Safety Goal Project
(45 FR 71023)

NUCLEAR REGULATORY COMMISSION

In the Matter of:

WORKSHOP ON FRAMEWORKS FOR DEVELOPING

A SAFETY GOAL

SECOND PLENARY SESSION

DATE: April 2, 1981 PAGES: 1 thru 47

AT: Palo Alto, California

ALDERSON REPORTING

400 Virginia Ave., S.W. Washington, D. C. 20024

Telephone: (202) 554-2345

UNITED STATES OF AMERICA

NUCLEAR REGULATORY COMMISSION

- - -

PUBLIC MEETING

WORKSHOP ON FRAMEWORKS FOR

DEVELOPING A SAFETY GOAL

SECOND PLENARY SESSION

Rickey's Hyatt House

4219 El Camino Real

Palo Alto, California

Thursday, 2 April, 1981

The meeting was convened at 8:10 a.m., pursuant to notice, with George Sege, Project Manager, Office of Policy Evaluation, presiding.

PRESENT:

Mssrs. Beyea, Bradburn, Bridenbaugh, Burstein, Charnoff, Cochran, Darby, Eisenbud, Hutt, Joksimovic, Kouts, LaPorte, Lave, Levine, Lewis, Malsch, Lowrance, MacLean, Maxey, Mazur, O'Donnell, Okrent, Paige, Perrow, Salisbury, Sheldon, Slovic, Starr, Tamm, Wald, Zebroski, et al...

P R O C E E D I N G S

MR. SEGE: Good morning. Welcome to the second day of the NRC Safety Board Workshop. We are starting out the day with a plenary session devoted to interim reports from each of the three panels, for the purpose of letting each panel know what is happening in the other panels, and perhaps generate some cross-fertilization of ideas for the continuing discussion.

After the reports of the panel chairmen, there will be opportunity for other participants to offer comments and suggestions and viewpoints, particularly with respect to subject matter that is dealt with in panels other than their own, so that the further discussions today could progress in directions that would be helpful towards a good interfacial consideration of the subject in the plenary session tomorrow.

Before turning to the report of the first panel, Walt Kato has a couple of announcements to make. Walt?

DR. KATO: Good morning. Now that you all know who you are, would you please turn your namecards around so that the other panel members can identify you. Thank you very much.

I have got a couple of announcements. The recording firm, the reporters, have indicated that transcripts are available if you wish, however, the charges for the transcripts will be to you, not to the workshop, so I just

wanted to warn you that if you sign up for transcripts, the bill will go to you or your company, and so I am just so warning you.

DR. JOKSIMOVIC: How serious are the consequences?

DR. KATO: I don't know. I asked the reporter, but he wasn't sure either. There is an amount mentioned on the sheet, but he is not sure -- that the is the standard charge that the ACRS transcripts are given at, and he is not sure what they will charge. It is very expensive.

MR. SEGE: Thank you, Walt. I should have mentioned that there is also another item we have on the agenda this morning between the panel chairmen's reports and the discussion from the floor.

In response to our suggestion by Gerry Charnoff, as well as others, I asked Lester Lave to give a brief highlight report concerning the practices of other agencies with respect to safety goals, other safety regulatory agencies. This is based -- Lester's report is based largely on a study that he is in the process of completing under contract with NRC.

Now, I would like to ask Dr. Herbert Kouts to give the report of Panel A. Herb?

DR. KOUTS: Well, we arrived at a fair consensus on a lot of things, some of which are trivial, some of which are not, and some lack of consensus on other things, so I will try

to give the flavor of things in all categories.

We did agree that there ought to be quantitative safety goals in nuclear power. However, it was realized that it may be difficult to implement safety goals if they are constructed in poor ways, and it may be difficult at best to implement safety goals that are quantitative. Now, this will depend on the structure of the goals.

It will depend on the mechanism for determining compliance with goals. This may particularly at the outset of the establishment of quantitative safety goals reduce their value.

This is not a prediction. It is by no means clear that this will actually be the case, but it is something to bear in mind. It is a reservation to be kept in mind.

We agreed that the quantitative goals should have a qualitative overlay which should be understandable to the public at large, to laymen. One example of a qualitative overlay that was presented was nuclear power should be safer than any of of the competing ways of producing electric power.

This was by no means a unanimous view of what the qualitative overlay should be. There were some reservations on this, and there was in fact a range of views as to what the qualitative overlay should look like, but it was totally agreed that there should be a qualitative overlay, that the goal itself should be a qualitative goal, and the quantitative

goal should be a means of implementing the qualitative goal, interpreting it.

It was believed that we probably would not be able to jump right into implementation of a quantitative safety goal even if it is established, and we should not jump right into it. There ought to be a transition period during which safety goals are used as a kind of touchstone to establish the validity of currently existing regulatory practices, regulatory requirements.

It may in fact be a very long transition period, and there is some possibility that we never get out of this transition period, that this is how quantitative safety goals would be used even in the long run, rather than this being the touchstone to determine whether say an individual plant is licensable or safe or something of the sort, that this establishes the means by which a secondary set of requirements are established for making plants or safe or something of the sort.

It was noted that although we do need -- it is sensible to proceed with establishing safety goals, both of a qualitative and quantitative character, that this may not have very much impact on nuclear power plants until, oh, say about the turn of the century, considering the pause in construction which now exists, unless these safety goals provide some basis whereby plants currently being constructed

are backfitted with features that may be derived as a result of safety goals having been established, and if this is the case, if the real application of safety goals, the real importance, is to be attached to plants which are, say 20 years down the line, we will have to be quite forward-thinking in the way we construct these goals.

The purposes of safety goals were discussed, and we agreed on two reasons for having safety goals, just as a logical start, one to protect the public health and safety, and two, to make the regulatory process -- the words are more rational. Some said less irrational.

It was agreed that it was not a purpose of safety goals to make nuclear power more acceptable to the public, although it was also pointed out that if safety goals are established which do not -- which lead to unacceptability of nuclear power by the public, this would be a useless enterprise.

Understandability of the safety is clearly a required characteristic, but certainly not an objective either.

Goals themselves could incorporate conservatism in the quantitative limits which are established, though this is not necessarily the case. There is some feeling that as a matter of fact the argument has to be looked at very carefully to establish whether or not we do want conservatism

in the nuclear field based on a cost-benefit analysis.

The analytical methods whereby compliance with goals is determined should be completely realistic. That is, in calculations, do not introduce conservatisms into the calculational methods to determine compliance. If you are going to put in conservative bias, put it in the limits themselves, where it is clearly seen.

We did not arrive at any logical basis for determining the quantitative -- well, the limits to be put into quantitative safety goals. There were several suggestions. Some were partisan to one suggestion, some partisan to another.

Some of the suggestions were that nuclear power should provide no greater -- only a given -- no more than a given fraction to total risk of man's activities.

Another was that nuclear power, the risk from nuclear power should be less than the risk from competing technologies, and here there is also a view that competing technologies may not simply be those that produce electricity, because in the long term, if we think, say, 20 years down the line, these may be technologies that produce synthetic fuel, that produce space heat, and so on.

There was some belief that limits should be set as the result not of comparisons among technologies of any kind, but on a cost-benefit basis. That is, if you need the

electricity, then you should be willing to allow a certain dysbenefit as a result of it.

Finally, there was some view that none of these is really going to work, and what you ought to do is simply pick some numerical limits that everyone thinks are clearly acceptable to the informed public and perhaps some of the uninformed public.

At this point, we began to get into quantitative safety goals and the day ended. We were -- the day ended as we began to take up the first aspects of the ACRS proposal.

MR. SEGE: Thank you, Herb. Dr. Lester Lave, Panel B.

DR. LAVE: After some of the sharp words that were exchanged yesterday morning here, I think that those of you in Panels A and C would have had tears roll down your cheeks to see former adversaries getting together, congratulating each other on the positions they had taken on nuclear power, declaring that they individually had been wrong and sometimes had had impure motives.

I am sure that George was -- it was too bad that he missed the beginning of our session where we had a unanimous vote to praise the NRC staff for penetrating lucid documents, for a set of wonderful questions, unanimously applauding the panel titles, the memberships, and in general the set-up of the meeting.

DR. LEWIS: This is April 2nd.

DR. LA'TE: Indeed, my task as a chairman was made easy by the lucid discussion, and I have a set of slides to show you which reflects the lucidity. Now, let me try.

We decided that we had our own set of issues that we wanted to talk about, and so let me just go through those rather briefly.

We started out with the usual question of why is the public apprehensive, and got a whole set of feelings of why that is so, about whether the public was ignorant, of whether journalists were stupid, or had ulterior motives and so on.

I think that it was pointed out that by and large the public is not a set of dumb people who don't react -- or who react irrationally to things. The problem is one of information, and the cost of information, but even if we talk about technically educated members of the public, we don't always get people who are well-informed about the issues of -- such as nuclear.

One of the most important points made in our meeting had to do with whether the nuclear agency was paranoid, and there it was remarked the only reason was that your agency would regard itself as being singled out is that it doesn't know enough about what is going on with other industries.

Literally every high-technology industry around, and many that are not high technology are under public scrutiny --

DR. JOKSIMOVIC: They are all paranoid.

DR. LAVE: Pardon me?

DR. JOKSIMOVIC: They are all paranoid.

DR. LAVE: Well, remember, paranoia is believing that you have enemies and being wrong about that. I think that the answer is no, high technology and so on, they are not paranoid. There really are all sorts of people out there who don't --

VOICE: You know the saying, even paranoids have enemies.

DR. LAVE: We talked about how it is that the regulatory process can be enhanced, and that was really much more a matter of trying to get the right questions answered instead of worrying about why some single group was being, or some group was being singled out.

I think that the general comments reflected that the nuclear industry is not being singled out, that there really is lots of mistrust being heaped on technology in general, and one of the characterizations was to try and differentiate between the sort of inherent feelings people have about technology.

There are a set of technology optimists who regard

what has been going on in the last hundred or so years as being the best thing that ever happened to the human race, and believing that the faster we implement new technology the better off we will all be.

Opposing them are a set of technological pessimists who don't regard what has been happening as marvelous for the human race, who see disastrous potential all the time.

One way of sort of phrasing that issue was, is it necessary that all problems with a technology be solved before that technology is implemented, and you get quite different answers from the two groups on that kind of a question.

The second area that we were trying to look at was what should the regulators have in mind when they are trying to go about looking at safety, and here we were trying to sketch out what some of the relevant attributes were, and we have this sort of grocery list at the top, changes in life expectancy or changes in premature deaths, the difference between them really is the extent to which one aggregates and how one aggregates across individuals and across time, changes in morbidity rates, changes particularly in disability or chronic illness, a general category of nonefficience or levels of consumption, various kinds of noneconomic or aesthetic values such as species extinction and so on, institutional changes, civil liberties and so on, and the amount of

resources that are required to be devoted to the regulatory process per se, that is, one of the things that nobody seems to want to do, which is spend an indefinitely large amount of resources on the regulatory process itself.

And of course, life gets more complicated still than that, because there is no succinct easy way to talk about, for example, 10,000 cases where life is shortened to some extent, and so one has to look at distributional effects in some more detail.

For example, which of these effects occur now, which of them occur later, and how does one make the now and later commensurate somehow. Domestic and foreign effects, effects on the rich versus the poor, effects on the old versus the young, effects on the public versus various occupational groups, or effects on individuals versus groups, effects by race, by gender, or by region, where in general, the question is, who gives and who gets.

I want to emphasize that while we posed this general set of questions, we immediately said that any regulatory body that was required to consider all of these explicitly was never going to make a decision, that this was simply paralysis, but that this was the range of issues that one had to have in mind, and that if any regulatory body, for example, focussed only on these first three categories, of health matters, and neglected the rest, then they were going to

arrive at poor decisions, so that this is a sort of a partial list to keep in mind as one is thinking about decisions, even if one is going to do these things implicitly.

We thought that one had to look at some of the advantages and disadvantages of quantifying safety goals, and I carefully redid the notes so it looks as if there are more advantages than disadvantages, even though our discussion showed the reverse.

One of the major advantages is that it allows comparisons to be made with other kinds of technologies, with other situations in life.

If one is able to sketch out some risk levels, even if they are not very precise, then you can define other kinds of situations. By analogy with food and drugs, one of the advantages of quantification is that it enables the agency to require detection limits that are less than what is scientifically possible.

One of the mistakes that FDA and Congress made in the early days was saying that their goal was no detectable amount. The problem is that analytical chemists were far smarter than anybody ever dreamed, and it is really a bankrupt practice to try and talk about no detectable amounts, since somebody will find ways of detecting levels that you never dreamed possible.

The third was that quantification permitted the

notion of a diminimus risk. If you didn't have quantification, you didn't know what a diminimus risk was. And then we get into some more general advantages of quantification. It permits sort of the reporting of experience and of learning, so that one has some idea over time what your track record has been, whether it is good or bad, and what should be done about it, and permit consistent rather than ad hoc decisions.

It may be that in the English common law, that a marvelous body of doctrine gets built up over time, and you are talking about a modern regulatory agency handing out a relatively small number of decisions. It is not quite clear that they are going to get any consistent set of decisions out of this without some quantitative framework.

Another notion was one of the educational value of quantification both inside the agency and outside, to tell people rather precisely what you had in mind, without reifying the numbers, that it permits one to be fairly precise about what was intended.

A track record for evaluation, and then one of the discussions we had between two lawyers, and I must say I am very suspicious whenever two lawyers agree that something is marvelous because it will lead to less litigation. The notion was that by establishing some sort of generic rule that we would probably have less litigation in the future. If we ever had two economists together that this change was

marvelous because it would mean less employment of economists, I need hardly say there would be a lot of suspicious non-economists around.

In terms of some disadvantages, the first and principal disadvantage is that there is a decided tendency to ignore matters which can't be quantified. No matter how often one says that the following notions are important, but can't be quantified, people start falling in love with numbers, and tending to give short shrift to those matters which can't be.

Sort of an ancillary problem associated with that is that once one starts getting into reams of numbers, it is very easy to kind of camouflage value judgments. You simply stick things in there and nobody sees them again.

One of the problems is that many of us are goal-oriented, and if you set up a number such as getting to some risk level, then there is a fixation on that numerical goal, people move heaven and earth to get there, and it may turn out that you wanted to have a little bit of a more of a balanced approach to it than simply trying to meet exactly that.

There was an issue about whether the analysis was complete. There is this danger again in setting numerical goals when you set those numerical goals as to whether they were encompassing all that you wanted to encompass.

Another issue down here is that the quantification

will have the effect of pinpointing certain shortcomings, and probably lead to a sort of a worse trouble in the short run, that is, that once these things are put down on paper, the agency's flaws will be revealed for all the world to see, and it will be some time before some response can be made to those kinds of flaws.

One idea here was whether it was possible to quantify qualitative changes, such as war. That is, were there very important aspects that were being ignored in all this, and then a final notion was that at least if one looks at the reactor safety study, there is a clear tendency for the goals to be hidden.

That is, that one of the advantages -- or one of the attributes one would like to have of the system that the goals and the process be transparent for all to see, all to argue about and appreciate, and the more complex one makes the calculations, the less transparent these things are going to be.

As Panel chairman, I carefully tried to survey what everybody knew, and then tried to find the topics that nobody knew anything about, and it turned out that the marvelous topic knew anything about was epistemology, and so of course we spent a long period of time talking about that.

The idea was, what can really be quantified in this area, and how certain can we be about the nature of the

quantification. That is, are we going through the medieval exercise of how many fairies can dance on the head of a pin or is there actually something that can be measured that we can have confidence in, and this is one where we didn't resolve things particularly well. I think we are going to have the joy this morning of going back and spending the whole morning talking about it again.

The idea is that with respect to routine operations, that we are pretty good with quantification, figuring out what the consequences are out to many decimal places, but when you start getting over to accidents, then there are a number of difficulties.

Some of the criticisms of WASH-1400, if you remember, had to do with whether the events were actually independant, whether combinations of events had been considered, and then whether there were important unrecognized events, as inevitably there are.

And we had some ideas about whether you could actually calculate what everybody would accept as upper bounds, which was low enough so that it would be meaningful in some sense.

As I say, those issues were not resolved very well.

And then, some of the matters that we intend to take up. We want to spend some time on implementation of quantitative goals. The first notion here is one of trying to

look at narrow technical goals versus goals that are understandable to the public, how you manage to make the two commensurate somehow. That is also part of the notion of how it is that you get some wonderful global set of goals that we all agree with, and then translating them down to workable goals, making sure that what is learned then at the level of working with these goals gets translated up again into high level goals, that is, making sure that there is consistency among the various levels.

And then, we want to spend some time talking about how it is that you get the system to operate as designed if we manage to settle on all these goals, and how it is that the goal process itself will be affected by the extent to which you believe you can get the various systems to operate as designed.

And then the final area that we want to get into is talking about the process for what I will call identifying and involving stakeholders, that is, who does have standing here, and how much standing do they have to protest or have their views known about various parts of the process?

The principal notion here is one of trying to find felicitous processes, that is, processes that manage to involve the right people in the right way, so that you manage to get issues resolved, get compromises made rather than having people go into highly dysfunctional behavior, and the

notion here is one of trying to look at comparative advantages among various groups.

Well, those are notions that we haven't gotten into yet. I think I have probably talked too long, and I will stop.

MR. SEGE: Thank you, Lester. Dr. Paul Slovic, Panel C.

DR. SLOVIC: Well, I think we did reach one point of consensus regarding the ACRS document, and that is we agreed, I think, with their statement that management of risks is as much a sociopolitical problem as a technical one. After that, the going got a little bit rougher.

We spent quite a bit of time at the beginning trying to get our bearings on issues such as the distinction between goals and rules, and there was some feeling that we should focus on goals at this time, and that rules were -- quantitative rule development was a bit premature.

However, the subsequent discussion seemed to confound the two again, reflecting what I think is a continuing uncertainty as to how goals would be operationalized or implemented.

We spent also some time discussing the development and the philosophy behind the ACRS guidelines in order to familiarize ourselves with that a bit more so that we could proceed then to discuss the sociopolitical issues relevant to

the ACRS proposal.

I think basically our discussion did center on the general question, which is what are the main economic, ethical and sociopolitical issues in the formulation of nuclear power safety goals, and secondarily we did address the question how does the ACRS proposal deal with these issues.

We are still formulating the answers to these questions, but some preliminary version of it is as follows:

Among the major concerns that were identified were the following:

First, there was concern that quantitative goals might capture or dominate the decision process, driving out valuable qualitative standards or procedures. For example, there was concern that the great uncertainties in quantifying certain factors and the problems of verifying that these quantitative criteria were met might lead to sort of number games and hand-waving and designed to give the illusion of satisfying these criteria at the expense of emphasis on some sound tried and true principles such as defense in depth.

The counter to this concern was the view that quantitative goals would really be only part of the process. They would be designed to supplement these other procedures rather than to replace them.

We talked a bit about the problem of public

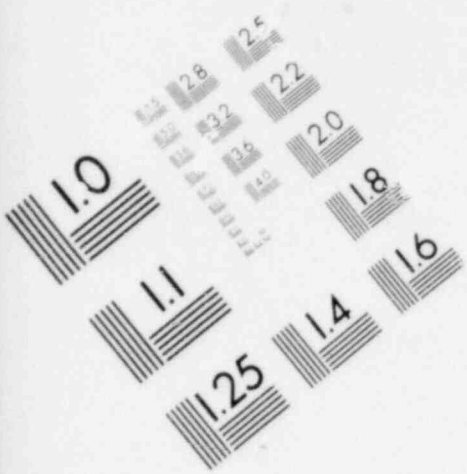

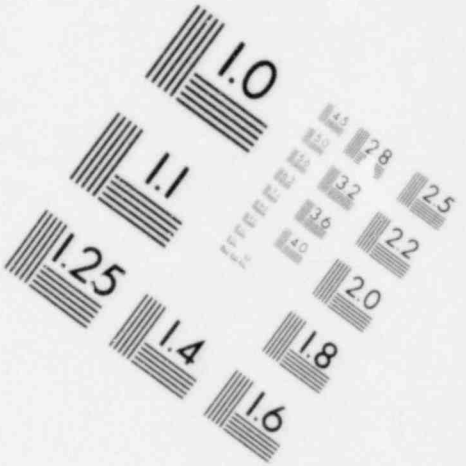

IMAGE EVALUATION TEST TARGET (MT-3)

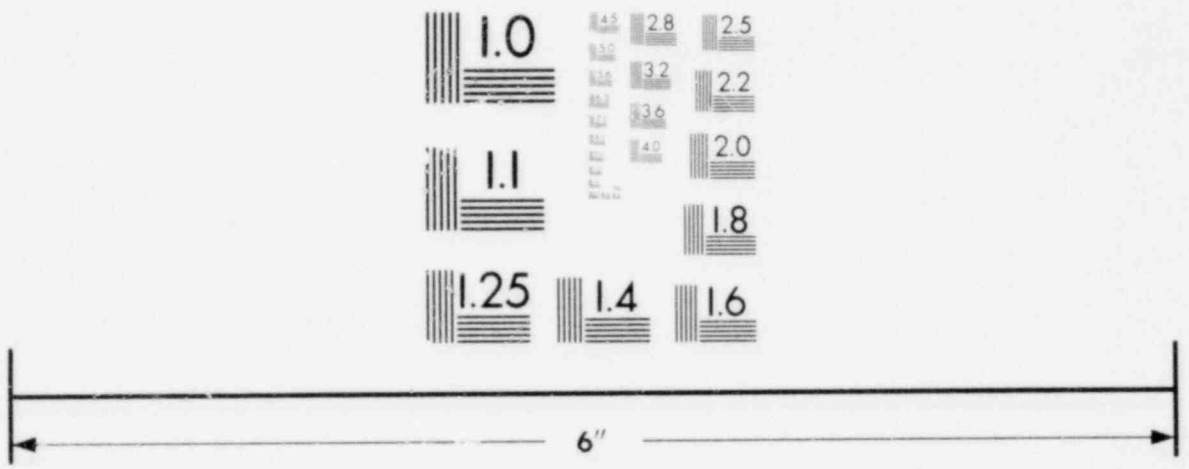

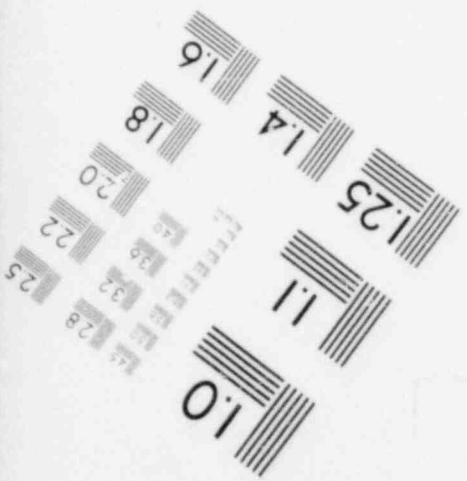

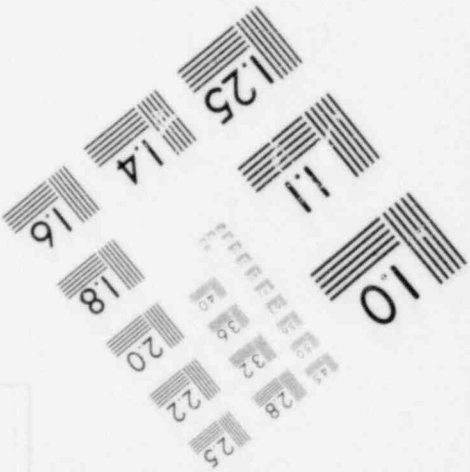

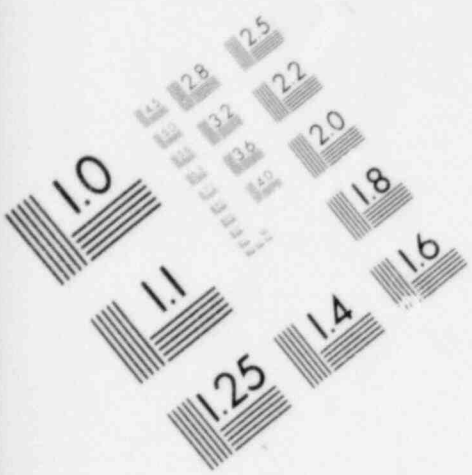

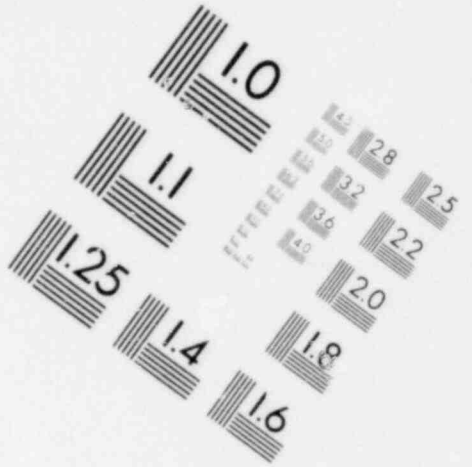

IMAGE EVALUATION TEST TARGET (MT-3)

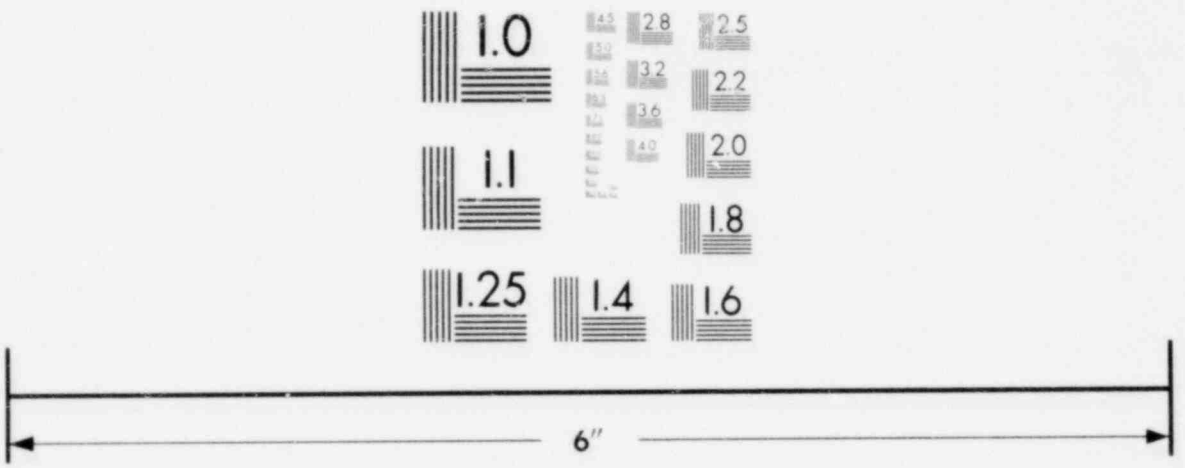

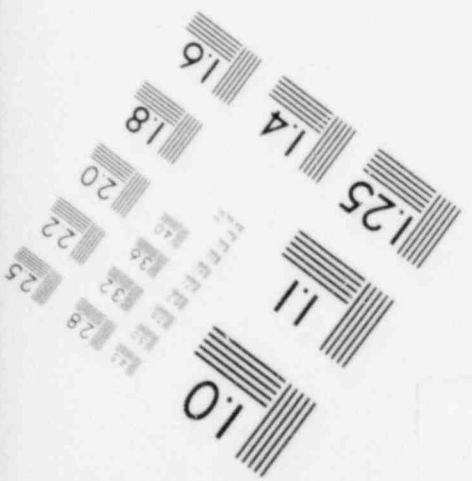

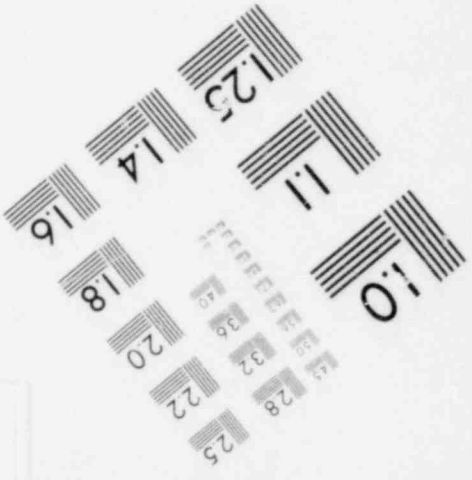

acceptance of these goals. I think we probably should work on this some more, but there was some -- at least three different views on the impact and relevance of quantitative goals to public acceptance.

Some felt that rational coherent goals would enhance the public trust, which is so important to the nuclear enterprise.

There were others who felt that actually the quantitative goals might reduce confidence and trust. For example, there was a feeling that the public would have difficulty understanding these goals, and would actually have more confidence in principles such as defense in depth, which sort of have a more deterministic and reassuring element to them, that the public somehow had a better feel for those kind of principles, rather than the complex quantitative goals that they don't understand and don't believe can be verified.

Finally, there was a view that in a certain sense these goals could proceed independent of specific concerns about public acceptance, that is, that the public should be consulted where issues of values were relevant for incorporation into those goals, for example, attitude towards catastrophic loss of life, or the importance of that is an issue where societal values should come into play, but that basically these goals were to be used as guides for decision-

making on the part of bodies such as the ACRS, and that they could proceed -- they would be valuable in that light alone, regardless of the impact on public acceptance.

Well, I think there was a lot of -- as you can see, there are different positions, and there is no agreement on that point.

Secondly, we considered the problem of scale, which seemed not to be reflected in the ACRS document, and that is, does it make a difference whether you are designing goals for a system with 50 reactors or 100 reactors, or perhaps a thousand reactors? There was a feeling that it did make a difference. For example, there was concern that there might be some slippage in the design, licensing, monitoring or emergency response capabilities as a function of large scale, and that these problems needed to be addressed.

There was also a feeling that scale would -- this is a related problem -- that scale would affect the availability of trained personnel, that another related issue is the sort of mix or dependence on mix of power supply systems, the dependance on nuclear power, that if the dependance got too great, doesn't this produce a certain vulnerability to problems that arise, so that if there were industry shutdowns, we would be in a quite difficult position, and what is the relevance of this for quantitative criteria?

Another related issue that has to do with scale is that with the large scale, large number of plants, you have the more likely possibility of individual serious accidents which could have -- which carry with them major societal and industry costs, just the sheer frequency of these accidents would increase, and somehow this should be factored into quantitative goals.

We also spent some time discussing ethical issues, also these two are not directly addressed in the ACRS proposal. There was a general feeling that it is desirable to compensate those on whom risks are imposed. For example, if one moves to siting in a relatively unpopulated area, in order to lower the expected loss of life, given an accident, then somehow these people should be compensated for the imposition of this risk upon them.

We separated these compensation or equity issues into two categories, those dealing with the spatial distribution of risks and benefits in the current generation, and those dealing with the intertemporal or intergenerational problems, passing risks along to the future, and there was general agreement that while there are some principles for dealing with these intergenerational equity issues, they are much more complicated even than the -- you know, there was at least some hope of dealing with the equity issues in the current generation, the spatial equity problems, but the

intertemporal problems would be very difficult to really incorporate into any criteria at the present time, not to say that they should be ignored, but that there are immensely complicated issues having to do with how do you -- you know, how does the future value certain risks that you will impose upon them in order to give them certain benefits, how do you make those decisions?

There were discussions of principles such as trying to leave the future with a menu of opportunities that is very much the same as the -- as exists for the current generation and so forth.

There was a discussion of discount rates, that is, can you or should you have a discount rate on costs or lives lost that are imposed on future generations. Some felt that a zero discount rate was really the most justifiable, meaning that you just simply count lives lost in the future the same as lives lost in the present.

There were others who felt that one could discern situations where you wouldn't want to have a zero discount rate, or where you act as though there not such a rate, you know, where you would act to save lives now more forcefully than to save them in the future, but that even then there was no specific or constant rate that really could be applied, so I think the general consensus was that this is an important, relevant, but immensely difficult problem to deal with.

Just in passing, there was a comment that the ACRS proposal seems to neglect genetic risks. It seems to assume that latent cancers are an appropriate surrogate for this. This was brought up, questioned, but we did not discuss it in any detail.

There was also brief mention given to the problem of how you deal with old or existing plants versus new plants, vis-a-vis quantitative criteria, and there was concern that there should be a distinction here and that the criteria that apply for new plants should not necessarily apply for the existing plants. The concern was that if there was a sort of across-the-board applicability here, then that might lead to a tendency to weaken the criteria for new plants, in order not to impose severe costs on existing plants.

We did not yet discuss the level of safety. I think that is an important issue which we will get into this morning.

Thank you.

MR. SIEGEL: Thank you, Paul. Now I would like to call on Dr. Lave to present the highlights of his findings concerning the practice of other agencies.

DR. LAVE: Rather than try to go agency by agency, I am going to show you a set of decision frameworks most of which are currently in use and to be proposed, and tell you what the agencies are that use these. In general, there are

relatively few pure cases of any one of these frameworks being used.

First, a framework that is used for the vast majority, that is all but a few health and safety decisions, is market regulation or individual decision, some combination of letting the market do what it wants, trying to make sure that information is available, for example, about saccharine, and then letting the courts handle cases of wrongdoing of one sort or another.

The second framework is no-risk, or the Delaney clause used by the FDA for food additives used by -- or imposed on EPA by Congress, for the primary air quality standards. Basically the idea is that no level of risk is deemed to be acceptable, and at least within this narrow class, the agency is to attempt to get risks down to zero.

That framework proved to be one that FDA did not like, and FDA originally testified against the Delaney clause and has found it uncomfortable to live with it. In particular in looking at sodium nitrite, a suspected carcinogen, the FDA proposed that it be allowed to consider the health benefits to consumers of having sodium nitrite around, versus the health risks of having it around. They called this framework risk-risk analysis, and it was the opinion of the attorney general that the language of the Delaney clause was clear, and that they did not have that option, that they were dealing with

no risk.

Congress has instructed EPA to use technology-based standards in various cases. For example, in the water area, Congress tells EPA that the standards shall be best available control technology, or one or another of those. Those wind up being a set of engineering judgments. When they are implemented in practice, they don't mean what the plain language says. They don't mean best available technology. What they mean is impose a technology which is stringent up until the point where the engineers either have reservations about whether it will work in practice, or the industry can no longer afford it, and without dumping on other disciplines, I must say that engineers making judgments about the profitability of companies doesn't let me go to sleep very easily at night.

The fifth framework is risk benefit analysis, and I will be careful what I say about that in this group. It is used, at least the name is used by a number of agencies, FDA with respect to drugs, EPA in the TOSCA regulations, Toxic Substances Control Act, although I sometimes get put on the program at meetings as an expert on risk-benefit analysis, I don't know what it is, and it must be that I just didn't talk with enough practitioners.

The sixth framework is cost-effectiveness analysis, developed out of the Department of Defense, where you

remember Eisenhower's Secretary of Defense, Wilson, said that his objective was to get the most bang for the buck. The idea is to try and maximize some goal given a budget constraint. Cost-effectiveness analysis might also, for example, come under the Nuclear Regulatory Commission's ALARA criteria, where one puts some dollar value on a saving of a man-rem, and then makes sure that all efforts which could save a man-rem for less than that cost are done.

One way of implementing the cost-effectiveness framework is the regulatory budget, a proposal that came out of Charlie Schultz' Council of Economic Advisors. The idea was to have an implementation budget where the amount of cost that each agency could impose on the private economy, those then would -- that would be a separate line item debated by Congress and set by Congress.

The final framework is conventional benefit-cost analysis, which requires doing all these things that people get intense reservations about. That is, not only finding out what the health effects would be in quantitative terms, but translating them over into dollars, and then adding up the benefits versus the costs.

I want to point to all of you that there is a government agency which always has used explicit benefit-cost analysis with a full explicit value on life, namely the FAA, and in every FAA safety proposal they go through and they

quantify lives lost, property damage costs. They then make everything commensurate in dollar terms, and my question of these people was, you don't really mean to say you have a dollar value on premature life, and their answer is yes, they do. They get asked about it every year at appropriations time by Congress, but so far, that is what they do.

The more general notion that I want to bring out here is that you people in the nuclear business are really luckier than almost anybody else in the government regulation of health and safety, because more is known about the health effects of ionizing radiation than about virtually any toxic substance around.

If you think for a second about what an FDA Commissioner should be doing when he receives a telegram that says we had 36 salmonella that reverted when we poured this substance on them, and we have ten consumer groups that want us to ban the substance immediately, you see that the kinds problems that are faced by other agencies are much much more difficult than anything having to do with ionizing radiation, and the thing that is interesting about them is that in fits and starts and with more or less controversy, the agencies have found ways of trying to live with the kinds of problems they have, and many of those techniques are of great interest, I think, for nuclear, for trying to find what compromises seem to work.

MR. JOKSIMOVIC: I don't see OHSA anywhere.

DR. LAVE: OHSA believes that its statute is a no-risk statute. OHSA up until this year believed its statute was a no-risk statue. The Supreme Court said they were wrong.

DR. JOKSIMOVIC: I don't know that that is true.

DR. LAVE; Well, if you go and read their statute, you could make the argument either way. Some of my OHSA friends tell me the Supreme Court made its decision up out of whole cloth. The people on the other side said that no, the agency had always been wrong, that they had always had a dumb general counsel. I have talked with both the critics and the general counsel, and I don't think either of those is true. It is just that it is a very ambiguous statute.

MR. SEGE: Thank you, Lester.

We are now ready for comments by participants on the reports that you have received, and comments for consideration by the panels in their further discussions.

I thought that we might start first with comments that focused primarily on Panel A, and proceed to Panel B, and then to Panel C and then comments that don't particularly relate to a particular panel, but cut across the different topics, to the extent that we can.

When you make your comments, would you please speak loud and clear, and identify yourselves to the reporter

who can't see all of the nametags, and this identification is particularly important if you are sitting behind somebody else's nametag.

Any comments with respect to Panel A? Mr. Malsch?

MR. MALSCH: I had one comment, and it was because of a statement that was made that I think I disagree with, and that is that you ought to set the goals very conservatively but then accept realistic implementation measures, and what I am concerned about is that if we establish very conservative goals, and then provide for realistic implementation measures, that would provoke intense controversy and concern as to whether the goals are met, as opposed to establishing realistic goals, let us say, lesser stringent goals, but then provide for very stringent implementation measures, and thereby diminish controversy over whether the goals are in fact established, just a comment that I had.

MR. SEGE: Thank you. Any other comments? Mr. Temme?

MR. TEMME: My comment is essentially the same, but I would like to add one more thing to it, I think. I think the objective of making the calculations realistic is probably not acheivable, because each of us has his own definition of realism, and that leads to a lot of controversy, as was pointed out. I don't think it is a

necessary objective, either. I think the real need is to specify what the calculations are when the goal is stated, and we can have that argument, but I think -- don't think we ought to get caught up in the argument of whether or not the numbers are realistic.

That is motivated, it seems to me, by one of two reasons. The first is that as engineers, when we are trying to meet goals, we like to think that we are performing cost-effective tradeoffs, optimization, if you will, and if we are using complex analytical models in which are embedded in a very uneven manner conservatisms, we can be completely misled by our calculations.

And so it is desirable from that perspective to have realistic calculations of accident sequence probabilities and so forth.

The other motivation seems to be that if -- one that appeals less to me -- and that is that if we put a lot of conservatisms into our risk estimates and they come out looking high and then we are accused of trying to kill a lot of people or something of that nature, and we become defensive about that, the other side of the coin is that we can also do calculations which produce infinitesimally small numbers, and then we are told who are you trying to kid, and I don't think we should be overly influenced by those kinds of reactions in what we do. I think -- as I have said before --

I think there is an incentive to do real calculations, to try to do a balanced job of designing an operating plant, but I think the real objective should be to just specify the calculations rather than to specify that they be realistic or something else.

MR. SEGE: Thank you, Mark. Any other comments with respect to Panel A? Mr. Hutt?

MR. HUTT: As I heard the panel report, they have reached the conclusion, I think it was twice stated, that nuclear energy should be safer than competing energy sources. Did I mishear that?

DR. BEYEA: That is not a conclusion. That is one point of view that was expressed.

MR. HUTT: I would simply point out that I think the vast majority on our panel concluded quite the opposite, but then I don't know that there is any -- not that it should be more risky, but that it should be -- there should be a uniform safety standard for competing energy sources.

DR. KOUTS: That would be an alternative that could be talked about, too.

MR. SEGE: Thank you, Peter. Yes, Mr. Burstein?

MR. BURSTEIN: I think that in Panel A, we tried to deal with the purpose of a safety goal, which as Dr. Kouts expressed earlier in a summary, had two facets, one to protect public health and safety, and the second to make the

regulatory process more rational.

If the imposition of a safety goal is accomplished, it is going to require some means of perhaps demonstrating acheivement or compliance that utilizes probabilistic risk assessment or some other kind of approach.

If that is superimposed upon current demonstrations of safety compliance, as I have heard in some other groups, and as we discussed in our own panel, as an additional regulatory procedural requirement or intent, then I for one want no part of it, and I think we are wasting all of our time, particularly for something, as we said, which may not be applicable to a generation of nuclear plants until some significant time in the future.

Unless the development and implementation of the safety goal serves to substitute for some of the current regulatory activity, that does not serve to make the regulatory process more rational, but on the contrary, it makes it more complicated, difficult, costly, and confusing. I think that is a very significant and important issue.

MR. SEGE: Thank you, Saul.

If there are no further comments on Panel A, we could -- yes, there is one further comment, Mr. Derby?

MR. DERBY: I would like to set forth an opposite view to what was just expressed. In Panel B what we discussed or part of the discussion centered on intuitive decisions that

were made on an ad hoc basis in the implementation of current Nuclear Regulatory Commission standards.

In the sense that a quantitative goal would formalize that intuitive ad hoc decision process, it would supplement the present standards. I for one think that is an advantage, and makes that soft process of regulation a little more regulated, to use the word, but also make explicit declarations of what goes into that soft part. That seems to be where the issues arise. That seems to be where the delays are seen, so I would like to offer that opposite point of view to the fact that replacing what is there is the only goal, or the only purpose of a quantitative standard. In fact, I think there is a supplemental advantage for additional formalization.

MR. SEGE: Thank you, Steve. Mr. O'Donnell.

MR. O'DONNELL: I would like to offer a viewpoint on this also. I think, speaking for myself as an industry representative, the problem with the current regulatory structure, is not so much that the current requirements are burdensome. I mean, they are being met. The problem is one of predictability and change, which is at present uncontrolled in any systematic manner, and I would view the application of quantitative safety goals as something to control that change, and I think that we are not looking to replace existing requirements, including concepts such as

defense in depth, and throwing them out and introducing something completely new, but the quantitative safety goals would be a means of measuring what is the level of safety that is provided by this existing set of regulations, and if it does in fact give a level of safety that is within these top level goals, we should then be focussing on the need to change those requirements on some quantitative basis, and specifically the use of the cost-benefit aspect of the goals, to make those changes, those decisions on faith, and I think that is the most constructive use of the quantitative safety goals, not to throw out what we have, but to introduce predictability and systematic decision-making into the need for change.

MR. SEGE: Thank you. Mr. Levine.

MR. LEVINE: I guess this is a very important matter that is being discussed now, and we discussed it extensively yesterday. I sense that there is some confusion about this matter. People seem to be regarding the fact as, we either have existing requirements or we have quantitative safety goals, or we have both, and I think none of those is right.

I think the idea of having quantitative safety goals is -- well, one has to talk about how one will use them. I don't think they should be used in the licensing process, but they should be used in the regulatory process. That is,

they should not be used in the licensing process to determine whether a particular reactor meets them or doesn't meet them, but they should be used in the background to look at the rationality of existing so-called deterministic requirements, to make them more rational, and those are the requirements that should still be used in the licensing process, at least for the foreseeable future.

This came up in Panel C, in fact, if I can say a word about Panel C, where in talking about the public acceptance of goals, the statement was made that the idea of defense in depth would have more appeal to the public than a set of numbers, again, I don't think it is either/or. No one is going to throw out the idea of defense in depth. No one is going to throw out all the safety engineering design requirements that now exist, which can't be covered by PRA or quantitative goals.

So it is not a question of either/or. It is a question of supplementing our existing NRC requirements with another tool to help them make it more rational, but not to throw anything out, or to add anything in particular.

MR. SEGE: Thank you, Saul. Let us move on now to Panel B. Any comments on the -- in connection with the Panel B report? Mr. Levine?

MR. LEVINE: I just have a general question on Panel B. I took some notes which make it look like all they

did was discuss quantitative goals, and I wonder if they are going to consider whether or not we should have qualitative goals, or what they should look like, or anything like that.

MR. SEGE: Would Dr. Lave please respond to that question?

DR. LAVE: Certainly we are. I guess that we had thought of that -- yes, we are.

MR. SEGE: That was a good unequivocal answer, thank you. Professor MacLean?

DR. MACLEAN: I had three very particular comments, well, there are two particular, one more general, just things that I happen to disagree with.

One was, it was mentioned as comparing the quantitative, the qualitative values or goals or whatever, it was listed as a disadvantage of quantitative goals that attempting to quantify qualitative goals, that it has a tendency, perhaps to lead to doubt. I think that is true, I don't knock it, but I think it is an open question whether the problem is merely to ease things out or whether there are some things that in principle cannot be quantified, and I don't know if I am disagreeing with you there, but I -- okay.

Secondly, it is listed as an advantage of quantitative goals that they can yield consistent decisions. I found this puzzling because a qualitative decision principle will also yield consistent decisions. The

measure of consistency itself is either quantitative or qualitative, so I think that claiming that as an advantage for quantitative goals is question-begging.

And third, more generally, I think, and after having thought about this myself for some time, and trying to write some things up on this, I think that talk about public apprehensiveness in terms of technological optimists and technological pessimists is more confusing than it is helpful.

I am thinking about it, first of all, I don't know a single member of Friends of the Earth who doesn't own a better stereo than I own. I have never heard a greater technological optimist than Barry Commoner talking about solar power. I have never heard a greater technological pessimist than my colleagues in the nuclear engineering department talking about solar power, and in general I think it is a red herring.

MR. SEGE: Thank you, Doug. Any other comments on Panel B? Professor Lewis?

DR. LEWIS: Just one, since we are making fun of technologists, I would like to register a general but very, very minor comment, in one of Lester's slides. There is a tendancy to always put the word "narrow" in front of the word "technical," and I would like to register an objection to that tendency.

MR. SEGE: The objection has been noted. Anything else on Panel B? Dr. Zebroski.

DR. ZEBROSKI: I guess it is commonplace, but I would like to supplement the chairman's excellent report with the observation of --

MR. SEGE: Excuse me, could you speak louder? People can't hear you, and the reporter can't hear you.

DR. ZEBROSKI: I would like to supplement the -- the chairman did an excellent job of summarization, but I think several interesting although commonplace points are also very important on the risk-risk aspect. I think two or three of the panel members brought up the observation that if you look on a 10 or 20 year time scale, the risk of human misery and death from social chaos and local or world war runs a very high probability. One panel member said .1 chance, and another panel member said 1.0 chance of such events affecting the human condition in the future.

The omission of something which is hard to quantify, which is the impact of adequacy or inadequacy of domestic energy supplies on world tensions and war, it is very difficult to handle quantitatively. I know of only one attempt, at MIT, but it is a very important part of the equation, which has always been left out, and is still being left out, and perhaps should be considered at least qualitatively.

MR. SEGE: Thank you, Ed. Dr. Eisenbud?

DR. EISENBUD: I would hope that at the conclusion of this workshop we don't give the implication that the subject of the risk goals, or safety goals has been introduced de novo and is a concept that is going to arise out of the recent work of the ACRS and the NRC staff and this workshop.

I detected that implication in some of the discussions that took place yesterday. I believe that there is a double chain of quantitation now, and some risk goals, or safety goals that have been introduced and gradually improved or expanded as the technology changes, and I hope that this, I think at some point I think this workshop has to address the really two chains of quantitation leading up to the safety goals.

One is based on the fact that there exist certain risk coefficients that describe the amount of health damage per unit of radiation exposure, and that information, which comes from various sources, national and international, has then been incorporated in Parts 20, 50 and 100 of the NRC's Code of Federal Regulations, into what amounts to safety goals.

They describe what the goals should be for normal operation in Part 20; in Parts 50 and Parts 100, they describe what would be acceptable in the event of an accident, and my

point of view as a somewhat distant observer is that what Parts 50 and 100 do, Part 100 more specifically, is deal with the question, and here comes the second chain of quantitation, of what the probabilities are of reaching the limits given in Parts 20 and 50, and there is an implicit acceptance of what the probabilities should be, the probabilities have been described in WASH-1400, as modified by the Lewis Report and others, but I think all of this constitutes a body of information which at the present time provides the safety goals on which the NRC policy and regulatory procedures are based.

And I for one have not detected yet what the relationship is of this workshop to the existing body of information and body of practice. I was trying to -- should we test the present criteria to see whether they are adequate? Have we decided that they are not adequate and therefore have to be made more strict?

In short, I don't think that we should emerge from this workshop without some continuity from the present body of practice as it now exists to whatever we recommend in our report.

MR. SEGE: Thanks for the comment. I am sure that the remarks that you have made will be taken into account in the panel discussions today. Mr. Malsch?

MR. MALSCH: Yeah, I just wanted to make a comment

in response to that, that there is sort of an implicit goal in terms of normal releases in Part 20 and Appendix (i) to Part 50, because that is based upon the old FRC radiation detection standards which were based upon a generalized kind of risk-benefit analysis, and then the overlay on top of that that at least it should be kept as far below those as reasonably acheivable, which is based explicitly on a cost-effectiveness analysis.

But in terms of accidents there is no probability goal at all established in 10 CFR Part 100. The language in the regulation refers to the word "credibility." That has been debated back and forth for years. Earlier decisions lead one to believe that credibility meant conceivability. Later decisions imply that some sort of probability is involved.

But you will search in vain for any official NRC pronouncement as to what is the quantitative dividing line between a credible accident sequence and an incredible one. And thus, the particular focus of the workshop on accidents, which is where there is less with respect to probability goals in the NRC's regulations as compared to almost any other area.

MR. SEGE: Thank you, Martin. We are running somewhat behind schedule, so I would like to ask that any further comments on Panel B should be made only if the commentor considers it to be of considerable urgency to make it this

morning rather than reserve it for possible consideration tomorrow in the longer discussion.

The criterion should be if the comment would be likely to influence the further deliberations of Panel B, it should be made now. Mr. Levine?

MR. LEVINE: I hate to do this, but you just said the words that made me feel I had to comment. In looking at the list of disadvantages to quantification, that it ignores what can't be quantified, the value judgments, et catera, et catera, et catera. This is true of any model that anyone makes, whether it is quantitative, qualitative or what. We all do it in our daily lives. We all are modelmakers and we all ignore factors that we think can be ignored. What we are looking for is consensus around a thing that can work, and that will always have compromises of one kind of another.

And I find most of this list of disadvantages very distasteful, and maybe they can use that comment in their further deliberations.

MR. SIGE: Well, thank you. Professor Perrow.

DR. PERROW: In that spirit, I have a helpful comment, too, for the panel. I am nonplussed by their conclusion -- I hope they explore it in much more detail, because we haven't been able to in our panel, Panel C, the conclusion that quantifying these goals is fine if there is no problem, that it is routine operation, but when you have

accidents quantifying the goals that deal with accidents is not really possible. I recognize the combinations and dependence and so forth, and I would say that that is what the thing is all about, and if not, I am very surprised, so I hope they devote the whole day to that.

DR. SLOVIC: I am sorry, we didn't say it wasn't possible. We said that there were lots of difficulties, that it was relatively simple, straightforward, to look at routine operations, and that when you got accidents, the problem got immensely more complicated, and I was trying to spell out some of the reasons why it was.

We have not come to closure yet in trying to define what is knowable here. Perhaps we can spend some time on it.

DR. PERROW: Okay, the preliminary list is good. The three things you have listed are useful.

MR. SEGE: Thank you, Chuck, and Lester.

Let us turn now to Panel C. Any comments on Panel C? Paul, you are not getting any comments at all. Yes, you are. Mr. Temme?

MR. TEMME: There was the suggestion at one point that the public should be consulted about their values, and perhaps it isn't a fair question, but my question is, how do we consult them? What do you do to consult the public about their values?

MR. SEGE: Dr. Slovic, would you care to respond to

that?

DR. SLOVIC: You can either look at ongoing behavior and try to infer important concerns and attitudes from that, or you can consult with various groups representative of different public interests, try to get some feelings for it and then decide -- but, for example, in the ACRS documents, there are questions, there are some philosophical value decisions embedded in those criteria, which I think are representative of attitudes, for example, the risk aversion coefficients.

But I don't think there is any perfect way of coming up with the public values, in fact there is no such thing in a sense, but we have to, you know, it would be implicit in whatever standards are set, I believe.

MR. SEGE: Thank you, Paul. Professor Okrent?

DR. OKRENT: I would like to make a request that each of the panels, including the one I am on, try to come up with either a position or positions or points of view on should there be risk aversion somehow included in what the NRC does, whether it be quantitative or qualitative, and if so, how, and should there be some kind of an ALARA used in design with regard to accidents? And if so, could they suggest how?

MR. SEGE: Thank you, Dave. Are there any other comments, either on Panel C, or comments that should be taken

up at this time, that may cut across various panels? I don't see any hands, so I guess we are ready for adjournment of the planary session. The panel discussions will be starting in approximately 20 minutes at ten minutes to ten.

(Thereupon, at 9:39 a.m., the planary session adjourned.)