things to do, assuming that we still have a goal of trying to develop some kind of possible approach to quantitative risk acceptance criteria that we could present to the full Committee, that they might buy, that they could present to the NRC for consideration.

I think I have indicated that the Commissioners themselves have now indicated that they feel a need to move in this direction. And I think we heard from Mr. Von Thun that there are others who are moving in this direction. We see that, for its problems, EPA is developing approaches.

Who wants to start the discussion?

DR. WILSON: One thought to me is that the biggest single gap at the moment -- and I don't know how to fill it -- in thinking of how one might use -- there are some cases where one can obviously, explicitly or implicitly, use risk assessment, is -- and it is both a gap and a characterization, and incentives -- is after you've thought about reactor safety, how do you make sure? Because almost certainly, somewhere in there there is something cheap you can do to make reactors safer, and on any cost-benefit ratio it will be cheaper.

I don't know what it would be. If I knew what it would be, I would have told somebody what it would be, and it would be done.

And the question is -- the interesting question is -- I was thinking of an analogy on the radiation things of the 20s, when all of the medical people were saying: Here are medical X-rays. The diagnosis we can do with them is so beneficial that it hardly matters to us what the risk is; and the benefits are so huge, we should go ahead and do them. And there were one or two people whispering from behind the closed doors, where they had been pushed off to the side: Yes, but you can probably get the same benefit with much less risk by using sensitive film and not covering off the parts of the body, and all of those things we know about now and we all know of.

We've got the exposure factor. It took a long time, about 30 or 40 years, to get that point across. And I think the major thing we're talking, we're trying to get across, is how to get that aspect of looking at those unknown things.

DR. OKRENT: Well, I would suggest that you could have an ALARA and I have said an AGARA, as good as reasonably achievable, criteria. In other words, you may have to meet some limit. Whether it's a non-acceptance limit or an acceptance limit, I don't know. But after that, you could say you would still have to make such other improvements.

DR. WILSON: But that is if you think of them. It is the forcing -- I mean, some of the improvements of Three Mile Island, for example, that the Kemeny Commission report comes out with are so obvious to us now. It is unclear to us -- I mean, if anybody really thought about them as clearly as we now think about them, they certainly would have got done. But nobody did.

And I think that is the biggest single gap.

DR. LAVE: Well, Dave, I think that the problem is -- or at least the problem as I see it -- is that asking the question of what is an acceptable risk criteria is not going to be a fruitful way of proceeding; that part of what Dick was just saying and part of what Toby Page was telling us a little while ago, and part of what Skovic has, and I think that very nice piece that was included, is that there are a whole set of other concerns, that it is not quite as simple as deriving a single number.

I mean, here in the last few minutes we've been talking about uncertainty and how it is that one could reduce the uncertainties surrounding any risk. We've been talking about feedback mechanisms. We talked at the first meeting about verification of risks, that is, if you have some estimate, what in the world does that really mean? Can anybody really believe it?

We have been talking about benefits. And it seems to me we have a whole nexus of problems that are just too complicated to be expressed in what is a single risk number. I guess if I were being a little bit playful, I would probably answer the question you posed to Toby by saying: Well, I guess that 10^{-20} is a risk that I would accept, even though it was of no benefit to me, if I was absolutely sure that it was no greater than 10^{-20}. Now, are you happy because I have given you that number?

I mean, I think not. I think that the number is surely going to be high enough so that the costs that are associated with achieving it are going to be out of all bounds.

DR. OKRENT: Well, to answer your direct question, I would say no, your proposed number I would find to be unacceptable.

By the way, I did not say there should be a single number. I don't think you heard me urge that there should be criteria that involves a single number. And I am not proposing that there be criteria that are independent of benefits or whatever.

I am asking, though, if one is going to try to develop risk acceptance criteria which are quantitative in nature and which include whatever other aspects they should, because that's part of the framework in which they are cast, how should we go at it and what is a possible approach, or however you want to put it?

So I choose not to set up that straw man of a single number and knock it down. I could, also.

DR. LOWRANCE: I gather when you say "risk acceptance criteria," what you mean is criteria that would be useful in managing the future of nuclear power, or at least the reactor side? Is that what you really mean by, quote, "risk acceptance criteria," close quote?

DR. OKRENT: Well, it obviously could mean different things to different people. We heard Mr. Von Thun earlier mention that they are beginning a process of trying to lay out for top level of management in what was the Bureau of Reclamation, which now has a new name I cannot remember, what their best estimate is of the risks. And they also have comparative risks of different approaches and so forth.

So when they make a decision, this is part of the information. But they clearly included the fact, at Jackson Lake, that this was not a body of water with no benefits. There clearly were benefits associated with this. And so, they in fact were buying risks under any of the approaches, even the one that was nominally not a risk for failure mode one, but there still was a failure mode two.

I think in fact one can argue from today 'til tomorrow, as it were, about how hard it is to develop risk acceptance criteria. The regulators, the people who build things, are still going to build things that are imposing risks, or standards that impose risks. And my own feeling is it is better to have them laid out as clearly as you can, either as Von Thun said or as we are seeing the EPA is trying to do, and, let me say, as the NRC should try to do. And they can say: Look, this is what we're trying to achieve. And they may have to say: This is what the best estimates are, and these are the uncertainties, and here are the residual risks that will exist if and when they meet the design criteria the way we think they should be. And it is presented as the total package of what we mean when we say these reactors can be operated without undue risks to the health and safety. And if you the Congress think that is wrong, tell us what you think should be different; or if you the public, or however it is.

It would seem to me that would be preferable. In fact, it might have averted a lot of the difficulties that exist today. Not all and not some of the technical issues, but you could still omit some of the technical factors.

I don't want to mix two different questions. How do you treat design errors, about which I've been worried for many years. It's not that I think they are unimportant, but I don't want to mix the two things.

e-16

DR. WILSON: There are about two or three places where it seems to me that one should be able to develop over a period of time places in NRC where risk calculations should be useful. First, it is not clear to me, if you develop a standard for anything, that a risk calculation is necessarily explicitly, although it should be done implicitly.

I mean, we talked about these standards EPA is developing. Elizabeth Anderson never mentioned that this was calculated solely on the mathematical equation going from a risk calculation to the standard. And in the same sense, the radiation standards, of course, are the oldest ones. And in a very real sense, all of the criteria which the EPA has talked about have already been identified by ICRP in setting the radiation standards, of compatibility with the background, and a rough calculation of risk on the linear hypothesis, and small increments -- all of those things are in there.

So, although -- and they are not explicitly mentioned in the standard-setting pieces of paper or -- in fact, the only words on the application of them are these words "as low as reasonably achievable, economic and other factors taken into account," and all of those nice phrases that get added onto the standard.

So, I think it is important to emphasize that a

pv HEE

risk calculation procedure should be backing up a standard and sort of be implicit behind it even though it is not explicitly used in the standard.

The other places where some calculations of risk, calculations of benefit, and maybe even comparison of risks and benefits might be important, are the continual tasks that the NRC gets on what it does when it suddenly gets a new piece of information which is important for the question of whether the present 72 reactors should stay on line or maybe a subset of them, those General Electric boiling water reactors or even those boiling water reactors which don't have the special jet pumps.

Now, every now and again a piece of information comes up which one hadn't expected. I was thinking, particularly because it affected us locally and affected my personal electricity bill, of the shutting down of reactors this spring because of earthquake hazards. Now, earthquake hazards are not something very amenable to risk assessment, of course. Nonetheless, we could have made a sort of rough estimate on it, and I am not sure one was done.

But I tried to think about it a little bit myself because while I was suffering the pain of increased electricity bills, the question came: should one, for example, shut down after two weeks, give them two weeks? I mean, what is the risk of — granted, it is a thing you really want to look at this question carefully, you want to put incentives to have someone look at it quickly, is shutting down the whole reactor too much of an incentive or just a "fix it and give us some information in two weeks' time or we shut down." Presumably, that could be addressed by risk calculation.

DR. OKRENT: I agree with you. In fact, there -- in our letter of May 16, in its second sentence, said, referring to quantitative safety goals, "This could be helpful, for example, in developing criteria for NRC actions concerning operating plants."

So, we are conscious of the possible use of such criteria. And I don't think the NRC staff now have a yardstick. So, even if they had a yardstick, it might encourage them to see are they able to measure against the yardstick, which you cannot always do. But in that particular case, they might have been able to.

DR. WILSON: Well, again, it is the uncertainties in the calculation that have to be brought in. In earthquakes, the uncertainties are huge. But nonetheless, it was clearly a case that the retrofitting -- you know, it is again — the retrofitting clearly is not because they won't — couldn't stand an earthquake at all. It is because if the calculations were wrong, they wouldn't stand quite such a big earthquake as they were designed for. And things of this sort could all go into it.

So, I think those are quite important financially, and the things I am most interested in, however, I think, we ought to try to find methods of using a risk calculation to simplify the regulatory process. And that, I think, is perhaps one of the most important things — not necessarily major regulatory processes, but — and that then has to be done in what one would, as scientists say, in a form where there is a negative feedback so the incentive is automatic.

One of the problems with one place where it is brought in in the standard after the slowest practicable hearing was the statement that "if you can reduce the exposure for" -- what is it -- "a million dollars a man-rem, then you should do so."

There are no particular incentives in there. It didn't replace any of the other low-level radiation things. It was added to.

One might, for example, talk about such criteria if on the following basis — I am not sure it was the right one, but it was the one I thought about for that particular -- it is going back on something which has been done in a similar criteria one might say instead of setting 10 millirems per year at the site boundary for the criteria, set an absolute maximum of 100 millirems combined with a mandatory regulatory measurement of calculation of what the

pv HEE

dosage is, a mandatory charge at the base of a million dollars per man rem to give automatically not only a demand that he reduce but a financial incentive — something of that general approach.

But is a combination. It is placing the mandatory standard which can produce a lot of work on that regulation. If you go down to 10 millirems and forcing the tech specs and you would reduce the amount of work, and we would have a simple procedure, which is simpler for a variety of reasons. One, it is the only regulatory -- part of it is subsequent to the event and not before the event. Those are the only general ideas that — I am not sure that is workable in that particular case.

DR. OKRENT: Well, I like the idea of a risk tax which, in effect, is what you are saying.

DR. WILSON: Right.

DR. OKRENT: I don't know whether that is part of a framework for risk acceptance criteria where there is a question of risk management. In my own mind, they are related, but in a slightly different box. But I am open-minded.

I am trying to encourage you all to think positively. I think the community or too long has been trying to point toward all the difficulties in developing an approach, and I think it is time to --

DR. WILSON: Well, I think the important thing about any positive statement is this: that I am always worried when someone says this is too uncertain to do a risk analysis -- because that is usually the case -- where I would like to see someone attempt a risk analysis more than any other case.

DR. OKRENT: I agree.

DR. WILSON: It's not because the number is going to be reliable or it's not going to be extremely uncertain. But until someone has attempted a risk analysis, set the set out, I am not even sure he has thought about the problem.

DR. OKRENT: I agree.

Bill, do you have anything?

DR. LOWRANCE: I am concerned with knowing where to take all of this. I have been troubled today as I was in the previous session. It seems to me that the most difficult aspect of all of this or one of the most difficult parts is not at the early end of how one describes the physical risks and comes up with some probabilities and consequences and one tries to bound the uncertainties and so on, but the "so what" question -- how does that then imply social action and management decisions or cutoffs or upper limits or lower limits or whatever.

And you just heard Dick say that systematic approaches should be used more in developing and backing regulatory decisions, say, by the NRC. That sounds fine.

But then the question is: for any newly discovered hazard — let's say, within the reactor — at what point, at what marginal return, do we then ignore the new signal or at which point do we start taking action? We just danced around that question all day and, yes, the Bureau of Reclamation is doinng systematic risk analysis, but it is not really looking at the "so what" point, at what point you stop building those big dams.

And as I said in the last meeting, I think if people really knew what the hazards of those big dams are and when a big one finally goes — I mean, a really big one finally collapses and kills one of the expected communities, 100,000 people or so -- that whole complexion is going to change. We're no longer going to be able to sit back and say, "For years the Bureau of Reclamation and others have been doing their analyses, and it is all working out well and we ought to do that in the nuclear area."

I don't know if any major technological area has ever done this and done it in an open way so there would be scrutiny and fairly informed knowledge -- not by the general public, whoever that is -- but by the decisionmaking public. I am just not sure how we got onto that point.

DR. OKRENT: There is a member of our audience.

MS. HALLER: Agnes Haller, of Babcock & Wilcox.

I came across this study not too long ago that was done by the Canadian Atomic Energy Commission. It is a report — whose number I don't happen to have with me — regarding the comparisons of energy sources and their risks. And it traces everything -- coal, nuclear, photovoltaic, any source that we might be considering for major production use -- through the mining and source, obtaining the source of the material to construction to production, all the way to waste disposal.

And I would like to recommend the report to you as a possibility for background information.

DR. OKRENT: Thank you. I think you are probably referring to what is called the "INHAB Report," around which I guess there is some controversy about the results.

MS. HALLER: I am sure there are questions, but it is a methodology that I thought would be useful.

DR. OKRENT: Dr. Shinozuka.

DR. SHINOZUKA: If I may, I would like to talk about problems associated with hardware in the nuclear problem. This causes a discussion away from acceptable risk criteria, but I think I would like to lead to that end of my comments.

Could we -- the power plant consists of electrical systems, structural systems, warehouses for the electrical system and the control system. The mechanical-electrical system obviously might become an initiating event of a very many number of possible event trees. At the same time, it will induce in some cases structural failures.

Now, when we look at the structural system, it might fail due to natural hazard, notably by earthquake. And by doing so, that is structural failures and that in turn induces failures of mechanical systems.

If you look at control systems, including the engineered emergency safety systems, if they fail, definitely they will induce all kinds of problems.

To make a long story short, all of these will definitely increase the possibility of radioactivity release, if not definitely lead to it.

I see three basic problems which ha e not been really dealt with carefully, dealing with these systems:

Number one, the interactions of these failures are not well understood at this time. For instance, if an earthquake hits, what is the probability that the control system fails, which will in turn increase the probability of failure of certain mechanical and structural systems? These are not well -- this problem has not been very well addressed.

Number two, human factors, particularly human errors in terms of design situation or operating situation. This is an important aspect which we really have

pv HEE

to systematically deal with. And I don't think I have seen really a systematic approach to deal with human factors.

Finally, uncertainties in evaluating particularly structural and mechanical system behaviors are really, really wise. I think we really should make an effort -- and when I say "we," I mean the general engineering profession -- should make an effort to, if not of course eliminate, to reduce the, let's say, confidence associated with such uncertainty.

I think all of these problems must be looked into systematically before we can really come up with estimates of probability of, let's say, radioactivity release of a specified amount, which are needed, I suppose, for the ultimate risk assessment we are talking about. These are some, I would say, practical implementable suggestions I can offer right away with the engineering community.

Of course, probably we will have to have some research funding to receive these researches, to see that they be implemented. But I think that that is really important. That is what I would like to say.

DR. OKRENT: Well, let me not comment at the moment. I agree with much of what you said.

Bill, you had your hand up.

DR. LOWRANCE: I would just like to raise my own continuing concern, and that is the question: was the Three

Mile Island accident a surprise to anybody who had been involved in the WASH-1400, Rasmussen study, and all of the modifications made since, the Lewis follow-up and so on? Was that accident fairly within the assessment fully at hand within the industry and the regulators? Or was it some sor of total surprise?

My guess, to answer the question a little bit, is that some subparts of it were surprises to lots of people and were not fully anticipated by the formal analysis; but overall as a huge machine most aspects of it seem to have been within the boundaries within the limits that had been drawn around the system before -- in the before-the-fact analyses.

So, then, I ask myself, "Well, so what? Do we need to do any more studies if we are getting better and better at the Rasmussen-type studies? That is, how will the machine behave and then adding operator error and better management of accidents and emergencies and evacuations and things of that sort? Or is there something else really fundamentally wrong with all of this?"

And it may be then that that is, when I am driven to sort of thinking about the effect of an accident in this case where no one was killed outright, is that what we ought to be worrying about? And how do we accommodate that balking effect, the fact that it is terribly disruptive

pv HEE

because it started down a very serious branch of the fault tree? Is that effect something we ought to be worried about more? And how do we do that? How does formal analysis help us with that?

This is just a large question I am just trying to raise that is somehow wandering around in my mind.

DR. WILSON: That brings up something that I did not include in my list of issues. But I think it can help. And that is: as a partial comment on that and also a partial comment, I think certainly I was expecting something the size of Three Mile Island to be coming up any time. I was hoping I was wrong.

But I must say I have read Rasmussen's report, and, of course, it wasn't exactly what was predicted by it in any way, but many people in the industry had not expected it. I mean, they really thought Rasmussen was rather pessimistic. And so, in that sense, it was a shock.

The other thing is also true: that as I think about it, what did the Rasmussen report do, and how could it be used? It was used incorrectly by two sets of people: incorrectly by the Commission and the industry in saying reactors were safe; incorrectly by intervenors who were saying it's a bunch of nonsense. But really, it was a rough assessment of where the industry was at.

But what it was not used as was a set of things which reactors are that safe if you follow these procedures. And you look at this very carefully.

So, there is one set of things where I hope something must have been done where the risk analysis is in fact being used and it should be used. It wasn't before Three Mile Island. And it was recommended by the Lewis committee and everybody.

And so I think we should continue to realize it both in itself and all its ramifications, and that is if Herbert Dekamp and his staff had read the 82 sequences of the water reactors that Rasmussen went through and said, "Well, I wonder how this Westinghouse sequence applies to the Babcock & Wilcox reactors," it is inconceivable to me that they wouldn't have realized that this particular sequence is a thousand times more probable in the Babcock & Wilcox reactors. It's not inconceivable to me that in that case they'd do one of two things: certainly, alert people; and, if not -- and badder still, it was not done until May, which was just the set point of all of the things, so it didn't happen so frequently.

Those aspects of risk analysis seem to me in management of industry, seem to me more important. That approaches the question I was saying that one of the points of doing risk analysis is to guarantee someone has thought through the problem. And in fact, one of the things that

Three Mile Island showed is that the particular staff had not thought through the problem, and this I had hoped: that if you have a thing like the Rasmussen report, 82 sequences, it is not that much work to ask everyone of the trained operators in the simulator training to think in terms of what do you do if that happens.

CR 8493
HEER
t-18 mte 1

DR. LAVE: That's a good point. But I think there is one thing there that might be unfortunate. Three Mile Island does not serve to disprove the Rasmussen calculation, but neither did they serve to prove them. That is, they are simply consistent with them, and one cannot say, aha, since Three Mile Island was consistent with them, therefore Rasmussen was right and various other people were wrong.

But let me try and be positive. Using that and what the other people have said to try and be positive -- and Dave, I'm sorry. I was reacting to some of the questions you were posing to speakers, which is why I was trying to ask you whether you had stopped beating your wife.

DR. OKRENT: Well, I did not feel it fair to push any of the speakers except Toby Page on the question of quantitative risk acceptance criteria. But I feel it is fair to push you all.

DR. LAVE: Well, let me try and do that. Trying to react in the spirit that you posed the question, I think that we know enough now to know that the framework in which that question has to be posed is more complicated perhaps than had been realized by the NRC at the beginning of the study. That is, it has several aspects about it which we now know to be important. And le. me just put a couple of them down.

That, first of all, estimating risks without being able to verify the levels of those risks is almost a meaningless exercise; and that probably at this point the most important aspect of risk assessment for the nuclear reactors has to do with some attempt to verify what those probabilities are, and not simply to make them larger or smaller or whatever it is.

Secondly, we need to know much more about the residual uncertainty once the risks have been estimated and verified to the extent possible, and that magnitude of residual uncertainty has to do both with the unverified part of the risk, since most of the risk will necessarily be unverifiable -- but we're also looking beyond that part.

The third part is what Dick has just been stressing, which is we need to worry about a different framework in which nuclear reactors are operated, so that reactor operators are motivated to think about whether it is worthwhile to go to the Rasmussen sequence or some other sequence, or worry about whether the level of training is correct or whatever it is; that is, that the incentives are properly placed so that somebody other than the NRC Commissioners are worrying about all of this in detail.

However smart the NRC is, it can never approximate the amount of smartness represented by all of those people out there in the field operating reactors every day, who, if properly motivated to worry about the right kinds of problems, are going to do a much better job than the NRC staff could ever possibly do.

So you want to get the incentives right. And then the last part of that is that any notion of acceptable risk is necessarily tied to some notion of the benefit that you are getting from looking at that risk, which not only has to do with the distribution of the benefits and the distribution of the risks, but also the overall benefit. And in this case, that really comes down to what is the cost of a kilowatt-hour foregone or the cost of a kilowatt-hour generated by an alternative to nuclear power, such as coal.

And I think only by getting estimates of all of these things and by looking at the parts of those estimates that one can verify, can one begin to get this answer of what then is acceptable, what should design or operational criteria be.

DR. OKRENT: I must say I agree that the points you raise are relevant. I think we discussed them before at previous meetings, well before the ACRS got involved in this specific activity. And we are very conscious of the problem of uncertainties. In fact, to me the problem of the uncertainties and your inability to be confident of what we can predict, unless you're willing to talk about relatively high risk levels, is a big impediment in this whole thing.

Maybe I'm wrong in trying to divorce what I will call risk management from risk acceptance, because obviously you all think it should be there. I am not opposed to it. Remember, Whipple and I wrote a paper in this area. I thought that maybe life would be simpler if we left it out. But what you tend to be saying is that you shouldn't be leaving it out, you should have it in there somewhere. And I am perfectly willing to be open-minded in that regard.

I should note, we started a little bit of our own effort, although it is a small effort. In fact, David Johnson and Bill Castenberg have been asked to try to give their estimates as to what extent did they think we could produce estimates of the individual risk from a nuclear plant, with what degree of uncertainty, or something that they would be willing to defend, even subjectively. But that is a small effort.

So, having said that, I want to come back to you in a minute and say, now what do we do. But let me call on Toby Page.

DR. PAGE: I wonder if this would be an acceptable dichotomy, that when you are worried about assessment of risk in sort of a descriptive sense, in trying to figure out what the probabilities are -- that sounds sort of like what you were pressing me on and other people, in terms of trying to describe the nature of the risks as they are. But when we begin to ask questions about what is acceptable or what is a reasonable risk, what are the criteria for risk, then we're really beginning to get into the normative questions. And when we get into these questions, the amount that we can change the risk by having incentive systems and institutions this way rather than that way becomes an important part of calculus, to decide what we're going to live with.

So that's one of the reasons why some of us are sort of fighting the idea that what you're calling risk management can be divorced from acceptable risk. If you said, let's try to divorce risk management from an assessment of risk as it exists today, then maybe we could sort of go along with that sort of dichotomy.

DR. OKRENT: By the way, I favor what I called an ALARA criteria, and I also, as I indicated, really favor something that is the equivalent of a risk tax.

But let me note, there are different ways of trying to get to levels of risk. One is to say, well, I will try to give the licensee an incentive. But the NRC could say, we will try to build nuclear plants underground. Maybe that way we will reduce the risk. They might be changing the cost by 50 percent, and I think that should cover it. That is a lot of money.

Society has to ask itself, is that where it should spend that much money, because you multiple it times a billion times X times -- so it is not something that I think the NRC should do lightly, if it were to do it, or the Congress.

By the way, sometimes, to put myself in a little bit of perspective, I do another kind of calculation about effects of radiation, namely, what is the effect of nuclear war. And I come up with the effect on the U.S. of 100 million early casualties and 100 million delayed, in Rasmussen parlance. And if you change the probability of such a war by a thousandth of a percent, if you increase it, that is still 10^8 over 10^{-5}. So it is 10^3 people expected value per year.

That is still a pretty big number, in fact. And who is there to argue that he knows what changes the probability of a nuclear war by one-thousandth of a percent? Certainly even oil supplies, you could argue, fall in that category readily. So it is a tricky bit when you start trying to put everything in perspective.

DR. WILSON: I always like to argue in terms of bounds, because so much of the talk people have is -- can be bounded. The comments that Dr. Lave had and criticisms I have had on risk benefit analysis is no good, you're just kidding us, and so on, but you can't do this, that and the other -- and they put up straw men and destroy them, because they are picking something which Lester Lave never said, and he never said the numbers were highly reliable. They said -- he said the numbers are this sort of magnitude if you calculate them this way. I very much doubt whether that's zero and whether that's half a million.

And if you start putting bounds on it, you get things which people can't disagree with. I think when you do something of the same on acceptable risk -- I mean, you talked about 10^{-20} and whether that was acceptable. Well, if I was being subjected to risk, it certainly is more acceptable to me to accept 10^{-20} than to pay for a 15-cent stamp to complain to my Congressman. I mean, that is a very simple risk-benefit calculation. I mean, much more -- $100 million per life is what I gain on that one, even with 15 cents per stamp.

So if one puts it in that ridiculous form, it becomes clear that some low level, like 10 to the minus -- I don't know. In this case, I think probably 10^{-10} could be argued in a fair form, in which everybody would agree, when you put it in a ridiculous form.

I even got Samuel Epstein to agree on the risk of carcinogens which is acceptable. But one goes by pointing out that saccharin -- one shouldn't bound waitresses serving in the cafes, because the risk is 10^{-10}. Well, 10^{-10} is what a critical risk analyst calls absurd, because it is small. But I think it is worth doing that, because that does give you a boundary point which is non-zero, and it immediately puts some of the things where you can immediately start thinking at a certain stage. And then you have to do your thinking between a range of 10^{-10} and 10^{-2}, And 10^{-2}, you say, look, if those reactors aren't safe at 10^{-2}, well, you certainly will not build them.

DR. OKRENT: You know, society imposes risks of 10^{-2}, knowing them. They are doing it in Los Angeles now.

DR. WILSON: Per lifetime.

DR. OKRENT: No, per year, on seismically substandard buildings, pre-1933 brick buildings with no seismic protection at all. I think the probability, even after allowing for you being there half of the time and so forth -- and the people don't know. They have not been notified.

DR. WILSON: Well, we had 10^{-3} in Canberra Island, as we heard in the previous meeting, and that's a pretty reliable number.

DR. GRIESMEYER: It is most likely not any smaller than 10^{-3}.

DR. WILSON: I think it is in Everett, Massachusetts, where the LNG ships come in there, and the population density is bigger.

DR. OKRENT: Well, in any event, we need somehow to decide in some relatively short time what would be fruitful steps to take next. Now, it may be that the National Academy of Sciences group will hold a workshop in this area. I don't know. We probably won't know until next year.

DR. LAVE: Don't hold your breath.

DR. OKRENT: I'm not going to hold my breath in any event.

So again, I urge you to think positive.

DR. LAVE: Well, without volunteering to do it, I think one could set out a document which lists the structure of what would be required to get some answers to risk regulation and risk management, and then worry about how it is that you could fill in that structure in order to then get something that was workable.

DR. MARK: I wonder, Lester, if you could say that again? Are you thinking of trying to -- you spoke of a structure -- make a pitch for plans you might make and the ways in which risk could be estimated, made available to the licensing or regulatory process, rather as Mr. Wilson was saying?

DR. LAVE: Yes, I think one could specify a structure which could wind up being a how to do it. Step one is blind, step two is blind. We need to know exactly these kinds of things, and these are how you go about knowing them, and these are the steps you have to take as you get from A to B.

DR. MARK: And these are the reasons why you should do it?

DR. LAVE: Oh, yes.

DR. WILSON: This is sort of a manual for the people who are trying to do it?

DR. LAVE: Let me be careful in the reaction. I think the manual is the end product. I think for right now what one could do would be to set down perhaps three pages of what the outline of that framework looks like, and then from there start trying to figure out how it is that precisely which pieces of information are needed for each part of the framework, and then precisely how is it that you go out and get those pieces of infor.ation.

So that I think it is not something that could be written full-blown at this point, but I think one could put down the structure and start asking the questions at this point, which I take it is what the four of us are supposed to be doing.

DR. WILSON: I think the one thing that I find -- different people work differently. But I find I can't ask the questions in any abstract way, without ever having tried to look at the specific problem. I always find me full of crazy examples about this, that or the other. But this is the only way I know how to think about it. And we could certainly set up examples of how one might address the question in certain specific examples that you have in the past.

I mean, this earthquake hazard thing; how one might have gone about that decision of whether to shut down immediately, later, or something a little more logically than it seemed. It may have been done, for all I know. I mean, I wasn't -- I have got good lines into the NRC, but they're not that good.

DR. MARK: You haven't found any logic.

DR. WILSON: Yes. In general principle, logic is not to be found anywhere in Washington.

But the point is that one could presumably look at that and find some example and say that should apply to this, and try and build on that. That is the way I would go about doing that.

DR. LAVE: I think that's fine.

Let me say that I too like to have specific examples. But I like to see if I can put down a general framework first and then find out whether the general framework makes any sense by testing it with specific examples.

DR. OKRENT: I will be happy to see an approach coming from each end.

Let's see. Toby Page?

DR. PAGE: Well, just sort of picking up this thing we said a little while ago, in this general framework it would be nice to see a chapter on how you go about evaluating previous risk assessment efforts. What has happened is that a lot of probability assessments are made, and then -- I don't know. I mean, I am an outsider, so I don't know. But my feeling is that they are not systematically evaluated and have a reality test run upon them. Maybe I'm wrong, but it seems to me that one very simple thing -- and maybe this cannot be done, I don't know -- is that if you take a report like the Rasmussen Report, which must have thousands of specific probability estimates buried here and there -- would this part fail and that part fail -- you go down this path and that path. Each one of these by themselves can't be tested.

However, if we have a bunch of these which are assumed to be independent, then we can add them up to a variable which you can observe. And by doing so, we can have two tests: One is whether or not the aggregate of these individual trials really adds up to what you would expect to see as an expected value of the binomial, because each binomial trial has a different probability to it.

And the other is, then you check the independence assumption, which is a critical assumption for a lot of these error trees. So you sort of get a feeling as to whether or not these independent assumptions make sense and whether or not these individual assessments are consistent with the observable kinds of variables.

DR. WILSON: If I may comment on the first thing, I would mention I put a thing on David's pile at lunchtime.

A research fellow and I are just writing a book on some of these things and have a chapter on just that. But it is in draft, and there is a copy on your pile there. And this is the thing I was going to send to you in the mail. So in a sense, that is an attempt at some of these questions. And the conclusion is that most of the attempts in the past have been full of problems, largely because I think they haven't asked the questions, and they have not been complete as to what they were aiming at in the end.

But I think that sort of thing, maybe we should -- that was a rather general thing we were coping with. But for the NRC, what we want to go through here is those aspects which are particularly applicable to the NRC.

DR. OKRENT: By the way, let me know -- I have been an advocate of the proof quality control under risk assessment. So I have to support your general thesis.

I guess I don't want us to get into a position of doing what NSF has asked the National Academy to do for them. They have asked the National Academy to help them out with a long-range research program. If that is all we can do for the NRC, I would say we have been a failure. Maybe that is the best we can do, but I would consider it a failure.

So in that context, I don't know how we deal with the problem that risk assessments, not only nuclear but on whatever, have been subject to uncertainties, systematic and otherwise, generally in the past. And a few have received careful review in detail.

DR. LOWRANCE: I would just insert my feeling that the NRC and some related agencies have really done more assessment than we seem to be giving them credit for. A lot of what has been done in my mind is not all that good, but there have been an enormous number of assessments, both in specific site proposals for reactors, industry, Government studies of all sorts on design problems and operator errors, and siting, and philosophy of siting and design, and it goes on and on.

Some of it has been done other places than NRC. I think we ought to give some credit and realize there is a lot of work to be drawn on; and then groups that have tried to pull it all together, have made some progress, and have come up with some kind of middle of the road assessments, as I understand it.

The Academy's CONIA study is coming out now with its overall reports in 23 different pieces or something of that sort. And maybe you want to speak to the basic assessment problem on the nuclear side. But I think there is more work that has been done than you are giving credit for. The Lewis panel in examining the WASH-1400 study I think was a good exercise.

DR. WILSON: That wasn't a risk assessment per se. I'm all for it. It was a review of -- a very valuable review of the Rasmussen Report. That makes the Rasmussen Report that much more valuable.

e-18

DR. LOWRANCE: It was a critique, a constructive critique.

1573 267

WHEE

DR. WILSON: I was wondering: as a matter of mechanism, the several ideas that I have now, I think, are things where NRC should and could be using risk assessment, which they don't at least appear, from my perspective, to be using, where they probably are on the back of Saul Levine's envelopes.

I could write down in a few pages those items and send them to you, and you could circulate them to other people.

DR. OKRENT: Well, I think that would be of interest. I would say it is not directly attacking the objective. I would like to see what these things are to see whether in fact there are some places where perhaps they should be that nobody has recommended to them.

DR. WILSON: What do you think is the objective?

DR. OKRENT: To try to develop some approach to quantitative safety goals and improve management.

DR. LOWRANCE: What are some examples, David? For example, what are a couple of things that we could urge someone to do?

DR. OKRENT: Well, you heard what George Kinshen proposed. He thinks he has proposed a set of quantitative safety goals for nuclear reactors. He has published this in the Institution of Civil Engineers, or something like this. There have been other proposals for quantitative safety

pv HEE

goals. You may think, "Well, they can't be posed that way. They have to be posed in some other way." What is the way in which they should be posed?

DR. MARK: David, I think I have been a little unclear. I am not sure if the rest of the group is. You have very frequently used the word "acceptable." That, of course, throws everything off the track, because "acceptable" -- is there such a thing? When you say "quantitative safety goal," that is a rather different kettle of fish. It is acceptable to Kinshen only, probably.

DR. OKRENT: You see, you can put out quantitative safety goals and say, "We plan to be using these." You don't have to call them "risk acceptance criteria." If somebody thinks this difference in working is beneficial, it doesn't particularly bother me. I think you are talking about the same thing.

DR. LAVE: David, I don't want to be churlish. But I was just trying to object to this notion that there is a schedule of that sort. I don't think there is a schedule of that sort like Kinshen laid out or a 10-9 or some other criteria. But it is necessarily more complicated than that.

DR. OKRENT: I wasn't saying it shouldn't be more complicated. Somebody asked me what do I mean by "quantitative safety." Well, I said, "Here is an examle that one man proposed, and he gave reasons -- brief, but reasons -- why he thought those were appropriate."

DR. WILSON: You see, Lave's constant comment is you can't make any logical sense out of risk calculations until you consider the benefits. He is constantly beating me on the head with that, and, of course, he's right. And that is why I bring up my comment, "Well, sometimes the benefit of not having to spend a 15-cent stamp is all the benefit I can itemize, but it is still a benefit." And he completes the logical structure.

I think that is one of the important things, because you can then stop somebody's arguing with you, and I think -- so I hadn't completely realized that that was Kinshen's aim, but, of course, very little of what we talked about today was directly approaching that. I mean, nowhere did Elizabeth Anderson really discuss how they picked up the 10-6 per lifetime, which I think the average on that really would start -- or the origin of that really started with the FDA and it --

DR. OKRENT: The primary focus of the talks today was to ascertain as we could what was going on in some other federal agencies that were working in some way in the area of risk assessment. I never seriously expected them to try to address the question of what should be a quantitative safety question.

DR. WILSON: That's a very serious problem for them, by the way, because EPA and FDA get into fantastically large problems at the moment.

DR. OKRENT: Well, I didn't really expect them to get into this problem with regard to the NRC except possibly the man who was from the insurers, and he was possibly the furthest of all of the speakers that we heard.

So, the purpose, as I say, was to see what was going on. I found it interesting to see what, for example, the Bureau of Reclamation was doing and where they stand now. I think that was a pretty giant step from where they were six years ago.

DR. LAVE: But it just seems to me that the next big public stink is going to be over those dams because of the vagueness of the numbers. I mean, there is an air of precision about them which is completely false.

DR. OKRENT: Well, he stated, in fact, to us that the uncertainties with regard to structural failure were larger than those of seismicity, in general. Those of seismicity are large. And he stated that also. The site he happened to deal with was one of those that you might say is a more fortunate one. You have some faults you can use to get a handle on what is the probability of a pretty severe earthquake for lots of sites. You are not in a position to use some strong nearby faults, and you are less able to say what is 10-4 and 10-5 per year earthquake.

DR. SHINOZUKA: Since this is the first subcommittee meeting I have participated in, I must confess I am completely confused. Are we accepting the kind of approach of structural reliability -- and that is what the dam safety program dealt with — that kind of approach we are going to accept by way of reaching the consensus as to acceptable risk criteria and the part of the analysis we have to perform?

When I talked about — when I talk about this, my overriding concern is the real uncertainties and difficulties we face when we have to come up with numbers. Sometimes I feel it is hopeless, but we have to come up with some numbers. And that person from the Bureau of Reclamation must have had the same problem: that "The liquefaction -- my god, that is a terribly serious problem." And I think that the variation would be 100 percent more.

So, my overriding concern is how we are going to resolve this difficulty now as a subcommittee. We are not addressing ourselves to this problem and just assume that somehow we can arrive at reliability figures, then establish risk criteria, taking the cost-benefit analysis into consideration. Is that what we are doing here?

DR. OKRENT: I think one could try to answer your question in different ways. One is to say you should be able to develop what constitutes proper safety goals or risk acceptance criteria or whatever term you wish to use for nuclear reactors without having quantified what you think it is for the current designs, but from other considerations, and then, if the current designs meet that, good; and if not, you have to change the designs to meet it. That is one possible approach.

Now, another approach and one I am sure that is taken throughout the regulatory business is: they find at some point to which they can go or some point where they are, they think and they decide what the thing looks like, and then they try to rationalize: does this seem to be okay, and if it is okay then they can develop criteria that match.

DR. SHINOZUKA: My feeling is I think we could talk about which design is safer, but I don't think we can place absolute values of structural reliabilities. That is where my difficulty comes in.

DR. OKRENT: It may depend upon how reliable you think things have to be. That's the only way I can put it.

Well, I said earlier -- and I agree with you -- the uncertainties are not going to be small.

DR. SHINOZUKA: So, supposing we agree upon acceptable risk interpreted in terms of, let's say -- probably many people don't agree -- but suppose we interpret that in terms of probability of failure. Can we then build or construct a structure which really can perform with that kind of probability?

I don't think that is guaranteed. I think we will have really difficult technical problems to build that kind of reliability into the structure by way of design.

DR. OKRENT: Again, I am going to have trouble answering in the abstract. 10-10, yes, we will have trouble. It will be possible. 10-3, maybe. 10-4 -- what degree of confidence do you want?

DR. SHINOZUKA: Well, I am simply trying to understand what we have to do.

DR. WILSON: Well, maybe one of the things that -- I think it is an important question because -- are we basically trying to set a level which is then for, you say, acceptable risk for a safety goal which, when explained to people, makes them reasonably happy? I mean, when I say "explain to people," I mean to politicians and economists and everybody.

So, then comes the question: what at the moment seems to be making a lot of people unhappy, and when one looks at it it is not the expected value of the risk people calculate, it is the possibility of very large societal consequences which they -- so large they can't envisage them, and I keep getting at that on all sides.

So, I think the safety goal can actually be -- it must be adjusted. I think it is probable that a less safe reactor would be more acceptable if it could be more easily demonstrated that it has a certain safety level or if it could be shown that the consequences of a failure are much less severe.

I think this is not usually thought of: "I can't prove it's safe, but it's the best thing I know how to do." Sometimes something is not the best, but you can prove it a little more readily, and I think this is an important question which goes contrary to most engineering judgmet.

DR. OKRENT: Well, you included two things, though, into your comparison. And I think that certainly the question of severe accidents is a factor in the discussion of nuclear power in what I would call the "risk." It is not just using the expected value, but you could build that into the risk acceptance criteria if you thought that is what should be the goal of the NRC.

DR. GRIESMEYER: That was suggested last time.

DR. OKRENT: It is possible there is nothing about criteria that automatically says "use a single number, use expected value," or what have you.

DR. GRIESMEYER: It also might be that --

DR. OKRENT: Or you don't allow for benefits or you don't provide incentives.

DR. GRIESMEYER: The framework wouldn't even have to be necessarily insufficient condition: it may just be a necessary condition for acceptance. Because it may be very difficult to come up with a framework a priori that adequately deals with the whole decision process.

But there may be things that you can say are not acceptable risks, so use them as a hurdle and then continue with your decision process.

DR. OKRENT: Well, can I ask: would at least the three of you who have been fairly active in the general area one way or another for many years, would you be willing to sit down in some corner and think about it and try to come up with what you think represents a possible next step or a possible general step or a framework or whatever it is, in whatever way you see it?

And maybe Lester would come up with something that resembles the framework he was talking about, with an idea of how you would pursue it. And I think that would be real nice, in fact, to have.

And if Wilson came in with an approach that was partly the same or partly from a different perspective or whatever --

DR. MARK: And make it clear at what level he would regard the 15-cent stamp is worthwhile.

DR. OKRENT: Well, more than that. Because that part of it, I think, we have a feel for.

But if you could try to do this in the next month -- and I know that is asking quite a bit -- it would really be useful.

DR. WILSON: It can certainly be done by January 1 by me.

DR. OKRENT: Well, when I say a month, that is obviously a round number.

DR. LOWRANCE: An NRC month.

(Laughter.)-

DR. OKRENT: Don't put it that way.

By the way, I am going to try to get Paul Slovick in joining us in the flesh.

DR. LOWRANCE: I really would like to enter the mildest complaint that we have today a meeting of consultants and not a meeting of the subcommittee, because I find this a little troublesome, because if this is to be a two-way interchange, that hasn't been achieved. And although we have had one faithful member and yourself here all day, it is hardly a committee meeting. Surely, it is not a quorum.

DR. OKRENT: Actually, though, David Johnson and Mike Griesmeyer are supposed to be working in this area full-time, and Bill Castenberg is working in it part-time. So, you do have the people who are trying to think actively in it. Kerr would ordinarily have wanted to be here, but he had a conflict.

What developed was the Commission thinks it needs comments from the ACRS on what it is supposed to do about the Kemeny Commission and some things like this, and so there are several members who are sitting in another room all afternoon. I think they are still there.

DR. MARK: And all of that blew up after all of this was planned.

DR. OKRENT: Yes. But if you would be willing to do that, I think that it would be, to me, a logical next step. If there is any way in which you can provide further information that could be useful, let me or Gary Quittschreiber know.

And I would like to talk -- sometime maybe I will get you on the phone, Dr. Shinozuka: I would like to discuss some of the things you have in mind and explore them in some more detail, if we can.

And we will also arrange to see that you get documentation. I don't know what Gary Quittschreiber has sent you so far, but we want to make sure you have the necessary background information.

Is that okay? All right.

pv HEE

MR. QUITTSCHREIBER: We have not received comments from any of the consultants on your frameworks.

DR. OKRENT: Well, if you would be willing to, Griesmeyer and I would appreciate it. You were given something that Mike Griesmeyer and I generated in November, and if you have any comments, we would appreciate it.

Well, thank you all very much. I will adjourn the meeting.

(Whereupon, at 6:25 p.m., the meeting was adjourned.)

* * *

1573 279