UNITED STATES
NUCLEAR WASTE TECHNICAL REVIEW BOARD

PANEL ON STRUCTURAL GEOLOGY & GEOENGINEERING

PROBABILISTIC SEISMIC AND VOLCANIC HAZARD ESTIMATION

March 8, 1994
San Francisco, California

BOARD MEMBERS PRESENT

Dr. John E. Cantlon, Chairman, NWTRB
Dr. Clarence R. Allen, Chairman, SG&G Panel
Dr. John J. McKetta, Member
Dr. D. Warner North, Member
Dr. Dennis L. Price, Member

STAFF MEMBERS PRESENT

Dr. William D. Barnard, Executive Director, NWTRB
Dr. Leon Reiter, Senior Professional Staff
Dr. Victor Palciauskas, Senior Professional Staff
Mr. Russell K. McFarland, Senior Professional Staff
Ms. Linda Hiatt, Management Assistant
Ms. Donna M. Stewart, Staff Assistant

ALSO PRESENT

Dr. Keiiti Aki, Consultant, University of Southern California
Dr. Robert Budnitz, Consultant, Future Resources Associates
Dr. C. Allin Cornell, Consultant, Stanford University
Dr. Michael Sheridan-Consultant, State University of New York
Dr. William Melson, Consultant, Smithsonian Institute

I N D E X

PAGE NO.

INDEX (Cont'd)

P R O C E E D I N G S

DR. CLARENCE ALLEN: Could you take your seats, please, and let's get underway.

Good morning and welcome to the meeting of the Panel on Structural Geology and Geoengineering of the Nuclear Waste Technical Review Board.

I'm Clarence Allen, the Chairman of that panel, and let me introduce the other board members who are present; John Cantlon, Chairman of our board, and Warner North will be here presently we hope; Dennis Price, John McKetta and Ed Cording may be in tomorrow morning.

In addition, let me introduce some of our staff people who are here; Bill Barnard, Executive Director of the board; Russ McFarland, Victor Palciauskas and Leon Reiter. Leon, as a matter of fact, is almost entirely responsible for setting up this meeting, providing the speakers, and we thank him for those duties.

I should also point out that sitting on the far side of the table here are a number of consultants to the board that you will be hearing from for the most part later in the program; Bill Melson, Bob Budnitz, Allin Cornell, Michael Sheridan and Keiiti Aki.

The last time we met on seismic issues was more than two years ago, and the last time we met on volcanism was about a year and a half ago. A lot has happened since then, particularly in the area of hazard assessment. During the past year, the DOE and its contractors have produced, and are about to release in final form, two documents that assign significant roles to probabilistic hazard assessment in the Yucca Mountain program.

Our meeting during the next two days will be devoted to this topic; that is, probabilistic, seismic and volcanic hazard assessment, or in shorthand PSHA and PVHA.

We have decided to discuss both earthquakes and volcanism at this meeting. The structures of the probabilistic and volcanic hazard analyses are similar and face many of the same questions. Seismic analyses, of course, have a longer history, and many more have been done in the United States and around the world, although there's a history of volcanic hazard assessment at the Yucca Mountain project itself.

In addition, there is now evidence that there has been some physical coupling of earthquake and volcanic activity in the past in the vicinity of Yucca Mountain. Until now, they have been largely treated separately by the DOE.

With respect to seismic issues, emphasis at this meeting will be placed on the future use of probabilistic analyses and its validity in the Yucca Mountain program. There has, of course, been some criticism of probabilistic approaches to seismic hazard assessment.

With respect to volcanic issues, emphasis will be placed on the validity of the assumptions by many on the DOE side that probabilistic assessments won't change much in the future, the implication being that, at least with respect to certain aspects of volcanic hazard assessment, the Yucca Mountain program has already reached the point where enough is enough.

The board is particularly interested in the significance of any calculational differences in hazard. I might point out also that it's not the primary purpose of this meeting to debate the use of probabilistic approaches versus determination approaches. Clearly both have their place under certain circumstances. Rather, we wish to concentrate on probabilistic approaches, their strengths, their weaknesses, their future trends, specifically as related to the Yucca Mountain project.

Following are some of the questions that the board would like to be addressed within the next two days, and this list of questions has already been made known to the speakers and perhaps to many of you in the audience:

What are the objectives of PSHA and PVHA in the Yucca Mountain project? How will PSHA and PVHA be used in critical suitability, design and licensing decisions? What are the specific probabilistic criteria that will be used in decision making? If they are not now in place, how will they be generated and how will they be approved? How will PSHA and PVHA be used in programmatic decisions, such as priority setting? Are the existing or proposed methodologies sufficient to meet the objectives and criteria? What are the current and ultimate roles of expert judgment in these assessments? What is the role assigned to deterministic hazard assessment in the Yucca Mountain project? What would be the effect of increasing the time period of concern for post-closure performance from 10,000 to 100,000 years or more? Increasing the time period, of course, could be one of the recommendations from the NAS committee on Yucca Mountain standards that is now carrying out its deliberations.

What are the lessons to be learned from the use or the lack of use of PSHA in the siting, design and licensing of critical facilities such as nuclear power plants and other engineered facilities?

With respect to PVHA in the volcanic hazard analysis, how valid is the conclusion that estimates of volcanic hazard at Yucca Mountain won't change much in the future? What kinds of discoveries could cause them to change? What is the likelihood of these discoveries and the ability of site characterization to reveal them? What are the criteria for determining when enough is enough in both PVHA and PSHA?

Have outside investigators supported the way their PVHA estimates have been used in the Los Alamos report? What role, if any, will--what will be the role, if any, of the proposed geomatrix PVHA in the Yucca Mountain project using expert judgment? How well integrated are the seismic and volcanic efforts at Yucca Mountain? How much integration is appropriate or necessary? What are the differences between the probabilistic seismic hazard evaluation at Yucca Mountain for ground motion and for fault displacement? What are the differences between PSHA for pre-closure and post-closure? What are the significances of non-homogeneous and non-Poissonian models in PSHA and PVHA for Yucca Mountain?

And based on current knowledge and models, what are the most critical geological, seismological and volcanic studies that need to be undertaken at Yucca Mountain?

That's a long list, and I'm not sure we're going to get the answers to all of those questions, but that's our purpose.

Today we're going to start off the meeting with updates on the seismic and volcanic investigations. We have asked the speakers to give emphasis to those findings that have the most impact upon hazard assessment. We will also be hearing about a new integrated structural model for the Yucca

Mountain region.

Following these updates, we have asked Allin Cornell and Bob Budnitz to give two general presentations on probabilistic hazard approaches in their applications. The rest of today will be devoted to seismic hazard issues. We will first hear from the DOE and its consultants, followed by comments from seismologists and other interested parties. We have asked Keiiti Aki of the University of Southern California to sum up the seismic section by giving us his prospectives on the issues.

Tomorrow we'll use the same structure to address volcanic hazards. In this case, we have asked Mike Sheridan of the State University of New York at Buffalo to sum up the volcanic hazards. In the middle of the afternoon tomorrow, we will convene a round table made up of all the speakers to discuss both seismic and volcanic hazards, what has been presented in the past two days, and answers to some of the critical questions that the board has raised. We will also entertain questions from the audience, and several people have already been lined up to speak.

So let's get on with the meeting, and I'll remind you the meeting is being recorded. So everyone who uses the microphone, including board members, consultants, please be sure to identify yourself before speaking into the microphone.

So our first speaker this morning is John Whitney of the United States Geological Survey, who will give us an update on seismic investigations at Yucca Mountain. John?

DR. WHITNEY: When the Department of Energy wrote its topical report discussing the approach for seismic hazard methodologies at Yucca Mountain, a strong emphasis was put on the fact that there would be a significant database with which to assess seismic hazards at Yucca Mountain. And so the tectonics program in the U.S. Geological Survey is really devoted toward collecting data that will be useful in both assessing fault displacement through the potential repository block and seismic hazard analysis that's primarily directed at ground motion assessment.

The list of the questions that I got was very similar to what Frank Perry got, which was basically to discuss the findings in tectonics, and for us, that's really the last two years when the program was restructured and we were allowed to collect data again in the field, and emphasis on investigations and results that have the most and the least impact on seismic hazard analysis; but future investigations will have the most and least impact on seismic hazard analysis, and to make sure that we put this into the context of pre-closure, post-closure, surface and underground, ground motion and fault displacement aspects.

For us, we consider the pre-closure and surface facilities issues pretty much equal, if not quite, in that the post-closure and underground are activities that are also related to the 10,000-to-100,000 year period.

I'm going to try to start off with our most important findings of the last two years. We now have a complete inventory of Quaternary faults at the site. We've produced a Quaternary fault map that is now in press of the Yucca Mountain area. We now have a map in press that shows all the Quaternary active faults within 100 kilometers of Yucca Mountain.

We've completed fault behavioral studies on the primary, what we feel are the most important faults at Yucca Mountain in the immediate area; the Bow Ridge, Solitario Canyon, Windy Wash, Paintbrush, Stagecoach Road faults, Bare Mountain fault and the Death Valley faults, which are outside the site area. We'll have significant results that will be completed by the end of this September, our fiscal year.

We've completed the Midway Valley study, which is an assessment of faulting at the proposed surface facilities and near the ESF, at the ESF.

We've completed a 10-year GPS survey over the region. We now have 10 years of geodetical leveling data as well. We have an analysis of the Little Skull Mountain earthquake that's just been completed and some of its aftershock sequences, and that analysis of aftershock sequences will probably go on for quite awhile.

We have an initial assessment of relevant earthquake sources for the region, which we'll go into.

DOE has also just about completed a preliminary probabilistic seismic hazard of the ESF at Yucca Mountain. So we've actually gone through an exercise within DOE to look at what the real hazard was right there at the ESF, and it did come up with a couple of points that were different from the assessments that were made in the mid-80s. We'll talk about that.

And this year we'll also complete a preliminary tectonic model of Yucca Mountain.

The most important future studies for seismic hazard assessment are trying to gather data that we really don't have at the moment. It's not really refinement data.

We need seismic reflection profiles across Bare Mountain, Crater Flat, Yucca Mountain and Fortymile Wash to help us examine questions of fault geometry in interconnectiveness of faults. We want to complete the detailed mapping of faults within the proposed repository block. That's an ongoing effort within the site geology group at Yucca Mountain.

A very important study that started this year is an analysis of fault movement on the Ghost Dance and Sundance faults to determine whether or not there's any Quaternary displacement on the faults, the bedrock faults within the block.

We would like to refine the ages of paleoseismic events so that we have the best recurrence data and fault slip data that we can have for the analysis.

We would like to complete paleoseismic investigation of relevant earthquake sources for which we have no information at the present time that appear to be a contributor of ground motion to the site.

We are going to start the ground motion modeling of these sources next year.

We would like to refine our knowledge of fault geometries. If the seismic reflection line doesn't give us what we need, that we perhaps will try other geophysical techniques.

We want to assess the possible connections between faults. Is there a fault interconnectedness that will tell us something about the behavior of these faults?

We will hope to improve earthquake locations by completing the digital upgrade to the Southern Great Basin seismic network.

We hope to complete the modeling of local site effects on ground motions, and we will be refining the tectonic models as these data sets come in.

What studies do we do that have the least impact on seismic hazards? Well, in the SCP we said that we would look at the tectonic geomorphology of the region. We would look at folding in the Miocene rocks, that we would look at lateral crustal movement. And another large program that's just starting up is looking at basically the different tectonic effects on different aspects of hydrology and rock properties within the mountain.

Well, the only activity that we're actually doing at the moment is No. 4; we're beginning to look at tectonic effects, and these probably won't be used directly within seismic hazard analysis.

The first three activities, we really aren't doing them, and we believe that, at least in terms of looking at the hazard from faults within 100 kilometers, that the program that's going on right now will collect all that data. And so we don't really feel that we actually need to start these studies up.

So we're really quite focused to collecting data that's relevant to seismic hazard analysis.

Now, on your right is a figure that, for those of you who have been following this program you've probably seen for at least 10 years, of Quaternary faults of Yucca Mountain. This probably goes back to before the SCP, about seven faults that have been identified in the early '80s. In the fault map that we have just completed, we can now break the faults into three different classes, faults that we definitely know offset Quaternary units, and that's what these red lines represent. That's the actual fault segments which we know cut Quaternary deposits.

There's a second class of suspected Quaternary faults, and that's usually where the fault comes up against the bedrock ridge and doesn't--it's actually offsetting bedrock, but it's usually an extension of one of these Quaternary faults. And then we have bedrock faults for which we have no evidence of Quaternary offset.

And actually the number of faults that display Quaternary offset did not change, but we have more segments actually at the present time. And so one way we're looking at the behavior of these faults is by trenching them. It's a real classic approach, and just to give you an idea of the volume of information that we'll have, we have 26 trenches on the faults right at the site that will be either in some stage of completion or will be complete by the end of this year. There are another 10 trenches which did not yield any tectonic information, either lineaments were trenched or segments of the fault that were not active had been trenched, and they're not included in this list.

I'll give you two quick examples of the kinds of paleoseismic information that we have. In the Bow Ridge fault on your left here, we have a record that it goes over

200,000 years with approximately five events in it, the last event being somewhat older than about 70,000 years. Offsets on the Bow Ridge, which is a rather short fault, of the five or six kilometers, only the offsets are on the order of 10 to 20 centimeters. The recurrence interval is between 60 and 100,000 years of those four events.

We have a rather spectacular fault exposure over on Busted Butte, 60 meters of exposure, and we have a total of net slip of over five-and-a-half meters vertical, and probably as much as seven meters of total accumulative slip over about a 700,000 year period.

We have three stone lines and three buried soils that were offset, as well as colluvial wedges along the main part of the fault that we could discriminate.

The two upper soils, which we thought would be around 100 to 150,000 years old, turned out to be older than we anticipated. The upper soil in green is about 300 to 350,000 years old, and the youngest soil in brown at the top is over 200,000 years on the downthrown side, and about 100 to 150,000 years on the upthrown side, which tells us that there was--the second to the last event that created that scarp there was somewhere between 200 and 100,000 years old. The soil reformed on the upthrown side.

We have one, possibly two, events that are younger than that 100,000 year soil, and we have a TL date of about

35 to 40,000 years of the unit that's unfaulted.

So the slip rate on this fault, the Paintbrush Canyon, which we've determined is one of the primary sources of hazard at the site, is about .01 to .02 millimeters per year and has a recurrence interval of about 40 to 60,000 years. The recurrence interval is very similar, but the displacements are--displacements on the Paintbrush Canyon fault range from 20 to 1.2 meters, and they average about 60 to 80 centimeters per event.

So just a preliminary summary of the paleoseismic data on the faults at Yucca Mountain themselves, fault lengths vary from 5 to 20 kilometers. The number of events generally ranges from two to five for the past 100,000 years. The displacement sizes are about 10 centimeters to a meter. Recurrence intervals range from about 20,000 to 100,000 years; slip rates from .001 to .02 millimeters per year.

That range has been from the fault work that was done in the mid-80s, we were primarily in the .001 to .008 category, and with the new work we've done, we have more evidence of getting back to about a hundredth of a millimeter per year, and that's consistent for several of the faults. I think we're getting to a point where we're getting convergence on rates now.

On one of the faults, the Windy Wash fault, at the southern end we did a shallow seismic reflection line to see how much offset there was on a 3.7 million year old basalt, and what we were able to show was that the basalt was offset about 95 meters vertically and about 100 to 110 if you add a left oblique component to that for a total offset.

So that showed us that, or demonstrated to us, that the long-term offset for about three-and-a-half million years is about .03 millimeters per year. So that rate is very similar or just slightly faster than these Quaternary rates. So what we're seeing is a long-term consistency in offset rates at Yucca Mountain. We're not seeing an increase in Quaternary activity. It's either constant or slightly decreasing, which should help in the predictability of the faults.

In the study of regional Quaternary faults, as I said, we now have an inventory of these faults, and they've been very useful to Silvio Pezzopani and Dave Schwartz in assessing relevant earthquake sources, which we'll get into.

Two of the specific studies that are being done at the moment are on the Death Valley fault system and the Bare Mountain fault system. And just to show you how important the study of these regional faults are, in the mid-80s, the Bare Mountain fault was considered to be the primary source for ground motion at Yucca Mountain, and the original estimation was that there was a Holocene event on the Bare Mountain fault that had a recurrence interval of about 20 to

150,000 years, a slip rate of almost .2 millimeters per year.

However, we have just completed a study, a trenching study, on what appeared to be the largest scarp in Quaternary materials at Yucca Mountain near Tarantula Canyon. We did not find any evidence of Holocene offset, and, indeed, we only found one or possibly two events in the whole trench.

So we have decreased the slip rate significantly in order of magnitude on the Bare Mountain fault. And we've looked at the ESF results. You'll see that that drops the Bare Mountain fault as being a significant source for ground motion in terms of a hazard assessment.

In Death Valley, we go the other way. Published estimates were for slip rates of about .2 to 2.5 millimeters per year, recurrence intervals of about 1,700 to 3,700 years between events. However, the recent work that the Bureau of Reclamation has done has shown that the recurrence interval may be as low as 500 years, in the range between 500 and 2,000 years per event, and the slip rate is as high as four to eight millimeters per year.

So this becomes a far more significant source for ground motion for low level frequencies.

We have quite a number of faults that really haven't been studied at all in the region, and that's one of our larger tasks ahead of us. This here gives you some idea of the length of the Death Valley/Furnace Creek system, and these little crosslines here are our best guesses for fault segmentation at the present time. That's very preliminary.

The study that was completed last fall, and the final reports are being completed as we speak, are for the Midway Valley study, the assessment of possible Quaternary activity near the proposed surface facilities, and a trench that was something like 360 meters long was put across the reference conceptual site. It crossed several of air photo lineaments and suspected faults, as well as a trench was put on the projected northward projection of the Bow Ridge fault. This down here is Trench 14.

And in this trench, what we found were two zones of fractures that had a North 15 East trend to them, whereas you can see our actual lineaments had a northwest trend for the most part.

So the zones of fracturing did not really--were not reflected in the surface at all, and that follows the fact that these fractures did not come to the surface. They were actually only found in middle Quaternary age deposits; that is, over 130,000 years old, and they did not extend up into late Pleistocene or Holocene deposits.

There was no vertical offset that was--vertical separation found on any of these fractures, and there was no evidence of lateral separation either.

So we have concluded that there were no significant faults, that is faults with greater than five centimeters of displacements during the last 100,000 years at the reference conceptual site.

Another positive aspect of this study is that the fault that is found on the east side of Exile Hill may serve as--in the study of it, may serve as a calibration fault for an intrablock fault that may be correlated to the behavior of say, Ghost Dance and/or Sundance faults.

So what do we do with all this data? In this past year the relevant seismic source program has started up, and using the data collected by the Bureau of Reclamation for the 100-kilometer region, Silvio Pezzopani and his colleagues have assessed maximum magnitudes as best they could, given the fault parameters that they had to work with. And where the data is in parentheses, these are basically estimated because there is no data. So these are--I think there are 26 or 27 relevant sources, going all the way out to about 97 kilometers.

All these have been characterized, and this table is continually updated. This is the fourth version of this particular table that you have.

Silvio has plotted these sources in their magnitudes against their distance from Yucca Mountain, and then looking at the attenuation or the peak acceleration relationships that Boore, Joyner and Fumal have constructed for--we think it's, if I remember, it's somewhere between 40 and 50 instrumented earthquakes primarily from California, you can see that the peak acceleration for Site A bedrock sites, if we look at what we call significant faults from the NRC at .1g, we have about half, or a little over two-thirds, of our primary Quaternary faults become relevant earthquake sources. If we move out to the 84th percentile, then we include quite a larger number of the faults in the outer areas there.

So this is how we are evaluating our data in terms of ranking the importance of the faults to be studied and helping us to select which ones need to be trenched at this point in time. And, of course, this information is extremely valuable to our ground motion modelers.

One thing that we are questioning is whether or not the peak acceleration is really an adequate measure of the damage potential at Yucca Mountain, both for surface facilities and underground, and we actually believe that using spectral velocities that span the frequency bands of engineering significance is probably a better way to assess relevant earthquake sources, and actually we can't do that until we have some feedback from the engineers as to what kind of structures will be designed.

This here is the peak ground motion acceleration combined attenuation for the sources at Yucca Mountain, or the ESF actually, and this comes from the DOE ESF/PSHA study, but it shows, and this model is composed of several attenuation models, it's combined, that the hazard up to 3 or 4,000 years is totally dominated by the background earthquake. And between 2 and about 6 or 7,000 years, you get a very small amount of input from the primary faults at the site, the Paintbrush, the Solitario Canyon and Fatigue Wash faults.

When you get down to 10,000 years, you begin to pick up a fair, a significant component of hazard from the Paintbrush at .4g and the Solitario at .3g.

So as you move, is that the period of concern increases, the hazard then becomes more dominated by the local faults at the site.

In terms of modern deformation, we completed a ten-year survey, trilateration survey, over about a 50-kilometer radius in the Southern Great Basin, and the amount of strain that was recorded is basically insignificant. In fact, the amount of strain is actually--these microstrain units are actually lower than the precision units over that distance there.

One thing that they were able to pick up, they picked up the Little Skull Mountain earthquake event and calculated based on the published moment magnitude and assumed a rupture area of about five kilometers, a total slip of about .6 to .7 meters for the main shock in Little Skull Mountain earthquake.

The Southern Great Basin seismic network was transferred from the U.S. Geological Survey to the University of Nevada at Reno, and we are going through an upgrade of that network so that it will become digital. We'll be able to look at, and we'll be able to record smaller events and obtain better locations for these epicenters and focal mechanisms for events in the area. It's a fairly sophisticated system, one of the best in the world when it's completed.

This is the earthquake, or the seismic catalogue, from 1978 up to the main shock at Little Skull Mountain. The small circle there is around Yucca Mountain itself, and as you can see, the seismicity in the Southern Great Basin is very sparse, and this has been commented upon by many people.

The Little Skull Mountain earthquake event was, from a seismic hazard standpoint, a very positive event.

As has been published recently, or last year, most seismologists and geologists interested in tectonics are now convinced that the Landers event did trigger seismic activity in several areas in the edge of the Great Basin and in California north, and the Little Skull Mountain earthquake appears to be one of these events, which has extremely low microseismicity and has for--that's been its characteristic primarily since 1978. We had an increase of foreshock activity. We got into the tens of microearthquakes for about 24 hours before the main shock, and then after that we had hundreds of aftershocks after the main shock. In the first six months, there were about 3,800 aftershocks that were recorded.

Work that has been completed by Kent Smith and his colleagues at UNR show that the event which UNR Seismic Lab believes is a 5.8 magnitude, the information from the National Earthquake Center is a 5.6. Their best solution is that the earthquake took place at about 11.77 kilometers depth on a fault dipping to the southeast that is sub-parallel to the Rock Valley fault system. And so the seismologists at UNR believe that this might be evidence for a slip partitioning on faults in the Yucca Mountain area.

The Rocky Valley fault system is a left lateral fault system, and this solution is consistent with that interpretation.

The aftershocks primarily are concentrated between six and ten kilometers, and actually, there were several structures in the immediate area that movement also took place on.

The amount of information in aftershock data collected by UNR really provided a great database for assessing site effects at Yucca Mountain and for ground motion modeling.

And just in terms of fault displacement through the repository--well, before doing that, I think I'll run through--the data that we hope to collect this year from the intermediate seismic reflection profile we hope will give us the fault geometry of the Bare Mountain faults and hopefully the Solitario Canyon, Windy Wash faults. Do they merge at depth? Do they shallow significantly at depth, i.e., is there a detachment fault under Yucca Mountain, and if so, at what depth? Out best hope is to basically image that contact of the volcanics against the paleozoic carbonates.

Chris will talk a little bit more about tectonic models and how they've evolved, but this is one of our keys that we hope that we'll have significant refinement of this year for seismic hazard assessment in terms of ground motion, fault geometries.

DR. ALLEN: John, two minutes.

DR. WHITNEY: Right. To assess fault potential through the repository, we're working together with the site geology group under Rick Spengler and they are looking at--they are creating these three dimensional models of Yucca Mountain, and they are doing detailed fault mapping at a scale of about one to 250, which is giving us an inventory of faults, secondary faults, fault splays, fracture patterns that will be used in that assessment of faulting through the repository.

And the ESF itself within the next year will go through both the Bow Ridge and the Drill Hole Wash faults, which will give us a chance to examine these faults at depth, and especially to see what's in them.

So to summarize where we're at, I've put together this little list of where I think we're at for a database to do seismic hazard assessments at Yucca Mountain for ground motion and faulting through the repository.

Geologic mapping is nearly complete. The regional work is fairly well done. It's just the work at Yucca Mountain over the repository block that needs to be completed. Site fault characteristics, I think we're about 85 per cent complete there. It's primarily a documentation exercise we have to go through in the next year. In the regional faults, we're nowhere near that secure in our knowledge. I think we've got about 40 per cent of the information that we need. There's probably at least a half a dozen faults that need to be studied.

Geophysics, fault location, I think from the aeromag and gravity that we know where the faults are. The 65 per cent confidence on subsurface geometry is basically going to come from the seismic line that hopefully will be run this summer.

Tectonic models, I think we know the bounds of our tectonic models, and we have some preferred models at this point that Chris will talk about.

In terms of modern deformation, we have a completed GPS survey. We hope to put another one across the Walker Lane. We have quite a bit of geodetic data. We have done a comparison of historic level lines that's not complete. We have in situ stress data from the early '80s, which we may add to if one or two of the geologic holes are bored on Yucca Mountain. And we will complete a revision of the historic earthquake catalogue by the end of this year, so that should be available. And the catalogue for the modern activity, of course, will be available.

Modeling for site effects, I think about 70 per cent of that work will be completed by the end of this year.

The assessment of relevant earthquake sources is early about half done and, of course, has to be revised with new information as it comes in.

Ground motion modeling is sort of getting off the ground, and our development of the seismic hazard analysis is also just beginning. The 15 per cent kind of represents the topical report and the fact that we'll complete a study plan this year.

DR. ALLEN: Thank you, John. One quick, but perhaps provocative question. As a result of the Northridge earthquake and other recent large earthquakes in California, I have heard a number of people, including some geophysicists, arguing that the study of surface faults is becoming increasingly irrelevant to the understanding of seismicity and the quantification of seismicity; arguing instead that somehow surface faults were some form of damage, the result of shaking or something, that was basically unrelated, or at least not directly related to the earthquake at depth.

Do you have any comment on that increasing skepticism of the relevance of geology?

DR. WHITNEY: Well, I think the geology, the difference in the structural environment between the Northridge area and Yucca Mountain is quite a contrast, and the basin and range faults have quite a bit more predictability in terms of their behavior, in their normal behavior as well as their geometry, although we're still working out geometry.

But in terms of having blind faults and blind thrusts in the Yucca Mountain region, I don't think tectonically we have that kind of environment at all.

There is going to be some discussion about amplification site effects, I think. The UNE work shows that there is site amplification of about 1.7 to 2.2 that could either be controlled by topography or some property, physical properties at Yucca Mountain. But I don't think that we are going to deal with structures that we can't model.

Furthermore, the background earthquake, which is--which actually is the primary hazard for nearly half of the 10,000 year pre-closure period, should include these aerial structures or the ones that we can't--we have no surface evidence for.

DR. ALLEN: Other questions from the board? Staff? Consultants? Keiiti Aki?

DR. AKI: You showed this strain accumulation from GPS measurements. How does this strain in depth compare with geology?

DR. WHITNEY: With the geology?

DR. AKI: Yes. Have you tried to compare?

DR. WHITNEY: Well, the characteristic that drives the hazard assessment at this point in time at Yucca Mountain for the near source faults, the faults within five, six kilometers of Yucca Mountain, is the very long recurrence intervals. When you have tens of thousands of years between earthquakes, the amount of strain that is accumulated over a ten-year period, it certainly isn't out of character to not be able to see that accumulation over a ten-year period. If you were, of course, to move over to Death Valley, you'd have another story. But for these faults with very, very low recurrence intervals, we're not seeing strain accumulation.

DR. ALLEN: Okay. Thank you. I think we must move on here. Before I introduce the next speaker, let me point out that we have with us today five representatives from the PNC of Japan, which is a research group on high-level radioactive waste disposal in Japan, and I would simply like to welcome them here today.

Further, I should point out that since I introduced the members of the board, Warner North showed up, and he is also with us this morning.

DR. NORTH: I apologize for the traffic delay.

DR. ALLEN: The next speaker will be Frank Perry from Los Alamos National Laboratory, who will bring us an update on volcanic investigations.

DR. PERRY: After reviewing briefly some of the areas of progress in the last year, I'm going to spend the bulk of the talk talking about the Lathrop Wells volcano, which is 20 kilometers south of the proposed repository, and it's the youngest volcano in the region.

In the past year, we spent a lot of time wrapping up the Lathrop Wells studies, and we've come to a number of conclusions about its history, which I think has important implications for risk assessment for Yucca Mountain, mainly in that we believe it gives us some spatial controls--on the location of future volcanism.

So I'll spend some time speaking of the evidence for Lathrop Wells having a long history of polycyclic volcanism.

Just to briefly remind everyone, the volcanism studies are divided into three main areas, three study plans. It's a lot smaller program than the seismic studies. The three study plans are characterization of volcanic features, which provides geochronology data, petrologic studies and field studies, probability of magmatic disruption, and then physical processes and effects of magmatism should it intercept or come near the repository. And characterization is the basic data feed into these other two study plans.

Some of the areas of recent progress, the regional geochronology is well under way. We have both Lehigh University and the New Mexico Bureau of Mines under contract to do $^{40}Ar/^{39}Ar$. They are mainly dating the centers older than Lathrop Wells back to about five million years, and so far we've dated about half of the centers that have been active since five million years ago. We've also dated the one aeromagnetic anomaly that's been drilled near Armagosa Valley commercially, and I'll show the dates on those things.

We're also proceeding the geochemical and geochronologic sampling for the rest of the centers in the Yucca Mountain region, including Buckboard Mesa.

As I mentioned, the work at Lathrop Wells is in a wrap-up phase. We've concluded that it is polycyclic, has erupted in four main eruptive episodes covering a time span of about 100,000 years. This has involved a minimum of six to eight magma batches, and the importance of that is that it means there's been six or eight separate diking episodes into the shallow crust concentrated at Lathrop Wells.

Currently we're using sanidines enclosed within tuff xenoliths that are in the lava flows at Lathrop Wells to refine the geochronology there, and this is pretty much our last major effort to add any more geochronology information at Lathrop Wells.

Greg Valentine has gotten started on his magmatic effect studies. He's completed field studies at Paiute Ridge on the test side and Alkali Buttes in New Mexico. These are analog centers; Paiute Ridge, to look at the effects of dikes intruded into tuff, and Alkali Buttes, there's a number of eruptive styles there, and he's looking at the amount of wall rock incorporated into the different eruptive episodes of this center as an analog for incorporation of waste.

He's also--he got sensitivity studies from modeling liquid and vapor flow in the unsaturated zone in response to a magmatic intrusion.

This is an example of some of our new Argon/Argon dates. These are results from Crater Flat that we've gotten in the last year. Our results are the open symbols, and it compares dates that the geological survey got in 1982. These are the results for the 3.7. In '82, conventional potassium

Argon indicated an age of about 3.7. I think the individuals were about 3.6 and 3.8. We've gotten three new analyses, Argon/Argon from these centers, and they all come in right about 3.7.

We've also done the Armagosa Valley aeromagnetic anomaly, and it comes in about 3.8 million years old. And at this time we conclude that this is a part of the 3.7 million year episode because the ages are so close.

We've also dated Black Cone and Little Cones in the million year cycle at Crater Flat. The previous dates were just a little bit over a million years, with fairly large errors. Our results, four dates from Black Cone and one from Little Cone on the basalt, show that they erupted at right about exactly a million years ago, and you can see no difference between Little Cone and Black Cone.

We've also gotten a sanidine separate from New Mexico Bureau of Mines. This is from one of the Little Cones, and it gave a high precision number of about 905,000, plus or minus 10,000, and that's within error of the Argon/Argon basalt date from the flow at Little Cone.

So what we see from this is that in 10 years, using a different method and a higher precision method, basically the numbers don't change. We still get 3.7 for the oldest cycle in Crater Flat and a million for the youngest cycle, but the precision is a lot better. It's a factor of two or more better, so that we know these dates are higher precision.

So we think that dating the other older centers in the area is going to be fairly straightforward from these results, and we don't see it as any type of problem.

Now, I'd like to start talking about polycyclic volcanism and begin by emphasizing the difference between a monogenetic and a polycyclic volcano.

When we came into these studies, we assumed, like just about everyone else, that small volume basalt volcanoes, like you see in the Yucca Mountain region, are monogenetic, meaning that they erupt during one episode over a period of weeks to several years, and although their plumbing system may be complicated, what it is basically is that one dike intrusion episode bringing one magma up to erupt. And once this eruption is over, the center's effectively extinct, and there will be no further eruptions at that center. So it's a fairly simple type of volcano.

In the last few years, we've been gathering evidence that some of the small volume volcanoes in the area are actually polycyclic, meaning that they erupted in several discreet eruptive episodes over periods of tens of thousands of years. This would necessarily involve several generations of independent dike formation and probably different magma batches, and I'll talk about the evidence for that later.

The scale of these two types of volcanoes is the same, but the polycyclic volcano is a much more complex volcano. And part of the reason this wasn't really widely accepted in the community is that people felt this small a volcano couldn't be this complex, but we are seeing this complexity at Lathrop Wells.

What we've concluded at Lathrop, based on field and geochronology studies, is that there have been four main eruptive episodes covering a time span of about 100,000 years. Geochemical evidence, which I'll go over, indicates multiple, independent magma batches. And evidence of Holocene eruptions, which I'll also review, indicates that the center can be considered to still be within its polycyclic lifetime.

And the implications for volcanic risk assessment are one, that the effect studies must consider multiple eruptive episodes, as indicated here in the schematic of the funding system of the polycyclic volcano. Because this is a repeatable pattern at one location, it provides a constraint on the location of future volcanism. In the case of a monogenetic volcano, once the volcano has erupted, it becomes extinct, and any future volcano in the area won't necessarily form a new volcano at some unconstrained location.

So this is a volcano that has no pattern, and here we have a polycyclic volcano that does.

With this in mind, disruption probability calculations, which assume a random distribution within particular volcanic event zones, can be considered conservative.

And last, you know, considering the history of the Lathrop Wells volcano, the most likely volcanic event in the Yucca Mountain region during the next 10,000 years we believe will be another eruption at the Lathrop Wells center.

This is pretty much our final map of the Lathrop Wells center. The field studies are pretty much complete. What we've concluded is that it did erupt in four eruptive episodes. The oldest episode is shown here in blue. It's the southernmost flows, and one flow to the north. It erupted from several north to northeast trending fissures, which are marked by these scoria mounds, which are in general fairly well eroded. Some of these showed dikes. There's been enough erosion to expose the underlying dikes.

Helium indicates that these have a minimum age of about 80,000 years. The southern flows where helium was done is shown by trenching and field studies that these were covered by a minimum of about two meters of tephra from the second eruptive episode. So these flows were covered, which attenuated the acquisition of the helium signal.

DR. ALLEN: Frank, excuse me for interrupting, but for the benefit of consultants or others who may not be familiar with the region, can you just say what the relationship is geographically between this cone and the repository site or--

DR. PERRY: Yeah, this is about 20 kilometers, pretty much directly south of the repository, a little bit southwest.

So, again, we feel that the helium gives a minimum, and we think these flows probably approach 100,000 years or older.

The second episode produced the most voluminous flow to the east of the cone, shown here in green, also erupted from northeast trending fissures. There's some other events over here. We're not sure what they fed, but they, from chemistry and field relations, they appear to belong to this episode.

It also produced a voluminous fall sheet, which is up to two meters thick. That's shown in the spotted green pattern. This is the most likely--we found this deposit as far in place, in stratigraphic context, as far as three kilometers north of the center. This is the most likely ash that's found in fault exposures in the trenching studies that have been done near Yucca Mountain.

The third episode produced the main cinder cone and a small flow to the north of the cone; again, from northeast trending fissures from the elongation of the cone. We have no evidence that the cone itself produced a voluminous fall sheet, which is kind of a surprising result, but I'll go through some evidence of that also in a minute. Helium ages indicate an age of the cone of somewhere between 40 and 60,000 years.

And then the last episode shown in red are these very small tephra deposits south of the cone. It's about two or three small volume tephras that overlie in some places the cone deposits that are separated by soils, and thermoluminescence ages indicate ages younger than 9 to 4,000 years.

And one thing we've been doing a lot in the last year is using chemistry to constrain some stratigraphic relations and also for petrologic models, and the way we've been looking at differences between--in chemistry between these four eruptive episodes is to construct a series of spider grams. These are by element, about 17 trace and major elements, all normalized to an average Lathrop Wells composition, which is about 99 trace element analyses.

So what we're looking at by normalizing, we're just looking at differences in chemistry between different eruptive units.

So this is an example in black showing a flow--this flow here from the oldest eruptive episode, and in red, this flow here in the peach color from the third eruptive episode.

And what we have is four analyses from the oldest flow and three from the youngest flow, and I just want to show what kind of differences we can pick up.

The total spread of each pattern, say the black or the red, includes both the--reflects both analytical precision of the analyses and also any internal heterogeneity within an eruptive unit. So you can see that they're fairly reproducible. There's not that much heterogeneity within a flow, and for most elements, it's fairly reproducible.

So you can see that these two flows are quite distinct in their chemistry in elements like thorium, strontium, phosphorous, the middle rare earth and titanium.

So we use these differences to constrain petrologic models, which relate these different eruptive episodes, and we've also used a lot to constrain some of the field relationships that are a little bit tricky, and in some cases, eruption dynamics in the case of the cone erupting and what type of distal fall sheet that are produced.

Here's an example showing that if you have enough samples, you can use fairly small differences in chemistry and get some useful information. What this is in black is an average of 15 of those patterns from the main cinder cone. And in the open symbols, an average of eight analyses from that fall sheet from the preceding eruptive episode. This is the distal fall sheet, which is the most voluminous scoria fall from the center.

And what we can see is that in the case of thorium and titanium, just doing a student t-test of the means, is that the differences are statistically significant. In these cases, there's about a 1 per cent and a 2 per cent probability that these means come from the same population.

So what we conclude from that, and also in this particular relationship, trenching and field studies, is that the fall sheet, which at first we thought came from the cone, which is the most likely source for it, didn't come from the cone and actually came from a preceding eruptive episode. So this type of information will help us when we try to correlate to ashes which are exposed in the trench and try to assign an age to those ashes in the trench to help constrain some of the fault recurrence rates.

And it also tells us something about eruption dynamics because we have a cone which apparently didn't produce a very voluminous fall sheet.

Just to summarize all the chemical differences for the four eruptive episodes, the top frame summarizes the first three eruptive episodes, oldest in blue, the youngest in peach, coded to the map. Again, you see significant differences, so we have a unique geochemistry tied to each eruptive episode. On the bottom, the same three at a different scale, showing how the first three eruptive episodes compare in chemistry to the youngest episode, which we think is Holocene.

In the youngest episode, we see some very different chemistry, and this has really cemented our conclusion that these youngest episodes did represent primary volcanic events. They weren't reworked from any older material because there's--you know, from all these analyses, physically there is nothing older that these chemically could have been reworked from. They're very high in rubidium and thorium, also the heavy rare earth elements.

So this really kind of finalizes our conclusion that these youngest events at Lathrop Wells were, in fact, new volcanic eruptions and represented new magma intrusions into the crust.

Now, I'd like to go to--on the slides. Okay. What I'm going to show briefly in four slides is the evidence and the chemistry of these youngest eruptions. What we'll be looking at is this area south of the main cone, and we'll be looking specifically at these two tephra deposits, one which directly overlies the distal edge of the main cone, and this one that sits above that in some sand units. Is that focused?

So what we'd be looking at is this area here. These red deposits here are the distal edge of the main cone. They're the upper part of the outer cone slopes, and we'll be looking at tephras that lie above this separated by soils.

The next slide?

This is just a closeup of that deposit. You go from red and grade into black here, and that's the uppermost part of the main cone deposits. The silty layers above that include soils and tephras of the youngest deposit, which we have evidence for being Holocene.

Next slide. Thanks.

This is a closeup of the deposits that overlie the cone. Way down here you can see a little bit of black where we dug a hole. This is that uppermost layer of the cone slope. There's two soils developed, one in the top of the cone deposits and then one--there's a tephra unit in here, which is so infiltrated with carbonate dust, we haven't been able to analyze. There's a soil developed in that, and then overlying that is this tephra deposit, which on volcanological grounds we would always argue was primary. It has a planar top and bottom, is sorted how you would expect a primary deposit to be sorted.

The chemistry of that--what this is, is the same type of plot comparing the chemistry. The lower two patterns here are the upper part of that cone deposit, the red and the black unit. Then this pattern in red was this unit here, the hydrovolcanic unit, which is that one that's very unique in its chemistry.

So what you have are two very different tephras in terms of chemistry, separated by this soil here, and this was the first evidence--this soil was the first evidence that led to the idea of this being polycyclic. Now with the chemistry, we're confident that it does represent a new eruption separated by time.

We have thermoluminescence dates on this soil within the upper cone soil deposit of 9,000 years, which dates the emplacement of this overlying tephra. We have a thermoluminescence date of 4,000 years on this soil, which would date the emplacement of this hydrovolcanic unit on top of that.

Next slide.

Then within the sand above that unit, this is that other red unit I showed on the map. We've recently discovered in May this other tephra deposit which sits within sand. It's slightly cross-bedded and reworked; again, planar top and bottom.

The chemistry of that, this is compared to the distal cone slopes of the main cone in black, and red is this deposit. It's very similar to the cone. It's very different from that hydrovolcanic unit again. It can probably be distinguished from the cone in terms of thorium content. It also has a slightly higher Mg number.

Our work with this is pretty preliminary, but at this point, we feel this represents the youngest eruption from Lathrop Wells, and at this time we have no date on this, but we're considering doing a thermoluminescence date on the material underneath this deposit.

Okay. You can turn off the projector.

Now I'd like to talk just briefly about evidence for multiple magmas of Lathrop Wells, and this is important because each different magma separated by time must have been in place by a separate episode of dike intrusion.

One of the most important constraints on distinguishing different magmas is this observation that the Mg numbers of the magmas are very much the same for all magmas we've--for all the lavas we've analyzed. This is 121 analyses, and they sit at a value right about 54.

What Mg number is, it's a measure of how evolved they are from a primitive basalt that's produced in the magma. A primitive basalt would have a number of about 70. So these are quite evolved. This involves 20 or 30 per cent fractionation to get down to this number.

And in light of the chemical variations I've shown you, there's really two ways we can think that this could be produced. One, you have separate magmas coming up, and there's some type of density filtering going on where they can only, you know, send and erupt at the surface after there's been a certain amount of fractionation and they've reached a certain critical density where then they can go on and erupt. Or, you may have only one magma involved, but you have complex processes going on where possibly recharge is going on to a magma, and you're buffering the Mg number. You have enough input of primitive magma coming in that you have a buffering going on where the Mg number reaches a steady state value at about 54.

And the approach we've taken to look for different magmas is to look at Mg number versus several different incompatible element ratios. This is thorium/potassium. Again, we see these units at the same Mg number. This goes from the first eruptive episode and increases steadily as you get to the third eruptive episode. These are all the major flows in the cone at Lathrop Wells.

For thorium and potassium, they're both highly incompatible in any fractionating phase in a basalt. So if you were fractionating, you wouldn't change the thorium/potassium ratio. It would stay the same, and you would just decrease the Mg number.

So these differences, systematic differences you see in thorium/potassium, must be related to different magmas, and what we've concluded from this, that there are at least four different magmas involved for these different eruptive episodes.

You see the same type of thing for lanthanum and samarium, but in this case, there are ways to fractionate lanthanum and samarium because lanthanum is more incompatible than samarium. If you had a large degree of pyroxene fractionation, which fractionates those two elements, it's conceivable that you could get a spread like this if recharge was going on to buffer the Mg value at a certain value.

So what we've done is set up a series of equations to model recharge assimilation into a magma, or affecting a magma, and this is an example where we have a high amount of recharge going on relative to crystallization. And if I can just show you a couple of panels here.

What this shows, this is Mg number versus magma mass, and it shows for a sufficiently higher recharge, you can buffer the Mg number at a certain value. The recharge is set here to buffer at a value of 54.

And in the case of lanthanum and samarium, this is the real data here. You can produce an evolutionary path for a magma where you reach a steady state in Mg number, but still continue to evolve a lanthanum/samarium ratio. So, but in this case it involves 75 per cent fractionation of pyroxene for the whole assembly. So it's a high amount of pyroxene fractionation, and I don't think that's realistic. For other things it doesn't fit so well. For lanthanum/samarium versus lanthanum, you still can't get the extreme lanthanum/samarium fractionation you get for only a small amount of lanthanum enrichment.

So in detail, I don't think this model works. For elements like thorium, potassium, which are both highly incompatible, you can't get any fractionation. You can never produce something like this. Even with a large amount of pyroxene or any other type of fractionation, you still get a fairly flat trajectory.

So we've done this model to look at more complex scenarios, but the evidence is still that there are multiple magmas, and even complex processes can't explain in detail what's going on at Lathrop Wells.

If we use the stratigraphic model we've come up with, the four eruptive episodes, we also see some systematic changes through time and certain trace elements. This is bieruptive episode from oldest to youngest. We see increases in thorium, potassium, lanthanum and samarium to some extent, decreases in titanium, and increases in thorium. We're still working on this, but we think these are related to processes in the mantle, either changes in the amount of melting through time or depletions in the source as you extract out different increments of melt.

We've also been using some of the major element analyses to look at the same type of thing. This is a normative plot, looking at the amount of silica saturation. Under-saturated lava is on this side, and saturated on this side, going from nepheline to hyperce normative. This is, again, eruptive episode. The first is slightly nepheline normative, and as you go through each eruptive episode, they become progressively more silica saturated.

This may be--we think, again, this is due to mantle processes, may be due to any combination of amount of melting changing, the depth of melting, you know, the pressure at which it's melting, or the volatile content in the source.

For Black Cone we see basically the same type of thing. We see two different flows of Black Cone that are related in a geochemical way that can't be explained by fractionation from one batch. So we see the same pattern at Black Cone, and our conclusions there are that it's also polycyclic.

So if we look at the region, what we think is that polycyclic volcanism may be pretty typical for the Quaternary. We see evidence of polycyclic activity at Black Cone and Crater Flat, also Red Cone. Gene Smith has done work there, and I think his conclusions are that it's also polycyclic.

At Sleeping Butte, we have some evidence that it's polycyclic, but we still need to go in there and really do some more work, and then, of course, at Lathrop Wells.

If you count the magmas at these centers, assuming that the ages all come into the same at Crater Flat, the ones we've had so far, everything's coming in at about a million years. Assuming this pattern holds up, in some ways we think of Crater Flat as really be a distributed polycyclic center that's just spread out along some structure. And if you count from chemistry the separate magmas that would be involved, it's about seven for Crater Flat. At Sleeping Butte, about two. The two centers are different in their chemistry. There's a possibility of one younger eruption from one of the centers, but we haven't confirmed that, and then at Lathrop Wells, a minimum of about six.

And so looking at the history of Lathrop, it has this 100,000 year pattern of repeated volcanism, and it's been maintained into the Holocene, assuming our evidence for these Holocene eruptions is correct. It indicates to us that the most likely eruption in the region will probably be another eruption at Lathrop Wells.

This is a block diagram of the region based on that map.

What we see, then, if each magma represents a new diking episode of intrusion into the shallow crust, what we see is a strong pattern in the diking activity in the last 100,000 years. So what this portrays is what we infer for the last 100,000 years what the diking episodes have been. We have multiple episodes at Lathrop Wells, possibly one at Sleeping Butte, but we have to do some more work on that.

So the point is that in the last 100,000 years, from what we can tell, diking episodes have been very concentrated at a particular place. They're not random. So the type of calculations that Bruce Crowe does where he looks at a random distribution, a possibility of random distributions for any future event within a certain defined event zone, those types of calculations are conservative because actually you see clustering, in this case, away from the proposed Yucca Mountain site.

And these are what we think the important future work is. One, we'd like to get an overall evolutionary model for the Crater Flat zone, which is this zone of volcanism from Sleeping Butte down through these aeromagnetic anomalies. We's like so that using chemistry and geologic constraints, we'd like to get an idea of what the magma production pattern through time is for this zone from five million years to the present. The question being, is magmatism waxing or waning?

We'd also like through this time span to see if we can see systematic changes in volatile content, looking at eruption dynamics, that type of thing, and also the fractionation depth for different assemblages of minerals to see if there's some change in magma chamber depth through time, which may be related to magma flex through time.

And we feel that this is important, that it provides a necessary physical framework for all the probability models and effect studies. Magmatic effect studies will, of course, continue. We'd like to refine this --what the mechanism and the duration of a polycyclic episode is. From Lathrop Wells, we get the idea it's at least about 100,000 years. If Crater Flat can be considered polycyclic, it couldn't have been more than 50 or 100,000 years duration because we get about the same for all the Argon/Argon dates and the errors, plus or minus 100,000 years basically. So that whole duration would have to be hidden in that Argon/Argon error.

We need to, of course, wrap up geochronology. We need to correlate ashes in the fault trench to these eruptive episodes at Lathrop Wells, and the approach would be a geochemical approach to try to fingerprint the ashes in the trenches.

At some point we believe it's necessary to finish the volcanism drill holes, which have never been started, but there's four anomalies in Armagosa Valley and also one in Crater Flat identified by aeromagnetic data. One has been drilled commercially, this one here, and is a basalt, and we dated that at 3.8 million. But we think it's important to date the others and rule out the possibility of any Holocene or Quaternary intrusions.

And, of course, Bruce will continue with the revised probability studies. One of the things he wants to focus on in the future is how this idea of polycyclic volcanism and its facial predictability affects his numbers.

That's all.

DR. ALLEN: Thank you, Frank. Questions from board members? From staff? And from consultants? Keiiti Aki?

DR. AKI: I see this fissure orienting northwest. Is this consistent with the stress pattern? Stress is more like--

DR. PERRY: Chris is going to talk in some detail about that. I'd really prefer him to go through his talk because he'll address that specifically.

DR. AKI: You seem to have a model associated with each center, but can't you think of the model, just fissure going through all these zone?

DR. PERRY: There is--I mean, Chris--as far as a unifying structure?

DR. AKI: Yeah, your model shows a very distinct channel, vertical channel--

DR. PERRY: Right.

DR. AKI: --associated with each center. But don't you think it's more realistic to have fissure continuous?

DR. PERRY: No, I don't think the centers are connected in any way by one dike structure, anything like that. They're probably related by structures in the crust that are somehow influencing where the magmas rise, but there's no direct magmatic connection between the different--

DR. AKI: Your chemical evidence supports this?

DR. PERRY: Yeah, if you look at all the centers along there, they're all very different chemically.

DR. ALLEN: Mike Sheridan?

DR. SHERIDAN: Frank, how important to the volcanology component is an integrated model of the geological aspect of volcanism from the generation of the magma transport towards the surface and then eventual eruption? I see that you have compartmentalized all aspects.

DR. PERRY: Yeah.

DR. SHERIDAN: But there doesn't seem to be an integrated model for volcanism.

DR. PERRY: We think it's important. I guess we haven't explicitly said that, but that is something more or less unifying, everything we're doing. Greg Valentine is involved from more of a physics and magmatic processes and what's going on as far as melt generation and that type of thing. I mean, we feel we have to tie all these things of polycyclic volcanism, how things evolve through time, back to what was going on in the mantle for us to feel confident that we know what's going on.

DR. SHERIDAN: It seems to me that a model that takes into account promulgation of magma towards the surface and cooling of the magma as it approaches the surface would be an important aspect to tie into these geochemical indicators that you have.

DR. PERRY: Right, yeah. We think, you know, all these are fairly small pools of magma. We've considered that in the light that none of these we think could have been long-lived magma bodies. And you have long separations between episodes. So that fits with these being totally discreet magma pulses because they're such small bodies.

DR. ALLEN: Yeah, one final question from Bill Melson.

DR. MELSON: Frank, you mentioned the densities of the magmas as being possibly one control in the compositions; that is, you're reaching a certain density and then it moves upward to some zone of neutral buoyancy, which may be the surface or may not be.

DR. PERRY: Right.

DR. MELSON: But I'm wondering, if you look at all the densities of the Crater Flat volcanism, the lavas, do you see a clustering of densities that suggest, in fact, it's a mechanical control on composition more than these other processes?

Let me add one other question, then.

DR. PERRY: Okay.

DR. MELSON: Given that these magmas, once they rise, perhaps commenting a bit on what Mike Sheridan was getting at, is at some point vesiculation will occur, and as long as a tectonic picture such as stress, these things then will rise; in other words, the buoyancy will go crazy and they will rise very rapidly. Is there some indication what depth --if that occurs first of all, and if so, at what depth such vesiculation might take over? If there is a cluster of densities--

DR. PERRY: Yeah.

DR. MELSON: --where would that correspond to, say, within the upper crust?

DR. PERRY: We're not sure really at this point. We haven't explicitly modeled what densities these are getting at. We just have observed that the Mg number, which is probably density controlled, do cluster. We see higher Mg numbers, say in the oldest Crater Flats cycle. All of these tend to be very evolved, about what you see at Lathrop Wells, but they're significantly higher, in the 3.7 cycle, and we'd like to compare those to see how that relates to a difference in density and is that the control.

We're doing some CO_2 measurements on one of the lavas at Lathrop to try to get a handle on what the CO_2 content was, to see if there could have been some deep exolution involved, and also looking at water content per shallower exolution.

DR. ALLEN: I'm afraid we're going to have to move on. Some of these questions we can debate later or in person.

Thank you, Frank.

The next speaker is Chris Fridrich of the United States Geological Survey, who will be talking about the integrated structural model of the Yucca Mountain region.

DR. FRIDRICH: Do we have a light pointer? No?

DR. ALLEN: Bill, do we have a light pointer? The answer is no.

DR. FRIDRICH: Okay. I'm going to present a tentative tectonic model I have developed based on recent geologic mapping around Yucca Mountain, and then I will discuss the implications of this model for seismic and volcanic hazards estimation.

Could we have the first slide, please? Let me just move this out of the way.

Okay. This is a generalized geologic map of the Yucca Mountain region. Paleozoic rocks in the big uplifted Bare Mountain. The tan color here is the silicic volcanics between 15 and 11 million years old, mostly. This is a repository area, and in the blue we have the basalts.

I've been mapping in the volcanic rocks taking off from the mapping that Bob Scott did of Yucca Mountain, going west over to Beatty, and going through this tail of Bare Mountain linking to Yucca Mountain.

The major tuffs in Crater Flat were erupted concurrent with the major pulse of late Miocene extension in this area. Hence, the field relations in the tuff record the tectonic evolution of this area. And the things I've been looking at are things like tilting, faulting, thickness changes and vertical axis rotations and how they change up section within individual areas and regionally to try to get a time space evolution of the whole thing.

If I could have the next slide, please?

Okay. This is a view of Crater Flat from the north side, looking to the southeast. This is the Bare Mountain Range front coming along here. You probably can't see them, the four little cinder cones. Red Cone and Black Cone and so forth are out there, Funeral Mountains, Panamint range. And so you have this big range front and then this whole system of little fault blocks facing it.

Next slide, please.

This is an angular unconformity within the tuff section, which I'm just showing as an example of the type of thing I'm documenting.

Here is the Tiva Canyon tuff, which is 12.7 million years old. It forms a hogback here, which is buried by a buttress unconformity of the Ranier Mesa tuff, only a million years younger. So here we have an angular unconformity of about 20 degrees between two formations only a million years apart.

The angular unconformity that this--it represents an event which occurred between 12.7 and 11.6, and this was the major pulse of extension out there in Crater Flats. Since then things have dropped off pretty much, almost exponentially to the present.

Next slide.

Okay. Next I'm going to go through two definitions. First, a structural domain I define as an area in which all stratigraphic changes, all structural changes are gradual and systematic, such that the domain constitutes a logical whole.

A logical corollary to that is the definition of a structural domain boundary, which is a zone across which an abrupt fundamental change occurs in structural style, per cent extension, and/or timing of deformation. And usually it's more than one of these.

Now, other people might define these things differently. I think what's important is consistency.

Next slide, please.

These are the structural domains of the Yucca Mountain region. Yucca Mountain is this multi-fault block domain coming down here like this. It lies in the eastern part of the Crater Flat Basin. The western boundary of the Crater Flat Basin is the Bare Mountain Range front fault, which actually continues to the north into the volcanics until it runs into the caldera complex to the north where it both dies off to the north and it's cut off.

The Tram Ridge uplift and the Bare Mountain uplift, which together constitute one domain, separate the Crater Flat Basin from the Bullfrog Hills highly extended domain to the northwest, and the younger, shallow Armagosa Desert Basin to the southwest.

To the north, the Crater Flat Basin, the faults within the Crater Flat Basin decrease in throw, basically pinching out in the moat of the Timber Mountain caldera complex. And so that whole northern boundary of the basin is kind of pivoting open in that there's a strong increase in the percentage of extension to the south.

The northeastern boundary of the basin is a right lateral strike slip fault, which separates Yucca Mountain from the much more extended Chocolate Mountain domain, and other faults related to that cut northern Yucca Mountain.

To the east we have a buried domain boundary separating Crater Flat Basin from Skull Mountain and Rock Valley. The timing of extension was very different over to the east, and that's the basis of this boundary, and it's also based on geophysics.

Next slide, please.

Okay. Now, I'm going to talk about the major internal features of the Crater Flat Basin, which are first the major range-front fault on the west side of the basin.

In the first kilometer or two, we have a lot of synthetic faults, other faults that are down to the east, but really very quickly it goes to a pattern where almost all of the faults are down to the west, basically facing into the major range-front fault on the west side.

And so the basic form here is a half graben by definition. You have a major fault on one side, lots of little antithetic faults facing it all across the basin.

This feature right here, which I've shown in the red, is a rollover. To the east of this rollover, the stratal dips are all to the east. To the west of it, the stratal dips are to the west into the major range-front fault.

Next slide, please.

In addition to those standard extensional features, there are a number of different features in the Crater Flat Basin which indicates strike slip shear. These are, first of all, that almost all of these north trending faults in the basin have a component of left slip, so they are all left oblique faults. Even though that the amount of left slip is usually small, it's very pervasive.

Two of the boundaries of the basin show right slip. The Yucca Wash fault is almost purely a strike slip fault, and the Bare Mountain fault at its southern end has at least a small component of right slip in addition to it being a large normal fault.

Yucca Mountain itself shows oroflexural bending. Basically, at the north end of the mountain there's no evidence of oroflexural bending, but when you come down about two-thirds to three-quarters of the way down, we get up to about 10 degrees of oroflexural bending in the Tiva Canyon tuff, and by the time we get to the southern tail, it's up to 30 degrees of oroflexural bending. That's vertical axis rotation.

One other evidence of strike slip shears is that we have scissors faults, and most notable being the Solitario Canyon fault. This fault decreases in throw in normal offset to a fulcrum point, past which it actually becomes a reverse fault.

Reverse faulting within an overall extensional province can be rationalized in the context that if you have vertical axis rotation like this and two different fault blocks rotate to a different degree, you will get a very localized zone of compression between them.

One thing I forgot to mention is that the pattern of normal faulting in the basin, basically in the northern part of the basin, is radial about the caldera complex, and then you see this prominent curve of the faulting as you come down Yucca Mountain.

You notice here that the strike of faults goes through a major inflexion point right down here. That inflexion point in the major strike of the normal faults within the basin correlates with the paleomagnetic evidence that there's a sudden increase in the degree of oroflexural bending.

As I said before, the vertical axis rotation goes from zero to about ten, and then ten to thirty. And so right here, where we have this change in the strike of the faults, there's a change from a basically weak, oroflexural bending to very strong oroflexural bending, where not only is the degree of, the amount of vertical axis rotation greater, but the gradient in vertical axis rotation at the southern tail of Yucca Mountain is very high.

Next slide, please.

This is aeromagnetic data over Yucca Mountain. Just to position us, this big fuzzy area is Bare Mountain, which appears that way because carbonates are not magnetic. These four little bits are the cinder cones out in Crater Flat.

The aeromagnetic data shows the patterns of faults in the basin because the Tiva Canyon and Ranier Mesa tuffs have opposite magnetic polarity. And so the very strong angular unconformity between those two units creates these ribs wherever there's a major fault that was active in that period, and virtually all of the major faults in the basin were most active in the period between eruption of those two tuffs.

One of the things you'll notice is that on this diagram in the aeromagnetic data, we can see this inflexion in the strike of faults on Yucca Mountain that I was talking about, and but what's most significant is that we can project that inflexion in the strike of the faults to the west into areas that are covered by alluvium because the alluvium is shallow enough that the aeromagnetic signature of the faults still shows up.

For instance, down here you can see that the major faults are striking northeast, but then up here they're striking to the north.

And so we can project this zone of--this boundary between the zones of weak and strong oroflexural bending on Yucca Mountain all the way across the basin, and I would propose that it's a northwest boundary, going up about like that.

Next slide, please.

And so I would summarize the major features of the basin as follows: Basically that we have the major range-front fault on the west side. It's a half graben where we have this whole system of antithetic faults facing that range-front fault across the basin. These faults decrease in throw to the north until they pinch out, so the basin is pivoting open on the north side, and this basin opened by virtue of dextral shear along the southwest trending zone on the southwestern boundary of the basin.

And it's this oroflexural bending, and probably the small of right oblique slip on this fault, is what allowed this basin to open. Basically this is a strike slip shear zone, a very diffuse and distributed zone of strike slip shear, but that's what allowed the basin to open.

And all of the information that I have indicates that the timing of the formation of these three features, the activity on the range-front, on these faults, and on the strikes of the vertical axis rotation in this strike slip shear zone were all the same.

Next slide, please.

This is a diagram which schematically shows the extensional evolution of the Crater Flat Basin that I've documented, where these are the ages of the major stratigraphic units that I've used to constrain the evolution of the basin. And what you can see is there was a small amount of extensional activity back in the period from 14 to 12-and-a-half million years, and then a huge pulse in extension right between eruption of the Tiva Canyon and Ranier Mesa tuffs at about 12.5 million years.

Since then activity in the basin has dropped off almost exponentially to the present, basically just increasingly feeble faulting activity.

Now, in the second column here, what I've done is I've made that a linear scale, and two things I should add as caveats, we actually have almost no information from 10 to 14 to 4 million years, and so that's basically an interpolation. Also, the existing data suggests that the activity, rather than being really smooth like this, is actually kind of episodic, that it's waxed and waned in various pulses. And moreover, there appears to be a coupling between the pulses of seismic activity and the pulses of volcanic activity. For instance, the 10 million year basalts are inter-layered with rock avalanche breccias.

Next slide, please.

This slide and the next slide present really detailed data on the tectonic evolution in the area of the space time pattern, and I'm just going to summarize it very quickly. What it shows is, and in the first period here, going back to 14 million years, the activity started from the east, and it basically migrated to the west. This was the major pulse of tectonism.

Next slide, please.

And then going to younger and younger periods, the place where the major tectonism was occurring kept being moved further and further to the west until it just migrated out of the area. Now, there continued to be tectonism after that around eight million years, and this is poorly constrained. We had some younger basins cut across.

Basically in the Pliocene and Quaternary, the pattern of activity has been that the mostly north trending normal faults of Yucca Mountain and Crater Flat Basin in general have been reactivated. These faults all formed at about 12.5 million years, but they're being reactivated for some reason. But it appears that the basin is still behaving as a half graben.

Next slide, please.

Okay. So my major conclusions about the structural model are that Crater Flat is a half graben, but has many strike slip features.

The entire Yucca Mountain region is segmented into domains, which makes sense with the fact that it lies in the Walker Lane belt.

Extension occurred in distinct belts that migrated from east to west in the region between 14 and 9 million years. The faults that are active now formed at 12-and-a half million years.

Next slide, please.

Okay. My next step will be to try to place the volcanic--the basalts in the context of the structural model. This diagram shows the four major episodes of basaltic volcanism in Crater Flat. We had basalts at 10 to 11 million years, about 3.7 million years, 1 million years and about 100,000 to the Holocene, being Lathrop Wells cone.

The thing I want to point out is that the vast majority of the basalts in Crater Flat lie in this zone of strong dextral shear that I was talking about along the southwestern boundary of Crater Flat Basin.

Moreover, almost 90 per cent of the total volume of basalts lie at the intersection of this dextral shear zone, and the extensional axis of the basin, which is this rollover from eastern to western stratal dips. And two of the episodes of volcanism were actually aligned apparently along this extensional axis of the basin.

So I would suggest that there is actually a very strong structural control in the basin on where the basalts are coming up. It's not just random at all.

The only other major occurrence of basalts in the basin in up here in northern Yucca Mountain, there's a very small cluster of basaltic dikes that are along the extension of the Drill Hole Wash fault and the Solitario Canyon fault.

Next slide, please.

And this is a detailed view of that. This is where the repository area is. These are where these small basaltic dikes are in northern Yucca Mountain. I believe that these dikes are related to the right lateral strike slip shear zone that cuts through northern Yucca Mountain.

The point, though, is that all of these dikes are very small. There doesn't appear that there was any significant edifice built up because there's no plug or anything. They're very skinny dikes, and there's no surviving edifice certainly. And it all occurred at 10 million years, and nothing has happened since.

Moreover, this structural zone of right lateral strikes of the shear appears also to have been inactive since 10 million years ago.

Next slide, please.

Okay. Now to step back and try to put this into the larger context, the occurrence of basalts in the southwest Nevada volcanic field as a whole, this is Crater Flat down here, showing the distribution of basalts there as I talked about, and then this is the other part of the Crater Flat volcanic zone that Frank Perry discussed up around Sleeping Butte.

The thing I want to point out is that our paleomagnetic results show that this also is a zone of very strong right lateral strike slip shear, and, in fact, Mark Hudson, who's done the paleomagnetic work, believes that this zone of strike slip shear is linked with the one along the southwestern boundary of Crater Flat. Hopefully we'll be testing that idea in the coming year.

There is one other occurrence of basalts in the strike slip shear zone, and that is these basalts in the left lateral Rock Valley shear zone done here.

In addition to strike slip shear zones, basalts in the southwest Nevada volcanic field are clustered along the ring fracture zones of the caldera complex in the middle of the field, Buckboard Mesa and so forth. You can see they really fall very well on those structures.

In addition, there are some other outlying basalts that lie along specific extensional structures, the Nye Canyon basalts, the Paiute Ridge basalts. I have to admit that that structural control is not as strong a case as the others because their extensional structure is everywhere out there. So anyway, it's not as strong as the other case.

Next slide, please.

So to sum up, the three major structural controls in the volcanic field as a whole appears to be caldera ring fracture zones, strike slip shear zones and extensional structures. But as I said, we have to take that with a grain of salt.

Next slide, please.

The basaltic clusters active today both occur in northwest trending right lateral strike slip shear zones, and as I've discussed, they might actually be the same shear zone. Both of these clusters that were active in the Quaternary were also active at 10 million years and in the

Pliocene.

The Yucca Mountain repository area lies completely outside of these zones of recent activity. However, there are basalts in northern Yucca Mountain in another right slip shear zone, but all of the indications we have are that this area has been dead volcanically and seismically since 10 million years.

Next slide, please.

Just briefly, I want to touch on the detachment fault model, which was the preferred tectonic model for the Yucca Mountain region before I got onto the project.

Recently there have been a number of different types of data that have dealt blows to this model, being geophysical data, seismological data and recent geologic mapping, which does not support the predictions of this model.

Next slide, please.

I'm just going to discuss one of these, and this is a slide that John Whitney had showed. This is the aftershock pattern associated with the 1992 Little Scull Mountain earthquake. What you can see is that the aftershocks defined a plane that projects up to the surface. There actually is a fault at the surface that lines up with the plane defined by these aftershocks. And so what this tells is, is that the surface faults apparently are planar structures that go down through the upper crust to the brittle-ductile transition. and that's very hard to reconcile with a detachment fault model.

Next slide, please.

Okay. To sum up, the implications of the structural model that I'm proposing for seismic hazard assessment are that the faults that were active in the Quaternary formed at about 12-and-a-half million years, and the chances of a new fault forming through the repository I believe are nil. Secondly, that the rate of extension has progressively declined since 11-and-a-half million years ago. However, activity probably is somewhat episodic, rising and falling, and there appears to be a coupling between the rises and falls in seismic activity and the rises in the episodically of the volcanism.

The implications for volcanic hazard assessment are that the Quaternary eruptions have been confined to a narrow zone that does not include the repository area. Hence, if you're going to include structural control in your volcanic hazards estimation, it appears to me that it would decrease the chances of magmatic disruption of the repository. The one thing that might operate against that is the dike zone in northern Yucca Mountain. However, as I said, it appears to have been completely inactive both volcanically and structurally since 10 million years ago. So it's hard for me to believe that that's a significant thread.

Thank you. That's it.

DR. ALLEN: Thank you, Chris.

Questions from the board? From the staff? Consultants?

Could I ask one question about detail? At the scale of your mapping, you showed the Drill Hole Canyon fault and the Solitario Canyon fault as not offsetting each other. When you get down to greater detail, which one of those trends is more recent?

DR. FRIDRICH: The Solitario Canyon fault is the younger fault. It cuts across the Drill Hole Wash fault. And so that's why the Solitario is not offset because the right lateral movement occurred first, and then the normal movement cut across, and the strike slip fault being vertical shows no real apparent offset where it's offset by a normal fault.

DR. ALLEN: The reason I ask is because that also bears on the question of the relationship between the Ghost Dance fault and the Sundance fault.

DR. FRIDRICH: I think that it's the same type of--I think that the same thing applies, yeah.

DR. ALLEN: Any other questions before the break?

Okay. Thanks, Chris, and let's have a break for exactly 15 minutes. We'll come back at 10:40.

(Whereupon, a break was taken.)

DR. ALLEN: The next speaker on this morning's program is Allin Cornell, who will give us some general comments on probabilistic approaches. Allin has been among the real leaders of seismic hazard assessment. His 1968 paper in the SSA Bulletin has long since been famous. He is a professor at Stanford University. He also runs his own consulting firm. He's a member of the National Academy of Engineering. He's been president of the Seismological Society of America. We look forward to his presentation, even if he can't spell his first name correctly.

DR. CORNELL: I'm a structural engineer, and despite the title of my predecessors' presentation about structural models, you're now in for something completely different. Leon asked me to talk about the broad background of how we got into this position of trying to characterize natural hazards in probabilistic and on certainty terms for use in engineering design evaluations and decision making, and to give some perspectives from that sort of broader view that is not necessarily focused on seismic and volcanic problems, not necessarily focused on Yucca Mountain.

As an engineer, a structural engineer, I became very interested early in my career on safety of structures. That naturally led me to probability. It also very quickly led me to realize that the loadings are the primary source of our randomness and uncertainty and potential troubles. And dealing with loads, immediately led me to have to deal with the scientists involved with the natural phenomena that lead to these natural loads, whether it be the seismologists or meteorologists, whatever.

So I'm sort of an engineer who hangs out with scientists, and I know almost nothing about Yucca Mountain. Together that puts me in a unique position to tell you exactly how to do your job.

I also start out with a very strong bias. It's written down there at the bottom. In fact, this is a new example of a multi-media presentation you'll see. What you've got here are a combination of the overheads and your notes, and they're in bold, which is what you're supposed to be able to read from back there, and in small print, which you're not necessarily supposed to be able to read. I hope it doesn't distract you; that's the negative side of this. And it gives me something to look at to remind myself what I wanted to say, and it gives you something to take home.

So don't necessarily try to read the small print. If you can read it, I should make it smaller next time. So let me know in the feedback section.

I've tried rather faithfully to follow the outline that Leon proposed because I thought it was a good one, and the list of questions that the group is supposed to address here over the next days I thought were excellent ones, and I share Clarence's concern that we'll probably not have unique concrete answers to all of them, but they're the ones we should be asking. And the sort of five or six, seven topics that you'll see here are precisely the ones that Leon proposed we talk about.

The first is what are these products? The product is presumably, depending on what the hazard is, some scale or measure typically, but it may be a vector. For example, peak ground acceleration, wind speed, whatever, and typically an annual probability of exceedence as a function of the level of that indicator, some effect variable, as I call it here.

And secondarily, but perhaps more importantly, an uncertainty band of some kind reflecting the degree of confidence, however you'd like to call this, in the estimate of that annual frequency of occurrence. The forms of this output may be different. It may be in terms also alternative scenarios of different things that can happen with their estimated frequencies and your uncertainty of the frequencies of those alternative scenarios.

And the uncertainty analysis may include looking at sensitivity studies, confidence bands, as I've indicated here, and so on.

I think an important aspect of the second part of this is so called epistemic or a knowledge-based, knowledge-related uncertainty; that is, something associated with what the limitations of our current scientific knowledge about the phenomena are. How those project onto this hazard curve, is that these are a current assessment of knowledge, and as we've seen in the previous talks, when a scientist studies something very hard, that state of knowledge evolves, and, therefore, these epistemic uncertainty bands evolve. It's the nature of them that they are not constant, and the question is only for a particular engineering decision application, when have you decided you've got them as narrow as you can afford to make them in the larger context of the decision process.

I think that this notion of presenting the results of your scientific investigations in a format of a hazard and an estimate and an uncertainty band, although it has indeed been driven by the users, whether they're regulatory users or engineering decision users, should I think be a natural way to report the output of science. I suspect that the notion of providing a concrete end product in a finite amount of time is not something that's natural for the scientists in their activities, but it does have to be done.

The objectives of the scientific process, let's say by which we arrive at those end products, as I said, I believe should represent good scientific practice. It should not be inconsistent with scientific practice, although it may be a new way of practicing such things; that is, the notion of coordinating, communicating, describing among yourselves uncertainty of the data, the alternative theories, your degree of confidence in those theories at any given time, the identification of factors which might be critical to the end result, all of those that deserve further investigation. Describing those alternatives and the information about them in probabilistic terms, it seems to me should not be an unusual or unexpected type of thing to be doing, and hopefully, it's, in fact, a useful process.

The idea of then combining those uncertainties, that is uncertainties in, for example, occurrence processes, recurrence processes and uncertainties in what are the effects of a given event on the structure or on the ground motions, for example. That requires some combination of information which has now been expressed probabilistically. It requires the use of some kind of probability theory. The idea that those pieces of information can be put together into this end product is, again, something that seems to be natural and good science.

The fact that the communication must ultimately be among yourselves and this end product must ultimately be scrutable and so on is, again, part of the process, and it seems to be a useful and not unexpected thing. It's something that I think we, as users, should be able to expect from you.

The third step, that is communicating these hazard results and the uncertainty to other people, which may be specialists in their own right or a review body or decision makers, politicians or engineers, is where things can begin to get a little dicey because as we'll see, the problems we're talking about are complicated. Putting uncertainties on top of them makes them--gives them at least another dimension of complication. And trying to reduce the presentation to something that's easily communicated could be very difficult.

So this interface problem may be one of the parts of the process that turns out to be one of the most cumbersome. The question of transparency, for example, can the reviewers see what you've done when it's already been integrated and multiplied and compounded a few times is one of the difficulties that we face in introducing this sort of combination of probability and uncertainty assessment on the problem.

The final step is a step that I'm saying I think the process should avoid, and that is something which I think is not scientific. It doesn't mean that it's not science. It doesn't mean the scientists aren't involved, but making value judgements is not part of science, and which this in turn means such questions as how safe is safe enough, cost benefit analyses that lead towards decisions, engineering implications of what will be the implication of an event which is beyond the design basis. That's not science.

All of these involve priority setting. It is because ultimately a decision maker has to set priorities, allocate resources, that we need your results that are probabilistic context with uncertainty bands, and we would like you, thank you, to stop there as scientists. If you want to join in the discussion of how to make the decisions, that's fine, but I contend that's not a scientific exercise.

Okay. For example, the final example here was this notion that comes up again are the questions, "When is enough enough," or "Is enough enough?" I'm never quite sure what enough is enough exactly means. But we all know what it means here. "When do you stop spending money looking for something else?" is, indeed, in this category because it involves prioritization and resource allocation. And what the scientists can bring--that it can be formally analyzed, and the decision theorists want a call of pre-posterior analysis, I suspect.

And what the scientist brings to that problem is what the likelihoods are that he'll find different outcomes when he carries on a proposed experiment that he's come to you to ask more cash for. And what we can also ask is: what will the impact of different findings from that experiment be on his best estimates and uncertainty bands on the best estimates. And the combination of those two things, coupled with a formal analysis, can help lead to an answer of the question, "When is enough enough?"

A little bit of background. Probabilistic characterization of design loads or design criteria for engineering purposes grew throughout this century. Structural engineers have used wind loads of 100-year return periods and snow loads and flood loads since early in the century. Many of these early models were rather direct empirical kinds of statement. You plot a few data points taking the annual wind speed. You plot it on appropriate probability paper and cast a straight line as far as the engineer wants it. And they usually stopped at something like an annual probability of 1 in 100.

And that was put into an engineering design process with load factors or some allowable stresses, another set of big conservatisms.

More recently what we've seen are much more structured models about how to develop probabilistic models, where instead of empiricism, we get much more of the science into the problem and the physics, and as been said in some of the comments this morning, what we're doing the physics for is to help structure the models of what now have become probabilistic models.

Examples of events which drove this kind of specification of probabilistic frequencies and uncertainties include, for example, Wash 1400 in the seismic safety area. A friend, Hal Lewis, wrote a report in response to that study which said we have to be very careful about drawing uncertainty statements on these technical and scientific inputs to these problems.

Today the question has come up, how widely is this used? Today in engineering practice in all countries for all fields for all types of natural hazards, probabilistic methods are absolutely the norm for the use of establishing design basis, whether we're talking about offshore structures at wave loads, whether we're talking nuclear power plants and probabilistic input, seismic input to them.

Some exceptions remain, and we know some of those. The flood people for high dams still like to talk about probable maximum participations, probable precipitations, probable maximum floods. Bob I think will talk a bit about those in just a moment. They argue about where those are in the probability of the main. We've had National Academy reports on this. These remain in, I would at least say, a state of flux.

As I said here, the higher tech fields, as we've gone to things, for example, like one billion dollar offshore structures, some of the nuclear power plant studies and so on, the evolution has also been towards getting out of this area of looking at, say 1 in 100 year loads with big load factors, and have gone to trying to characterize the load at the 10^{-3}, 10^{-4} level; that is, higher loads at lower probabilities, because that's where the action is, and that's where the needs are from safety perspective.

Many of you may have read Sunday in the New York Times the article about Yucca Mountain, New York Times Magazine. A very nice one, but the man there talks about the famous drunk who loses his car keys in the dark alley and looks under the street light. That's famous because I always use that example. Everybody here that knows me knows I use that example. And the point of view is that you don't look at 1 in a 100 year return periods for safety problems because the problem is in the $10^{-4, -5}$ region in the dark alley, and you're much better looking over there, no matter how dim your match is, or how weak your flashlight, you're much better looking there where the problem is than under the street light.

So we're going towards these low probabilities, tough as it is.

As I've indicated before, we must focus, because from the engineering point of view, natural hazards problem, the phenomena, the interesting problems that threat the uncertainty lies in the loadings. We need to focus there and not in the structural systems by contrast. There are some sidelines on that that I won't get into.

But what recent experience has brought, now we're talking about the last 20 years, to this exercise, is this notion of trying to quantify the uncertainty about the probability or the uncertainties about the frequencies, this epistemic uncertainty as I prefer to call it, your lack-of-knowledge uncertainty.

And this is where we're now starting to struggle and see the implications, good and bad, of trying to go through that exercise. What it does bring to science is the opportunity to not have to come up with unique answers, unique models and unique numbers that you know you can't in your heart of hearts defend with total confidence, even though the regulator may want you to say that.

It gives you the opportunity to put in alternate models and your degree of confidence with them to express your uncertainty explicitly.

I contend that should be--I hope you find that useful.

The basic structure of the models we're usually talking about--Clarence, remind me of the time. I started about?

DR. ALLEN: Pardon?

DR. CORNELL: Remind me of the timing here. I forgot to take a starting point watch.

DR. ALLEN: You started at 10:42.

DR. CORNELL: 10:40, okay. Halfway through, sure. Okay, good. Just so about--

DR. ALLEN: You've got about 25 or 27 minutes.

DR. CORNELL: Okay, good, no problem. On target.

And this question is what are the basic structure of the usual kinds of models we're looking at in natural hazard assessment. Most of them fitted, whether we're talking about tornadoes or storms, hurricanes at sea, earthquakes or volcanoes, we end up trying to come up with, first of all, a recurrence model at the time in space, a temporal-spatial recurrence model of something of which is effectively a point in time in space; that is, in some time space scale, it's effectively a point. That is the duration of the earthquake is small compared to the design life of the structure. The location of the source of the earthquake is relatively small--in dimensions we're usually talking about, et cetera.

And so we have some kind of XY plane, and history is going to give us points on this plane which events happened, and it's going to give us the order and the dates in which they happened. Typically those events are then what are called marked point processes. Associated with each of these events is some scaler or vector, we would recognize or model as a ramdom vector of source characteristics. The obvious one is a magnitude of the earthquake, but it could be magnitude stressed throughout the length, at a whole vector potentially of whatever you use to describe the source in your scientific model.

So the first step is a recurrence model in time in space. It means beginning to talk about whether these events are homogeneous in space. We've heard that discussion on the volcanoes, or whether these average recurrence rates are not homogeneous; that is, clustered in space, as a diagram like this might suggest, relative to our site.

On the temporal side, we begin to ask the same kinds of questions. For example, is the process Poissonian? If so, is it homogeneous in time? Is the rate relatively constant, non-homogeneous? Is it growing? Is it decaying, as the last model suggested? Is the process, indeed, Poissonian itself, or do we find clustering in space, clustering in time, or the reverse, some kind of pseudo cyclic behavior as a characteristic magnitude model would tend to suggest to us?

So these models all tend to be of roughly this type, and so there is a benefit of sort of a common modeling approach that's taken to natural phenomena, particularly of the extreme type. And the key point is that these models, the probabilistic models that are available are as--I hate to use the word complicated, but they're as complex as you need them for the physics of your science.

And what is it? They should be as complicated as necessary and as simple as possible, but they should indeed keep all of the physics, and there's no reason why the probabilistic modeler should put any limitation on your physical models.

The other side of the coin is he's going to demand from you a lot of information about the parameters of these models and characterizations of them that you may feel hard pressed to make estimates of, but we would--I think you would admit they are the essence of the problem.

Each element of this model, then--pardon me, the second step of the model is some kind of effect representation. That is if an event of a given size, whatever, given characteristics, occurs at a given location in space at a given time, what will the effects be on the structure, which in most cases is, again, relatively localized in space? And without going through details, what usually ends up with some kind of summing over these possible sources of events, places where they can happen, something to do with the recurrence rates or mean rates of occurrence, something to do with the duration of the interval of time looked at, something to do with the likelihood that, for example, ground motion or wind speed will exceed a certain level, conditional on what the size and location are, and then integrated over possible alternative values with the relative frequencies of these size levels and distance levels, et cetera.

So some general kind of form of probabilistic integration of the randomness and the event location, size time, et cetera, comes about. If the models are not Poissonian, some of these steps become a little more difficult, but the key point is that there's a very common structure to virtually all of these natural phenomena, and if you look at tornadoes, you see they look the same way. If you look at hurricane models, they look the same way. If you look at North Sea storm models, they look the same way, et cetera.

Each element, then, of this model needs to be characterized; that is, there's a size, scala random variable, the magnitude. You need to give us a probability distribution on it. It may be in cases like Yucca Mountain that some of the things that are unimportant are the relative frequencies of very small events, but, in fact, it's the relative rate of occurrence of the large events that matters, and so you focus on near upper bound magnitude events, et cetera.

And so where exactly you should expect feedback from your engineers and decision makers as to what portions of these distributions need the most attention and characterization, and, in fact, that's an obvious outcome of projecting these results forward into this hazard analysis to see which parts of that hazard curve are most sensitive to which parts of the input, the typical kinds of sensitivity analysis that should be an interactive reciprocal cyclic kind of approach.

Still back on that basic structure thing, which I pulled off the screen here, and it went where? Oh, yeah.

So it is a characterization of each element, and here's where we often get into alternative models of the characterization; that is, consider a model of volcanoes which is homogeneous in space, consider another one which is clustered. We're not--we can't be absolutely sure that one or the other governs, and so alternatives show up. And that shows up finally at characterization of the uncertainty, not only in estimating the parameter values, but their uncertainty, and we'll come to that next.

So within these models, which are now probabilistic, we have a vector of parameters that need to be estimated, mean max of a magnitude, mean rate of occurrence in the next few years, the slope of the decay of recurrence rates in time, et cetera.

Many of these parameters may vary in space as we've heard. As we've heard, they may vary in time. So the models begin to take on a level of complexity that still in the probabilistic context, stochastic modeling context, maybe it makes the numerical analysis a little bit difficult, but it should not be a barrier. That is it should not be a barrier to calculating this hazard curve from whatever level of complication of physical stochastic model you want to construct. This should not be an issue.

Where we start to bump up, incidentally, against our deterministic design basis, friends, may very well be in estimating, for example, limits on some of these distributions. If we believe the magnitude distribution stops somewhere because it's limited by the length of a fault, then we may both agree that this maximum magnitude, maximum possible magnitude, is an interesting number to us, and we both agree, might agree, that we don't know exactly what it is, and there's a high degree of uncertainty in it, where the deterministic approach stops as saying, that's the only number I'm interested in. And we're going to argue and agree in some kind of decision-making, non-scientific process about what that maximum magnitude is, where the probabilistic approach would go forward as to try to put an uncertainty distribution on that maximum magnitude describing what your current level of degree of confidence and knowledge is.

That comes to the second part of the whole general structure of these models, which, as I said, is the kind of thing that's come up much more recently in the last 15 to 20 years, and that is explicit quantitative uncertainty assessment on the uncertain parameters of these probabilistic models, rates, upper bound magnitudes, co-efficients on regression, attenuation laws, et cetera, et cetera.

And this is the tough part. This is the one that should be in principal relatively easy; that is, the objective is simply to put in the same way you've been doing for years standard errors on the outputs of some tests. But now it's going forward into putting standard errors or distributions more generally to reflect statistical and current level scientific knowledge on all the parameters of this probabilistic model we talked about.

So the reality is that this becomes very complex. For the kinds of physical models we now have available, physical structural models of the processes we're talking about, varying in time, varying in space, non-homogeneous, non-Poissonian, scaler descriptors of the sources of these things, the complex theoretical, physical attenuation laws, not just dumb empirical regressions, et cetera, et cetera. The process of recognizing that each of the parameters, and you now may have 10 to 20 or more parameters, some varying in time, some varying in space, suddenly becomes very complicated, and unfortunately, opaque, and unfortunately, not familiar to very many People. Especially the people that are responsible to putting information into that process and reviewing that process and that's the really--that's the tough part of what we face in this exercise right now.

And why does this become complicated? Something such as the mean rate of volcano occurrence in the next 10,000 years, suggesting maybe it's falling off with time. So we have to--the probabilistic model says the mean rate follows linearly in time or exponentially in time.

So you've got a couple of parameters suggesting that exponential fall-off, but the fall-off itself is now, in fact, a random process. And to describe a function in time about which your uncertainty becomes a random process, you have to have its best estimate at any point in time, your uncertainty at any point in time, your correlation between any two points in time.

So what looked like a pretty simple thing going in now becomes something where your random process theory, however much you had of it, comes into play, is specifying, operating on and understanding the output of.

So it's non-trivial, and this uncertainty analysis which puts an additional dimension on this whole stochastic modeling of physical processes is the thing that has really kind of put, unfortunately, kind of a cloud or curtain somewhere in the process I'm afraid. And it's something we have to work on very much I think.

What the benefits here, of course, are, as I said before, it permits the opportunity to retain alternative models in the science; that is, it's not necessary to go forward with a unique model, but to retain the fact that you're not absolutely sure about what they are. And this I think is critical and, in fact, beneficial to the scientist I would think.

The other part of the problem is this notion of maintaining diversity; that is, if indeed experts' interpretations to create models become an important part of this exercise, what we know characterizes, it seems especially the geological sciences, is diversity of opinion about what these models are. They take pride in this, and thank God they do. But what it means is there's also a responsibility on the scientist's part and the decision-maker's part to recognize that diversity and do something with it other than push it under the rug and say, there's a consensus among science that this is the way it is. And what the uncertainty assessment gives you the opportunity to do is keep it and do it, but it also makes it very hard to communicate.

These steps of eliciting uncertainties where there's a whole new field--I guess I'll call it a science. It's at least a social science--in eliciting the uncertainties from experts, from technical people, from scientists. It's a tough job, but somebody's got to do it.

I don't know how to do it. Fortunately, we have people that make a living doing this thing, and we think they're doing a better and better job of it. It's tough. In many cases, it's going to be new for the scientists to be asked to do these things, and don't forget that the poor guy who's the elicitor, the science is new to him. And this means there's a very difficult, and our experience says very time-consuming, job of getting this communication going between the scientist and the uncertainty elicitor. I'm not sure what these people often call themselves today, but that's what their job is, to pull out from you with relative degrees of belief on alternative models, on alternative interpretations of the future trend in volcano rates.

So this is the difficulty that comes up, this notion of uncertainty assessment, aggregation among experts, et cetera, et cetera. It's an opportunity, it's a responsibility, but it's tough, and as we said, it leads to these questions of a lack of transparency and understanding the results. This is a common criticism of probabilistic seismic hazard analysis with uncertainty bands as used in the nuclear regulatory environment, for example, or DOE critical facilities.

But it's necessary, as I said, and it's important that we all work on it, and I would suggest that it's important that people on both sides of the fence work on it;

that is, the reviewers have to put some effort into finding out what all of this means also.

Some examples of use, very quickly, offshore structures, an area we are working in. All of these same kinds of models are used there in characterization of, for example hurricanes in the Gulf of Mexico. I believe the experience base has been good. It's followed the kind of evolution we talked about. That's starting off with design levels with probabilities of 1 in 100, even though the failure rates and target safety levels are in the range of 1 in 1,000 and 1 in 10,000 per year, and evolve towards procedures now which take this second level. That is they begin to push the dim flashlight into the alley and look not only at the rare events and the small probabilities, tough as they are for the scientists, but just as tough for the structural engineers, how is my structure going to behave, not when it's down in the elastic rubber band area, but when it's up in a highly non-linear, near-failure condition. It's an added responsibility on the engineer's predict behavior, too.

And this is, again, brought on by a more realistic regulatory environment and the needs of our safety analyst friends to carry forth this work into larger probabilistic risk assessments.

This has led to--I think it's led, this need to go into these small probabilities in other fields, as well as this one, into a lot of interesting new science. We see paleoseismology, paleo flood analysis. We see in hurricanes in North Sea looking at hind casting of what the waves must have been in an event in 1902 when the pressure drops were the following as this track came across the Gulf of Mexico.

And we see the need to address very strongly questions of space time exchangeability. If I haven't got a long history, can I exchange--looks at other places in space for analogies, and so on and so on. So this driving towards rare events and small probabilities has led, I think, to these kinds of issues.

Another area, of course, is the application of these exercises in seismic safety in the nuclear power plant area, let's say particularly focusing on the eastern United States. This is an area which is on the whole, I would say, a success story, although it's not been without its rocky bumps along the path. But I think among questions resolved by that process would be, as I alluded to earlier, this attempt that we went through in the '60s and '70s to try to find the unique seismo-tectonic zonation of the eastern United States. And blood was spilled on many tables trying to make those characterizations and to make the lines very fine because your power plant might or might not have been on one side of those lines.

But what we have today, despite Ellis Krinitzky's criticisms, are, in fact, you know, pieces of wallpaper with different floral patterns that represent different experts of judgments as to what these seismic source zones might be, and, unfortunately, those are alternative models, and they represent the state and diversity of current opinion. But they are carried through, and the arguments will follow as to the basis for those zones. That's fine, but the idea that everybody has to agree finally on unique zonation has disappeared, and I think that's a benefit.

Some issues and problems. As I've said, alluded to earlier, in these problems, if we do look where the action is and where car keys are in the alley, these are rare events, and it implies that we have to bring all the relevant information, scientific and interpretative information, to bear the problem. It means we need to go for these space time exchanges, as I've suggested, interpretations, et cetera, and to combine these sources of information intelligent ways. That is, the preferred approach here is usually to not just take an empirical extrapolation of flood data or wind speeds, but to desegregate the problem into its physical pieces in the relevant physical parts of the model. Make models and assessments about each of the pieces and let probability theory in a logical way put the pieces back together so that you end up making predictions about 10^{-4}, not by extrapolating from 100 years of data, but by putting 10^{-1} assessments on three or four pieces, and then combining them.

One of the final problems of this is, of course, it sort of lies out of classical statistics. I mean we grew out of classical statistics, but it's virtually impossible, for classical statistics brings very little to bear on assessing these 10^{-3}, 10^{-4} events. It brings a lot to bear on assessing the individual pieces, and then you put them back together.

Another issue here, as I've alluded to again, is that multiple disciplines evolve, not only within the science, as we have seismologists, geophysicists, aeromags and everybody involved here, but also because there has to be communication to engineers and to elicitors and to regulators and reviewers, and this takes time. It takes cross-training. It takes time to develop communication about this. I would suggest that probability is a common line, universal line reached to do this. But it's kind of like Esperanto around the turn of the century. If everybody would learn how to speak Esperanto, we could do away with all the languages in the world. The problem is not very many people, even today, 100 years later, speak it very well, and we're left with English, weak as it is, that's the closest approximation.

The results, of course--a final issue, the results are often very--use a very visible arena. Yucca Mountain is an obvious example. This is a contentious, litigious, et cetera, environment, and that has a lot of implications about the degree to which these results have to be defended, the degree to which you'd like to be able to say they present a consensus, but it's some kind of a new definition of consensus when it's put together this way.

As I alluded to before, this probabilistic analysis involves models which are no longer trivial. They were in the early days, they aren't anymore, and the degree to which the physics is a fundamental part of the probabilistic model implies more and more complexity on the part of the probabilistic models, and not necessarily everybody has been well trained in his scientific career to look at these things.

And that means building the models itself gets a little bit--becomes not as familiar an exercise as we'd like it to be. That's changing with time, of course. But as I've alluded to, the really tough problem is putting this uncertainty description on top of the probabilistic models because what became--what were simple parameters become random variables. What were simple functions become random processes, and everything gets not really one-dimensional, but multi-dimensionally tougher, faster.

I was asked to comment on Ellis Krinitzky's criticisms. I have read his article in the engineering literature, Civil Engineering magazine, which he called "The Hazard of Using Probabilistic Hazard Analysis". I have not read his papers in Engineering Geology. I'm an engineer. I only read those papers that have come highly recommended to me by my engineering scientific friends. That's the case, and their comments are that these are heavy-going.

The Civil Engineering article talks about the hazard of using this, and as I would suggest, everything that's useful is hazardous, and this is a couple examples that are on the board. And the question is, of course, what are the alternatives? We're talking about siting critical facilities, hazard analysis of natural phenomena. The only apparent alternative on the table is what is usually referred to as deterministic analysis or deterministic design basis. And I think if the science is evolved along the lines we're talking about to incorporation of alternative models, as opposed to collapsing to single ones, the deterministic basis simply hasn't come along to help that out.

I would also remind you that setting a deterministic design basis is, again, not a scientific problem. It's a valuated resource allocation problem, and science is part of it, but it's not a scientific problem.

Finally, Yucca Mountain specific issues to make some comments on. The long time frame, the 10,000 or perhaps maybe 100,000 year notion, has, of course, a number of implications. Some of them are technical; for example, issues like Poissonian versus non-Poissonian clustering in time, et cetera. This has a number of implications. For example, if it implies with the kinds of slip rates we're talking about that there may be more than one near-maximum earthquake on one or more of the features that are critical to the facility nearby the facility, it means that the focus clearly has to be on multiple recurrences of near maximum events and not on multiple, multiple, multiple purposes of small events.

It means things like segmentation, which involve, for example, parts of faults breaking, and occasionally lots of parts of faults breaking, is less of a problem I believe. It means we have to more or less--there's a high likelihood that you're going to get a multi-segment event presumably. It also brings to question issues in conventional deterministic procedures; for example, use of the "maximum magnitude" together with an 84th percentile ground motion. If that is appropriate for a situation in which there's only going to be one such event of the future, is the 84th percentile still the right number when there are going to be multiple such events? It's not clear, and so I think even deterministic bases need to be reviewed in the context of these things also.

I would say more importantly is what the implication of long time frame means for in terms of careful thinking about the criteria, the statement of the criteria. We hear of numbers like 10^{-1} or 10^{-2}, and 10^{4} years, and so on. The question, is this really different here from, say 10^{-6} in one year?

Structural engineering practice, even though we deal with lifetimes typically of 50 or 100 years, economic lives of our facilities, states life safety concerns in terms of annual risks for very long-debated and good reasons, which I won't go into, but it's the way to do it. It gives you out of conundrums associated with building one structure every five years versus one for a hundred years, and their having different safety bases, which they surely should.

But if these processes were stationary, for example, then presumably we could be looking at the 10^{-5} risk instead, annual risk, as distinct from something which sounds a lot different, which is a 10^{-1} risk in 10,000 years.

If this degree of non-stationary, should it exist, is not very large, factors of 3 to 10 over the period of time, 10,000 years, potential non-stationary, that may not be important, given the kinds of uncertainties we already have in some of these rates. But I think where we really need work is in some kind of feedback between the decision makers back to the scientists about the questions of the sensitivity of statements you're trying to make about what's happening in

10,000 years.

And here we get into questions such as how are resource allocation prioritization decisions made in principal, and this principal involves some kind of risk cost benefit for society's resources. And the society means this generation and the next generation, no question about that, and lots of generations in the future. But it also means that there has to be, if you're making an intelligent risk assessment process, risk process of priority allocation, some kind of discounting. And once that discounting takes place, the impact of what the situation looks like 1,000 or 10,000 years from now on today's decisions is less. Sorry, that's the way it is. And that means if you want to do the best job for your progeny 10 generations from now, maybe you don't want to spend so much money, maybe you want to put it into other technology which improves our health care. And if you don't do it now, they're not going to be as well off somewhere else by if you waste the money here.

So this feedback of the decision process back to what the science means, I think it's something that has been missing. We've heard a lot about top-down decision making, or top-level, top-down processes. I haven't seen much of it, but, again, I don't know a lot about Yucca Mountain. It sounds attractive. I haven't seen a lot about how that's impacting the decisions as to when enough is enough, and it seems to be that's the essence of the problem.

Another question of Leon's was the facility involves radioactive waste. So what? To me, that means, wow, it's an important problem, and, therefore, we better do a state of the art job in terms of the science, and I think a state of the art job in terms of the science means a probabilistic hazard analysis with uncertainty and alternative interpretations, as opposed to trying to find a unique one that we all agree on, and that's why we're here, and that's what we're talking about, of course.

It means also that those scientific assessments and their coupling into a risk statement or frequency statement with uncertainty bands has to be communicated to the engineers. It has to be reviewed by the reviewers and dealt with by the decision makers, and they may have to do some hard work, too, as a result, to make sure that they're up speed with reviewing procedures that are done this way.

Volcanism versus earthquakes. As far as I can see, given the kind of structure of the model we talked about at the beginning, they are equivalent problems from the point of view of--the approach to them from a probabilistic point of view. There truly are differences, and I plead a great deal of ignorance on the volcanic problem, but from the point of view of a probabilistic model, they're equivalent problems.

Finally, I would ask the question whether--because

Clarence said we have to talk about deterministic alternatives, right? If I criticize the deterministic people for not thinking about probability, I've got to think about the alternatives, too.

We have at least two models, I would contend. Let's say critical facility analysis today mostly means nuclear power plants and the past 10 or 15 years of doing deterministic design bases. And one is the eastern United States, which is sort of a low-seismicity, long-history case, and the other is California, which is a relatively short-history, but high-deformation rate case, and Yucca Mountain is neither one of those, as I understand it. And the question is which of those two models is right for Yucca Mountain, if either, and how do we differ? And Yucca Mountain is both the question of the seismology and the question of the time window in which you're looking.

For example, if indeed the seismic deformation rates are 100 to 1,000 times less, does it mean we can take the 10,000-year window and divide it by 100 to 1,000 and say, this is just like a California problem with a 10 or 100-year economic life of our facility? There's a time exchange problem for you, in which case it would argue that a California deterministic design basis procedure ought to be about right, which is where you usually take sort of a max credible magnitude, some judgment of it, and use an 84 percentile ground motion.

So I've solved the deterministic problem, if that's what you want.

Thank you.

DR. ALLEN: Thank you, Allin. Very provocative. Unfortunately, we're not going to get any lunch if we don't open up the discussion here. So I think I'd rather go on.

DR. CORNELL: Of course.

DR. ALLEN: Do you have one short comment?

DR. NORTH: I'd just like to make the short comment, for those in the audience who did not attend or do not know about the workshop that was held on elicitation of expert judgment, I think Glen Hoffman in the audience can tell you where you could find a copy of the final report on that workshop. It addresses many of the very deep and provocative issues that Dr. Cornell has set forth briefly in his excellent summary.

DR. ALLEN: Okay. Thank you, Allin. I'm sure we'll come back to some of these questions in the session tomorrow afternoon.

Our final speaker in the morning session is Bob Budnitz, who has been involved with nuclear reactor safety for many years. For several years he was with the NRC, where he was director of the Office of Nuclear Regulatory Research. He's currently president of Future Resources Associates Incorporated, and he currently chairs the National Academy of

Sciences Committee on remediation of buried and tank waste, and he is also a member of this NAS committee on the technical basis for Yucca Mountain standards that I mentioned earlier. He's also been involved in WIPP. Welcome, Bob.

DR. BUDNITZ: Well, Allin, I'm a scientist who hangs around with engineers.

Many of you may also know that there's a project going on for the last year, and another year to go, to try to develop an improved methodology for probabilistic seismic hazard analysis. It's co-sponsored by EPRI, NRC and DOE, and it's a seven-member committee and a whole lot of technical support developing what we hope will be an improved methodology PSHA. And I chair that. Allin Cornell is on it, and Kevin Coppersmith, who is here, is on it, and several other people that many of you know, Dave Boore, Lloyd Cluff for example, and that's my most recent exposure to probabilistic seismic hazard analysis. But that's not what I'm going to talk about today.

This is going to be a systems perspective, and what I hope to give you is a perspective about how probabilistic hazard analysis and probabilistic facility analysis generally works, what the problems are, and in particular, how the analysis fits into how you regulate or assure safety in a probabilistic framework.

Now, I'm going to start with a simple problem. I want you to imagine some external hazard. It might be a tornado or a flood or an earthquake, and this is a particular hazard that has a maximum size. For example, there are no 800-mile-per-hour tornadoes, so it has a maximum size. This is a very--and here's a hazard curve for it with some annual frequency. I'll show you a wind hazard curve later that looks like that.

And now I have a single component. It might be a valve, or it might be a small building that's made out of steel that can withstand whatever this maximum size is. It's very simple for you to figure out that the way to assure that this thing is absolutely robust against that hazard is to make its fragility curve or its capacity higher than wherever that cuts off. That's the easiest thing in the world. And, of course, for those of you who are not familiar with a fragility curve, this is size in some figures of merit. This might be wind speed, or some way of characterizing an earthquake, or whatever, and this is the probability of failure of this gadget as a function of size, and at certain size, it fails.

By the way, we have data like this, for example, from shake tables for earthquakes, and this isn't a step function because not all gadgets that are identical actually fail identically because they aren't actually identical, even though you think they are.

So this is a fragility curve which shows in a trivial way that it is possible, at least in principal, to design a single thing to withstand with high assurance some external threat that has a maximum size.

Of course, you have to know that well, and you have to know this well, which is a story which I'm coming to.

But, of course, the world isn't that way. Most of our hazards don't have a maximum size. Most earthquakes, for an example, we don't think at least in the regions of interest here that there's a maximum size, or at least Tarzana was 2g, right?

And what that means is, even if you have a single gadget, unless it's very strong, you can't design for absolute certainty. All you can do is say I have some goal that I'm going to try to design for and do that.

So for example here, you might say, gee, what I'd like to do is make sure that this gadget, it might be a pump, or it might be a nuclear power plant, has a high assurance of being better than 10^{-6} per year.

So you pick off the hazard. You find out what size that is, and this might be a tornado or an earthquake, and you just make sure that that fragility curve looks like that. Actually, I could have drawn it so it started here, and I would have high assurance of 10^{-6}. If you want 10^{-7}, you've got to make it stronger.

Now, there's another problem which everybody understands. We have uncertainty in these curves, and we have uncertainty in these fragility curves, even though I drew it with a shape like that, because the uncertainties are actually quite broad. There are all sorts of reasons why there's uncertainty.

What that means, of course, is if you want to have a certain level of assurance, you have to understand the uncertainty as well, and you have to make sure that this thing is strong enough to meet that.

That's the simple way of understanding how a hazard and a gadget, and again, it might be a valve or it might be a nuclear plant, interact and how the risk of failure, whatever that risk is, can be determined by working out the fragility curve. Again, for a valve, you can put it on the shake table. For a nuclear power reactor, you have to do analysis as well of structures and tanks, some of which you can't put on a shake table, and some which you can. And you have to know the hazard.

Now, if we could regulate probabilistically, which we're not doing, for example, total nuclear power, I'll come back to that in a minute, and if we know the hazard curves, and if we knew these probabilistic of failure, like this is a function of size, and if we could characterize size properly for all the hazards, you would do that, and you would know what your target was and you'd know when you met it, and that would be a terrific world.

By the way, the world is that way for some of our external hazards for some facilities, but that's generally not the case for most of the important things that we work on. Furthermore, it's a trick to define fragility or capacity for complex systems, and I'm going to tell you about that in the next slide.

You see, I want you to imagine that this is a nuclear power plant or a refinery, and it's an earthquake we're worried about, and the fillet has four components, just four components, A, B, C and D.

Now, Components C and D fail together. They don't fail alone. This is their fragility curve. They fail together at around 10^{-5} per year earthquake, however big an earthquake that is. And what happens is when you get a 10^{-5} per year earthquake that comes along, that's the fragility curve for the plant, and that means that the probability of failure, if that's the simplest model of all, is 10^{-5} per year, you're going to get an earthquake that big, and it's going to be the failure, and that's going to be the probability that the power plant is going to have an accident.

And if that's all you knew, then you could define very well, and that's all there was, the fragility or the capacity of this nuclear power plant.

The trouble is it's not that simple. I want you to imagine the same nuclear power plant has two other components, A and B. B is extremely strong for earthquakes. It never fails, but it's out for maintenance some of the time. And the other thing is A and B. A is a--maybe you can't see because it got rubbed off. A is a seismic failure, but it's much weaker. It occurs at a much smaller earthquake, a 10^{-3} per year earthquake. But when A fails with that earthquake, you still don't get any trouble unless B fails.

Now, here's the point: If B is out 10^{-1} of the time; that is, 35 days a year it's out for test and maintenance, a tenth of the time, right, it's just not there when you want it, then the overall failure is 10^{-4} per year, because it's 10^{-3} at the time you get the earthquake, and 1 time in 10 why the thing ain't there, and you get a core damage accident.

If B is 10^{-2}, the failure is 10^{-5}, which, by the way, is the same as that other one. And if B is 10^{-3}, that is it's only out 8 hours a year out of the 8,000, a third of a day out of the whole year for test and maintenance, then that multiplies out to 10^{-6} per year.

Now, the question I want to ask you is, what is the seismic capacity of this gadget? And I insist in this scenario it's totally not well defined. It depends on B. And that's a lesson that I have spent 10 years trying to convince seismic PRA people, including seismic structural people and regulators about. They don't seem to understand that in this situation there is no unique seismic capacity for the plant, whereas if this never happened, if B never failed, and that one wasn't there, we would have a seismic capacity. As plain as that.

And that's an important lesson I'm going to come back to at Yucca Mountain because, in fact, the seismic capacity has a meaning only in terms of non-seismic processes. And by the way, the volcanic and non-volcanic processes. There are all these other things that interact with what the earthquake does, which tell you about the figure of merit. And the figure of merit isn't the seismic capacity anyway. It's this thing down here, which is the core damage frequency, or at Yucca Mountain, the probability of some release that you don't want.

And that concept, which I'm going to come back to, I think is a complex one, especially at Yucca Mountain where it isn't a pump that's failing in an earthquake.

Now, I'm going to give some examples of hazards just to show you that, in fact, there are cutoffs. This is a wind hazard at Indian Point Nuclear Power Station on the Hudson River north of New York City. And without arguing there are hazard curves with different probabilities because of different--I won't go into the details. And there are hurricanes and tornadoes. But you can see that there are no 800-mile-per-hour winds at Indian Point.

The other thing to tell you is that Indian Point is designed for 320 miles per hour, which is around 10^{-7} per year, if you believe these hazard curves. 10^{-7} per year is the design basis for the Indian Point containment structure, and as far as I'm concerned, and everybody in the business understands this too, that's not a risk at Indian Point. You don't have to worry about winds in the containment at Indian Point. There are other hazards, but not that. This is an example of something that effectively has a cutoff against which you can design the whole facility.

By the way, other things in the yard fail when you have winds that size, but the reactor itself can survive because it has enough things that are protected, and so winds aren't a problem.

I'm going to give you another example, though, that's quite different. This is a plot obscurely shown, but I have to explain it. There are 110 nuclear power stations in the United States, and they sit on 69 or 71 sites, something like that. And the SSE is the design basis earthquake. This is the probability of exceedence, accumulative plot of the probability of exceedence of those

100 reactors. A few of them, the SSE is actually not much higher than 10^{-7} per year. These are the EPRI hazard curves mean.

The vast majority of them, the SSE is in the range of 10^{-4} to 10^{-5} per year. A few of them, the SSE is several 10^{-4} per year. You see, the SSE wasn't picked in a probabilistic basis. It was picked in a completely different basis, and now that we've done hazard studies, like the EPRI hazard curve, hazard study, you find out that the SSE that was picked for all of our 100 odd power stations, nuclear power stations, varies all over the world, from 10^{-6} or worse, or lower, to higher than 10^{-4}, almost 10^{-3}.

By the way, the Livermore hazard curves are plotted this way--excuse me, this is median, not mean. The Livermore hazard curves plotted this way follow almost exactly on top of this, plotted this way, except there's a factor of three difference.

So the median, the 50th percentile one is about 10^{-4}, instead of 10^{-5}.

And I raise that because in my next slide I'm going to comment about it.

Now, as I'm sure everybody knows, nothing in the licensing of the current 100 odd nuclear power stations in these arenas was done in a probabilistic basis. It was all done in another basis. And the question that's worth asking is well, how did it come out?

You see, nuclear power plant licensing, as I say, is not probabilistic. It takes a traditional deterministic approach. It uses something called design basis earthquake that was arrived at in a way that I can't tell you about here without going into a diversion. It has a design basis wind. It has a standard project flood for the first 25 power stations, and for the next 80, it was the probable maximum flood that was used. And these turn out to be a few hundred year things.

And then it uses the standards and codes and design rules, the ways of inventing margin against those design bases, and you rely on those margins to get you where you are.

And when they started this process, they had no idea what the probability of core damage was going to be. But today we know. I have on my shelf 40 full-scope PRAs, and by the end of two years, all 110 plants are going to have PRAs that include all of these external hazards. And so now we know how well these judgments came out, and I'm going to tell you roughly how they came out.

And I'm also going to point out that the NRC has established as a policy goal that the body of plants should have probabilities of a large release of radioactivity in a range of or better than about 10^{-6} per year.

Well, from these PRAs that I have on my shelf, for earthquakes, that should say about 10^{-3} that that thing is scrubbed out. For earthquakes, the design basis ranges from about 10^{-3} to 10^{-6}.

The probability of a large release in the PRA, the outcome, is the range of 10^{-5} to 10^{-6}. Actually, it's better than 10^{-6} for the best of them, which means that there's another factor of 100 between the recurrence of the design and the worst of these, which is actually pretty good. That is the plants where the recurrence is about 10^{-3} per year, and only every hundredth of those produces the bad outcome, the other 99 out of 100 don't. That's that factor of 100.

For winds, as you saw, the design basis is 10^{-7}, and the outcome is too small to worry about. For floods, many of those floods are in the range of several hundred year recurrences, although some of them, by the way, are high and dry where you don't have to worry. They're 200 feet above the Susquehanna River, for example. But a lot of them are in the range of several hundred years.

Nevertheless, for floods, the core damage frequency is 10^{-6} or better, which means that they have a very strong, robust behavior against those floods even when they happen, even when the project flood is exceeded.

For internal fires, just to give you a bench point, at around 10^{-3} per year, some damage, some fire damage happens, but not a complete disaster, and the fire--show that you get the very bad accidents around 10^{-5} or 10^{-6} per year.

Now, there's an important point I want to be sure to make here, which is that all of these numbers have a hell of a lot of expert judgment in them, as well as analysis and data and models. The PRAs have as much expert judgment as data and models in the way it affects the outcome, and because of that, they're not very solid numbers. But nevertheless, the lesson of all of this is a pretty good lesson, separate from the details of the numbers.

The reason I--and by the way, what the actual recurrence is here, for some of these are full of expert judgment, too, you know about the seismic we're talking about, and it's certainly true of the others as well.

Now, why is it, going back to the nuclear plants--this is my previous one--not licensed probabilistically? Why is that? The reason is because the regulators traditionally, and I don't blame them at all, don't trust the probabilistic approaches. They don't believe that regulating this way is the right way to regulate. And that makes perfect sense to me because these aren't good enough. You're not sure you captured everything, and so they've relied on the traditional methods of picking a design basis, making sure it's strong against it, defense and depth, redundancy, diversity and so on, and because the engineering was as good as it was, the outcomes are actually pretty good, which is nice to know. We are in the range of 10^{-5} or 10^{-6} per year.

But I want to insist that for today's nuclear power plants, there's no specific role for the PRAs, for the probabilistic analyses and licensing. They've been used to examine the plants to find out if there are weaknesses. For future nuclear power plants, they're going to be used as a check. They're asking all the applicants for the new power plants, if there ever are any, the new PRAs, as a check, but that's not going to be a licensing criteria.

And the acceptability of all of this is based on expert judgment, and there's nothing wrong with that. It has to be that way even for something where we have thousands of years, of reactor years of experience already, they're running every day, there are hundreds of people watching them all the time, and how can it then not be so for Yucca Mountain? How can it not be so? Of course, expert judgment is going to have to be there if it's there for nuclear power stations, for which thousands of reactor years exist.

So I want to insist that expert judgment is an intrinsic part of this whole thing.

Now, let me turn to Yucca Mountain and talk about the applicability of these lessons for Yucca Mountain, and just to start with a thought, I want to talk about the standards for a minute.

Now, what are the standards for Yucca Mountain, and I want to insist that I have no idea. And since I'm on the Academy committee that's supposed to be recommending that, it's a bit disingenuous, but, in fact, we're still working it out. And any of you that have ever been on one of those committees, I just couldn't say if I could, but anyway, I have no idea. I have no idea, and I don't think anybody knows yet whether it will end up being a dose based standard or an individual risk based standard or a release fraction standard like the old Part 191 that was remanded, or how it's going to be cast, whether it's going to be 1,000 years or 100,000 years. I just have no idea.

But for the purposes of this discussion, I'm going to postulate something. I'm going to postulate that the figure of merit, whatever it is, is performance-based, rather than deterministic where you have rules, and then you go and see whether you met the rules. And by the way, the Part 191 that was remanded, but by the way, it's in place for WIPP, is a performance basis standard. I'm also going to postulate that it's a probabilistic standard, like the old Part 191 was, but I don't know what the figure of merit is going to be. I want to insist, though, inevitably they're going to be expert judgments even in working out these probabilistic analyses, never minding judging whether the standard's been met.

And now with those postulates, I'm going to go and discuss what I think the issues are for these external threats. I want you to imagine that a performance based--well, performance assessment, a probabilistic performance base assessment is done. What does that mean? It means--and let's forget about earthquakes at the time. Suppose there are no external hazards. It sits there undisturbed for this whole time without any volcanism or earthquakes or anything like that. That's just Postulate No. 1.

Analysts have to work out, and by the way, they have started to work out, although it's a difficult process, how the thing performs. So it is a function of time that casks finally may degrade and the radioactivity decays, and then the groundwater does this or that, and there may be infiltration. And finally, may or maybe or may not some radionuclides are transported someplace, and then finally, somebody may get a dose, or there may be a release.

And I'm going to show that in a very stylized way of saying that whatever your figure of merit is, a dose or a release or something--suppose the figure of merit is, the cumulative release over 10,000 years, which by the way was the old Part 191. Then somebody then does this analysis. This is just a probability density function for it, and it looks like that. And this is a CCDF, which is nothing but the probability of--one minus the probability of exceedence.

It's the same curve turned around.

Now, let's suppose that analysis has been done for Yucca Mountain, but not for earthquakes or for volcanism. We want to ask the question, well, what does earthquakes have to do with this analysis? Well, let me now postulate for you that the same analysis is done, but for a couple of different earthquake sizes, so that the black curve is the original one I just showed, and that's no earthquakes.

Earthquake 1 is the earthquake that's 10^{-3} per year. By the way, I don't know which earthquake that is, and you just heard a talk saying that we don't know that well enough yet, but we ultimately would know with some accuracy what Earthquake No. 1 looks like. And there may be, whatever this figure of merit is, there may be some higher release. There may not be. It may look like the black curve.

And then we're going to take Earthquake No. 2, say it's a much larger earthquake, a 10^{-5} per year earthquake, and by the way, I've only now postulated just one earthquake here, and you get some other release, and the CCDF looks like that.

And, of course, that analysis has not been done, and then ultimately, you know, if it's the old Part 191, there's some figure of merit in here, you know, that step thing, and you have to see whether you exceeded it or not. And if it's in here, you lost, and if it's over there, you've got your license.

I mean, this is all very stylized, but I'm trying to make a point. And by the way, another way of saying it is you might have a limit in here on some dose and you exceed it or not. There are various ways of expressing it, and I can't speculate at all as to how that's going to come out.

Now, that analysis has not been done. By the way, it's not been done at WIPP, but because WIPP is in salt, it's thought not to be a problem for earthquakes, but it's certainly not been done here.

I want to postulate for you that we need a seismic performance assessment, and what I mean by that is suppose this is the undisturbed case, no earthquakes, and it's below the unit. I'm going to use probably density rather than CCDF space, but it's the same concept. And these are different earthquakes, Earthquake 1, Earthquake 2. These might be 10^{-2} per year, 10^{-3} per year, 10^{-5}, 10^{-6} per year. They're just larger earthquakes however defined.

What I'd sure like to know is whether not much changes, and this is some fragility curve which measures when you get into trouble. But ultimately you get the same earthquake, and you know it, you're in trouble. And so it goes along, and maybe the curve looks like this. Or maybe it looks like this. Or maybe it looks like this. Even the largest earthquakes don't get you into trouble, or maybe even small ones do.

By the way, I don't know what trouble means. It's this figure of merit I didn't define very well, but we need to know that. And the reason we need to know that is if, in fact, even the largest earthquakes don't cause a release of concern for the standard, then we're wasting our money worrying about earthquakes. On the other hand if .01g earthquake kills you against the standard, we're also wasting our money. We've got to go find another site, right?

We need to know that, and we don't know it. And until we know that, we don't know how much to characterize or what. Of course, it's driven by the standard. We don't know what the figure of merit is. We don't know whether it's going to be doses, individual risk or release limits. We don't know whether it's 1,000 years or 10,000 years.

And, of course, there's a further complication, which I'm going to come to on the next slide, which is do we run each of these little earthquakes singly and see what it does? Do we do a weighted sum? If the Earthquake 1 is--if it's a 10^{-3} earthquake, and we expect 10 of them in 10 millennia, do we run 10 of them, Poisson-wise? I don't know. Nobody has thought that through yet in terms of how the standard is going to be, and certainly nobody in the project has done a performance assessment with 10 earthquakes of 10^{-3} size, or is it a 10^{-5} earthquake where we expect only one-tenth of one, a chance of one in ten that we'll even have one in 10,000 years? Or is it a 10^{-6} earthquake that we'll have one chance in ten in 100,000 years?

Has that been done? Until it's done and until that interaction between how the thing behaves, which, by the way, means affecting hydrology or changing the paths to the environment or the casks will disintegrate, or it will change the chemistry, or God knows what else, until that's been done interactively between the behavior of the system and what the earthquake does, different earthquakes, different numbers of them, we don't know whether earthquakes matter at Yucca Mountain. We don't know now, and we won't know.

And that's my lesson, and it goes back to this question here, to the nuclear power station. You see, if there's never a problem with B--the seismic capacity is here. If B fails almost all the time, and if it's out one-tenth of the time, why there's really a problem here. And the seismic capacity of this reactor is different depending on what B is about, and that's the response of the non-seismic parts, which is related to this issue here about what happens with this performance assessment, and I insist that we don't know. And, of course, we don't have a standard, but that shouldn't be an excuse for not having got started on this analysis because the analysis has got to be done one way or the other and understood all these phenomena separate from whatever the standard is.

So let me just--I'm coming to the end here. Let me finish with a question I really don't know about. This is an open question for me, and of course for you, and for the Department and for EPA and NRC. Do we design for the largest earthquake in some time period, even if it only has a chance of 1 in 100? Do we design for it, or do we assess it and see whether we have to design for it? Do we design for some risk index, which is a weighted sum of these things? And I ask the question, is the design tied to the performance assessment, in which case you do the performance assessment for all these different things, and then you change the design if you've got to, or is it only reactive? You do the design independent of it, and you go and look and see what you've got.

And we know the answer to that. It's going to be interactive. There's too many billions being spent here for it not to be interactive.

But, in fact, this is a necessary iteration in which the behavior of the system for these earthquakes or volcanoes or whatever has to be assessed, and then if necessary, the design has to be looked at and changed, tied to the performance assessment until you meet the standard, unless, of course, the fact is that the bottom curve is right. Even the largest earthquake won't cause a release that matters, in which case we're all wasting our money. Or, unless this curve is right, even the smallest earthquake kills the project.

And just going back to my last slide, I think that all that's being done in seismic on this project has to be linked to the performance assessment, which, by the way, is linked to the engineering, which hasn't been settled, and the seismic hasn't been settled. So there's this great big iteration that isn't going on that's necessary and which I think is the fundamental lesson that this project can learn from the nuclear reactor business where that coupling between the seismic and the non-seismic failure, wherever it is, was shown to you before, that B that's a non-seismic failure, that tells you about what the actual seismic issue is.

Thank you very much.

DR. ALLEN: Thank you, Bob. You're finished ahead of time, and so we do have time for questions.

DR. BUDNITZ: Well, you can ask Allin questions.

DR. ALLEN: Let me ask the board members if they have comments or questions. Warner North?

DR. NORTH: Warner North. I think your questions are right on. I think one of the frustrations, I will speak for myself as chair of the Risk and Performance Analysis Panel. I've been here five years, and I'd love to see these questions answered. We have repeatedly urged on the

Department of Energy's program that they use performance assessment in an iterative way to assist in the basic decisions such as design, given the repository goes ahead, to guide the program.

And there's a lot unrealized potential in that area, and I think you've done a very good job of giving us a well-focused example of how one should proceed to try to realize that potential, get these questions answered, and use them as the basis for setting priorities and deciding how much more we really need to know in one area of the scientific investigation.

The interaction with the design questions as we move to a new baseline for the ESF and for a potential repository if the process proceeds in that direction, these issues really need a lot more analysis and thought communicated in public than they have yet have.

DR. BUDNITZ: I agree with that. That was easy for me to agree with that.

DR. ALLEN: John Cantlon?

DR. CANTLON: Yes. We started off with a set of standards that were to choose among sites, as opposed to evaluating the total system, which had a technology of a repository. And in a sense, we're caught in a process in which DOE had a great deal of time to begin to put the system together as a real system interacting between characteristics of the site and the characteristics of some kind of an engineered repository.

DR. BUDNITZ: Oh, I understand that perfectly well, but I want to remind everybody in the room that had some aggressive lawyers not sued EPA in 1986, we would have had the standard. It would have been in place all these years. We don't have it--we haven't had it since I don't know when, '88 or so. But, and in fact, Part 60 was written before 191, and it's going to have to be--there's a whole lot of stuff going on, but that doesn't excuse, and Warner just said it as well as anybody, it doesn't excuse going ahead with the analysis, which is going to have to be done in any event to support whatever happens.

And I can assure you, without talking out of school about the NAS Committee on the Yucca Mountain standards that I'm part of, that we're not thinking about site comparisons. Of course, our mandate is to recommend the technical basis for standard at Yucca Mountain.

DR. NORTH: Right.

DR. BUDNITZ: Which it's in the legislation.

DR. NORTH: Yes.

DR. ALLEN: Questions from the consultants? I'm sorry. Bill Melson?

DR. MELSON: Bob, in terms of designing, let's say in a reactor, so we'll withstand a certain magnitude earthquake, there's some experience in engineering and others that allows you to do this.

DR. BUDNITZ: Yes, sir.

DR. MELSON: I think Greg Valentine's work that Frank Perry talked about, he's trying to model how the heck you deal with, for example, intrusion of a tunnel by a body of magma.

DR. BUDNITZ: Yeah.

DR. MELSON: These are things we don't have a whole lot of experience in, where the engineers and people like Greg have to get together and allow that possibility to be and design something that can handle that. But the difficulty is there's a lot of uncertainty in how you do that, more so than say, the background of seismicity.

DR. BUDNITZ: Oh, yeah, I, of course, not only understand, but I'm puzzled by--imagine you have a five in a year recurrence earthquake, and it happens 20 times in the first 10,000 years, and gradually changes the water flow. Well, I'm not a hydrologist, but I don't think that's an easy thing to model, either here or at any other place where there might be a natural analog that you could use. It presents a formidable challenge unless it's very strong against it, you know, unless you can--and if that's so, it would be great.

I mean, that's what's nice about WIPP. The salt seems to be--the fact that it's there almost tells you that it seems to be robust against that, and that's what makes that particular against these insults--it makes that particular site analyzable almost by default. I use the word default not as a pun, but just as a fact.

DR. ALLEN: Michael?

DR. SHERIDAN: Yeah, I was just thinking, Bob, that your argument seems to be a very strong one for sensitivity studies of different models to determine if a model is prohibitive or permissive.

DR. BUDNITZ: That's fair enough, yeah.

DR. SHERIDAN: And there's a range in between so that your recommendation would be to examine multiple models with some sort of CCDF; is that correct?

DR. BUDNITZ: Yeah. My friend, Chris Whipple, keeps arguing for a 20,000-year cast, which, of course, would satisfy a 10,000-year criterium by itself. I'm not going to argue that here.

DR. SHERIDAN: Unless a volcano erupted right through that.

DR. BUDNITZ: No, I understood it. I understood. But the fact is that I'm not going to argue that here, but there are engineering approaches that you could conceive of, which override the necessity for analysis. And by the way, the salt site for seismic is another example of a site that overrides it, but there are engineering approaches which could do it. Of course, you know, it's only a few billion dollars, and none of us want to spend that money frivolously. It's really--it's a totally non-trivial problem of interaction between design and figures of merit.

DR. ALLEN: Other questions? Since Bob purposely timed his talk to encourage questions, let me turn to the audience and ask if there are any questions? Leon Reiter?

DR. REITER: Just a comment, Bob. I think your comments about the seismic not having the seismic PA to get the insights, a couple comments on that. I think that you'll see, particularly tomorrow, that there's perhaps a lot more insights being developed in the volcanism, volcanic hazard.

DR. BUDNITZ: I know. I know that.

DR. REITER: And you'll see some good arguments being presented by the people in DOE, which would suggest, at least in certain aspects, enough may be enough. So I think you can see some of that.

DR. BUDNITZ: I understand that. I think that's not true for earthquakes it.

DR. REITER: Yeah, with earthquakes, I think--you know, we had a meeting on seismic vulnerability, and I think the board became convinced that, at least for the suitability of the site, the focus is really not so much in ground motion, which is a relatively easily designed against, but more in fault displacement.

DR. BUDNITZ: Yeah.

DR. REITER: And there has been a lot of looking at to how much fault displacement is needed to look at the cask. So I don't think that--I'm not sure the situation is as bleak as--

DR. BUDNITZ: It may not be the casks. It may be other natural processes that are modified by succession of small ones or some larger ones. As I said, water flow or infiltration, combined with I don't know what else that might happen through other processes over this very long time. And by the way, 100,000 years is a lot higher than 10,000 in that regard.

DR. ALLEN: Okay. Let me turn to the audience and ask if there are any comments or questions.

Everyone must be hungry. Thank you, Bob, very much, and we'll reconvene at 1:30.

(Whereupon, a luncheon recess was taken.)

AFTERNOON SESSION

DR. ALLEN: Okay, the first speak on the afternoon program is Richard Quittmeyer, the Woodward-Clyde contractor to the DOE.

DR. QUITTMEYER: Okay, I've been asked to talk about a methodology to assess seismic hazards at Yucca Mountain that the DOE has been preparing a topical report that describes this methodology, and the first thing I'm going to do is present an overview of how the topical report fits into the overall seismic hazards program at DOE. Then I'll discuss the objectives that methodology was designed to meet; also spend a few moments discussing the design context in which the methodology is going to operate, and then go into discussing the various components of the methodology itself and summarize at the end.

This figure is an attempt to show the overall seismic hazards program for Yucca Mountain, and the relationship of the probabilistic seismic hazard methodology that the topical report describes to the overall program.

One aspect of the program which is ongoing is the collection and analysis of data to the site characterization activities, and these are described--these studies are described in study plans.

The topical report that I'm going to discuss deals with methodology to assess both vibratory ground motion and fault displacement hazards at Yucca Mountain. We also, though, envisioned two additional topical reports to describe other aspects of this program.

The second topical report will deal with determining the hazard levels appropriate for risk consistent seismic design. This will involve determination of performance categories and associated performance goals for SSC's at the geologic repository operations area, the seismic design criteria used to design a potential repository, risk reduction factors associated with those design criteria, and finally the level of hazard that's appropriate for design for each of these performance categories. And this will be, as I said, the topic--the subject of a second topical report which we hope to begin soon.

A third topical report will deal with how to develop the seismic design inputs, the seismic loads or fault displacements that are used in the design process.

You can see that the ultimate customer of the seismic hazards program is the seismic design process, and the assessment of the containment performance of a potential repository. This is where we ultimately want to get to. These are the steps along the way, and I'm going to be talking about this first step, the assessment of seismic hazards.

The objectives of the methodology to assess seismic hazards ultimately go back to the regulations. What we need to do is to provide the information that will allow us to design the potential repository, the geologic repository operation for seismic safety, and design it for waste containment, waste isolation.

We also need to design it to ensure that we can retrieve the waste during the pre-closure period, and we also need to describe and assess and evaluate features that might affect the design and performance, or the potentially adverse conditions.

Giving just a little bit more detail, to do that, there are some things that the methodology needs to do. It needs to assess the hazard from vibratory ground motion and fault displacement hazards. There are faults at and in the vicinity of the potential repository, so we need to address both types of seismic hazards.

The repository will have facilities at the surface and below ground, so we need to deal with the hazard in both of those situations. And there are also different time frames that we need to assess hazard for. There's a pre-closure time period. That mostly deals with design aspects. And there's a post-closure time frame that mostly deals with performance assessment aspects. And also we want our methodology to be such that it facilitates the regulatory review and decision-making that the results of the assessment need to support.

So next I'd just like to talk for a few minutes about the design context in which this methodology will operate.

When I think about seismic design and what the purpose of it is, it's to ensure that society is not exposed to unacceptable risks related to the occurrence of earthquakes. That's why we do seismic design.

If we think about it that way in terms of a risk, we're led to the fact that risk is a function of both the frequency of occurrence of an event and the consequences of an event. And if we're going to try to carry out seismic design for this purpose, we then need to factor frequency of occurrence of the event, in this case earthquake ground motion or fault displacement, into our assessment of the hazard.

Over the past several years, a performance goal-based design process has been developed, or has evolved, that links the consequences, frequency of occurrence, the design criteria, and the hazard level for design in a logical framework. And I'll just spend a second talking about the performance goal-based design process.

So the performance based design process is designed to give a design, and in this case, we're talking about seismic hazards, but the concept would apply to all the other natural phenomena hazards.

Take a look at the structure, systems and components and categorize them according to the consequences of their failure. And then establish performance goals for each category with the goal that risk is constant across the performance categories.

The SSC's that, if they fail, have more adverse consequences, will have more stringent performance goals than components that, if they fail, don't really have very large consequences.

Another aspect of this is establishing the design and acceptance evaluation criteria for each performance category. And these are the details of how the engineers design the systems for the various categories. And the conservativeness of these criteria result in a risk reduction from the performance goal. And the more conservative the criteria are, the larger the risk reduction.

And coming out of these three steps then, the hazard level that's appropriate for design is just related to the performance goal times the risk reduction factor. And application of this concept, of this design process, to Yucca Mountain will be the topic or the subject of the second topical report in which we'll establish the performance categories and goals as they apply to the Yucca Mountain situation.

Now I'll talk about the methodology that's described in the first topical report. We've adopted a probabilistic methodology, and that's primarily for three reasons. There are three aspects of the traditional deterministic approach in which we feel weaknesses are accommodated in the probabilistic methodology.

The first of these is incorporation of the frequency of occurrence of the earthquakes, of the hazard. The second is that we, within the probabilistic framework, we can explicitly incorporate the variability in the data and inputs that go into the assessment. This includes both the randomness of the earthquake process and the diversity of interpretation that results of different scientists looking at the available data.

The probabilistic assessment is required, is needed to support the long-term performance assessment and also the probabilistic approach is needed to support the performance goal-based seismic design process.

Okay, as I said on the first view graph, the methodology relies on established, generally accepted data collection and analyses. We probably require--or we do require more information to carry out a probabilistic assessment than we would for a deterministic one, especially because the frequency of occurrence is factored in. We need to know that.

The methodology also feeds back to the data collection by identifying through sensitivity analyses the types of information that have the most influence on the outcome of the assessment. We can use sensitivity analyses to identify which uncertainties carry through the analysis to provide the most uncertainty in the answer, and then direct resources trying to reduce those uncertainties.

A preliminary seismic hazard assessment for the ESF design, for instance, has identified that the background earthquake is very important, and DOE has allocated the resources to complete the historical earthquake catalog during this fiscal year in order to support a better understanding of what the background recurrence rates are.

The various components of the methodology are; to identify the sources and characterize them, to assess the frequency of occurrence and the maximum magnitude of the earthquakes associated with each source, and then, for the next step, is dependent on whether we're assessing the ground motion hazard or fault displacement hazard.

For ground motion, we need to assess the attenuation and levels of ground motion, and if we're looking at fault displacement hazard, we need to understand the amounts of fault displacement and the distribution and space of those displacements as a function of magnitude.

Once the inputs are developed, then you integrate over the data and the uncertainties and carry out sensitivity analyses to develop a more complete understanding of the hazard.

The methodology is based on a growing experience base which will be discussed in more detail in Kevin Coppersmith's talk a little bit later.

The experience is almost entirely with vibratory ground motion, but fault displacement is a very--is a similar phenomena. It's a time dependent phenomena and it has uncertainty in the spatial distribution of fault displacement. So it can be treated in a similar manner.

Now, I'll just discuss a little bit more the various components of the methodology.

Seismic source characterization. Here we're just trying to provide a spatial description of the sources of future earthquakes. For Yucca Mountain, we're primarily interested in fault sources, although there are--we do use volumetric sources to characterize activities such as background earthquakes which don't cause surface rupture.

Seismic source characterization will also include, for the Yucca Mountain situation, the identification of, or assessment of underground nuclear explosions.

Recurrence for fault sources will be based primarily on the geologic and paleoseismic data that's being gathered during site characterization. The models of recurrence that are employed, poissonian or characteristic, will be based on what the data shows.

Recurrence at Yucca Mountain will also potentially have to deal with issues such as temporal clustering. And, again, it will be the interpretations based on the data that's developed out there that tells us whether that is an alternative that we'll need to include in the analysis.

For volumetric sources, geologic and seismic data will form the basis of recurrence estimates. Seismic data here will play a much larger part.

In the methodology in the evaluational sources, we also need to identify maximum magnitudes for the various sources that are identified. For fault sources, we'll be using empirical relations that relate magnitude to various physical parameters of the fault sources, the length, rupture area, displacement, using again the data.

If fault segmentation is important, we will also incorporate that into our estimates of maximum magnitude. For the volumetric sources, the maximum magnitudes will be based on tectonic analysis on comparison to observations and tectonic regimes that are similar to Yucca Mountain. And we can use the magnitudes of earthquakes with observed surface rupture. The smallest magnitudes that have surface rupture is an upper bound on that.

For vibratory ground motion assessment, the next step is to develop the ground motion attenuation evaluation. We'll be using both empirical and numerical methods. The empirical and stocastic numerical methods will be the primary focus, and other numerical methods will be used primarily to provide information on near-field ground motion effects.

The ground motion evaluation will also include assessments of the various factors that can lead to site responses, the local geology, the velocity, shallow velocity gradients, topography. And of particular importance to the Yucca Mountain situation, attenuation of ground motion with depth. This will be important for evaluating hazards and designing the underground facilities.

If we're looking at fault displacement hazard, then we need to develop an evaluation of the amount and the spatial distribution of faulting. Again, we'll be looking to empirical relations between displacement and magnitude, and also empirical relations that describe the amount and distribution of secondary faulting. This will be, you know, particularly important for the faults that are in the vicinity of the potential repository site.

Using the probabilistic approach, we can also include the possibility of new faulting, even though the likelihood may be very small. That alternative can be included in the analysis if it is appropriate.

In developing all these inputs, a question has arisen as, you know, what is the role of expert judgment. And I guess I would start by saying that I think expert judgment is going to be used whether we use a probabilistic methodology or deterministic methodology. Expert judgment is necessary in interpreting the available data for all the various inputs to the assessment.

In terms of how exactly we'll do that at Yucca Mountain, the current concept or current approach will be to rely on the experts who are involved in the program and who have the most familiarity with the geology and the work that's going on out there, and to have them develop the data and evaluations that will be used as input and to describe and assess the uncertainties, the different alternatives that can be supported by the available data.

Once the data and the evaluation of the uncertainty have been produced, then the hazard assessment progresses by integrating over the inputs to produce a curve that shows the annual probability that either ground motion or fault displacement will be exceeded.

Propagation of uncertainty within the methodology can be done by either of two equivalent methods. The logic tree approach will define discrete alternatives for the various inputs and evaluate their likelihood. The Monte Carlo method will take continuous distribution description of the uncertainties and use a random sampling approach to incorporate the uncertainty into the analysis.

The final step in carrying out the assessment, and this is an important step, is to carry out the sensitivity analyses to provide a more complete understanding of what's going into the hazard assessment, of what the hazard assessment is telling us.

The types of analyses that will be carried out are looking at the sensitivities to different inputs, to the uncertainties in those inputs. We'll also be de-aggregating the results to determine at various hazard levels, what the strong contributors are to the hazard at those levels. And we can also do reality checks, comparing the seismicity that's calculated from the sources and recurrence inputs, and comparing them to the observed seismicity, for instance.

The sensitivity analyses will also be used to help us determine when enough is enough. If we determine that for a particular source or for type of data that additional reduction and our knowledge of the uncertainty will not produce--you know, that that type of information is not a strong driver of the final hazard, that will be information that management can use in terms of deciding where to allocate their resources.

So to summarize, the approach that we use, the methodology that we've described in the topical report as the probabilistic methodology will provide the results that we need for waste containment, performance assessment and to support the seismic design process. It explicitly incorporates the frequency of occurrence of earthquakes, the variability, the randomness and uncertainty in inputs.

It includes the contribution to the hazard from all the sources, and will provide a basis for design and licensing decisions that's based on safety performance goals, compliance with waste containment performance goals, and provides extensive documentation of data interpretations and the sensitivity analyses from the assessment that will facilitate both regulatory and decision making processes within the project.

Any questions?

DR. ALLEN: Thank you, Richard. Questions from either the consultants or the Board? Dennis Price.

DR. PRICE: Let me ask you something about this. It appears to me that this topical report is--maybe ought to have an "R" between the "T" and the "O," a tropical report because it seems to be paradise, with the tradewinds blowing, and I could almost sense myself sitting at the beach and enjoying this thing. It only contained methodology in your report or, number two, where you have seismic source characterization and evaluation; will it be some evaluation output in that report? Will there be some evaluation ground motion attenuation in the report? Or is this only a report of methodology, what you will do someday?

DR. QUITTMEYER: Methodology; the series of comparable reports are designed to ascribe the methodology. We want to get the NRC's acceptance of our methodology, and then we'll go out and apply it using the data at the site and develop reports describing the actual results of that application.

DR. PRICE: Okay. These subsequent reports, will they have results in them or something of substance? This is methodology, and everybody knows we have to have methodology. But I was just--where are the results and when will we see the results and when will these come?

DR. QUITTMEYER: The results of applying this methodology are now planned for FY 96. Is that correct, Tim? I believe so.

MR. SULLIVAN: Tim Sullivan. I'll address this in a minute here after Richard has answered a few more questions.

What Richard is describing in these topical reports is a part of DOE's issue resolution strategy. The concept there is that DOE and NRC hopefully can reach closure on the appropriateness of the methodology to assess seismic hazards. And that methodology will not need to be addressed again during the license application process. Rather, we'll focus on the results.

DR. CANTLON: Yes, you mentioned DOE's plan on use of expert judgment, and if I understood you correctly, you restricted the experts to those that are involved in generating the data and are intimately involved.

It would seem to me in the light of DOE's clearly established credibility problem, that you do everything you could to get some external experts involved in the expert judgment phase of your work. Could you comment on that?

DR. QUITTMEYER: Input from scientists outside the project could be involved during the process. If we need to go to a more formal elicitation of interpretations from a wide variety of scientists, we can certainly do that. We're not excluding that.

Our primary approach, or our first approach would be to allow the scientists working on the project to use the data that they're familiar with to try and include the interpretations from anyone, from those in and outside the project, you know, try to define that diversity of interpretation. You know, if outside review panels convince DOE that that's not sufficient, we'll certainly do more. What we're trying to do is get the true diversity of interpretation. If it's decided that doing that within DOE is not sufficient, then we'll have to take the next step.

DR. NORTH: I share my colleague, Dr. Price's, assessment that it's like tropical paradise. I'm very concerned that five years after I became a member of this board, with very extensive discussion of the seismic issue in the board's report and promises from the Department of Energy that they were really going to take our advice seriously and do interative performance assessment, that you are standing here at this point and giving us methodology which, in my judgment, lacks substance.

I don't know how you are going to do it. I don't know how you are going to use the data. As far as I'm concerned, what you've given us are a set of reasonable platitudes for how the methodology is going to work. And, frankly, as an outside reviewer of your program, I can't say I have any confidence in it until I see the details.

I'd like to put you and the Department of Energy on strong notice that I, for one, am very impatient about the lack of progress. I want to see the details, I want to see iteration one with numerical illustrations of how you are going to take the data, how you are going to assemble the expert judgment, and how you are going to give us an initial iteration on the issue of seismic risk that can be useful to those who, for example, are considering the design decisions to go to, for example, horizontal drift emplacement instead of vertical bore holes, as the program is currently contemplated. Because from what I can see, what's been presented in public so far, there is nothing that will be very helpful to them as they address those decisions. I think you need to get serious and get specific.

DR. QUITTMEYER: Okay. You know, just given a half hour to describe the methodology, I certainly can't get into all the details. We are, though--

DR. NORTH: We want to be convinced that those details exist. Where is the report? Where is the product? Where's iteration one?

DR. QUITTMEYER: The report will be out within the next two weeks, delivered to DOE. They need to review it. So probably within a month, it will be out to the general public.

DR. NORTH: Is it doing to give us some content, some specific illustrations of how you are going to deal with the data, how you are going to deal with the many problems of the implementation of this methodology which you have given to us in such general terms?

DR. QUITTMEYER: I will certainly describe that in more detail than I was able to present right here. The second document that DOE is working on includes a preliminary--a seismic hazard assessment employing this methodology to support the selection of seismic design basis for the ESF. That's a little bit--the completion of that report is a little bit farther out in the future, maybe another two months. You know, that will show an example of how this methodology has been applied.

DR. ALLEN: Other comments from consultants or staff?

DR. REITER: I think, like Dr. North and Dr. Cantlon and Dr. Price were talking about, was the need, the immediate need of not only a bottoms up approach, but a top down approach, and those kind of insights. That's really important. Looking at this as the first step and then looking at other things, other steps may be orderly in some manner, but it may be wasteful in not developing the insights.

I think there were some good examples of what efforts that DOE has done which are helpful. To give an example, several years ago, we had a presentation by Asa Hadjian, who's sitting in our audience now, about the surface facility, and looking at earthquakes in terms of surface facility and what can go wrong. And I thought that was really insightful, understanding what can go wrong in a surface facility and how different types of ground motions or fault displacements affect it, really affect the way one looks at that facility.

And the other example about the external--use of external opinion, I quite honestly think that putting it off may not be the best way. In fact, you have a wonderful example of what was done, not by DOE, but by EPRI and Kevin Coppersmith, which you utilized outside expert opinions to look at fault displacement hazard, I think gave a very useful and in many ways certainly perceived as a more unbias kind of evaluation than you would do just by internal experts.

So I think the material is there and I think you have some good examples, and it might be worthwhile to learn from those.

MR. SULLIVAN: Leon, DOE's plan is to--I'm sorry--Tim Sullivan, DOE. Our plan is to first develop preliminary conclusions based on the work that the DOE experts do, and then as appropriate, involve outside experts. Do you feel that's a flawed strategy?

DR. REITER: My personal view of the "as appropriate" worries me.

MR. SULLIVAN: But in addition to that, again, I think the other issue, the issue of--you're looking at the top down, even those insiders, Bob Budnitz really laid out and gave examples of nuclear power plants that I think are really important, and I think--has advanced a little more than you have in giving us some of the insights.

DR. QUITTMEYER: DOE is now developing the study plans that will work out the tectonic effects, the effect of tectonic type events, earthquakes, on processes that may affect the performance, things like water table, fracture permeability. So the studies are moving forward.

DR. REITER: Again, excuse me, I think that's a certain mentality about well, we'll wait for the study plans and then we'll understand. I think there's a need for some scoping, initial studies right now to help you gain the understanding of where you're going to go with all parts of the program.

DR. ALLEN: Bob Budnitz had a comment or a question. Please pass the gavel to Allin Cornell.

DR. CORNELL: Just within the context of the seismic characterization, I think it's also important that you start to train your scientists to respond to people who try to elicit these uncertainties. And if they haven't been doing that on a regular basis, they've got a couple years of experience to gain before they're going to be able to do that well. So that's a reason to get an early start on that, even within the context of seismic characterization.

DR. ALLEN: Bob Budnitz?

DR. BUDNITZ: Well, my comment Leon said in part, but I'll try to say it in a little different way. There are many different outcomes of a seismic hazard assessment, including displacement, velocity and spectral acceleration, all sorts of other things, as you know perfectly well, and you require constant feedback from the designers and facility to understand which of those are necessary for the repository itself. But you also need a different kind of--same kind, but a different perspective from the people who are modelling the broader repository, that is the ground water people and the transport people to understand which kinds of motions or accumulated displacements they need to know about in order that it might affect their part of the modelling and performance. Without that feedback, you may find that in years hence, and it isn't one, what you've produced isn't what they need. And I didn't see as much of that in your brief presentation as I would have liked. That doesn't mean you haven't thought that way, but it didn't emerge, maybe you are thinking that way, on both facility and what I'll call the environment.

DR. ALLEN: May we move on? Tim will be next up.

DR. QUITTMEYER: I can stop here.

DR. ALLEN: Thank you, Richard. And we'll turn next to Tim Sullivan of the DOE, who will be talking on the use of probabilistic seismic hazard assessment in the Yucca Mountain program.

MR. SULLIVAN: I'll ask Richard to turn these view graphs for me so that I don't forget to make any of the remarks that I have prepared.

My name is Tim Sullivan. I work in the project office in Las Vegas, and I used to be a geologist, but these days I interact mostly with project management staff and DOE in headquarters with schedulers and planners and bean counters and occasionally I go out in the field and look at the trenches.

On this first view graph, I've identified those questions that I will address today, some in full, some in part, and others I will still waffle on a bit.

The site characterization plan still forms the foundation for DOE's pre-closure tectonics program. In italics there I have identified the objectives of that program. I think it's important to keep in mind that by pre-closure, we're referring to that period of approximately 100 years, and our concern there is for the engineered structures at the surface and in the underground.

Recently, DOE has identified some changes in the SCP which have led to the topical report that Richard just described. The SCP adopted a dual deterministic-probabilistic approach to pre-closure of seismic hazards assessment. We have now determined that a probabilistic approach is appropriate, both for the pre-closure and the post-closure, and the topical report just described will present that approach.

After that is finalized, changes in the baseline documents, including the study plans, will be initiated and completed.

Pre-closure tectonics has been emphasized in John's talk and as well in my remarks today, because all of the data collection and analysis activities that support both the pre and post-closure tectonics programs are actually in the pre-closure section of the SCP, a total of some 18 study plans.

The result of the pre-closure tectonics program will be to provide repository seismic design bases for ground motion and fault displacement. And to do this, DOE intends to incorporate frequency and rate of occurrence information in developing these design bases by using a probabilistic seismic hazard assessment.

As described by John, the average slip rates on faults in the Yucca Mountain area range over more than two orders of magnitude, from less than--from several hundredths of a millimeter a year in the site area to several millimeters per year on the Furnace Creek fault. Thus, probabilistic seismic hazard analysis will allow for a combination of all of these sources in a single hazard curve or family of curves. These curves will be developed for ground motion and then separately for potential fault displacement.

The USGS is the participant organization that's responsible for this work. The first study plans that will describe the implementation details of the seismic hazard assessments are being prepared now. And as source characterization or site characterization data related to earthquake sources is finalized next year, the probabilistic seismic hazard analysis will be developed. And I'll take a broader view of that schedule here in a moment.

For the ESF, the current seismic design basis assumes conservative values for those portions of the ESF that have been designed and constructed, specifically the portal and the pad.

To support underground design of ESF, the technical assessment review that Rich referred to, and some results of which were presented earlier, will be completed shortly. This review of the current ESF design basis has resulted in the preparation of an initial simplified probabilistic seismic hazard analysis that use preliminary paleoseismic and historical earthquake data.

The results emphasize the importance of the contribution of an areal source referred to as the background or random earthquake to the overall hazard. As a result, the DOE and the USGS have prioritized development of the methodology for the characterization of this source.

For the repository, the current seismic design basis is .4g. That's peak acceleration at the surface. This is as described in the SCP. This design basis will be updated as the probabilistic seismic hazard analysis is developed. An outline of the program developed to meet this objective is presented in the next view graph.

The sequence of topical reports described by

Richard is shown on the left, and the technical activities described earlier by John Whitney is shown on the right.

In addition to design, advanced conceptual design, performance assessment--there on the lower right--is the user of the results of the tectonics program for post-closure evaluation of the total system performance.

When referring to total system performance, we're not talking here about the engineered system, whose lifetime is the pre-closure or perhaps 300 to 1,000 years. Total system performance refers to the performance of the natural barriers or the natural system itself for periods of 10 to 100,000 years.

The consequences of interest to performance assessment are releases that could exceed release rates at the system boundaries.

In order to perform performance assessments, we need to have characterized the natural system itself, much as to do performance assessments of engineered systems, the characteristics of the engineered systems need to be understood.

The first two total system performance assessments in 1991 and 1993 considered tectonic effects in relatively little detail. They focused on phenomenal case, the description of the natural system as we understand it now.

The result of both of those assessments emphatically indicated, to site characterization, that the data of greatest significance has to do with percolation in the unsaturated zone. DOE has prioritized its site characterization efforts accordingly.

The next total system performance assessment scheduled for '95 will consider tectonic effects in greater detail. In order to support this performance assessment, alternate tectonic models will be defined and described. Probabilities of initiating events or disturbing events will be provided, and models of post-closure effects on the waste package, water table elevation, fracture permeability and porosity and rock geochemistry processes will be provided.

I thought I'd just talk for a minute about site suitability. In January of 1992, the early evaluation of site suitability, which was the first of a series of evaluations by outside experts, provided the following results for tectonic hazards. These are based on 10 CFR 960.

The disqualifying condition is not present for tectonic hazards. This is a higher level finding, meaning that the conclusion will not change with the collection of more data.

The qualifying condition is likely to be met, but their conclusion was that was a lower level finding, meaning there is not yet sufficient confidence. However, DOE expects that the program of data collection and assessment that has been described, at least has received an overview description by John and Richard, will result in a higher level finding.

The panel also concluded that tectonic hazards can be accommodated by reasonably available technology. These conclusions seem to be strengthened by DOE's recent decision that the baseline waste package engineered barrier system will be based on horizontal in-drift emplacement.

This conclusion particularly applies to the potential for waste package disruption by fault displacement. The SCP waste package concept was a small, thin walled canister emplaced in a bore hole in the floor of a drift, leaving a 7 centimeter air gap between the bore hole wall and the canister. This concept does allow for the possibility that the air gap could be compromised or the canister disrupted by even small displacements.

However, emplacement of a six foot diameter multi-purpose canister in a 14 foot diameter drift seems to mitigate potential concerns about fault displacement. That is, the only consequences seem to be minor tilting and possibly some spalling.

The Board's question on design criteria seems to be a request for us to address what are appropriate hazard levels for pre-closure design. These will be established in the next topical report. The basis will be the ASCE guidelines, which Carl Stepp will describe tomorrow, DOE guidelines for other nuclear facilities, and previous nuclear power plant experience. Kevin will describe both the DOE guidelines and nuclear power plant experience in the next presentation.

I have identified several critical activities, the first of which is the assessment of the background earthquake source. As I mentioned earlier, this potential seismic source seems to be a major contributor to the hazard, particularly at lower annual probabilities of exceedence.

Complete paleoseismology studies. Rates of occurrence derived from data on average slip rate and average or event specific return periods from trenches are key ingredients of the probabilistic seismic hazard analysis. An important element of this data set is the development of a sound geochronological data base to date the offset deposits and individual events.

And, finally, subsurface information, specifically the down dip geometry of site faults. The close spacing of Quaternary faults at Yucca Mountain, spacings of a couple of kilometers, and possible evidence from surface displacement on several faults at the same time suggest that the subsurface geometry of the faults may be simpler than is indicated at the surface. That is, there may be fewer fault planes at seismogenic depths than indicated by the surface geology.

Subsurface data would then support resolution of tectonic models and could contribute important refinements to the rate of occurrence of past and future surface displacement events.

The SCP and study plans lay out DOE's strategy for data collection. The PI's then will determine if the data is adequate for assessment or characterization of seismic sources. The scope of the data collection effort and the variety of data collection techniques has been guided by past experience in sighting studies for critical facilities.

The main difference at Yucca Mountain is that the site area, that is, the repository plus the controlled area, are much larger than for typical critical facilities, necessitating more extensive site studies, although regional studies are comparable.

In addition, NUREG 1451 provides NRC guidance on the appropriate scope of data collection to support a license application. And that is incorporated in DOE's program.

In addition, DOE has a planning system in which the principal investigators, participant managers and DOE staff;

review annual plans for site characterization activities, agree on the scope of work, and chart the schedule on budget. Thus, DOE is directly involved in the planning and execution of all work activities after initial assessments determine if data is sufficient to support PSHA and design basis.

DOE expects that the preparers of the probabilistic seismic hazard assessment and ultimately the design bases for ground motion and fault displacement will specify the required data input during the development of the assessments, and as appropriate, feed that information back to PI's for additional data collection if needed.

Subsequently, uncertainty and sensitivity analyses will be performed that may indicate the need for more data or may indicate that the data is adequate. As assessments are finalized, we'll determine if additional data could reduce the uncertainty associated with the sensitive parameters that are significant contributors to the hazard, and conduct independent technical and peer reviews. Review is now and has been a part of the DOE process, and additional data needs could be defined through this process.

DOE will rely initially on their internal experts who have collected the data. We feel they are the best to analyze, interpret, evaluate uncertainty, and determine the completeness of the data set.

The DOE experts are qualified investigators from a variety of organizations with experience in research and applied seismic hazards evaluations. They include the USGS, Denver, USGS, Menlo, U. S. Bureau of Reclamation, Geomatrix, Sandea.

Independent technical review; the technical reviews by experts who were not involved in the work is a normal part of all Yucca Mountain project technical activities.

Peer reviews will likely be a mechanism for ensuring diversity of interpretation or completeness.

DOE believes that the probabilistic seismic hazard assessment is the appropriate methodology for pre and post-closure assessments, because it's consistent with PA and design needs.

DOE recognizes that past reactor licensing practice has been to use traditional deterministic hazard assessments, but probabilistic seismic hazard assessment has evolved to state of the practice, as Kevin will discuss here in a few moments.

And let me summarize here by briefly describing DOE's licensing strategy, which is to establish reasonable assurance, reasonable assurance that the MGDS, the Mined Geologic Disposal System, can isolate waste for the performance period. And I would note in 10 CFR 60, the section Technical Criteria, complete assurance--and I quote here from that section--"Complete assurance is not expected. Reasonable assurance is the intent." Although I'm reminded by people who have attended reactor licensing hearings that lawyers don't find that distinction as clear as I might.

Develop an adequate data base and a sound methodology. Those elements have been discussed. Document the methodology and submit it to the NRC. Methodology is being documented in the topical reports and the implementation details supporting the methodology will be in the study plans which will be provided to the NRC and other interested parties.

Develop design bases for ground motion and fault displacement. This is the ultimate goal of the pre-closure tectonics program.

And, finally, an important ingredient that has been referred to by the Board and others is to establish or describe the consequences of tectonic events. Ultimately, this will be done through total system performance assessment, as I've mentioned. But the NRC in public meetings and in a recent draft NUREG has urged DOE to come forward to the NRC and describe special design measures that may be required to accommodate tectonic hazards. At this time, DOE has identified no special design measures that seem to be required.

For the surface facilities, as John described, we have completed the Midway Valley studies. This was an early focus of the tectonics program. The characterization parameters laid out in the study plan for identifying a suitable location for the repository surface facilities have been met through the geologic studies described in that study plan. A final report on that work will be available later this year.

And with regard to the underground, as I discussed earlier, the current waste emplacement system seems insensitive to earthquake hazards, which leads me to my final remarks. Further evaluation of the sensitivity of the surface and underground facilities should guide DOE in determining the level of sophistication, detail or alternatively the simplicity of the probabilistic seismic hazard analysis that we are now contemplating. And I read sophistication as equivalent to resources expended.

We need to find a way to look at the costs versus the benefits of detailed probabilistic seismic hazard analysis. This will be a part of the topic of topical report Number 2, and will evolve with the further evolution of the multi-purpose canister design concept that is currently being assessed within DOE, and also should evolve from total system performance assessment Number 3.

DR. ALLEN: Thank you, Tim. Any quick comments from Board of consultants?

DR. NORTH: Again, I find this extremely dissatisfying because of what you haven't told us. I feel what you've given us is a very ill-defined statement of what DOE is doing, a set of DOE conclusions, and the basis for it is nowhere evident. I haven't seen a systems viewpoint. I haven't seen any evidence that the PI's guiding the data collection and the people that are going to do the modelling leading to the probabilistic seismic hazards analysis even talk to each other, let alone whether the PI's are doing the right kind of decision making.

I don't feel you've presented any evidence that your process is working. You have a series of documents that you're preparing, but how far you've come, what you've learned, the insights, the basis for considering, for example, whether the 0.4g that apparently came from the SCP years ago makes sense. You haven't given me any reason to believe that there aren't significant problems in the move to in-drift emplacement where we are going to have rather large and heavy things in a tunnel, supposing with an earthquake, those things could move and run into each other.

Maybe it's a simple concept and maybe that might only happen during a short period of years before closure, but I don't think you can ask on behalf of the Department of Energy trust us, we know what we're doing, it's all great. Because, frankly, the people don't trust you, and if you don't do a better job of presenting the basis for a program that makes sense that is indeed trustworthy, you can expect, I think, very serious criticism from some of us on this Board.

DR. ALLEN: Other comments or questions? Leon?

DR. REITER: Warner, just in all fairness to DOE, Clarence and I attended a tectonic workshop where the people who were trying to do the seismic hazard analysis and the investigators are beginning to talk to each other. So I think at least that part of the process they've begun. I'm not going to say anything about the--the consequence is something else. But at least part of the process they have made, at least from our perspective, have made progress.

But, Tim, I want to ask you a question. About two years ago, we discussed seismic vulnerabilities. The DOE proclaimed that its philosophy was fault avoidance, in placing waste, they're going to avoid faults so as to avoid the possibilities of fault offset.

What you've proposed now here is you say that we have a system of drift emplacement and where, I think you said, earthquake hazard no longer becomes a consideration. Does that mean now that you're going to change that? And particularly I'm trying to assess the impact of what might be with discovery of these new faults, such as the Sun Dance fault, which now seem more prevalent.

MR. SULLIVAN: I don't recall that DOE's position was that seismic hazards would be mitigated by fault avoidance. I'm sure Keith is going to address that in a little while. I do understand that the DOE was intending to avoid any faults of potential engineering significance, and identified the Ghost Dance fault as one of those.

DR. REITER: But I guess based on your--and now you're saying you have how many feet of freeboard? Does that mean that the Ghost Dance fault and the Sun Dance fault, even if they prove to be active, let's assume that they don't generate any of the meters of offset needed, would that mean that you're going to abandon what was then told to us, the philosophy of fault avoidance, or is that still part of your philosophy?

MR. SULLIVAN: Well, those faults that pose, you know, a potential risk to system performance will need to be addressed. I don't think I have a good answer to your question, Leon. I don't recall the DOE making that blanket commitment.

DR. ALLEN: Okay, we've got to move on, but I should give you a chance to respond to Warner if you wish to.

MR. SULLIVAN: Maybe we can cover that in the roundtable.

DR. ALLEN: In the roundtable, okay. Very good, thank you, Tim, and thank you, Richard.

We'll take a brief respite here from Yucca Mountain, I think, and Kevin Coppersmith from Geomatrix Consultants is going to talk about PSHA case histories, which I assume are mainly elsewhere.

DR. COPPERSMITH I've been asked to talk about some case histories of the use of probabilistic seismic hazard analysis particularly in facilities other than high level nuclear waste repositories, and I will do that. In the view graph package, I'm sure are more case histories than I'll have an opportunity to go through. It doesn't matter. The point that I'm trying to make is what have we learned in the course of these analyses. And those lessons learned, of course, can have some applicability to the methodology and application at Yucca Mountain.

I'm going to look at some cases where we're talking about critical facilities like nuclear power plants, but also critical facilities like major lifelines, and in this case some of the San Francisco Bay area bridges, a dam or two, and some DOE facilities.

One concept here, of course, is that many of these studies are graded to the type of facility that you're looking at, and often one of my underlying themes will be that there is a consideration of risk in this process in identifying criteria for hazard levels, and I'll try to highlight that as we go along.

For those that are from the Bay Area will know that in the last three or four days, we've had the biggest outburst of pollen I think ever recorded. Out where I live, pollen counts are higher than they've been in 20 years, about 20 miles east of here, and I'm not a pollenologist, but I'm highly sensitive to pollens.

This is the purpose of my talk. I want to look at some of the--

DR. ALLEN: Maximum credible pollen outburst?

DR. COPPERSMITH Or minimum incredible. That's right. It's a thousand year event.

I want to talk as we go through the regulatory context for these studies, which does vary. In some cases, these are in contexts that are not under heavily driven regulatory frames, but financially driven frameworks.

The use of the study for decision-making, why was it done in the first place, how is it going to be used. How does earth sciences make its way into the assessment and how will uncertainties in the earth sciences data base actually make it in. The subject of the use of expert judgments, how was that used, and you'll see a range in the case histories in how that can be done. And comparisons of probabilistic and deterministic. Many of the agencies that carry out these studies rely on one or the other or both of these in making decisions about things like design or design retrofit.

Some of the applications of probabilistic seismic hazard analysis, many of these appear in the fine print in Allin Cornell's view graphs. And obviously I won't have a chance to go through those, but the theme here is that this tool of probabilistic seismic hazard analysis is now prevalent throughout many, many types of facilities, particularly for design purposes and for design evaluations.

Obviously we see the studies that have been related to nuclear power plants, both in the Eastern United States, large regional studies like the EPRI and Livermore studies, but site specific studies for the Western U.S. power plants have also been done for all of the plants, including Diablo Canyon and so on. I'll show some examples.

At the present time, a program called IPEEE, Individual Plant Examination for External Events, is looking at the beyond design basis evaluation of seismic margins for all the plants. Probabilistic hazard analyses have been used for all of those as well.

In the DOE, nuclear and non-nuclear facilities, both new facility design and review of existing designs are using probabilistic approaches. I'll show an example of that for the Hanford site.

Major bridges and highway structures, a few that I'm familiar with, Illinois, Oregon and Arizona, others are doing the same thing. In particular, site specific design review, Caltrans, rather than doing the entire state, for example, is looking at particular regions like the San Francisco Bay area or the Los Angeles Basin and so on, as a basis for making decisions about design, retrofit, and the costs thereof for their facilities.

In the dams, dams are one of the last holdout of deterministic approaches. And basically for design review, though, it's very common to see evaluations of probability of exceeding a particular ground motion level. I'll show an example of that. And often it is used as a check on deterministicly derived values.

Building codes; obviously the federal maps, seismic hazard maps, are the core of development of building codes and they're in the process of being revised almost continuously. And for many commercial facilities, particularly large buildings and so on, probabilistic analyses will be done where there's significant financial investment involved, particularly in places like San Francisco and Los Angeles. A high rise in downtown San Francisco will go ahead and look beyond building code requirements, develop probabilistic analyses, look at elastic and beyond elastic type design scenarios.

Look at Diablo Canyon, this is--again, these will all just be very quick thumbnail sketches of some of these case histories. This is the--the saga of Diablo Canyon goes on for a couple of decades. I'm going to talk only about what's called the long term seismic program that began in about 1984 and ended in about 1989.

During the course of that, we had some interesting earthquakes in the State of California, and we've had some since. The purpose of that study was to satisfy a condition on the operating license that said that you will re-examine the tectonics of the regions, the implications to earthquake magnitudes, and evaluate the seismic margins of the plan. And satisfaction of that licensing condition was the basis for conducting the probabilistic hazard analysis, a full scale seismic probabilistic risk assessment.

A scoping study was done early on in the program of probabilistic hazard to identify significant geology seismology geophysics issues right at the beginning on the basis of present knowledge. What do we know now, do an analysis, look at what is most sensitive and carry out a program of design to address those significant issues.

An extensive program of data collection occurred over a period of about three years with the NRC staff and its consultants, there was a lot of interaction, dozens and dozens of meetings, field trips throughout the process.

The assessments that actually went into the probabilistic analysis represent consensus estimates of uncertainty by a large project team. The team that actually conducted this work over a period of a few years consisted of about 15 individual geologists, seismologists and geophysicists, who together developed expressions of uncertainty and documented that expression.

Extensive NRC staff review and NRS consultants, including the USGS, people from the University of Nevada at Reno and so on. I think this is very important in this case, a highly regulatory environment with many, many interactions with NRC throughout. Both probabilistic and deterministic approaches were used in evaluating the seismic margins of the facility. So both were conducted.

In a nutshell, the geologic environment is one of coastal central California. This is the location of the plant site in this location along just to the north of San Luis Obispo. Obviously, a consideration is the off shore Hosgri fault, a large fault. An early considerations, early basis for the license condition had to do with the sense of slip on this fault. Many of the studies that were done were solely done to try to establish whether or not this was a stretch lift fault or a reverse fault, and in turn what are the implications of those differences senses of slope to the slight ground motions.

So obviously in the course of this over 1,000 kilometers of seismic reflection surveys were carried out. Unsure studies, particularly in the area of rain, Quaternary terraces were carried out to look for faults in the near plant area. Studies were done up in the San Simeon over 85 kilometers from the site basically to get an idea of the style of fault and rate of slope and what is the off shore equivalent of the Hosgri fault zone, a long process of data collection.

The approach here, of course, was to try to incorporate this information and its uncertainties into the hazard analysis. And just one simple expression of that is the logic tree that's shown taken from the long term seismic program final report for the Hosgri fault. And, of course, it's impossible to get into all of the elements here. I'll just show that some of the key components of the problem, like the sense of slope, studies had gone on that were aimed primarily at making that assessment and designed to give us the highest level of confidence in deciding on these different components of slope.

And, of course, when we come down to the end, there is still uncertainty and there's still the need to document the basis for the assessments on the tree. The implications, some key elements that have to do with things like the fault slip rate which drives the recurrence assessment were also the focus of the field program, focus of a lot of the discussions in the review process.

I thought an interesting element, this is a probabilistic analysis, but also had a deterministic parallel study carried out. This is an example of one of the components, one of the key parameters in the assessment. This is the assessment of the maximum earthquake, maximum magnitude for the Hosgri fault. As we talked about before, there is uncertainties in the particular parameters that go into making that assessment, and Allin talked about how these days we need to actually acknowledge and explicitly incorporate that uncertainty. And in doing so, this is the type of distribution that results in all the uncertainties in things like fault segmentation, fault rupture length, the amount of displacement and so on. And this expression of uncertainty of about a full magnitude unit in the maximum magnitude is not at all uncommon. In fact, it's something that we now see quite a bit in these assessments.

Shown just to point to a particular point is the mean value of magnitude seven. The deterministic maximum earthquake, the maximum credible earthquake in this case was a magnitude 7.2, which was selected for the deterministic analysis. It's not a worst case. It's, in fact, what we'll see in the deterministic never represents the worst of the worst. In fact, it's usually a negotiated value, and I think it's the problem of that negotiation that's taken so long in past licensing.

Let me jump a few hundred kilometers to the north up in the coastal Washington area, a nuclear power plant that is in a mothball condition at the present time. A good bit of work was done in the licensing of the seismic components of the plant, the WNP-3 SATSOP plant. The probabilistic analysis that was done here was to answer a question that had been asked. The deterministic SSE was developed as part of the final safety analysis report, and a staff question from the NRC said what's the probability of exceeding that SSE. A very simple question. I think Jerry Kenya asked it, as a matter of fact. The process of answering that question of course entails conducting a probabilistic seismic hazard analysis.

One of the key issues here at the time the original studies were done for the SATSOP plant were changes in the perception of the hazard related to the Cascadia subduction zone. At the time of the original licensing activities, the PSAR and so on, the Cascadia subduction zone was believed to be aseismic, to not be a source of earthquakes. So we had to incorporate this change in the perception of the hazard. Obviously, no large earthquakes had occurred. We had no new empirical data. We had many other geologic indications that in fact there may be a seismogenic potential.

This is a different approach. The approach here was actually the expert elicitation, the individual elicitation of 14 experts and the aggregation of their assessments and, of course, extensive documentation of that.

An advantage--one thing that we tried to do was to look at what's called component level aggregation. One of the key technical issues is what is the geometry of the subduction zone, and we can aggregate at that level and look at it. What's the maximum magnitude on the subduction zone? We could aggregate that component of the model and look at it as well.

So we knew the staff was interested in looking at these key components, and so we set up the models so that you could have that key component level aggregation and not just have aggregation across the bottom line, the final hazard analysis.

NRC staff was using this to assess, evaluate the conservatism of the SSE. And, of course, this study was done in about the late 1986 to 1987 time frame. And, of course, the use of probabilistic analyses for this type of insight is something that went on in many other cases, many other sites as well.

Diversity of expert judgment is a key issue and it was mostly because of this new tectonic interpretation. And I'll simply show--you have a map in your package that shows where things are. Let me just jump to some of the seismic hazard curves. These express the annual frequency of exceeding a particular ground motion or the probability of exceeding a particular ground motion, if you will.

And what's shown outlined in the solid lines, orange on the slide, is the overall distribution across all of the 14 experts taken into account all of their uncertainties, and that's shown as the 15th, the median, or 50th percent and 85 percentile hazard curves across all of the experts. And also shown are the median estimates of seismic hazard for each of the experts themselves. And you can see that some of the experts, their median estimates were significantly lower than the 15th distribution, and some that are significantly higher.

The point of this and the point of the next view graph is to show that in this particular case, there was quite a bit of diversity in interpretation. When we look at the aggregate hazard curve, which is the lower right, it expresses, at least in the 15th and 85 percentiles, across all experts, we have a good measure of the total uncertainty.

If we went to an individual expert and asked him, in some cases we would get significantly less uncertainty, in some cases comparable to or maybe even more than the aggregate. I think the point that we saw here is primarily because of very, very significant differences in the tectonic interpretation of what was going on. In the earthquake potential of the subduction zone we saw a very strong component of the total uncertainty which related to the expert to expert diversity of interpretation. And by capturing it in the total distribution, we are able to have a better expression of the total uncertainty.

The San Francisco Bay area bridges, I'll show just a couple of examples, and these are part of studies that are ongoing, will be used by--the front end, the seismic hazard analysis is used as a basis for design review of the Caltrans bridges. Ultimately this will go into a process of design retrofit if required, and decisions regarding capital outlay over the next several years within the state will be based on this type of analysis.

As we know, some of these bridges that include the Bay Bridge, for example, in San Francisco have tremendous consequences of disruption. We saw--created a loss of the Bay Bridge for a month, had severe consequences that we continue to feel in the Bay area. So they put a high premium on trying to keep these open.

Site-specific assessments, these are not regional, they're down to even differences in particular abutments of the bridges. There are differences in hazard at the west end versus the east end of the Bay Bridge, for example, as you move away from one fault towards the other.

Incorporation of fault-specific paleoseismic studies in the Bay area; some of this type of work occasionally gets funded, and we can of course incorporate that.

The way things were assessed here is by a project team essentially of four or five geologists and seismologists, and heavy consulting board review, interactions, multiple meetings with the consulting board. We're basically testing whether or not you have the hypotheses that are out in the scientific community into your assessment, if possible, forcing the grounds on uncertainty as wide as you can get through this multiple interaction with the consulting board.

Basically, the way Caltrans assesses its design ground motions is looking at probabilistic and looking at deterministic and making judgments about which are most appropriate to be used. And I'll show a couple of examples of that.

You have your map in the package that shows where these bridge sites are, so I'll save some time. One example of--for those that aren't familiar with looking at sensitivity probabilistic hazard, obviously the results of hazard are dependent on a bunch of things. One of the things is what structural frequency, what structural period you're looking at, whether or not it's very high frequency ground motion, peak ground acceleration out to structural periods as long as three seconds. Of course, for bridges we're often interested in motions down in these longer periods. And we can also see that the contribution to hazard that different faults make varies as a function of that as well.

We'll also see that it functions as does the level of probability. The hundred year ground motion, the 10 minus 2 ground motion can have a contribution from some faults that don't contribute at 10 to the minus 4 and so on.

But what I wanted to show here is that when we look at peak ground acceleration, high frequency ground motion contributions, we have contributions coming in this case from a zone of seismicity that is presumably identifying or outlining some local faults in the Berkeley Hills that are very close to this particular site, the Carquinez Bridge site, and they dominate the hazard, they're closest to the total hazard in this plot. Then we're into longer period ground motions, a higher dominance from some of the bigger structures like the Hayward Rogers Creek that can contribute more in terms of somewhat less frequent, but larger magnitude earthquakes.

Also I've outlined on here a little fault, the Franklin Fault, which probably is not active, but may be, and has some indications of possible activity. If it is active, the evidence for slip rate puts it way down. In other words, it has a very low slip rate, very long recurrence intervals, and the probabilistic hazard does not contribute much. Keep that in mind when we look at deterministic hazard, this is the controlling source, this small fault that happens to be near the site becomes a dominant contributor despite the fact that the frequency of occurrence is very low.

I'll show this quickly just so when you look at it in your package, it shows the contribution that different magnitudes make to the total probabilistic hazard, and we see it as a function of the return period, 100 years versus 1000, as well as structural frequency. We see different contributions from just different magnitudes. This was one of the insights that can be used for designers, for example, who are trying to see is my hazard driven by a magnitude of five earthquakes or driven by a magnitude of seven and a half. Of course, there are different design implications for those types of events.

But what I want to show in these response spectra are a couple of things. This is a comparison, a direct comparison of probabilistic and deterministic results for a particular site, in this case, the Carquinez Bridge. What is shown are equal hazard spectra, and I won't get into what those are other than they are ground response spectra that are equal in their probability of exceedence throughout the structural period range.

So if we look at this, this is the 100 year return period or annual frequency of exceedence ground motion response spectra. And as we move up into the 300 year and the 500 and 1000 and 2000 year, we see the way ground motions go up across all the structural periods.

We can compare that with deterministicly defined ground motion values. Deterministic means they have a maximum credible earthquake that is assumed to occur at the closest approach to the site. And those assessments are done both for median of the ground motion attenuation wall for the 84th percentile of the ground motion attenuation wall.

And we, as we know in nuclear power plants, watch to see the 84th percentile was often used as the deterministic ground motion value. When we look at that, the solid line here, this is the median deterministic ground motion, and we can look at it across the different structural periods. We see out in the short period or high frequency end, this represents somewhere between a 500 and about 1000 year ground motion. As we move into the longer period motions, this actually moves into a lower probability of exceedence. We move into levels out here that actually are of most importance to these bridges. 84th percentile shown by the yellow dashed line above sits out here somewhat above the 2000 year return period.

And the value actually selected for design evaluation sits out at about the 2000 year probability of exceedence level pretty much throughout the range of structural period. That is the value that was selected for design evaluation, and it obviously is up in a level of conservatism, but Caltrans feels comfortable with. A couple thousand years out in the structural period are the most important to the bridge.

So the decision was made that this is reasonably conservative and falls between the mean and the 84th percentile, is richer out in the longer periods to account for the fact that it is not as deficient as this local nearby fault would say, and levels of conservatism were added over here to make it richer and higher levels of conservatism.

DR. ALLEN: Kevin, you have about two minutes, according to my watch.

DR. COPPERSMITH Let me go to the lessons learned. I'll be happy to answer the questions of any of those who are going through the rest of the case histories. But I definitely want to touch on these.

I think it's been acknowledged that deterministic approaches do not take into account the likelihood of occurrence and the actual rate information, and typically do not include uncertainties as well. These usually tend to go hand in hand with the probabilistic approach. Therefore, design values are often contentious. MCEs can be argued about long and hard. If we perhaps put it more in the context of an uncertainty distribution, we could get past some of the contentions.

There's been increasing use in probabilistic approaches I think as people--all of us have understood them better, what drives them, what the important components are. We also now realize the importance of things like slip rates and paleoseismic information, and they themselves become the areas of research. And as scientists, earth scientists have developed that type of information, they find a way into the probabilistic approach. These recurrence related parameters do not find a way into deterministic studies.

Comparisons are often made between deterministic and probabilistic results. And when you do the comparisons, and there are others in your package, we see, number one, the deterministic is usually not a worst case. In fact, we see the probability of exceeding the deterministic case varies quite a bit. And I think this gets back to Bob Budnitz's comments that we see the SSE probability of exceedence across the Eastern United States nuclear power plants varying by two or three orders of magnitude.

Likely also we'll see that in the West and highly active areas, the probability of exceedence of the deterministic case might be--the deterministic might define a 200 or 300 year event in low activity environments that might represent a 10,000 year event.

So there is, in the comparison, we do see quite a bit of difference in that probability of exceedence of what is considered the deterministic value.

The incorporation of uncertainty now is something that we're accustomed to. To see uncertainty distributions about parameters is something that we can feel comfortable with. I deal a lot with earth scientists who go nowhere near probabilistic or statistical techniques. They're beginning to see and feel comfortable with uncertainty distributions.

The advantage and one of the burdens of probabilistic approaches is extensive documentation. You have to, in characterizing 20 or 30 parameters for a fault and various models, they need to be documented and often the documentation itself leads to a higher level of assurance by review bodies.

The ways of approaching the issue of capturing diversity of interpretations I think is still somewhat of an open issue. There are accepted approaches, and I had a chance to talk about a few, and there were others that range from a formal elicitation of expert judgment going through those procedures, to one of the development of a consensus type of assessment by a large panel, for example, of experts, to one of the development of a particular probabilistic hazard assessment and a lot of peer review and regulatory review. All of those at the present time have their advantages and are in operation in one form or another. I'll stop there. Thanks.

DR. ALLEN: Thank you, Kevin. Any quick comments or questions?

DR. NORTH: Warner North, quick question. Kevin, could you give us an idea of the time and resources needed to carry out the probabilistic analysis in the examples you've described? And contrast that to what a deterministic approach might have taken in time and resources.

DR. COPPERSMITH I would say in general, the probabilistic analysis will require 30 to 50 per cent more time and resources than the deterministic analysis, primarily because of the need to gather rate related information and the need to document and incorporate uncertainties.

Now, that would be just doing the analysis. In terms of regulatory review and contention, it's impossible to guess.

DR. NORTH: Could you just take the example of the Carquinez Bridge over the nuclear power plant in Washington and give us a sense how long did it take, how many people were involved and roughly how many person years did their work take?

DR. COPPERSMITH The Caltrans, we did work for seven Bay Area bridges developing these analyses over the course of about nine to twelve months, multiple meetings with the consultants and consulting board. I would guess when you add in our effort plus the consulting board plus Caltrans' effort, probably, oh, close to ten man years worth of work to carry this out.

The SATSOP case we actually had--it occurred over about a year and a half, 14 experts, workshops, formal elicitations, probably 25 man years, I would say.

DR. ALLEN: But certainly the circumstances are much different from, depending on the level of seismicity and so forth, a site right next to the San Andreas fault in Southern California, I could give you a deterministic assessment in two minutes. We could easily spend $50 million on a probabilistic assessment. So a lot depends on the--

DR. COPPERSMITH I think a lot has to do with the regulatory environment and the levels of assurance that are required. One example that's in the packet that I didn't get to is a dam up in the Portland area that's undergoing FIRC review, and this is one of hundreds of dams that are undergoing FIRC review. It's not under a particularly high regulatory pressure. The study was done for about $100,000.00 over about a three month period, provided levels of assurance to both the owner, the operator, as well as FIRC that there was sufficient seismic margin, and that was it. So I think there are scales. A lot of it is determined by the regulatory framework.

DR. ALLEN: Other questions, comments?

DR. CORNELL: Would you comment, Kevin, on the distinction between SSE determination to the deterministic basis in the Eastern United States and the West and the relationship to maximum possible earthquakes?

DR. ALLEN: Maximum dreadable earthquake.

DR. COPPERSMITH That's a leading question, because we just finished up a large study for EPRI on assessing methods for assessing maximum earthquakes within the Eastern United States. And it is, in the Western U. S., we typically use estimates of fault dimensions, rupture lengths and other constraints of what we feel are reasonable maximum scenarios for rupture. In the Eastern United States, things are driven much more by the largest earthquakes that have been observed within your source zone or ones that are felt to be reasonably analogous tectonically to your particular source zone.

From the standpoint of nuclear power plant licenses, it normally follows the idea of the largest observed plus an increment, an intensity unit is added to what's been observed. So the net effect, I think, is in the Eastern U. S., to have earthquakes that are probably more in the lines of a design type of events, and in the Western U.

S., there are more in the lines of maximum possible events. They average in the East about magnitude five and three-quarters. They average in the West on the order of six and a half to seven.

DR. ALLEN: I think we'd better be moving on. Thank you, Kevin. And the final presentation before the break is by Keith McConnell of the Nuclear Regulatory Commission, who can say anything he wants to. His title is Comments by the Nuclear Regulatory Commission.

DR. MC CONNELL: Thank you. The only guidance that we got from Leon when we called to discuss our participation in this meeting was not to be boring. And for a regulator, that's difficult, but we'll try. And we particularly want to leave more time for questions for him and Richard.

What my co-author and I would like to do, my co-author being Bakr Ibrahim, is to give you some of the staff feelings or beliefs on several of the issues that have been raised today and that Leon specifically asked us to address in our--in his outline.

What we'll try to do is give you the staff position as it exists on deterministic and probabilistic assessments. We'll run through some, at a general level, some acceptance criteria for data analysis, or when enough is enough from the staff perspective, and then we'll try to outline some of the investigations that we believe are critical for fault displacement seismic hazard analysis.

The policy guidance that we are using was given by Bob Bernero in a speech to the ASCE actually in 1992 and published in 1993, that both deterministic and probabilistic techniques will play a role in the analysis of fault displacement and seismic hazards.

To take that or a corollary to that, is that we would expect in our review of the license application, that both deterministic and probabilistic approaches would be presented, and the basis for that is that we would expect that someone associated with the program will do a probabilistic assessment and someone will do a deterministic assessment. And it's good, or perhaps in DOE's best interest to address both approaches early on instead of waiting till we get into licensing and start worrying whether somebody comes up with a maximum credible earthquake or maximum gullible earthquake.

Also in fiscal year 1995, the staff intends to prepare a staff technical position on the criteria and analysis needed for the development of design bases for fault displacement and seismic hazards. This is, in part, contingent on the DOE topical report that you've heard about. It's also contingent on the discussions that are ongoing regarding the revision to Appendix A to Part 100, which most of you are familiar with. All of this will play a role in deciding what is done with this technical position.

What we've tried to do with this slide, and it's somewhat redundant at this stage after the other presentations today, is go through some of what we consider to be the positive attributes to both deterministic and probabilistic approaches.

Of course positive is probably in the eye of the beholder, but I won't go through them all, but obviously from previous presentations, we know that the deterministic approach has the regulatory and licensing precedent in the licensing of nuclear facilities. It's relatively straightforward and it's transparent.

Probabilistic approaches, obviously the requirements, 10 CFR 60.112, requires, along with the EPA standard from which it's derived, a probabilistic analysis. And from what Bob Budnitz said earlier today, any future standard would also probably require a probabilistic analysis.

Also there is proper--or when it's properly implemented, a probabilistic approach explicitly includes uncertainty. Frequency and magnitude of earthquakes and displacement events are considered, and multiple estimates--

it uses multiple estimates and considers the range of possibilities.

What we would suggest by this is that these positive attributes provide support for doing both analyses or providing both analyses in the license application, or even prior to the license application so we can resolve any concerns before we get into the licensing process should it get that far.

Now, moving on to acceptance criteria or determining when enough is enough for fault displacement and seismic hazards, we've put down what we consider five of the minimum requirements for determining when enough is enough.

Specifically, the collection of data used in support of the analyses is sufficient to support the assumptions made. I think in our reviews of some DOE documents, we've developed concerns that some of the assumptions made and some of the most likely estimates made have not been thoroughly supported and, therefore, are challengeable. And so we would ask that the data collection focus on supporting the assumptions made.

With respect to fault displacement and seismic hazards, we also would expect that the positions provided in NUREG-1451, which is the staff technical position on investigations of fault displacement and seismic hazards, as well as the positions in the draft STP on "Consideration of Fault Displacement in Repository Design," which we, in shorthand, call the avoidance STP, have been satisfactorily addressed.

I would comment that I don't know that what Tim presented would satisfactorily address those positions in the second STP or not, but it didn't appear to. This STP basically says that you should avoid faults if possible. If you can't avoid them, then you can design for them, but you'd better come to the NRC and show us how you're going to accommodate design and performance issues.

MR. SULLIVAN: Didn't it say Type I faults?

DR. MC CONNELL: Type I faults.

MR. SULLIVAN: Correct.

DR. MC CONNELL: Another minimum requirement would be that expert judgment has not been used as a substitute for field or experimental data or other more technically rigorous information that is reasonably available or obtainable. This is not to say that we don't believe that expert judgment isn't going to play a significant role in the licensing of a geologic repository. What we are saying is--and I'm sure it's not the case with anybody here--that sometimes experts do make mistakes, and that the data analysis should be thoroughly supported, that the analysis should be thoroughly supported and should not substitute expert judgment for data collection if it's reasonably obtainable.

Fourth minimum requirement for when enough is enough would be that the analyses are transparent, and we've heard all this earlier today, sensitivity analyses have been performed, alternative models, both statistical and conceptual, have been identified and evaluated, and the results of the analyses of individual alternative models are explicitly treated. And, again, that addresses the transparency issue.

Fifth, that the analyses clearly reflect the uncertainty in the understanding of tectonic processes.

Ultimately, the final determination of when enough is enough will be an assessment of repository performance and full consideration of uncertainty. And I'd point out that the staff in its own IPA Phase 2 work has considered both seismicity or vibratory ground motion effect and vulcanism in a relatively rudimentary form, and we are using that input to develop our license application review plan and the site specific acceptance criteria in the license application review plan.

Most of these--I say most of these requirements are somewhat motherhood statements. The staff believes in motherhood. But the other concept that we're trying to get across with these criteria is that we would expect both the bottoms up and a top down approach to determining when enough is enough. In other words, you do use your IPA efforts to determine significance of various processes and events, but you also use your data collection efforts and your scientific analysis to tell you when enough is enough. So it's both bottom up and top down.

And, finally, we've tried to describe some of the key, or what are considered critical investigations that need to be done in addition to those that are already ongoing at the site. Specifically we believe that high resolution geophysical investigations to identify buried structures and the down-dip expression of faults at depth is necessary.

An appropriate model, tectonic model needs to be developed for earthquake locations that also addresses the apparent contradiction or discrepancy between fault plane solutions and the nature of displacement of faults expressed at the surface.

DOE should provide site specific information on surface and subsurface ground motion at Yucca Mountain and the development of an attenuation function for Yucca.

The identification of all Type I faults, and again that refers back to NUREG-1451 in which we define what a Type I fault is, which is a fault that "is subject to displacement and has or could have a significant effect on repository design or performance."

We also believe that there should be a determination made about the possible coupling of faulting and igneous activity, including structural control, and that there be a determination of stress and strain patterns in the Yucca Mountain region.

Now, just based on the discussions by DOE and some of the other earlier discussion, like John Whitney and Chris Fridrich, I would say that the NRC staff and the DOE are not that far apart in determining what's needed to input into fault displacement and seismic hazard designs--or fault hazards. And that's it.

DR. ALLEN: Thank you, Keith. Any comments or questions by consultants or Board? Bob Budnitz?

DR. BUDNITZ: Can you explain the rationale for asking the applicant--supplicant, excuse me, to do both deterministic and probabilistic analysis? In particular, are you going to ask that they be done double blind by different teams, or is it going to be the same team?

DR. MC CONNELL: I'd say we haven't gone to that level of thought process.

DR. BUDNITZ: Well, let me go through it with you. Let's suppose it's not done double blind. If it's done double blind, there may be value because somebody with a deterministic might find a hypothesis not captured properly. But if it's done by the same team, how could the deterministic ever come out anywhere but right in the middle? Let me propose for you that the deterministic comes out on the high end of the probabilistic range, it will be adjusted so it didn't. Whatever comes out at the low end, it will be adjusted so it didn't. So the deterministic can't possibly provide anything new if a full probabilistic has been done.

Therefore, I think it's a waste of time for the supplicant to do it. What they ought to do is you ought to ask them to do a probabilistic and you ought to commission, if you feel a deterministic is useful, which I don't, but you ought to commission your own deterministic double blind and see whether your own deterministic comes up with a hypothesis that would somehow make you feel the probabilistic wasn't-- you didn't have proper sense in it.

DR. MC CONNELL: Well, that may happen. It may be a part of our license application review plan that that sort of analysis would be done at the staff level.

DR. BUDNITZ: But I suggest that you consider seriously abandoning the notion that they waste government money-- excuse me--rate payer money on a deterministic that can't possibly come out except how I said. And there's a letter which I wrote to Andy Murphy in the context of Appendix B to Part 100 for reactor siting, which makes this argument in plain English, which I suggest you might go and read and use to abandon your notion that they ought to do both.

DR. MC CONNELL: Well, I guess it depends on who's going to waste the money or who's going to use the money to determine analysis. In the past, we put the burden on the supplicant to do the analysis, and we would then review it in our license application review plan. Now that could change.

DR. BUDNITZ: Yeah, but do you understand my notion that if the supplicant does both, how can the deterministic possibly be anywhere but in the middle? It's got to be. So it's a waste.

DR. MC CONNELL: Is that an issue of how the analysis is done or who does it?

DR. BUDNITZ: Unless it's double blind, it can't possibly come out any other way, in my view.

DR. ALLEN: Further comments? Allin Cornell?

DR. CORNELL: No.

DR. ALLEN: Others? Staff? Leon Reiter?

DR. REITER: Kevin, let me try and put you on the spot here--I'm sorry, Keith. The idea of fault avoidance and fault displacement is sort of a great burden that you're placing upon DOE to come and to argue with you--what's your perception of the fact that if it's indeed true they've gone from vertical emplacement with like 7 centimeters of freeboard to drift emplacement with essentially several meters of freeboard, doesn't that--from your context, from your view as a scientist, as a seismologist, as somebody who's concerned, doesn't that sort of--does that relieve them of a large part of this burden?

DR. MC CONNELL: I think the philosophy behind the staff technical position was one of common sense, and no matter what design you have as far as waste emplacement, the common sense would tell you you lessen the uncertainty by not putting those waste emplacement features across faults that you know exist. So we would say use that common sense approach. If for some reason there are problems where common sense needs to be overridden or it can be accommodated by design, then we would be willing to go with that sort of mechanism too.

DR. REITER: Well, let me sort of pursue this and take devil's advocate. Speaking of common sense, would common sense tell you that if I have several meters of offset, of freeboard, that it's a lot less serious problem than 7 centimeters of freeboard?

DR. MC CONNELL: Certainly.

DR. ALLEN: Assuming no backfill.

Other comments or questions? Okay, let's take a 15 minute break and we'll get together again at 3:45.

(Whereupon, a brief recess was taken.)

DR. ALLEN: On the program this afternoon, on your agenda, there's comments from the State of Nevada by Carl Johnson. Unfortunately, Carl could not be here today, but Dave Tillson will be here to present what Carl might well have said, or anything Dave himself wishes to add.

MR. TILLSON: Well, as you might expect, I was involved to some extent in preparing the comments, so I'm not entirely without blame.

I wish to make one comment, however, aside from these prior to reading Carl's speech, and that has to do with the WNP-3 probabilistic study and the question that Dr. North asked about the time that was required. There were some very large mitigating circumstances that dictated that that study did in fact take 18 months. But it was part of a much broader study that was going on.

I was the principal geologist both at the time of the construction permit in 1974, '75, and I also was the one responsible for establishing those studies in 1982 and '83. And one has to remember that in that particular case, we were trying to establish a position which we knew would not be completely evaluated. It was to be put on the shelf, so to speak. We were aware that we were not going to complete the construction and operation at that point, and so that similar study, whether it would take 18 months, at this time, that's questionable. I think it would take much longer, frankly.

We also had a considerable body of information that probably would not be available at Yucca Mountain if such a study was attempted today. So it was a good piece of work. I have no question about that. But I don't want you to be misled that that study in itself was holding up the process as such.

Again, I give Carl Johnson's apologies. He had a family problem that he had to take care of, and I will provide his comments.

The State of Nevada has commented extensively to this Board about the seismic hazard assessment of Yucca Mountain. In addition, the State continues to question the adequacy and efficacy of DOE's study plans for evaluating seismic hazards. The remarks today will not repeat those comments since they are already part of the public record, but will focus our comments instead on the issue of hazard versus risk and what we know and don't know about the potential hazards of the Yucca Mountain natural system and the engineered system.

There is a need to make a clear distinction between hazards assessment and risk assessment. The site characterization program, in our view, is supposed to develop the information necessary to do a hazards assessment of the site sans, that is, without engineered systems. This is work in progress and we feel there's still a long ways to go. At some point when sufficient information is developed to provide reasonable assurances that the site will be able to meet all regulatory criteria (10 CFR 60 and 40 CFR 191), to be specific, without need to resort to any untested engineering solution, the design process can then begin in earnest.

As you have already been informed, hazards describes the potential for natural related phenomena to occur, that is, such things as vibratory ground motion from near field sources, fault rupture, fracturing, volcanic activity, intrusions, ground water rise, geochemical processes, et cetera, et cetera. It is primarily a spatial measure. Occurrence of these phenomena either singularly or as a coupled process could result in adverse consequences. That is, they could cause the uncontrolled release of radionuclides to the accessible environment. That is the risk that we are talking about.

To satisfy the siting requirements, we need to first known the natural systems and all of the potential operative processes in order to define the hazard. Subsequent to the hazards definition, we can then start to conceptualize what the engineered system needs to be and the ways it can fail when subject to the hazards in order to establish the potential consequences. We're talking about an iterative type process. Once a conceptual design has been decided upon that minimizes the potential consequences, then a risk assessment can be made.

Risk, as you are aware, is the probabilistic expression of the product of the hazards and its consequences. The level of risk that will be acceptable will ultimately be decided by the government and the citizens of the State of Nevada. To reduce the risk to an acceptable level, whatever that turns out to be, will require either reducing the uncertainties in our knowledge of the natural systems and how it operates and/or changing the fragility of the engineered system so that it is less vulnerable to being affected by natural phenomena.

We know that we cannot engineer the natural system. We can only strive to understand that system to the point where it will be reasonably assured that we know what all the significant operative processes are, how these processes are spatially distributed, whether these processes operate separately or are coupled, and how these processes might change in time when the engineered system is disturbed by the occurrence of any of these natural phenomena. Once the natural system is deterministically defined with reasonable assurance, then and only then can we begin to decide whether an engineered system can be designed, licensed and constructed that will meet the public's requirement for acceptable risk.

Now, you see this view graph up. This is a very generic and general list of things that we think we know about the potential hazards of Yucca Mountain. Now, there may be others. We know that there are some very active faults operating in the geologic setting that includes Yucca Mountain. We know that there are faults cutting through and bounding the proposed repository block that can provide direct fracture pathways to the accessible environment.

We know that fracturing on the surface in the proposed repository block is pervasive. We know that there are active volcanic processes operating within the Yucca Mountain geologic setting. We know that there have been volcanic processes that have directly affected Yucca Mountain in the past. And we know that there has been a coupling of volcanic processes and seismogenic processes in the past, and we know that there has been hydrothermal alteration of the rocks in Yucca Mountain.

What we don't know about the natural system of Yucca Mountain at this point is the type, location and extent of active blind faults under and around Yucca Mountain, the distribution of fractures within Yucca Mountain, how the fracture permeability will change due to earthquakes on any potential blind fault, how the ground water system will change in response to movement on any of these faults, the structural control for volcanic processes in the vicinity of Yucca Mountain, whether there is an active magma chamber in the vicinity of Yucca Mountain, and how fluids move through the vadose zone.

What do we know about the engineered system? Nothing. We, therefore, have no idea what the potential consequence could be in response to some natural phenomena. We also can't be sure that any hazards assessment results being produced by the present ongoing process are relevant to the needs of the design engineer. These ideas have been stated before today.

What we don't know about the engineered system. We don't know how much and what kind of waste there will be. Is it going to be 77,000 metric tons, 86,000 metric tons, as we heard at the full Board meeting in January, 100,000 metric tons? How much of it is defense waste in addition to the 86,000 metric tons? What other types of non-spent fuel waste is being considered? Is the plutonium from the weapons disassembly being considered for disposal as high-level waste?

If you take a simple calculation which we saw on the basis of the thermal loading, you do not have enough space to get 86,000 metric tons into Yucca Mountain, particularly if you have to avoid any faults or fractures.

What the thermal loading strategy will be. How can the thermal loading strategy be finalized if all of the waste streams that would go into the system are unknown? How pervasive the faulting and fracturing is at the repository level. We know a lot about the surface, but not much about the repository level.

We don't know how to determine near field seismic ground motion from as yet to be identified sources. We don't know how to effectively translate near field seismic ground motion into repository design. We don't know how to test a near field seismic design. We don't know how to design and test seals to withstand vibratory ground motion, both far field and near field. And we don't know what the potential consequences of system failure is.

Finally, to preclude any questions from Leon Reiter, I want to put up this last slide. I want to close my remark with this quote from a best-selling author who is in this room. Yeah, this is his. It's a direct quote out of his book "Earthquake Hazards," which I'd strongly recommend reading, by the way.

"The seismic analysis needed to help prevent the public from being subjected to an unexpected release of radioactive waste from an underground repository during its 10,000 plus year lifetime is quite different from the seismic analysis needed to help prevent earthquake induced deaths and serious injury during the 40 to 50 year life of a nuclear power plant."

"The analysis for a repository must take into account great public scrutiny--and I want to put a paren in here that we're talking about the scrutiny of the public within the State of Nevada particularly, and those who are working on the project outside of the DOE--the hypothetical changes in the tectonic regime during the next 10,000 years, and the effects of earthquake on the buried waste containers and ground water flow, and the path of the radionuclide release to the environment."

And I think the last one is the one you should pay the most attention to. It's not the seismic design per se, but it's how those radionuclides might get out into the environment that concerns us the most. That's the end of my remarks.

DR. ALLEN: Thank you, Dave. Any denial from Leon Reiter?

DR. REITER: I just hope that everybody goes out and buys ten copies of the book.

DR. ALLEN: Any comments or questions from the Board or consultants? Bob Budnitz?

DR. BUDNITZ: Sir, I've never met you before. A lot of what you said made perfect sense. But something that you said sounded to me so incredible as to defy common sense. You said that the design ought to wait for the characterization.

MR. TILLSON: I beg your pardon?

DR. BUDNITZ: You seem to say in the beginning that all design work ought to not go on until after the site was characterized.

MR. TILLSON: I didn't exactly say that.

DR. BUDNITZ: Whatever you said, I could find. It's in your text, but I don't have it in front of me. But that's what I heard, and if that was so, that makes no sense at all to me.

MR. TILLSON: No, I didn't really say that. If it came across that way, I apologize.

DR. BUDNITZ: Good.

MR. TILLSON: I say that fitting the final design--in fact quite the contrary. We need some conceptual design that can be iterated as the hazards assessment proceeds.

DR. BUDNITZ: Oh, good. I agree with that.

MR. TILLSON: Yes, very definitely. At this point, the design that's being presented constantly changes. It constantly is moving and we don't have any idea. The first decision and the most important one we think is the decision on thermal loading.

DR. BUDNITZ: I apparently had misunderstood. It wasn't as incredible as I thought.

MR. TILLSON: Oh, no. I think some in the State would certainly like the DOE not to do any design work until they finish the characterization, but I don't happen to hold that view.

DR. ALLEN: Other comments or questions from the Board or consultants or staff?

DR. REITER: I can't resist this question. What we don't know about the engineered systems is how pervasive the faulting and fracturing is at the repository level. Is that an endorsement by the State to proceed with underground exploration?

MR. TILLSON: Yes and no. I think they should proceed with underground exploration, but I think that the size of the TBM should be about 4 5/8 inch diameter, and it should be a horizontal drill hole. Our concern, or the concern the State has about proceeding with underground exploration is that somehow that facility will become part of the final repository, and we're concerned that until the design parameters are much more closely fixed, that that may be a mistake to proceed too far.

DR. ALLEN: Well, as you know, the Board has expressed some similar concerns.

MR. TILLSON: Yes.

DR. ALLEN: Any other comments or questions? Thank you, Dave.

The next speaker is Steve Wesnousky, the University of Nevada at Reno, on how good is PSHA. Steve?

DR. WESNOUSKY: As Clarence said, my name is Steve Wesnousky. I'm from the University of Nevada at Reno. So you're all probably thinking which side of the coin I'm supposed to come down on this issue of probabilities.

I've been involved with seismic hazard analysis at various levels for about ten years, and I've paid attention to probabilities and I've used them. And I've come to a general conclusion--can I have my first slide? And that's it.

DR. ALLEN: Thank you, Steve.

DR. WESNOUSKY: Now, Clarence, you must miss the days when you wrote the proposals, papers, and you put out that idea and you got the review back and they said your ideas are good, Clarence. And what did the next sentence start with? Next side, please.

So I'm sorry I didn't make handouts for this for you guys, but I think you can remember this.

However, probabilities aren't without uncertainty, and I guess that's the issue I want to touch upon, just to bring some of the issues, and perhaps the topics or the points I'm going to be making are points of semantics of interpretation of what probability is.

Seismic hazard analysis you can break down pretty clearly. It's a very simplistic process when you put it on paper. You estimate the size of the earthquake. You estimate the frequency distribution of earthquakes, the different sizes, both in space and time. And then you estimate how the ground is going to shake somewhere as a result of that, and the ground breakage, if you will.

Now, some of the elements of that are really conducing to statistical analysis or probabilistic analysis. This is a couple of slides from a colleague of mine, John Anderson at UNR, of strong ground motion data. And the vertical axis is peak acceleration, and the horizontal axis is distance from the earthquake.

This upper slide is for small earthquakes which occurred there, of which there are an abundance of earthquakes, so there are an abundance of dots. And we can use various different ways to fit this to make predictive curves from this, from standard regression analysis to non-parametric techniques. And if we go to the literature and we look, there is a curve by a guy named Joiner, and there is probably a curve that Allin Cornell uses. And they all use different methods to fit this data and they all make different predictions. So even within the sense, it means that the interpretation of this is somewhat model driven.

In the lower one, we have the same plot for larger earthquakes, of which there are fewer. And this illustrates another uncertainty in hazard analysis, that there are fewer big earthquakes and there are fewer observations close to the fault. And what you see here is much less scattered, so the question arises is this scatter real or is it an artifact that we don't have a large number of observations. There's an uncertainty there.

But, again, this is one of the principal aspects of seismic hazard analysis which you can put standard probabilistic estimates to, give me an earthquake size, I'll use this data and I'll make some predictions, but I'll have to choose which curve.

The other aspect of seismic hazard analysis is estimating the size of the earthquake, and that also seems to be a relatively straightforward process. We know earthquakes occur on faults. Often we see them looking at the surface or after shock distributions. We know how long they were, what their dimensions were.

So we can go around the globe and we can plot up the size of the earthquake on the vertical axis versus the length of the rupture, and then we can go to a place like Yucca Mountain and say here are the faults, here are the lengths, let's use this to estimate the size of the rupture. But even here it's not that simple.

For example, I've separated the dots from solid to open, and the open symbols are earthquakes which occurred on faults that slip more slowly over the long term than those that slip at a faster rate. So depending on whether it's perhaps a fast slipping fault or a slow slipping fault, we're going to have to make a decision on which one of these lines to use with respect to mapping a given rupture length.

Now, it gets more complicated than this because some of my colleagues don't agree with this. So we have different models on how to approach this data, and that's the point I want to make, is that even when we look at this data, you're going to be making subjective decisions based on some expert's opinion on how the earth actually works.

It gets more complicated when we recognize that all faults, all earthquakes don't occur on faults that are easily seen, and moreover, near Yucca Mountain and in the Walker Lane, we find earthquakes like this one, the Cedar Mountain earthquake of 1932, which produced distributed rupture over a zone some 60 by 20 to 30 kilometers wide.

So then how do we use our standard regression analysis of these nice earthquakes which produce faults on surface ruptures very distinct in which we measured their length? Again, there's an assumption, and those standard regression curves aren't necessarily useful in this sort of analysis for this sort of earthquake.

DR. ALLEN: Excuse me, Steve. What are the green areas there?

DR. WESNOUSKY: Oh, I'm sorry, Clarence. That's a good point. This is a map of the faults and basically I've just shaded the regions at the time of the earthquake which were characterized by relatively pervasive fracturing. So this is a zone of surface ruptures, whereas many of the earthquakes we think about would be limited to a distinct line.

Well, if we look at Yucca Mountain, what we see-- and this is nothing new to my colleagues that are working in the probabilistic format, and they consider all these things--that if we look at Yucca Mountain, we see that there's a distributed pattern of faults, which are these black and red lines. Also the investigators tell us that within these faults, they find volcanic ash which came from these volcanic vents, suggesting that perhaps they all did rupture simultaneously in earthquakes.

So how do we estimate the size of the earthquake here? Do we take the individual fault lengths or do we take the whole zone and some other subjective informed expert estimate of what the size of that earthquake is going to be if it occurs here at Yucca Mountain? And this also plays a role into the standard method of estimating recurrence times of earthquakes, because usually what we do is we estimate how much slip is going to occur on the fault, and then we divide that by the slip rate.

And so if we assume only one fault, we get lots of small earthquakes during a short period of time, or if we sum up all the slips and pretend they all occurred during one earthquake and we divide by the average slip rate, we get a bigger earthquake with a longer recurrence time. So when we start taking this model, we get very large uncertainties. Okay? And that's what you want to quantify, is the uncertainties.

And then there's also the question of whether we're looking at crustal standard normal type faults or whether or not detachment faults occur here as well as whether or not they can produce earthquakes.

The reason we have these models is because there is not enough data to look at them statistically and probabilistically to determine and to say that one of them is correct.

Other aspects is how does the recurrence interval of the largest earthquakes take place on these faults. Here I put a plot of displacement versus time. And so for the perfect idealized case, earthquakes would occur periodically. So we have an earthquake time, earthquake, same slip, and it's a very regular process. Well, there are also investigators and there's evidence to argue quite strongly that it really doesn't work that way.

Then we have clusters of activity separated by quiescence. So when you dig your trench, you might not know--you won't know if you're here necessarily, or if you're here. So that brings then in the order of uncertainty and an assumption has to be made on what model you're going to use to estimate recurrence for the largest earthquakes.

Now, what about the small to moderate sized earthquakes? Now we're talking about a long period of time, and we're not talking just about seismic shaking, I don't believe. The fracturing can also play a role in this coupled system that exists, if that's what you're concerned with. What we generally do is because the historical record is so short, the instrumental period of recording is so short, that for the biggest earthquakes--and this is a plot of magnitude in this direction versus the number of events greater than or equal to a given magnitude.

So at this end, we use geology to tell us how frequently these largest earthquakes occur, and it's down here that we use the instrumental record. And we have to make some assumption about how the two connect, so we might have 15 or 20 or 30 years of data here to plot a line up from observation.

Well, generally what you need to do is extrapolate, but now we have geology, so we can couple these. And what we observe in nature is different sorts of distributions. In some sense, sometimes we see that it goes very linearly, what seismologists would refer to as a Gutenberg-Richter relationship, or other times there's a paucity of data in here, and then you see the Gutenberg-Richter relationship. And we do see this in nature. For example, here's two faults in California where I've coupled the geological data with the instrumental data, and Southern California I think has perhaps the most--the longest period of recording of any network in the U. S., and perhaps the world. We see the San Jacinto shows a very linear relationship, and here we see this very distinct bend or flexure which we refer to in the lingo as a characteristic earthquake distribution.

Well, it's critical to know which one, and in places where the instrumental recording period is short and the uncertainties in the geological data are large, whether we choose this type of distribution or we choose this type of distribution can result in order of magnitude differences in the prediction of these moderate to small earthquakes during the time period which we're interested in.

So, again, there's another model that we have to make a decision on, which is not necessarily based on the standard statistical analysis of prior observations to come up with the distribution of inter-event times or magnitude frequency distribution.

The point I'm trying to get at is really summarized in this one, and this is the one that I think deserves some discussion, is that we can have a whole suite of models, Model A, Model B, Model C, and it can go for any of these things, estimation of earthquake size, the shape of the magnitude, frequency distribution, the recurrence rate, the rate of slip on the fault, and we can propagate these through the seismic hazard analysis and we can estimate what's going to happen to some facility. But the point is is when we have these models, they're basically mutually exclusive at the site that you're working on. And what you come up with when you go through this analysis is what Allin Cornell was talking about, is the probability of exceedence curves.

So here I placed probability of exceedence, which can be over a given time of some event, whether it's an earthquake, whether it's how fast or how frequently a fault slips, and we can have different models and we can choose some probability level that we're interested in. And what the curves do for us, and it's very useful, is that they define a range of uncertainty.

Now, I think where the problem comes in is now what happens is you have to make some subjective expert opinion evaluation of which model is most correct. And that's where the lines between probabilistic analysis and deterministic analysis become blurred, and that's not coming through--I think that's one of the problems as it's stated, is that when we speak about probabilistic analysis, it's really not coming out that there is a tremendous amount of subjectivity. All right? And we cannot actually categorize the probability density distribution to characterize these models.

What we can do is we can ask the experts what they think is the best model, and what we get are probabilities that the experts think that these models are correct. So this might be a 10 percent of the experts and 50 percent and 40 percent, but only one of these can be right in nature. So one of the concerns is when you wrap in all of these models to your analysis is are you in some sense degrading or lessening the probability and the net outcome. Although you consider all the possibilities, you're actually lessening the probability because you're putting weight on the other ones which might not be true.

There's one other aspect or element in the hazard analysis that I want to bring up, and that's the element of time that you have to deal with. And this is a suite of California earthquakes and this is an anecdotal statement or argument. California is very active. Earthquakes occur probably ten times more frequently in California than in Nevada in the basin and range and, therefore, we know a lot more about them. And it was in the 1930's that some fellow said, well, all we have to do to do hazard analysis is look at the faults and that will tell us where the earthquakes are.

And then in the Seventies, this fellow sitting with you wrote another paper and said come on, guys, this is where the earthquakes occur. This instrumental data is not going to do everything for us. Let's map the faults, let's get the slip rates, all the buzz words, let's trench them, and that's how we're going to learn where the earthquakes occur, because the instrumental and historical records are deficient because they don't span a long enough period of time.

And so there was a flurry of activity, and these are all the active faults back at about 1980 that we knew about in California. And I'm sorry I'm not talking about Nevada, but these are all the faults we knew about and we pretty much thought we had things wrapped up. I mean, there was this warm fuzzy feeling; we're mapping the faults, we're taking into account the instrumental record, and seismic hazard maps were made--and I'm not going to show any of those hazard maps because I made them--based on that approach, and it seemed like a very sound approach. And it still is sound in its general manner.

But what happened during the ten years later, subsequent to that approach and development of maps? Basically it was the occurrence of these unexpected earthquakes in Loma Prieta, Coalinga, Northridge, Whittier Narrows, all occurring on structures that weren't even conceived of ten years ago in terms of input to seismic hazard analysis.

Now, all my colleagues will now say well, we account for those now because we know about them. And that's the point I'm trying to bring across to you, and it came up in my mind this morning when John Whitney said, well, we understand normal faults better than we do California faults. And I don't think he meant to phrase it that way, but in essence, you can have some false confidence because of a lack of data in terms of these faults and what we're estimating inputting into seismic hazard analysis. And that's the point I want to bring across on this.

Are you bored, Leon?

DR. REITER: No.

DR. WESNOUSKY: Should I continue? Because this is a very nice place to stop. Okay.

The point here is that you folks are dealing with bases per hundreds to thousands to ten thousand years. So there should be a severe or a significant element of conservatism in your approach, particularly in light of my prior comments that this brings about that you can do your probability trees, you can put in all your models and you might not even have the model that's correct in there.

Now, how about the uncertainty of the uncertainties? I'm going to stay in California because there's a nice analog. A number of years ago, 1988, I think Allin was involved with this, and there was a group of scientists, seismologists and engineers--by the way, I'm the seismologist that runs into the engineers--and they made this map. This was the working group for earthquake probability put together by the United States Geological Survey--the National Earthquake Prediction Evaluation Council, and basically they evaluated the probabilities of earthquakes along the San Andreas and other major faults, which I've put up here for you. So these are the conditional probabilities, and that's a little bit different necessarily than what we're talking about, but these bar graphs here basically show zero to one probability from 1988 to 2018 of the probability that this fault is going to break along this section of the fault.

So I want to just talk briefly about the Mojave section of the fault and the Parkfield section of the fault. Now, the basic principle behind these probabilities was an assumption of a model on how the earth fault worked, and basically the assumption was the time until the next earthquake will be equal to the slip that occurred here during the last earthquake, divided by the slip rate. I can even understand that one. You just put the dot on top. Okay? So you come up with a T, an estimated time of occurrence, the time from the last event, plus that interval that you've estimated here. And with that, they assumed a certain distribution around that expected recurrence time, and they plugged it into a log normal distribution and estimated a probability of .3.

Well, what is ignored in that type of approach, or was ignored and pointed out by a fellow named Jim Savage at the Survey was that estimate of the next time of occurrence has an uncertainty to it itself. And that was ignored and has been ignored commonly in these types of analyses.

And so if you take into account the actual uncertainty in the prediction which will propagate through because of your uncertainty in your slip rates and your uncertainty in what the amount of slip was during the last earthquake, you can actually do a simulation and ask what's the probability that the probability is a given amount. And then, sure, you come up with a bar graph that shows the frequency of which you could expect the given probability to be correct, and it's relatively flat, the argument being the uncertainties are so extreme that you could argue that perhaps they're not even significant.

Now, this is interesting because it also brings another element which I'm just kind of learning about. I've never been involved in the licensing domain at all, but here we've assumed a given model, and there was a group of scientists who said this is the way it worked. But there were no intervenors. Allin Cornell was involved with this, and I think I recall that you--no?

DR. CORNELL: I was going to wait--

DR. WESNOUSKY; He was involved with it, but he had misgivings about it. And I think Allin has lived in the consulting world that if one of the intervenors of Nevada or somebody came and said, Allin, do you agree with this approach, and he would have said no, and he would have been up somewhere saying this approach is not valid, or there is a better way to do it.

Now, the Parkfield approach--and this is going to hammer home somewhat on the previous idea that I put up in terms of models--Parkfield is a very famous place in seismological circuits. Time's up?

DR. ALLEN: Oh, no, you've got about seven minutes.

DR. WESNOUSKY: Oh, I'll talk slowly. Parkfield. Parkfield is a famous place in seismological circuits. It's the first place that the U. S. Government, the National Earthquake Prediction Evaluation Council actually anointed officially the prediction of an earthquake. And the basis for the earthquake was the historical records showed that since about 1850, there was a sequence of about six earthquakes which occurred very regularly in time.

Now, interestingly enough, the prediction was made based on a regression of this curve, which predicted the occurrence sometime around 1988 to 1992. What they didn't include was this event, because it didn't seem to fit. It's rather capricious, but it was done, and it was argued and there were people that didn't agree, but there was physical arguments to say well, maybe this was triggered a little bit by the last earthquake, so we'll just do our regression through here.

Now, if I calculate the uncertainties on this depending on these two fits, they look like this, and they're basically mutually exclusive models. Parkfield prediction--okay, this is a probability density function from--focused on about 1988, which was when the prediction was made for, was a very tight probability density distribution.

Well, again, Savage just pointed out, well, perhaps we've ignored an alternative hypothesis, and that's really the key word, is to examine the alternative hypothesis. And if we take into account the other events, all of a sudden our prediction window becomes much wider because there's more scatter in the curve, and it's defined by curves like this red one or this blue one. And I can talk about details of those, but the point is is that we're bringing probabilistic analysis into the licensing arena.

Does that mean that you're going to take something and say we're going to use a Gaussian distribution, we're going to use the 94th percentile and you're going to get that in the licensing package to be litigated and to be concrete? Or are you really saying, well, we're going to use our judgment here and we're going to take advantage of it and use probability to basically help us think about it, but when it comes down to it, we're going to provide you a piece of information that says this is what we think the response is, or this is what the size of the earthquake is. Because what I want to point out here is if it's the previous, it seems like a real can of worms to get into in terms of trying to define what those parameters actually are going to be. And you can't ask any questions.

DR. ALLEN: Thank you, Steve. Nevertheless, in spite of your admonition, we may do so. Do the Board or the consultants have any questions or responses?

DR. CORNELL: First, no, I was not part of the '88 working group. I was part of the '90 working group.

DR. ALLEN: I was a member of both, unfortunately.

DR. CORNELL: I think the '88 working group did a pretty good job, though I'd like to say I think you misrepresented it. They did indeed account for the uncertainty in the mean inter-arrival time, and the slip rates that went into it. And what they chose to do was report only what we would now today call the mean probability.

DR. WESNOUSKY: Yeah, I think that's fair.

DR. CORNELL: And what Savage did and what the '90 working group did was--but it was in the appendix of the report, because recall we have to present this to the user group, who in that case was the public, and we were a little nervous that they were having trouble with even the bar graphs that you showed, and we put uncertainties on the bars. So, indeed, we did express what the implications of the uncertainties in the parameters were. The key ones, as you alluded to, were estimation of the mean inter-arrival times because they're coming from uncertain slip rates or uncertain paleoseismic information, limited dated, et cetera, et cetera, and the one way that we would today, and as Kevin described this morning, do regularly is in fact put those uncertainties on parameters so that they produce--they would propagate and produce uncertainties on probabilities.

Now, these aren't uncertainties--unfortunately words become difficult here. As I said, we need Esperanto or something, but it becomes uncertainties on probabilities, and what we would do with this problem in that context, which is the next step up, not uncertainty in parameters now, but on multiple hypotheses as to the model, is part of what you alluded to, is you would poll the experts as to where they are, but you also let the individual expert take a look at the evidence presented by the different proponents of these models, and let him put relative weights on any--you know, each expert has the opportunity to put relative weights on the model based on his interpretation and judgment as to the evidence presented by the specialists that have perhaps generated.

That's one of the main reasons that Kevin talks about having peer review, is one way to do this. Sure, it's true that the peer review--members of the peer review panel haven't, you know, spent as much time kicking around the dust of Nevada, but they are presumably people with the capability, and they've looking at other such studies, and they come in and hopefully ask critical questions and maybe push another hypothesis out, but ultimately try to help and make sure that the evidence has not gone towards reporting, and in the end, as I'm afraid I might have interpreted your suggestion to be that, well, finally we've got to sit down with scientists and say we believe in Model 3.

The problem is that suddenly becomes--that becomes value judgment making, in my mind, because you've somehow said we should be conservative. Why should we be conservative? That's not a scientists decision to be made, in my opinion.

DR. WESNOUSKY: That's exactly right. But you're assigning probabilities to these, and you can't tell me that

that's a probability of how the earth behaves because--

DR. CORNELL: No, no. It's what I called an epistemic uncertainty, Steve. It's related to degree of knowledge. The 1000 years, we won't put our weight on any one of those. I mean, none of those models is right, and in 1000 years, we'll have another little band of models hopefully that will be closer together. But it's reflecting the fact that we have limited knowledge at any given time, and those--you know, some people call them red probabilities instead of green probabilities. Some unfortunately call them subjective as opposed to objective. Some call them epistemic as opposed to aleatory. You have a choice of words here, but the point is they do reflect a different thing. They aren't--

DR. WESNOUSKY: They aren't probabilities; they're judgments. It's semantics, but I think that's where--I'm sorry, Allin, to interrupt--but that's kind of where I want to point out is the problem, is that it has to be pointed out, I mean are you really doing probabilities.

DR. CORNELL: No, Steve, probability theory is something proposed by mathematicians that follows three axioms about--it turns out both frequencies of occurrence and subjective assigned degrees of belief follow the laws of probability theory. So that means you can call them both probabilities.

DR. WESNOUSKY: Okay. But it's--they're apples and oranges.

DR. CORNELL: You're absolutely right. We tried to keep them very carefully separated; that's why the hazard curve comes out, and if you give me what parameter value and what model, I'll get one hazard curve. But I have multiple models, multiple parameter values; I get a suite of hazard curves. They come out with weights properly propagated. I get uncertainties on the frequencies. That's maybe an easier way to think about it; the frequencies being your probability of how the earth works, and the uncertainties reflecting our limited knowledge about what those frequency--

DR. WESNOUSKY: That's exactly what this says. You can do the range of uncertainty, but it's the judgment--it's a totally different--

DR. CORNELL: You're right on track, and I think--what I'm saying is I think it's terribly important that we--that the scientists display all those models. And what we're trying to get away from, in my opinion, is the previous deterministic licensing practice or design basis practice where the argument comes over the scientists on both sides of the table having to say which of those is the right model. We can't do it. So that's not a question that we should be asking.

DR. ALLEN: We're going to have to be moving on, although I sense the way Warner is moving his microphone around, he has something to say.

DR. NORTH: No, I'll pass until later. I'm enjoying the discussion, which I think is right on the point. And the comment I will simply make is that this kind of work is not easy. It is as much art form as it is established procedure, and documentation so it is understood what the basis is for probabilities either of the apple kind or the orange kind is very critical. These points were made at length in the workshop on expert judgment that was held, what was it, a year ago November.

My serious concern is that I don't see any evidence that the Department of Energy's program in this specific area has learned the basics of how to go about doing this and demonstrated that they understand it by doing iteration one.

DR. ALLEN: Thank you. I have a hunch also that Keiiti has some thoughts on this, but he's going to be speaking here in a few minutes, so he'll have his chance.

Thank you, Steve, for a presentation that definitely was not boring. I even agree with most of what you said.

The next presentation this afternoon is by Paul Pomeroy on the same topic as Steve, presumably not with the same view graphs or slides. Paul is a seismologist. He is a member of the advisory committee on nuclear waste, which is advisory to the Nuclear Regulatory Commission. Paul?

DR. POMEROY: I have to offer you the usual disclaimer, that the statements that I'm going to make are my own and do not represent those of the Advisory Committee on Nuclear Waste, nor the Nuclear Regulatory Commission.

I do have the same title that Steve's talk has. I will give you--my answer is that it's very good in general, but--rather than however--in this specific instance, it depends I think on the way that the case is made for it and the scientific evidence that's associated with the way the case is made. I'll try to tell you something more about that while I proceed, but I really want to try to focus on a few actions that could conceivably move the regulatory decision making process forward.

I'd like to deal today with the part of the probability space that I characterize as a "degree of belief" probability space, and within that space, I have attempted to indicate a discontinuous spectrum representing increasing amounts of expert judgment associated with our probabilistic assessments, starting from the left-hand edge with little or no expert judgment involved, and ending in an area where nothing but expert judgment is involved.

All of you know that I've been concerned about expert judgment for many, many years now, and its use in this particular process. I am concerned that at the moment, as Warner has said so eloquently, that we don't see any significant progress in understanding how to incorporate in a more useful way expert judgment into the assessments to assist in the decision making process.

I believe that we have, as Kevin has pointed out earlier, a large number of studies that cover this spectrum, from non-expert judgment studies to ones that involve pure expert judgment. Contained within those studies I think there is a great deal of information on how to improve what we're doing with expert judgment, and I think--I feel very strongly that we should be looking at those studies, all of them, including the WIPP studies, to try to determine what it is that we can do in a better way.

Second item is that somewhere on this spectrum, PSHA and PVHA fall, not necessarily in the same place. I'd like to say from the regulator's standpoint, the regulator is looking for assurance, assurance that the probabilistic assessments are valid, that the time frame that the regulator is considering is the same, or at least covered by the time frame of applicability of the probabilistic assessment, and of course the regulator is concerned that the results have some consistency.

Finally, I'd just like to say that I want to make a point that each of these "degree of belief" probabilities as well as classical probabilities will be examined in intense detail in any licensing procedure. None of us have any conception, I believe, of how serious that intensity of examination will be. It behooves us, I believe, to be very aware that the concern is going to get progressively greater as we move along this spectrum towards complete expert judgment. And we've seen in last Sunday's New York Times the article that Allin reference, the criticism of people who are using expert opinion only, or relatively speaking only, in these determinations.

My main point here, though, is that we really need to do some critical research and compilation from the studies that currently exist regarding the use of expert judgment in decision making.

How do we expect expert judgment to be treated in a hearing process? And this is actually taken from a personal communication from a Dr. Warner North. It's a very succinct statement of the treatment of expert judgment, and I suspect the treatment of probabilistic assessments in a hearing process. They must be highly credible. And by highly credible, it means they must be well reasoned, they must be supported by available data. They must be consistent with the scientific literature and, in fact, in some cases this has been interpreted to mean they must be published in peer reviewed literature, and they must be consistent with the judgment of at least a portion of the scientific and technical community. I'd ask you to remember that last criteria in particular at a later point in the talk.

I'd like to say also that expert judgments are going to be judged on a number of other criteria that are included in this slide, particularly the identification and selection of experts, design and conduct of elicitations. And I want to stress particularly the question of aggregation of judgment since we've talked some about aggregation here today.

My perception of the NRC legal position is that you can aggregate expert opinion if you wish, but ultimately in a legal sense, what is admissible is the individual's testimony and expert judgment. And the aggregation will simply be de-aggregated and each individual's judgment will be explored.

All of these things, including the influence of the normative experts, those people who are experts on expert elicitation, will be evaluated in any process.

Let me turn briefly, since we've talked some about nuclear power plant siting, I'd like to talk a little bit about Appendix B and what Appendix B contains, because it's a representation of thinking of a number of people who are seated around this table, and it is a proposed approach for the future of power plant siting. So this is a proposed seismic siting criteria, Appendix B to Part 100. This is one relatively recent publicly available approach that's been suggested. This is certainly not the final version.

It starts with conducting a probabilistic seismic hazard analysis using the EPRI or the Lawrence Livermore technology or techniques, if you will. Then after some other intermediate investigations, it says calculate the site specific ground motion for the plant. There may be a strong deterministic element in that particular element of Appendix B. There certainly is in the next one. Bob Budnitz pointed out that he wrote a rebuttal that dealt with this next item. I wrote a rebuttal to Bob's rebuttal. Bob won.

DR. BUDNITZ: I think we both won.

MR. POMEROY: In a sense, we both did. We finally decided that the staff itself had to conduct an independent check of the probabilistic results using some sort of deterministic analysis. And as Allin is quick to point out, once you've done that, you automatically force the applicant into doing a deterministic analysis as well because he would be relatively unwise to enter any sort of licensing discussions without that in his back pocket.

And, finally, we all feel that the EPRI and Lawrence Livermore data bases of probabilistic methodology need to be updated every ten years or so. Eventually I'm going to get to the point here of telling you why this is an acceptable methodology.

There are many reasons why. Let me say that it's an additional process, obviously, between a fully probabilistic and a fully deterministic determination. I suspect that in the future in any case, probabilistic analysis will predominate, and I've strongly--I'm an ardent supporter of the probabilistic approach in this application using the probabilistic analysis for interplate regions where you don't know the causal mechanisms and you're dealing with 40 to 50 year lifetime structures.

Part of its acceptability relates to past experience, and I'd like to turn quickly to that if I can. A lot of us have experience with both the EPRI and Lawrence Livermore studies, probabilistic studies. I don't have any direct experience with the WIPP studies, but I know that they're, from the viewpoint of determining the usefulness of expert judgment, they're extremely important. They're also extremely important because, in my estimation, they're significantly ahead of the DOE, NRC performance assessments and probabilistic assessments.

Again, research should be conducted. I don't want to make that point too many times. I'd like to turn to the EPRI probabilistic seismic hazard assessment, because I believe it was something unique and unusual, and it contributed to the acceptability of probabilistic seismic hazard assessment.

It was a well planned, well funded and well executed massive transfer of technological information, not only to the large number of participants, direct team members, oversight committee members, advisory committee members, special studies committee members that were associated with the project itself, but also with a large number of observers and regulators that sat out in the audience.

The study is important not only for its PSHA results, which are important in and of themselves, but also because of the involvement of that technical community. And that involvement led to clear understanding, I think, of the methodology, purpose, limitations and usefulness of PSHA.

I believe that in some sense the EPRI study is a paradigm for a similar activity to inform and educate the technical community on the PSHA, PVHA and the other probabilistic assessments that we're going to do here. The paradigm follows in several ways, not only the massive technological information transfer, but in a critical way, the timing in which it was done.

The EPRI study did not first submit to the NRC a methodology topical report for approval. EPRI instead carried out this full scale probabilistic seismic hazard analysis with all its good points, and perhaps some bad points, and developed a consensus, or at least developed a presence of a portion of the scientific community that agreed with the methodology, in fact strongly supported the use of the methodology. That kind of a consensus or involvement of the technical community provides the regulator an assurance that he's got one of the key elements covered in accepting a methodology for use in this--in any licensing procedure.

I think that that might be an approach that would simplify the life of the regulators in this case.

I think we need also a public debate on the validity and applicability of PSHA, PVHA and other--all the key uncertainty areas for the time periods that we're dealing with here. And that kind of a discussion is--the responsibility for that discussion falls on the protagonists of the techniques and I believe it should be carried out. Whether you particularly like, for instance, the Krinitzsky criticisms, they do represent a viewpoint that needs to be discussed.

Finally, I don't think the interests of the country will be well served if all of these decisions on applicabilities of methodologies and assessments are postponed until the licensing hearings are underway. Resolution of many of these issues, at least temporarily, can be achieved in an approach similar to the one that EPRI used, and perhaps an even more massive approach would be that of a moot court hearing, which I have been advocating for some time now. This could be combined with or preceded by the informational transfer and the debate.

So I'd like to conclude just by saying that you can perhaps improve the process, the regulatory process and assist in the regulatory process in a number of ways. First of all, I think there should be a recognition by all parties that the underlying scientific bases of these technical judgments are going to be challenged vigorously in the courts and in the hearing process. We do need research on how we use in a better way the expert judgments. We really need to conduct the public debate in some form of these alternative idea regarding the use of PSHA and PVHA. We don't need to talk to each other about this problem. We need to talk to a much broader audience.

This is a highly visible and highly emotional public set of issues that we need to resolve, and we need to resolve these methodological and applicability issues in advance of any licensing procedure by I think something like a moot hearing approach.

And just so Dave Tillson doesn't get on me immediately, I will say of course we recognize that no issue can be resolved finally, but we can gain a tremendous amount of understanding by a full scale study, a bottoms up approach. Thank you.

DR. ALLEN: Thank you, Paul. Questions or comments from the Board or consultants? Allin, are you waiting to say something?

DR. CORNELL: No.

DR. ALLEN: Yeah, Bill Melson.

DR. MELSON: John, I just wanted to commend you on your comment about communication with the public about these matters. And I would say even these meetings and some of the probability discussions that go on, there is an assumption that everybody who understands the fact of what people are talking about, so I think one can even practice that here, not in your case, I think yours was very clear, but I think this is a critical issue that you've touched upon, which is the clear communication, the best one can do in complex ideas. I think sometimes the jargon gets very heavy in the area of probability, and it needs to be avoided. It's not a luxury that we can afford to have if we're going to be able to communicate with the public. So I'm really glad you made the point of trying to communicate what's going on.

DR. POMEROY: I agree. I think that's one of our principal problems right now. I don't see anything in the current structure that's going to allow that kind of interaction to take place, and I'm very concerned because I feel strongly that it should take place.

DR. ALLEN: Vis-a-vis what you said about technology transfer, during the EPRI study, I might just say that during the initial phase of the EPRI Eastern Seismicity study, I was on the advisory committee to EPRI, and after a number of these meetings, even though I was a paid consultant, my conclusion was I really should have been--they should have charged me to listen. And Carl Stepp, I'm still waiting for that bill.

DR. POMEROY: I think Carl deserves a lot of credit for that particular study.

DR. REITER: Paul, let me see if I understand something you said here. I don't want to misquote you. And that was your use of the EPRI study and the way they worked it in the topical report as a paradigm for Yucca Mountain. And essentially let me paraphrase it as the way I think you said it. Your belief is that the DOE is wasting its time at this point trying to submit a topical report to get some sort of poll of the NRC as to how to do something; rather, they should go out and do a study which involves the whole technical community, including the NRC, and develop a consensus. Is that correct?

DR. POMEROY: Yeah, that's the approach that I think is one way that we can move forward in this process. I don't say that there aren't other ways. I am personally, and this certainly doesn't, again, represent the NRC viewpoint, I personally don't think that approval will be quickly forthcoming. Even after the EPRI study, it took two years to achieve an approval by the NRC of the technical position outlining the PSHA study that had been undertaken by EPRI, in essence. I don't think that two years is anything like the representation of an appropriate time scale that it might take if you submit--if you simply submit a seismic hazard methodology, a topical report, at this point in time. I think it might take an infinite amount of time to achieve resolution, and I don't think that's a useful use of our resources. That's a personal opinion.

DR. REITER: Well, again, do you think--personally, do you think that the NRC could agree to such an approach rather than a beforehand agree to the topical report?

DR. POMEROY: I think the NRC might agree to that. I can't answer, obviously, but I think they might agree to it. I would like to see the--all of the participants involved in a very real way in this kind of moot court approach, this first round PSHA, PVHA, PCHA, all the other uncertainty probabilistic assessments that we're going to have to do in the licensing process. Let's carry that through. I have lots of ideas about that.

DR. ALLEN: Bob, you look like you're compelled to say something.

DR. BUDNITZ: Paul, you had a comment about the ownership of opinions being that of individuals rather than of a collection, and as I recall, you talked about admissibility in court proceedings being typically that of individuals. But there's a precedent for the latter; that is NRC accepted hazard studies about 1986 and 1988 as valid for the purposes of licensing proceedings and use in the regulatory process, and as far as I can tell, the Livermore study never had an author. And if I'm hurting somebody's feelings in the audience, I apologize, but the authorship was passing the buck. Everybody got equal weight and nobody had a chance to really own it.

MR. POMEROY: You're right on target there. They did indeed do that. I think that my comments were made, and I hope they were made in the context of this specific project, and I suspect that in the final analysis, when you got into a courtroom or a hearing situation, that you would find that the court would have a difficult time accepting the aggregate of this unknown quantity of experts. They really are going to want to know not only the identity, but also the qualifications of those experts and all the other things.

DR. BUDNITZ: So you would argue that if you had a process like the Livermore process, there ought to be one author who takes responsibility for it even if he hears everybody else's opinion?

MR. POMEROY: No, I wouldn't argue that at all. I would simply say that if the results are aggregated in any way, that they should be capable of de-aggregation because they undoubtedly will face that de-aggregation and the individual experts involved will be asked to testify as to their particular interpretations, their particular assignment.

DR. BUDNITZ: I understand you point. That is different than what I thought you were saying, and I want to make a contrary opinion. Of course that might work, but I think another thing that might work is if they found one person who said it's my report, I listened to everybody else, but it's my report, you know, Norm Rasmussen, and by the way, Saul Levine owned that report. 130 people worked on that report in 1973 to 1975, but it was known as the--actually, it should have been known as the Rasmussen - Levine report, but there were two guys that owned it and all of the judgments of the people down in the trenches ended up--they went up to Norm and he settled it, and that's a valid thing even though they didn't do all the work and even though much of it relied on the expertise that was beyond their individual expertise that they had to have from God knows where, and that will work, too. And it has a model that can--you know, and the guy can say I relied on Pete, but in fact I had to make a judgment call between Pete and Charlie, and it's my call. Is that a wrong model, Paul?

DR. POMEROY: My personal opinion--you know, I love to disagree with you--

DR. BUDNITZ: No, I mean, you understand the context in which I'm asking this. I'm sure this SSHAC committee, we try to sort this out as to what we might want to recommend for hazard analysis in terms of this integration.

DR. POMEROY: I think what you're going to have in a real situation, Bob, are groups, a large number of experts associated with the State, a large number of experts associated with each of the intervening parties. I think it would be disastrous in some sense not to have a broadly representative group of experts associated with any particular set of opinions, because the courts and the hearing process could only interview one person under your scenario and--

DR. BUDNITZ: Oh, yeah, I understand the point. I assumed that the Department would stand behind all this. I other words, it would become the Department's position, but it would have a single spokesman of a person of stature who could in fact speak for all the inputs he got.

DR. POMEROY: I would disagree with that. But let's talk about that some more in another context.

DR. ALLEN: Tim, since you suddenly appeared at the head table, does that imply you want to say something?

MR. SULLIVAN: Yeah, just briefly, Paul. I had the impression from your remarks that you thought perhaps DOE was going to await NRC acceptance of a methodology prior to proceeding with the development of the probabilistic seismic hazard analysis, and I wanted to make it clear that wasn't our intent. The NRC, if I judge Keith's remarks correctly, agrees on the necessity of developing a probabilistic seismic hazard analysis and we intend to proceed in accordance with the program that I showed there on the screen. We do hope prior to the license application, however, to reach closure on the methodology so that we can focus on the results.

DR. POMEROY: That's good. I'm glad to have that clarification, Tim. I would strongly advocate, however, that you try to carry out a broad based study similar to the EPRI kind of study, perhaps a full scale moot court approach to provide the technological transfer to provide the information to the public community and to build that consensus that's part of the acceptance process, I believe, within the regulating agency here.

DR. ALLEN: Thank you, Paul. I appreciate it. Our final speaker in today's program is Keiiti Aki, who is professor of geological sciences at the University of Southern California. He is the director of the Southern California Earthquake Center. He's a member of the National Academy of Sciences and he has been active for many years in the field of probabilistic hazard assessment. Keiiti?

DR. AKI: I was asked in the beginning that I'm supposed to come up. But this is the first time I was hearing about Yucca Mountain today, and I just can't summarize those political issues. So my view will be just mostly from a science view.

The first time I ran into PSHA is late Sixties when Allin was at MIT writing this '68 paper. And at that time, I was also working on--I was trying to compute seismic motions from the earthquake floor when the rupture propagates. This is the first time for that kind of configuration that we have tried. And I felt it would probably take a long, long time before PSHA can utilize this kind of new element in seismology. So I have a mixed feeling about this PSHA.

On the other hand, I felt that the kind of thing that I was doing is very specific to--and for the information that's available it's not very often, and it was--I thought PSHA is very important because it can include the effects of all possible aspects, not just one particular aspect. So this integration aspect, integrate all the possibility, I hope is a very strong point of the PSHA.

The second time I ran into PSHA was I was asked to chair the National Academy of Science panel on PSHA to evaluate PSHA, and it was the early Eighties. And half of the members were outside interests and the other half were hazard analysts, and it took three or four meetings--of course Leon Reiter was the driving force and it took three or four meetings before the outside interests accepted putting weight on the likelihood of hypothesis as a necessary evil, because engineers, hazard analysts must give answers today, and they cannot wait, and when in that situation, the best possible way is to do some kind of weighted--and this is really difficult for scientists to accept because any time--but we accepted it. And after we accepted that, we were able to write the report.

The third time that I became involved in the PSHA is when we wrote a proposal to the National Science Foundation, and this has been funded three years ago by USGS and currently supported by a little over $4 million a year, and there are 50 PIs involved from eight core institutions, including the USGS of Pasadena, and there are about a dozen participating institutions, about 50 PIs involved, and the goal is to integrate research findings from various disciplines in earthquake related science to develop a prototype probabilistic seismic--so this goal is a PSHA, and this was accepted by the National Science Board.

Through this three years of experience, I can summarize this saying that PSHA can follow the framework for integrating information from various disciplines engaged in the assessment of seismic hazards. Integration of multi-disciplinary work, PSHA can be very effective, and it's been mentioned by Allin.

Also the PSHA can promote interaction among the different disciplines, and it's not just the framework. It can really improve the understanding of the physical phenomena causing seismic hazards. So this is the framework, but the way you use it, somebody mentioned about this EPRI report, that's how you use these things and it makes a lot of difference and it can promote interaction among different disciplines and actually improve the understanding.

And, finally, PSHA can identify multi-disciplinary issues to be resolved for a better assessment of seismic hazards.

Today we heard about data from Yucca Mountain, geological data and others, and I was expecting some of these interactions taking place, but then we only had about methodology and almost nothing about the PSHA or nothing about this interaction that we might expect.

And what I'd like to show very briefly is what kind of things are happening in the center, Southern California Earthquake Center, to demonstrate this use of PSHA. Since the time is running out, I will be very brief.

This is the way we divide Southern California into 65 zones, and some of the zones contain the San Andreas fault, San Jacinto fault, for which we know very well from the paleoseismology and we can characterize probabilities. This mostly shows probability for the next 30 years. In other zones, we have many--and so depending on the zone, you'll have data available.

On the other hand, more recent development--one of the important elements of the Center is the GPS data. and from the GPS data, we have uniform coverage of this strain. And we distribute the measurements into the zones that I have shown here, and these are the numbers that Steven Ward, University of California at Santa Cruz, has assigned to different zones. And by the way, the zones which contain mostly--is one of the highest rates in terms of the strain data accumulation observed from GPS, only a few years of data. This was a surprise for all of us, that GPS data was useful for the hazard estimation.

We have basically three different kinds of data that characterize each fault zone, and there's a geologist measurement of the--and this strain measurement through GPS, and also we have a catalog of earthquakes in the past 150 or 200 years, and that can be assigned to--and we can look up the parameter which we use--which can be measured by seismic methods from--it can be measured from the volume of strain and also it can be measured geologically.

So this common parameter can be sort of integrated and it helps us to understand the nature of the data and also the nature of the--and here's part of Southern California, which includes Los Angeles and San Bernardino and in the distance, the San Andreas fault here. And these three figures here are the result of PSHA and a very simple parameter. We use .2g, peak acceleration. And this is a map--the 60,000 probability centers around San Bernardino and, as you might expect, high probability around the San Andreas.

This is when you just use the earthquake calculator and smooth it, you see fault zones and then--and then you see it spreads out as compared to the fault information corresponding to the paleoseismological data. It's rather spread out as compared to this. Here is the prediction from this GPS strain data, and as you see, it has higher hazard off the San Andreas fault.

Now, we tried to combine those three data sets into one most reasonable fault zone characterization, but we have some discrepancy among the different groups. One group--more or less this fault segmentation model, and for each of the source zones--then we account for moment rate in the way more or less similar to--in this particular model we use Gutenberg Richter quantitative.

What I'd like to show here is if you use the best of the combination of this geodetic, geology and predict what is the annual rate for the whole Southern California, and perhaps that is a function of magnitude, the top curve is the prediction, and all these curves are showing what is the contribution of characteristic distribution--but this predicted one is maybe sometimes almost three times higher than the observed in that past 150 years. This could mean that actually the past 150 years was one of anomalously low seismicity compared to the situation that you might expect from these geodetic strain accumulation and geologic fault information. But some group at Center doesn't like this interpretation, and they like to resolve this. This is one way of resolving it, is to make the earthquake of the San Andreas fault can have a much--say magnitude eight. If you allow that, these strains are being accumulated, measured by geodetic means, can be absorbed in this very large earthquake, which of course vary, but this will be closer to predicted.

This kind of comparisons among different data sets is giving rise to controversies that predicted earthquake rate is greater than historic earthquake rate, and the one possibility is change in seismicity, or it could be that the strain may be taken up as aseismic slip. Or it could be that the maximum magnitude may be eight or several earthquakes of magnitude six to seven, and some geologists say yes to this issue and geodesists in our group don't like this, and seismologists in Pasadena do or don't like it, because there has to be an increase in the seismicity the next 30 years or so.

But outsiders, first of all, are in favor of this use. They like to have more earthquakes in Los Angeles. So there is a very interesting serious conversation between the different groups, and this is, I think, because of our--

because the PSHA forces us to talk in terms of the same quantities that gave us common ground to compare their data.

So this one interesting graph here is showing cumulative moment rate, and this is a cumulative--means if the whole area has the same moment rate, and the moment rate you can think of as a seismic hazard, if everything is homogeneous, it would be a straight line connecting this. And this departure from the straight line, when this is sharper, this means most moment. So it means you have a lot of information about where the earthquake--and this is one of the results we got from one of our models, and they show this--our knowledge about geodetical distribution of seismic hazard, and this happens to be a zone containing earthquake and which is in the top 13 percent of the whole Southern California in terms of this. Steve Wesnousky included Northridge as an occurring unknown fault, but we anticipated this earthquake.

As you know, you can make all kinds of seismic hazard maps. This is just one example showing acceleration points through the 50 year exceedence for the 10 per cent. And we saw some high acceleration sort of west of--in addition to the San Andreas fault, and this was a comparison with a previous study by USGS that this region shows--and this is probably sheer luck that we seem to be giving more credibility because of this.

So in the last several years at the Center we are surprised that GPS data can be useful in such a short span, and we also found this geodesic data seemed to over estimate the historic seismicity. But all these things seem to give us the forecast for the future studies, like issues of people disagreeing about multi-disciplinary groups, the maximum magnitude issue and aseismic strain issue and if the seismicity can change over hundreds of years or so. These are very fundamental issues and PSHA helps us to forecast in a very quantitative manner on these issues.

I'm supposed to talk about Yucca Mountain. I think I have seen this morning the data covering all these geologic and also strain in the data, but strain was very, very small and it was, in the map shown, strain accumulation was negligible, within the error of the measurements. Except this Little Skull Mountain earthquake apparently created measurable strain which gave the seismic moment for this earthquake. So there is this one little earthquake which may be used in the context of hazard analysis to combine with the data sets. But I think this very low seismicity rate at Yucca Mountain and very long occurrence time make it very difficult to promote such an interaction as we have seen in California in a very short, short time span.

But we have a large number of scientists getting together, and we have an opportunity to discuss in workshops and meetings almost continuously, and this interaction is a very time consuming effort, and for Yucca Mountain I think we would need more and broader participation. It's probably more difficult, but if we did elect something in the direction that we have seen in California, we need a very large group of people involved. So that's my personal experience with PSHA.

DR. ALLEN: Thank you, sir. Any comments or questions from Board members or consultants? Staff? Leon?

DR. REITER: I want to revisit a little bit the panel on seismic hazard analysis that you chaired. Back in the Seventies, there was another panel and I think the title of their report was "Research Reactors." I think Clarence was on that panel.

DR. ALLEN: I deny it.

DR. REITER: One of the conclusions of the panel was that probabilistic analysis was not yet ready to be used, and I think your panel came up--and that was in the Seventies--you panel came up and said yes, we are ready to use it, and there have been some increases in our knowledge of that approach.

However, there was a small statement in that report that some of us, particularly the NRC, noted. And that said that for facilities where the likelihood of earthquake occurrence is less than 10 to the minus 3 or 10 to the minus 4 per year, you recommended that both probabilistic and deterministic analysis be conducted.

DR. AKI: I hope we recommended both on any case.

DR. REITER: I thought it particular to the very low probability. And I was wondering if you think that if you would write that report again today, you would still make the same recommendation?

DR. AKI: Yes.

DR. ALLEN: Budnitz was not on that committee.

DR. BUDNITZ: I'm not an earth scientist.

DR. AKI: The weakness of PSHA, as every realizes, is this sort of smooths out everything. So you lose what is the most important earthquakes affecting your hazard. And so the aggregation is the most important thing, and you have to de-aggregate. Once you do this PSHA, you have to de-aggregate and look at each individual event in the model. That's what we thought was an important element in this thing. So PSHA should be combined with deterministic analysis.

DR. ALLEN: Other questions or comments?

Thank you, Keiiti. We're virtually on schedule, but we're--I declare the session closed for today, and we'll start again at 8:30 in the morning on volcanic hazard analysis

Thank you all very much.

(Whereupon, at 5:35 p.m., the meeting was adjourned.)